University of California
Santa Barbara

Essays on Uncertainty and Risk in Economic Settings

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

David Scott Hales

Committee in charge:

Professor Theodore C. Bergstrom, Chair
Professor H.E. Frech, III
Professor Ryan Oprea
Professor Cheng-Zhong Qin

June 2022

The Dissertation of David Scott Hales is approved.

Professor H.E. Frech, III

Professor Ryan Oprea

Professor Cheng-Zhong Qin

Professor Theodore C. Bergstrom, Committee Chair

June 2022

Essays on Uncertainty and Risk in Economic Settings

by

David Scott Hales

I dedicate this dissertation to two people who have had great impact on my academic career, but who have rarely been visible to the outside world: to my wife, life partner, and best friend Stephanie Hales (Jänsch), for her countless hours of unfailing support, love, and understanding; and to my late mentor Stephen P.A. Brown, who inspired and guided me in my early academic years and who encouraged me to become an economist.

Acknowledgements

I gratefully and humbly acknowledge the tremendous guidance and advice of so many counselors, in support of the research that comprises the body of this dissertation. I benefited enormously from the helpful advice and guidance of my advisor Ted Bergstrom, and from countless other mentors including Stephen Brown, Frank Heinemann, Cheng-Zhong Qin, Tim Cason, John Duffy, Ted Frech, Uri Gneezy, Youssef Benzarti, Antoine Dib, Gary Charness, Rod Garratt, Sevgi Yuksel, Erik Eyster, Ryan Oprea, Rosemarie Nagel, Gabriele Camera, Shyam Sunder, Emanuel Vespa, Terri Kneeland, Jaime Ramirez, Paul Thistle, and Yi Zheng, as well as my coauthors and collaborators on past projects.

I am also indebted to the Harvard Center for Risk Analysis for help with their models of cancer risks, and to the organizers and participants of the 11th BESLab Experimental Economics Summer School in Macroeconomics, the 7th Spring School in Behavioral Economics, the 2019 International ESA Meeting, 13th Maastricht Behavioral and Experimental Economics Symposium and the 11th International Workshop on Theoretical and Experimental Macroeconomics for insightful discussions.

I also gratefully acknowledge and thank the Joachim Hertz Foundation for providing funding for a pilot experiment that preceded the research effort that culminated in Chapter 1, and to the University of California, Santa Barbara Economics Department for providing funding for the experiment that is the basis of Chapter 2.

Permissions and Attributions

The contents of chapter 1 and appendices A, B, and C are the result of a collaboration with Muhammed Bulutay, Patrick Julius, and Weiwei Tasch, and were published in *Journal of Economic Behavior & Organization* in 2021. The article is an open access article and is reproduced here according to the terms of that journal's license agreement (http://creativecommons.org/licenses/by/4.0)[1].

The contents of chapter 2 and appendices D and E are the result of a collaboration with Hazem Alshaikhmubarak, Maria Kogelnik, Molly Schwartz, and Kent Strauss[2]. This article is currently under review at a peer-reviewed journal for publication consideration.

The contents of chapter 3 and appendices F and G are the result of a collaboration with Mary Riddel, which were published in an article in *Risk Analysis* in 2018. This material is reproduced here according to the terms of that journal's license agreement[3].

Curriculum Vitæ

David Scott Hales

Education

2022	Ph.D. in Economics (Expected), University of California, Santa Barbara.
2017	M.A. in Economics, University of California, Santa Barbara.
2015	M.A. in Economics, University of Nevada, Las Vegas
2003	M.B.A, University of Hawaii, Manoa
1997	B.A. in Computer Science, National University
1994	Sc.B. in Physics, Massachusetts Institute of Technology

Publications

Imperfect tacit collusion and asymmetric price transmission, Journal of Economic Behavior & Organization **192** (2021) 584–599 (with M. Bulutay, P. Julius, and W. Tasch).

Predicting cancer-prevention behavior: Disentangling the effects of risk aversion and risk perceptions, Risk Analysis, **38** (2018), no. 10 2161–2177 (with M. Riddel).

Risk misperceptions and selection in insurance markets: An application to demand for cancer insurance, Journal of Risk and Insurance, **85** (2018), no. 3 749–785 (with M. Riddel).

Physical competition increases dehydroepiandrosterone (dhea) and androstenedione rather than testosterone among juvenile boy soccer players, Adaptive Human Behavior and Physiology, **2** (2016), no. 1 44–56 (with T. S. McHale, D. T. Zava, and P. B. Gray).

Abstract

Essays on Uncertainty and Risk in Economic Settings

by

David Scott Hales

This dissertation consists of three essays that investigate the role uncertainty and risk play in shaping outcomes in several economic contexts of interest.

In the first chapter, based on joint work with Muhammed Bulutay, Patrick Julius, and Weiwei Tasch, we experimentally explore the behaviors and response pathways of economic agents who are placed in a competitive market environment in with exogenous cost shocks, strategic complementarities, and large cross-price effects, but in which communication is not permitted. Here our subjects experience strategic uncertainty over the pricing decisions of their competitors, and the predominant response we observe confirms the asymmetric price transmission (APT) phenomenon that is widely observed in the empirical literature. Moreover, by varying the number of sellers across markets we evaluate the role competition plays in APT. We report similar magnitudes of asymmetry in markets with 3, 4, 6, and 10 sellers, but not in duopolies. Furthermore, sellers consistently set their prices above the best-response levels implied by their forecasts, particularly in periods following negative cost shocks. We interpret these pricing deviations as sellers' intentions to collude, and note that they mechanically drive the pricing asymmetries we observe, consistent with a hypothesis of tacit collusion. In the second chapter, based on joint work with Hazem Alshaikhmubarak, Maria Kogelnik, Molly Schwarz, and Kent Strauss, we investigate the role of strategic uncertainties, themselves shaped by the presence information uncertainties, affect subject willingness to engage in risk to pursue socially optimal outcomes. We employ an experiment to study the ways in which

mutual payoff information affects play in strategic settings. Subjects play the Prisoner's Dilemma or Stag Hunt game against randomly re-matched opponents under two information treatments. In our partial-information treatment subjects are shown only their own payoffs, while in our full-information treatment they are shown both their own and their opponent's payoffs. In both treatments, they receive feedback on their opponent's action after each round. We find that mutual payoff information initially facilitates reaching the Pareto-efficient outcome in both games. While play in the Prisoner's Dilemma converges toward the unique Nash equilibrium of the game under both information treatments, mutual payoff information has a substantial impact on the equilibrium selection in the Stag Hunt throughout all rounds of the game. Using a belief-learning model and simulations of play, we provide evidence that these effects are driven not only by initial play but also by the way subjects learn. We propose that strategic uncertainty is a probable channel through which payoff information affects play.

In the third chapter, based on joint work with Mary Riddel, we examine how public misperceptions of cancer risks interplay with risk preferences to influence behaviors that either increase or decrease cancer risk. While there are numerous previous studies that link risk perceptions with risky behavior, it is notable that none of these controls for risk aversion. Similarly, studies that control for risk aversion fail to control for risk misperceptions. We use a survey of 474 men and women to investigate the influence of risk aversion, risk misperceptions, and cognitive ability on the choice to engage in behaviors that either increase or mitigate cancer risk. We measure optimism in two dimensions: baseline optimists are those who inaccurately believe their cancer risk to be below its expert-assessed level, while control optimists are those who believe they can reduce their risk of cancer (by changing their lifestyle choices) to a greater extent than is actually the case. Our results indicate that baseline optimism is significantly and negatively correlated with subjects tendencies to engage in cancer-risk-reducing behaviors, and positively

correlated with risky behaviors. Subjects' control misperceptions also appear to play a role in their tendency to engage in risky and prevention behaviors. When controlling for both of these types of risk misperception, risk aversion plays a much smaller role in determining health behaviors than found in past studies.

Contents

Chapter 1

Imperfect Tacit Collusion and Asymmetric Price Transmission

1.1 Introduction

The phenomenon of Asymmetric Price Transmission (APT), that is, that supplier prices rise quickly after positive input cost shocks but fall relatively more slowly after similarly-sized negative shocks, has been repeatedly documented in the literature such that we can rightly describe it as a stylized fact.[1] However, while empirical evidence for the existence of APT is ample, identification of its causal forces is not settled. Many theoretical explanations have been proposed, but the empirical literature has yet to conclusively determine which of these are valid or are most influential.

Empirical studies of APT predominantly examine aggregate-level variables (e.g. inflation, concentration) proposed to be relevant in the theory literature. The focus on such variables occurs because firm-level determinants are either not directly observable,

[1] This is also sometimes termed as "positive APT" to distinguish it from the opposite phenomenon of "negative APT." In this paper we will simply refer to "APT" when we mean positive APT, except where doing so would create ambiguity. See Section 1.2.1 for an overview of the evidence.

or are not adequately measurable in panel data form. This approach yields helpful correlations between such variables, but the effort to identify causal relationships has met with only limited success, most notably in the context of firm-level underpinnings of the phenomenon. While the search for accurate firm-level data should certainly be continued, and where discovered used to further inform our understanding of pricing behavior, experimental methods offer a comparative advantage: testing theories that involve variables which are unobservable in the field (e.g. agents' information sets) lie outside the reach of empirical methods;[2] if however these same variables can be controlled through experimental design, we can overcome this obstacle to testing theory.

A question of primary interest is whether tacit collusion drives APT-like pricing behavior.[3] The field data does not convincingly exclude the possibility that market competitors secretly communicate, given the strong legal and even criminal incentives for firms to conceal – or avoid engaging in – such activities. This provides an obvious challenge for identification and motivates turning to the controlled setting of the laboratory, where we can directly observe competitor behavior and credibly prevent communication between sellers.[4]

An argument put forth by Borenstein et al. (1997)[6] is that a variation of the "trigger price" model of oligopolistic coordination originally introduced by Green and Porter (1984)[7] may explain the emergence of APT-type dynamics through tacit collusion. In their model, when positive shocks occur firms immediately raise prices in order to preserve profit margins; however, when negative shocks occur firms react adaptively,

[2]Meyer and von Cramon-Taubadel (2004)[4] and Frey and Manera (2007)[5] provide extensive discussions of methodological issues in econometric tests of APT.

[3]In this paper we will use the term "tacit collusion" to mean the phenomenon in which suppliers coordinate on prices above the competitive equilibrium level, through the channel of publicly visible pricing alone. Tacit collusion can also take the form of coordination on quantities below competitive equilibrium levels, but in this paper we will focus strictly on the role of coordination on prices.

[4]Furthermore, the laboratory may be the only environment in which we can reliably detect collusion, since the non-collusive prices or profits are unavailable without imposition of strong structural assumptions.

holding prices at pre-shock levels until they see convincing evidence that a rival has cut their prices. Rapidly lowering prices in response to a downward cost shock could be perceived as defection from a mutually beneficial regime of tacit collusion, thus inviting retaliation from other firms. In contrast, rapidly raising prices in response to an upward cost shock poses no such threat to one's competitors, and therefore incurs no corresponding risk of retaliation. Although their arguments are sound and are consistent with a deep empirical literature finding correlations suggestive of tacit collusion, Borenstein et al. (1997) conclude that they are unable to conclusively draw support for this hypothesis from their data[6]. As no other empirical study of which we are aware has accomplished this either, we thus find motivation to turn to the laboratory to examine the role of tacit collusion in driving the APT dynamic.[5]

A second question of interest is whether the number of competing sellers in a given market plays a significant role in the realization of the APT phenomenon. Notably, in his broad study of U.S. wholesale and retail markets, Peltzman (2000) finds a negative relation between the number of competitors in a market and the magnitude of APT observed[9]. As with any empirical study, however, this study does not exclude the possibility that explicit (but unobserved) communication between firms lies behind this result. Several non-APT focused studies of experimental oligopoly markets find that there is an inverse relation between the number of sellers in a market and the size of deviations from the Nash equilibrium (NE) outcomes (for example, see Huck et al., 2004[10], Dufwenberg and Gneezy, 2000[11], and Fonseca and Normann, 2012[12]). However, we are unaware of any experimental study that specifically studies the role of the number of sellers in driving the APT phenomenon. We therefore incorporate the number of sellers in our

[5]There are some studies that regress the estimated asymmetry with measures of market concentration, e.g. Loy et al. (2016)[8]. Counter-intuitively, the authors find that asymmetry decreases with higher concentration in German milk markets. However, it is difficult to associate this estimate with the causal impact of collusion on APT as their observed higher concentration index may stem from higher efficiency or product differentiation rather than from competitor conduct.

markets as a treatment variable in our experimental design.

To our knowledge, Bayer and Ke (2018) is the only experimental study that directly targets the topic of APT[13]. The authors' study employs a Bertrand duopoly setting in which sellers' costs either increase, decrease or stay constant at the halfway point of the experiment. With two extensions of this baseline condition, they further test the impacts of search costs and asymmetric information on APT. They find APT across all treatments, even in the absence of search and information frictions. They argue that the asymmetry can be explained with a backward-looking learning model: If a seller fails (manages) to sell the good in the period prior to the shock, it is more (less) likely that she will adjust her price downwards (upwards) in the following period. The authors' results support this regularity when the shock is negative, but not when it is positive. Hence, although this learning model may account for the downward rigidity, it falls short of explaining the asymmetry.[6]

While Bayer and Ke's study provides a useful benchmark to our own, our design choices differ substantially from theirs, as we pursue different research questions. Whereas we aim to assess the roles of cooperative behavior and tacit collusion on pricing asymmetries, they deliberately try to attenuate their impacts to isolate the role of learning.[7] In particular, in their experiment sellers whose stores are not visited by a buyer receive only limited information on the market price, due to the feedback structure. In our ex-

[6]Bayer and Ke (2018) also reason that following positive cost shocks sellers will reason that other sellers will all immediately raise their prices, and so they do the same, while following a negative cost shock sellers do not see any reason to cut their prices unless and until they subsequently lose sales. They cite factors such as bounded rationality as explanations for this behavior, but do not offer a more precise explanation of the channels through which the observed behavior emerges.

[7]Although Bayer and Ke (2018) exert effort to minimize the role played by tacit collusion with their study, their typed-stranger matching protocol significantly reduces but does not completely eliminate the possibility that subjects might repeatedly interact, and thus have the opportunity to establish reputation over time. By contrast, the perfect-stranger matching protocol, in which a subject is assured they will be matched with another only once in a session, does more credibly eliminate this possibility. Moreover, the duopoly setting of their study makes collusion presumably more achievable, since coordination is easier when there is only one other market participant. As a result, it is hard to assess the extent of the role to which cooperative behavior played in their study.

periment, we inform sellers of the average market price of the other sellers, as we want to create the conditions in which price signalling can be studied more explicitly.

In our experimental setting, subjects play the role of sellers and a computer plays the role of buyers. Each seller faces demand that linearly decreases with one's own price and linearly increases with the average price of others. We vary the size of groups across sessions as 2, 3, 4, 6, and 10, while calibrating the demand function to hold the best-response functions of each seller identical, across all group sizes. Thus, we isolate and study the impact of group size on behavior, while holding the price-based incentives facing individual sellers constant across markets.[8] Throughout our experiment, sellers experience a series of input price shocks – either large or small – that shift the NE price either up or down. Through this design, we are able (i) to test whether APT emerges despite the absence of market frictions and information asymmetries that are often theorized to be the causal forces behind pricing asymmetries; and, (ii) if APT does occur, to assess the impact of number of sellers on the magnitude of the resulting asymmetries. To our knowledge, ours is the first experiment that study the role of number of sellers in shaping APT, but also the first to study the pure number effect in a price competition setting.

Our contributions to the literature are two-fold: First, we document the prevalence of the APT phenomenon through experiments in which we possess strict control over the environment. In particular, our results indicate that the APT may emerge even in the absence of market frictions and information asymmetries that are often theorized to be the causes of pricing asymmetries. This suggests that in markets with three or more sellers, the presence of agents who attempt to coordinate on prices via price signaling may suffice for APT pricing dynamics to emerge. In our duopoly markets, however,

[8]This is akin to the "pure number effect" studied by Isaac and Walker (1988)[14] in the context of public goods game. See Hanaki and Masiliunas (2021)[15] for a similar approach in Cournot oligopolies.

our results suggest that coordination can be so successful that rather than the (positive) APT phenomenon, persistent pricing at collusive price levels, or even negative APT, may instead result.

Second, keeping incentives the same across differing group size, we are able to isolate and perform hypothesis tests on the pure effect of increased group size on APT. For markets with three or more sellers, we do not find significant differences in the magnitude of observed APT. Together, the results of our study support theories that highlight the role of tacit collusion on APT. We conclude that APT may be the product environments in which collusion is significant, but imperfect (i.e., unstable).

1.2 Related Literature

1.2.1 Field Evidence

Bacon (1991) provides an early empirical study suggesting that retail gasoline prices in the United Kingdom experience faster and more concentrated responses to crude oil price increases than they do to similar crude oil price decreases[16]. Bacon termed this phenomenon "Rockets and Feathers," and since this paper was published dozens of other researchers have detected the presence of this sort of asymmetry in a variety of consumer and intermediate goods markets.

Peltzman (2000) provides one of the most comprehensive empirical examinations of APT[9]. He conducts a broad study of pricing behavior of 77 consumer and 165 producer goods markets in the U.S., and he concludes that in more than two-thirds of these markets prices rise faster than they fall, in response to input cost changes. Peltzman also seeks correlations between various features of markets and industries, and the degree to which evidence of APT is present. Most notably, he finds that markets with fewer competitors

tend to exhibit more pricing asymmetry, while on the other hand markets with higher levels of concentration tend to be less likely to exhibit pricing asymmetry, as in Loy et al. (2016)[8]. Peltzman's study, however, does not provide an explanation for these correlations.

In an early survey of field evidence, Meyer and von Cramon-Taubadel (2004)[4] find that (excluding Peltzman's 2000 study), symmetry in price response is rejected in almost one-half of all cases in the literature. Their survey also shows that different test methods yield highly varying rejection rates (between 6% and 80%). Frey and Manera (2007)[5] and Perdiguero-Garcia (2013)[17] provide meta-regression analyses with more comprehensive and recent data sets. Both studies confirm that APT is very likely to occur but also emphasize the variation of reported outcomes. Their results show that this heterogeneity can be explained by certain characteristics of the data (e.g., data frequency) and of the employed econometric model. Most notably, Perdiguero-Garcia (2013) reports that the asymmetry tends to decrease in more competitive segments of the industry.

1.2.2 Theoretical Explanations

There is a growing body of literature on the theoretical accounts of APT, an unsurprising fact given that pricing asymmetries are not predicted by standard price competition models.[9] These studies propose explanations of APT mainly by introducing market fric-

[9] A notable exception is the case of Markov-perfect equilibria, and in particular the case of the Edgeworth cycle. In this phenomenon, firms undercut each others' prices successively until prices approach marginal cost; at this point, one of the firms decides with some positive probability to spike its price, and once this occurs the cycle is repeated, yielding each firm positive economic profits. Maskin and Tirole (1988) further show that these cycles provide a case where asymmetric pricing can be sustained in equilibrium[18]. However, the Edgeworth cycle model requires that firms make price decisions alternately; the model does not support an equilibrium when price decisions are made simultaneously or continuously. Moreover, the emergence of the phenomenon seems in practice to be limited to environments in which competitors rapidly and publicly change prices (see for example Byrne and De Roos (2019) for an interesting case in Perth, Australia petrol markets, in which a government mandate for retail suppliers to publish their prices daily seems to have facilitated the emergence of a weekly cycle of Edgeworth-like pricing dynamics that persisted for many years[19]). Thus the Edgeworth cycle model

tions, information asymmetries or boundedly rational agents into the underlying models. One reason there is such a variety in the way different studies explain the APT is because these studies typically focus on specific market structures (e.g., wholesale petroleum markets) and their idiosyncrasies. In this subsection, we review some of these studies in an attempt to categorize them as well as to highlight discrepancies.[10]

Borenstein et al. (1997) consider the role of search costs in facilitating APT[6]. They hypothesize that negative cost shocks in the presence of costly search provide firms temporary pricing power, which they then use to delay reductions in prices, yielding temporarily superior profits. Benabou and Gertner (1993)[21] and Yang and Ye (2008)[22] also develop explanations based on consumer search costs, but also on the volatility of input costs. They reason that volatility should reduce search incentives for consumers; producers, realizing that they face demand that is temporarily more inelastic, yielding them increased pricing power over the short term, respond by reducing prices more slowly. Reagan and Weitzman (1982)[23] and Borenstein and Shepard (1996)[24] propose explanations based on inventory costs, reasoning that it is relatively more costly for manufacturers and suppliers facing capacity constraints or sharply rising short-term production costs to deal with unanticipated increases in demand resulting from price drops, than it is to respond to corresponding drops in demand due to price increases. Ball and Mankiw (1994) consider a menu-cost model in conjunction with positive trend inflation as an explanation of APT[25]. In another study, Ahrens et al. (2017) show that the presence of consumers with loss aversion may explain why prices are more sluggish to adjust downwards than upwards in response to permanent demand shocks[26].

The various explanations and models described above provide differing implications for government policy: if APT occurs due to collusion, there may be room for regulation

arguably applies only to a relatively narrow range of market contexts.

[10]For more exhaustive surveys of theoretical explanations, see Meyer and von Cramon-Taubadel (2004)[4] and Brown and Yucel (2000)[20].

to improve economic efficiency; if however APT is primarily caused by the presence of inventory costs, asymmetric menu costs, or search costs, then regulation that controls pricing behavior may actually induce inefficiency rather than attenuate it. Given the robust evidence of the widespread existence of APT and its non-trivial magnitude and impact on consumer outcomes, identifying which theories best describe the asymmetric pricing behavior is key to determining effective public policy.

1.2.3 APT and Experiments

Despite the many possible explanations that have been proposed, the empirical literature yields only mixed evidence that is often inconclusive due to identification issues. This suggests there is room for further research to shed light on the phenomenon. We consider the advantages of experimental methods in isolating and studying causal determinants of APT.[11] In this subsection, we summarize the most relevant literature to our study.

There are two studies of which we are aware – in addition to Bayer and Ke (2018) – that conduct market experiments with APT-related results[13]. Deck and Wilson (2008) investigate gasoline markets and find that retail prices adjust asymmetrically to changes in station costs in zones with clustered stations, but not in zones with stations that are relatively isolated from competitors[29]. Cason and Friedman (2002) find weak evidence of APT in posted offer markets where customers incur switching costs[30]. While these studies examine their findings on APT, their experimental designs are optimized to investigate questions regarding the structure of gasoline markets (e.g., zone pricing, divorcement) and of consumer markets (e.g., switching costs), not to identify causes of APT. In particular, sellers' costs in both experiments follow random-walk shocks, which

[11] The usage of experimental methods in macroeconomic research is becoming more and more prevalent. See Duffy (2016)[27] and Cornand and Heinemann (2019)[28] for recent surveys.

may not be salient enough to detect APT. Our study distinguishes itself from this string of literature by examining APT with larger, persistent shocks.[12]

Apart from studies that directly target APT, price competition experiments that study the impact of group size on tacit collusion are also relevant to the current paper. Dufwenberg and Gneezy (2000) provide an early evidence for such a relation through an oligopoly game that corresponds to a discrete version of the Bertrand model[11]. They find that winning prices tend to converge to NE levels in groups of three or four competitors, but stay consistently high in duopolies. Morgan et al. (2006) find that increasing the number of sellers from 2 to 4 decreases the prices paid by some consumers (the ones informed about the entire distribution of prices) but not for others (the ones who buy with motives other than prices)[33]. Abbink and Brandts (2008) also find that there is a negative relationship between the number of competing firms and price levels[34].[13] Nevertheless, as in Dufwenberg and Gneezy (2000)[11], they find that collusive pricing is the modal outcome in duopolies. Fonseca and Normann (2012)[12], Orzen (2008)[35], Davis (2009)[36] and Horstmann et al. (2018)[37] provide further evidence that collusive prices are very likely to be observed in duopolies. Average prices approach considerably close to the NE in the baseline condition of these studies (fixed matching, no communication, symmetric sellers etc.) when the number of sellers is 3 or greater.

A finding common to each of these studies is that persistent coordination over collusive prices is unlikely in markets other than duopolies. This, however, does not preclude the possibility that players might manage to coordinate temporarily on high prices following

[12]Fehr and Tyran (2001) also employ large positive and negative shocks and report APT-like behavior in a price-setting game[31]. However, the authors do not analyze the phenomenon, nor do they probe its implications. In another related experimental study, Duersch and Eife (2019) consider Bertrand duopolies with zero marginal cost in either inflationary, deflationary or constant price environments.[32] They find that real prices are significantly lower in the inflationary environment compared to non-inflationary environments.

[13]Their results are particularly interesting since in their price competition setting, there exist multiple equilibria.

negative shocks. Experiments also indicate that increasing the number of sellers often leads to more competitive outcomes (in terms of price and output), which in turn should make APT less likely. Although, the meta-analyses of Fiala and Suetens (2017)[38] and Horstmann et al. (2018)[37] on oligopoly experiments indicate that there may not be a linear relationship between the number of competing firms and the degree of tacit collusion. On the one hand, Horstmann et al. (2018) argue that this result stems from the relatively small number of studies that provide pairwise comparisons and the lack of statistical power in these studies. On the other hand, Hanaki and Masiliunas (2021) consider, as we do, the fact that a change in group size simultaneously influences the difficulty of coordination (aka the "pure number effect") and the incentives provided to collude[15]. They find that the pure effect of group size is small if exists at all. Our study contributes to the literature through improvements of these axis, in particular, by changing the group size without changing the incentives provided to individual sellers in different markets in a price competition game.

1.3 Method

1.3.1 Pricing Game

We develop a price competition game with a linear demand model and employ this in our experimental markets.[14] In this setting, the demand facing seller $i \in N$ in period

[14]Linear demand systems have a number of advantages over alternatives and long been applied in the industrial organization literature. Such systems lead to closed-form best-response functions and Nash Equilibrium specifications, greatly enabling interpretation of empirical or experimental results. To the extent that non-linear systems of demand (e.g. the Almost Ideal Demand System of [39] or the Relative Love of Variation model of [40]) may be preferred on theoretic or other grounds, linear systems provide approximations with first-order accuracy.

$t \in T$ is equal to

$$q_{i,t}(p_{i,t}, p_{-i,t}; \delta, \gamma) = \begin{cases} \delta - \gamma \cdot (p_{i,t} - p_{-i,t}), & p_{i,t} \in [p^{min}, \overline{p}] \\ 0, & \text{otherwise} \end{cases} \tag{1.1}$$

where δ and γ are parameters of demand, $p_{i,t}$ is the price set by seller i and $p_{-i,t}$ is the average price chosen by the rival sellers in the same market (i.e. $p_{-i,t} \equiv \frac{1}{N-1} \sum_{j \neq i}^{N-1} p_{j,t}$) at period t and $\overline{p} = min\left(p^{max}, p_{-i,t} + \frac{\delta}{\gamma}\right)$. Parameters p^{min} and p^{max} refer to the price floor and the reservation price of the representative consumer, respectively. Variable $\overline{p}$ regulates the maximum price level below which (conditional on the average price of other sellers) $q_{i,t}$ takes on strictly positive values.[15]

The model of linear demand that we use, with own- and cross-price parameters equal in magnitude, is micro-founded by the Spokes models of Chen and Riordan (2007)[41] and Bos and Vermeulen (2020)[42], which are themselves adaptations of Salop's canonical Circular City address model (Salop, 1979[43]). In addition, this specification of demand can also be motivated by a limiting case of quasi-linear quadratic utility (see Appendix A for an exposition). Note that any demand specification with linear parameters on the price of each competing good can be equivalently represented in terms of own-price and the (weighted) average price of other sellers, as in (1.1). Thus the presence of average price figure $p_{-i,t}$ in this equation does not imply that when contemplating their appetite for good i that consumers consider the *average* price set by firms competing with firm i; rather, it implies that the aggregate effect of all consumers' individual demand for good i is based on the entire vector of prices, but can nevertheless be mathematically represented concisely in terms of own price p_i and and the average price of all other sellers, p_{-i}. This fact will be helpful in keeping both our analysis and our experimental

[15] In practice subjects chose prices that were revealed to be above this value in fewer than 0.2% of all cases.

design tractable, as we will shortly see.

Given the own-demand specification in (1.1), seller profits are calculated as

$$\pi_{i,t} = (p_{i,t} - mc_t) \cdot q_{i,t} - f, \tag{1.2}$$

where $q_{i,t}$ is quantity demanded from seller i as defined in (1.1), mc_t is marginal cost that shifts every $\overline{T}$ periods that comprise a round (denoted $r \in R$) and f is fixed cost. Sellers set their prices in each period simultaneously from a discrete set that is bounded as $p_{i,t} \in [mc_t, p^{max}]$, such that the price floor is equal to the marginal cost of that round.[16]

In the described finitely repeated game, we can express the maximization problem as

$$\max_{\mathbf{p_i}} \sum_{h=0}^{T} \beta^h \underset{i,t-1}{\mathbb{E}} \pi_{i.t+h} \quad \text{subject to} \quad p_{i,t+h} \in [mc_{t+h}, p^{max}]. \tag{1.3}$$

Sellers thus maximize the expected discounted sum of profits over T periods by choosing a vector of prices $\mathbf{p_i}$.[17] This in turn leads to best-response function

$$p_{i,t}^{BR} = \frac{1}{2}\left(mc_t + \frac{\delta}{\gamma} + \underset{i,t-1}{\mathbb{E}}[p_{-i,t}]\right). \tag{1.4}$$

The current period's marginal cost mc is revealed to the sellers prior to their pricing decisions thus is outside the expectation operator. Sellers form expectations of the prices their competitors will set during the current period ($\mathbb{E}_{i,t-1}[p_{-i,t}]$) by conditioning on all available information. The system of best responses for all sellers implied by (1.4) solves

[16] As a practical matter we needed to set a price floor of $p_{i,t} \geq mc_t$ to ensure subjects would not complete the experiment with a negative payoff.

[17] Although the model assumes that sellers choose a vector of prices for all periods, subjects only submit a price decision for the current period in the experiment.

for steady-state prices as:

$$p_{i,t}^{*} = p_{i,t}^{BR}\left(\underset{i,t-1}{\mathbb{E}}[p_{-i,t}^{*}]\right) = mc_t + \frac{\delta}{\gamma} \equiv p_t^{NE}. \tag{1.5}$$

This price level also corresponds to the unique stage-game NE (p_t^{NE}) and the unique subgame perfect Nash equilibrium (SPNE). Sellers may achieve the joint profit maximum (JPM) if they each set their prices to the maximum price p^{max}.

In this pricing game, neither own-demand nor own-profit depend on the number of sellers. These only depend on own-price and the average price of rival sellers. The best-response action is also independent of N for a wide range of expectation models, including rational expectations and adaptive expectations (Evans and Honkapohja, 2001)[44]. This feature assures that the incentives given to the sellers of different group sizes are matched and the market power of each seller, measured by the size of markup over marginal cost, is *ex-ante* equal. We consider this as necessary for ensuring a *ceteris paribus* comparison between the treatment conditions and to capture the pure number effect.

1.3.2 Experimental Design

Sellers interact repeatedly in the described pricing game for R rounds, which are each composed of $\overline{T}$ periods. Marginal cost mc_t fluctuates across rounds, modeling large exogenous cost shocks, but remains invariant during a round. Our experimental manipulations consist of varying the size of markets across sessions in a between-subjects design, and of varying the size and direction of shocks across rounds in a within-subjects design. We implement a fixed-matching protocol during a session.

The calibration of the experimental game is summarized in Table 1.1. The experiment consists of 5 rounds of 15 periods each, with a new marginal cost announced at the beginning of each round. The sequence of shocks is identical across all treatments:

Table 1.1: Experimental design parameters

General parameters	
Number of periods per round	$\overline{T} = 15$
Number of rounds per session	$R = 5$
Demand parameters	$\delta = 8.50$, $\gamma = 7.275$
Fixed cost	$f = 1$
Maximal/reservation price	$p^{max} = 3$
Varying parameters	
Group size across treatments	$N \in \{2, 3, 4, 6, 10\}$
Marginal cost across rounds	$mc : (0.90, 0.50, 1.30, 0.50, 0.90)$
Cost shock sequence	$\Delta mc \equiv \eta : (-0.40, +0.80, -0.80, +0.40)$
NE price across rounds	$p^{NE} : (2.07, 1.67, 2.47, 1.67, 2.07)$

marginal cost starts at \$0.90 in Round 1, drops to \$0.50 in Round 2, rises to \$1.30 in Round 3, falls again to \$0.50 in Round 4, then rises to \$0.90 for Round 5.

1.3.3 Procedures

Experimental sessions were conducted at the University of California, Santa Barbara's Experimental and Behavioral Economics Laboratory (EBEL) using the z-Tree platform (Fischbacher, 2007)[45], between September and December of 2018. A total of 245 subjects were recruited from the experimental economics subject pool of the same university, using the ORSEE tool (Greiner, 2015)[46]. Subjects were allocated to markets of size 2, 3, 4, 6 and 10, with a total of 36, 39, 52, 48 and 70 subjects assigned to each group size condition, respectively. This setup yields 59 independent markets for the analysis.[18]

At the beginning of each experiment, subjects are provided written instructions which are also read to them aloud by an experimenter. Subjects then proceed to take a short

[18]In one session (20 subjects), the data from the final period is lost due to technical reasons. All the analysis in the results section is performed based on all the available data.

comprehension quiz.[19] In the main part of the experiment, each subject plays the role of sellers and makes a series of 75 pricing decisions, whereas consumer behavior is simulated by computer. We also elicit subjects' one-period-ahead expectations about the average price chosen by rival sellers (i.e., $\mathbb{E}_{i,t-1}[p_{-i,t}]$). These expectations are not rewarded separately, to avoid creating hedging issues. Subjects are able to set a price between the marginal cost and the maximum price (of \$3.00), in increments of \$0.01. Once all subjects set their prices and expectations, they are individually notified by the computer of the average price established by the others in their market, reminded of their own price, and shown their own resulting payoff for that period. Subjects are able to track the previous values of these outcomes through a history box that is available in their screen (see Appendix B.3).

We notify subjects that a new cost shock will occur at the beginning of each new round, either an increase or decrease, of either \$0.40 or \$0.80. We reveal the magnitude and direction of each shock immediately prior to the first period of each respective round. At that time, we also hand out copies of a printed payoff table corresponding to the new marginal cost. These tables assist subjects in estimating the profits they will receive, conditional on the hypothetical prices they and others may set in each period of that round (see Appendix B.3).

Sessions lasted a total of 90 to 125 minutes. Subjects were paid \$18.66 on average (a minimum of \$10.89 and a maximum of \$28.50), which includes the \$5.00 show-up fee and \$3.00 for the completion of the optional survey (no subject declined this offer). The remaining payoff is determined as the average payoff of a randomly chosen round of the game.

[19] We reviewed answers for each subject and provided explanations where needed. See Appendix B for all experimental material.

1.4 Hypotheses

As the experimental design specifically avoids any of the features outlined in Section 1.2.2 (e.g., frictions, information asymmetries), standard theory suggests that prices react symmetrically to shocks. However, we may observe asymmetry if certain conditions are satisfied even in the absence of these frictions. To unravel these conditions, we need to investigate the strategic tension underlying our pricing game. On the one hand, individual incentives promote undercutting opponents' prices until the NE price is reached. On the other hand, sellers may generate higher profits if they successfully sustain coordination at a price level above NE.

We first focus on the incentive each individual seller has to undercut other sellers. Consider the following variable:

$$Inc2Dev(p_{-i,t}) \equiv \pi_{i,t}(p_{i,t}^{BR}(p_{-i,t-1}), p_{-i,t-1}) - \pi_{i,t}(p_{-i,t-1}, p_{-i,t-1}), \tag{1.6}$$

which expresses the incentive of a myopic seller to deviate from the coordinated price level from the previous period, $p_{-i,t-1}$.[20] This seller is myopic as s/he ignores the possibility that the competitors might also engage in similar reasoning. By substitution of (1.1), (1.2), and (1.4) we can rewrite (1.6) as

$$Inc2Dev(p_{-i,t}) = \frac{\gamma}{4}\left(p_{-i,t-1} - \delta/\gamma - mc_t\right)^2 = \frac{\gamma}{4}\left(p_{-i,t-1} - p_t^{NE}\right)^2 \tag{1.7}$$

The incentive to deviate from price $p_{-i,t-1}$ during period t thus increases quadratically with the difference between the previous period market price and the current period NE price. Thus, individual incentives promote convergence to the NE.

We now separately analyze $Inc2Dev$ in reaction to positive and negative shocks.

[20] We are indebted to an anonymous reviewer for proposing analysis using this approach.

Consider the case where a shock of magnitude $\eta \equiv |mc_t - mc_{t-1}|$ occurs in period t. Then, application of (1.5) allows us to express $p_t^{NE} = p_{t-1}^{NE} + \eta$ and $p_t^{NE} = p_{t-1}^{NE} - \eta$ for positive and negative shocks, respectively, and we thus define:

$$\begin{aligned} Inc2Dev_{i,t}^{+} &\equiv \frac{\gamma}{4}(p_{-i,t-1} - (p_{t-1}^{NE} + \eta))^2 \\ Inc2Dev_{i,t}^{-} &\equiv \frac{\gamma}{4}(p_{-i,t-1} - (p_{t-1}^{NE} - \eta))^2. \end{aligned} \tag{1.8}$$

Straightforward manipulation of (1.8) allows us to express the difference in the incentives to deviate as:

$$\Delta Inc2Dev_{i,t} \equiv Inc2Dev_{i,t}^{+} - Inc2Dev_{i,t}^{-} = -\gamma\eta \cdot (p_{-i,t-1} - p_{t-1}^{NE}). \tag{1.9}$$

This equation implies that when market prices are above (below) NE immediately prior to a cost shock, the incentive to deviate *following the shock* will be greater (lesser) if the shock is negative than if it is positive. When however $p_{-i,t-1} \to p_{t-1}^{NE}$, or when sellers are not myopic, the incentive to deviate converges to zero. We consequently consider the absence of APT as our first null hypothesis:

Hypothesis 1: *Prices respond symmetrically to (equally sized) positive and negative shocks.*

We can test this hypothesis by exploiting the exogenous within-subjects treatment variations in marginal cost.

An interesting implication of equation (1.9) is that if prices are already collusive prior to the shock, the relatively greater temptation to deviate from the coordinated price following negative shocks suggests we should observe negative APT as opposed to the positive APT that Borenstein et al. (1997) and many others have observed[6]. The arguments of myopic best-response and of tacit collusion – explained in the next hypothesis – therefore go in the opposite directions.

Our second hypothesis concerns the second force underlying this strategic tension: the prospect of sustained higher profits through tacit collusion. We use market power, measured through the size of markups over marginal cost, to study collusion (see Section 5.2. for the description of the exact measure). Markups should be invariant to the number of rival sellers, given that both the profit and best-response functions are independent of group size. Moreover, in the absence of frictions and the ability of competitors to communicate, the theory predicts a constant markup for all levels of marginal cost. However, if tacit collusion occurs, we expect to observe higher market power (i) in markets with fewer sellers, and (ii) in the periods occurring soon after negative shocks. For (i), we expect to observe persistent coordination more often where there are fewer sellers to dampen the strength of price signals. Moreover, it is arguably easier to sustain coordination above NE pricing when one has fewer competitors, as any one of them can undermine joint coordination if they individually fail to cooperate. For (ii), we conjecture that negative shocks boost the market power of sellers (at least temporarily), as such shocks may play the role of a coordination device for attempts of collusion. We can test this hypothesis by using the between-subjects treatment variations in group size, and within-subject treatment variations in marginal cost.

Hypothesis 2: *Sellers' market power is invariant to the number of sellers in the market, and is unaffected by the existence of periodic shocks.*

Finally, our third hypothesis concerns individual pricing strategies. The Rational Expectations Hypothesis (REH) of Muth (1961) admits the possibility of expectation errors at the individual level, but which should tend to cancel out in aggregate[47]. Also, after observing $t-1$ periods of price history, a seller may learn that the others do not behave consistently with the predictions of REH. Nevertheless, conditional on expectations, sellers should select the best-response action as this maximizes their profit. As we elicit subjects' guesses on the average price set by others, we can test this hypothesis without

assuming a specific expectation model. We can interpret any intentional deviation from the profit-maximazing action as an attempt to reach collusive prices.

Hypothesis 3: *Conditional on expectations, pricing behavior follows the best-response function.*

1.5 Results

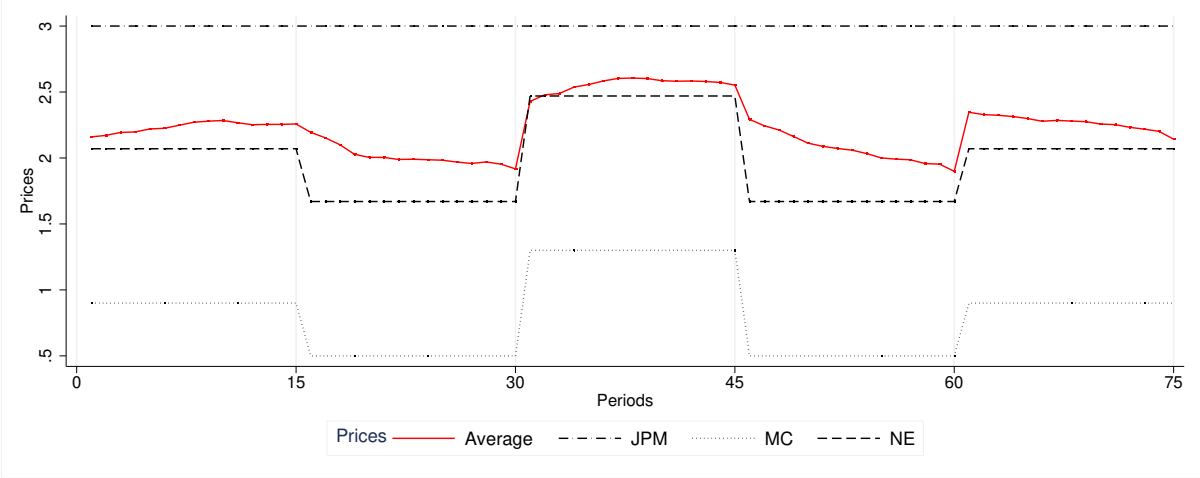

(a) Average of all market prices

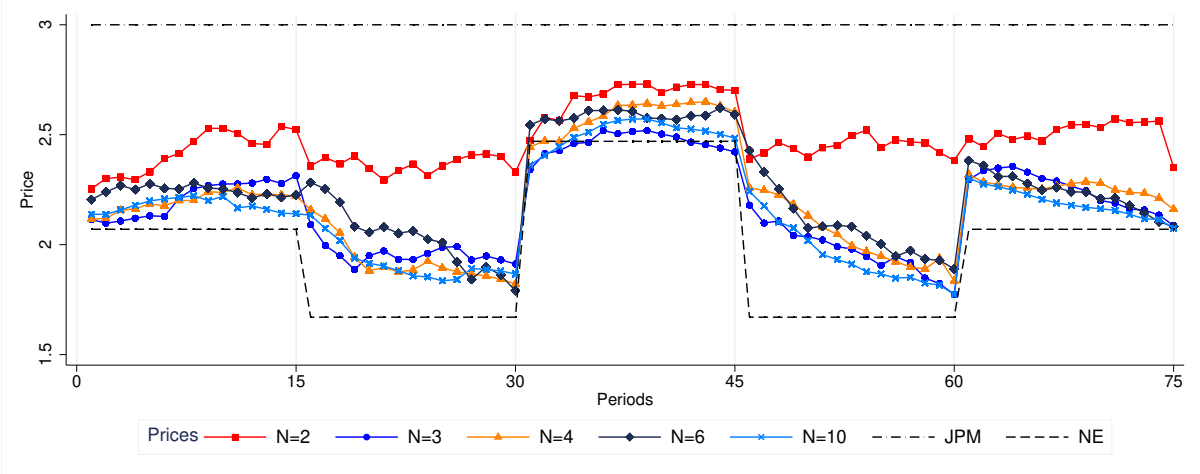

(b) Average market prices by group size

Figure 1.1: Average pricing behavior across periods and group sizes.

Figure 1.1 provides a depiction of the average price per period, as the average of all market prices and as broken out by group size. Here, market price refers to the average of all prices in market m (i.e., $p_{m,t} = \frac{1}{N}\sum_{i=1}^{N} p_{i,t}$). The reader can readily discern that for groups of size 3 and greater, average prices rise rapidly after positive cost shocks, while they fall more slowly after negative cost shocks. By contrast, for groups of size 2, it is not immediately obvious whether average pricing behavior is affected by cost shocks. A second observation that is immediately clear is that average prices are generally above the NE price, with deviations being higher following negative shocks than when following positive shocks. Overall, the visual inspection of the data suggests the presence of market behavior consistent with APT.

1.5.1 Estimation of the Asymmetry

We follow Peltzman (2000) and estimate the coefficients of the distributed lag model (DLM) to measure the magnitude of APT[9]. This model can be expressed as:

$$\Delta p_{i,t} = \sum_{k=0}^{K} b_{t-k} \cdot \Delta mc_{t-k} + \sum_{k=0}^{K} c_{t-k} \cdot ([\Delta mc_{t-k} > 0] \cdot \Delta mc_{t-k}) + \epsilon_{i,t} \tag{1.10}$$

where the change in output price (i.e., $\Delta p_{i,t} = p_{i,t} - p_{i,t-1}$) is modelled as a function of the lagged changes in marginal cost (i.e., Δmc_{t-k}). The indicator variable $[\Delta mc_{t-k} > 0]$ takes the value 1 if the change in marginal cost in period $t-k$ is positive and equal to 0 otherwise. The sum of interaction coefficients $\sum_{k=0}^{K} c_{t-k}$ reflects the magnitude of asymmetry and its persistence over K periods.

We estimate model (1.10) with Ordinary Least Squares (OLS) regressions in a stepwise manner. Figure 1.2 reports the estimated asymmetry for $K = 4$.[21] Estimates

[21] We report the full set of results in Online Appendix C.1. All estimations employ robust standard

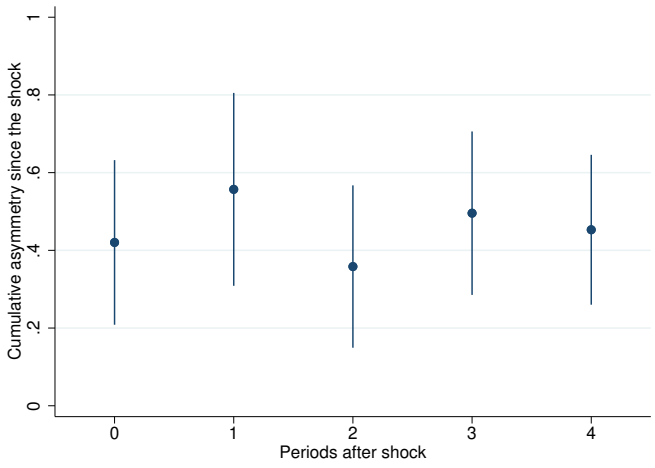

Figure 1.2: Cumulative response after K periods. Dots refer to $\sum_{k=0}^{K} c_{t-k}$. Lines represent 95% confidence intervals.

indicate that the APT is both strong and persistent. Immediate price reactions are 32.9 cents greater in magnitude for positive than for negative 80-cent shocks.

We now assess the reaction of prices to equally sized shocks between our treatment groups with non-parametric tests. We compare immediate pass-through rates of shocks that are β_0^+ and β_0^- calculated as:

$$\begin{aligned} p_{i,t+\tau}^{+} &= p_{i,t-1}^{+} + \beta_{\tau}^{+}\eta^{+} \\ p_{i,t+\tau}^{-} &= p_{i,t-1}^{-} + \beta_{\tau}^{-}\eta^{-} \end{aligned} \tag{1.11}$$

where η^+ (η^-) reflects either the large or small positive (negative) shock and $t-1$ corresponds to the period just before the shock. Note that aggregate market demand implied by (1.1) is perfectly inelastic, so that shifts in the NE price following cost shocks are exactly equal to the magnitude of the cost shock itself. Accordingly, we can denote

errors that are clustered at market level. We also include a set of indicator variables that are specific to each group size (i.e., $[N = s]$), the lagged change in the average price of rival sellers (i.e., $\Delta p_{-i,t-1}$), a three-way interaction term (between $[\Delta mc_{t-k} > 0]$, Δmc_{t-k} and group size specific indicator variables) and autoregressive terms amongst the set of regressors to check the robustness of estimates. The significance of asymmetry coefficients as well as their magnitude are robust to the inclusion of these variables.

the cases $\beta_0^+ = 1$ and/or $\beta_0^- = 1$ as incidences of "full pass-through" of cost shocks. The ratio of β_0^+ and β_0^- thus conveys information on the degree of APT in immediate cost-shock responses. A ratio of 1 would indicate the absence of APT.

Table 1.2: Asymmetry in the immediate pass-through rates

	All	$N > 2$	$N = 2$	$N = 3$	$N = 4$	$N = 6$	$N = 10$
Small shocks							
β_0^-	0.159	0.115	0.411	0.558	0.158	-0.144	0.0154
	(0.0676)	(0.0739)	(0.162)	(0.247)	(0.138)	(0.130)	(0.0985)
β_0^+	1.119	1.270	0.244	1.305	1.209	1.233	1.322
	(0.0672)	(0.0659)	(0.196)	(0.172)	(0.121)	(0.121)	(0.123)
p-value	0.000	0.000	0.411	0.028	0.000	0.000	0.000
Large shocks							
β_0^-	0.324	0.313	0.391	0.303	0.432	0.206	0.304
	(0.0362)	(0.0405)	(0.0734)	(0.107)	(0.0811)	(0.0706)	(0.0712)
β_0^+	0.639	0.718	0.181	0.537	0.779	0.945	0.619
	(0.0396)	(0.0400)	(0.111)	(0.116)	(0.0636)	(0.0796)	(0.0643)
p-value	0.000	0.000	0.035	0.202	0.013	0.000	0.006
Observations	245	209	36	39	52	48	70

The averages of pass-through rates by differing group sizes are reported. Below averages, standard errors are reported in parentheses. p-values correspond to the result of the Wilcoxon signed-rank test on the equality of pass-through rates for small or large shocks (i.e. $H_0 : \beta_0^+ = \beta_0^-$).

Table C.2 provides average value of pass-through rates for different aggregation levels.[22] First, we note that the hypothesis of full pass-through can generally be rejected. Exceptions consist of the small positive shock (i.e., $\eta^+ = 0.40$) and $N = 6$ for the large positive shock. Second, we test APT in the immediate post-shock responses by testing the equality of immediate pass-through rates for equally sized shocks as $H_0 : \beta_0^+ = \beta_0^-$ via Wilcoxon signed-rank tests. The pooled data and the data for groups of size greater than 2 suggest rejecting the null. For groups of size 3, we reject symmetry for the smaller but not for the larger shock.

[22]We report the results of an identical analysis for $\tau = 14$ in the Online Appendix C.2. Consequently, asymmetry in pass-through rates reduces but does not disappear entirely even 14 periods after the shock.

For duopolies, we see that the asymmetry is reversed; the average price response following the larger cost shock is significantly greater for the negative than for the positive cost shock. Reflecting on the differences in incentives to deviate calcuated in (1.9) from Section 1.4, and noting from Figure 1.1 that pre-shock prices were generally much higher than the NE price immediately prior to shocks in duopoly markets, we see that the relative incentive to deviate after negative shocks was much stronger in duopoly markets than in other markets. We explore the possibility that this relatively higher incentive to deviate in duopoly markets may have led to this divergent outcome.

Taken together with the estimates of the DLM, we reach the first two results of our paper:

Result 1.1: *Prices do not react symmetrically to equally sized positive and negative shocks.*

Result 1.2: *Price reactions in triopoly and larger markets are consistent with positive APT. Price reactions in duopoly markets are consistent with negative APT.*

1.5.2 Market Power

We now turn to our second hypothesis. We follow the literature in applying the Lerner index as the relevant measure of market power: $L_{i,t} = \frac{p_{i,t} - mc_t}{p_{i,t}}$ (Lerner, 1934)[48]. We propose that the difference between the observed Lerner index (i.e., $L_{i,t}$) and the "theoretical" Lerner index, that is the index that would be relevant if behavior was consistent with NE predictions (i.e., $L_t^{NE} = \frac{p_t^{NE} - mc_t}{p_t^{NE}}$), provides a measure of "excess" market power due to collusion. We further propose this as an appropriate measure of tacit collusion, as our price competition structure incorporates homogeneous goods and we control marginal costs. Thus, we do not suffer the identification problem of observational studies. Our

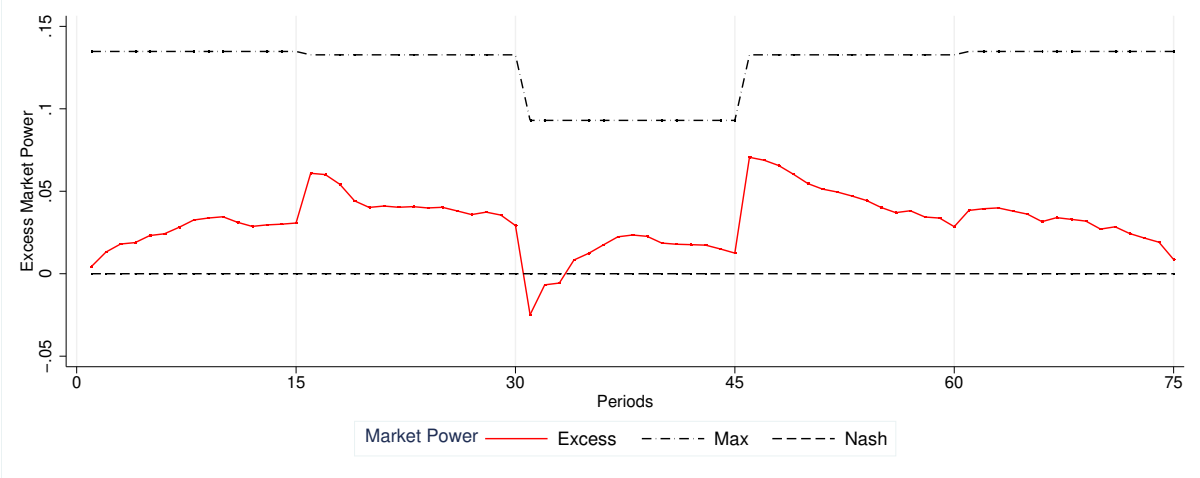

(a) Average of excess market power

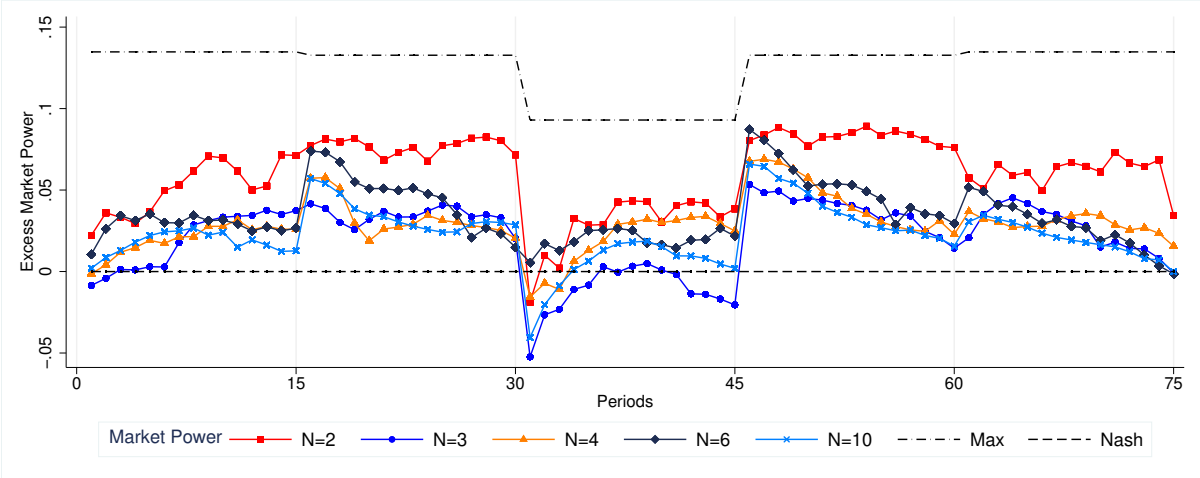

(b) Excess market power by group size

Figure 1.3: Excess market power across periods and group sizes.

In both subfigures, "Max" refers to the maximum excess market power that can be observed (i.e., when $L_{i,t} = L_t^{max} = \frac{p^{max} - MC_r}{p^{max}}$) and "Nash" refers to the case $L_{i,t} = L_t^{NE}$.

measure of excess market power can then be expressed as:

$$L_{i,t}^x = L_{i,t} - L_t^{NE} = mc_t \left(\frac{1}{p_t^{NE}} - \frac{1}{p_{i,t}} \right). \tag{1.12}$$

Figure 1.3 depicts the average of our measure of excess market power, by period

and treatment. Upon visual examination one can immediately see that excess market power generally lies above the theoretical "Nash" level, consistent with an environment in which tacit collusion exists. Also, this average measure reaches its highest levels during the second and fourth rounds, the two rounds that immediately follow negative shocks. Following the large positive shock at the beginning of the third round, excess market power falls so much that it turns negative for several periods. Following the small positive shock at the beginning of the fifth round, excess market power does not notably react.

We test the veracity of these observations by performing OLS regressions.[23] We consider the following specification:

$$L^x_{i,t} = \alpha + \sum_{s \neq 2} \delta_s \cdot [N = s] + \sum_{e \neq 1} \gamma_e \cdot [r = e] + \epsilon_{i,t}, \tag{1.13}$$

where the excess market power of seller i in period t is modeled as a function of group size- and round-specific indicator variables. According to our null hypothesis, the model with no independent variables should fit the data as well as this model.

Table 1.3 reports the estimates in a step-wise manner. In model (5), we truncate the data to the periods where shocks shift the marginal cost (i.e., periods 16, 31, 46 and 61) and replace the dependent variable with the change in excess market power as $\Delta L^x_{i,t}$. This allows us to interpret the estimates of round specific indicator variables as the immediate effect of cost shocks on the tacit collusion in model (5).

First, we reject the null hypothesis in all specifications except (2) at a confidence level of 0.01 with the F-test. The fact that the constant α is positive and significant in model (1) indicate the overall presence of tacit collusion. Second, the coefficients of round-specific indicator variables in model (3) indicate that tacit collusion is higher

[23]The non-parametric counterpart of this test is reported in Online Appendix C.3

Table 1.3: Excess market power

	(1)	(2)	(3)	(4)	(5)
Constant	0.032***	0.060***	0.025***	0.054***	0.004
	(0.004)	(0.011)	(0.004)	(0.011)	(0.007)
$N = 3$		-0.039*		-0.039*	0.018
		(0.015)		(0.015)	(0.020)
$N = 4$		-0.031*		-0.031*	0.028**
		(0.014)		(0.014)	(0.010)
$N = 6$		-0.026		-0.026	0.047***
		(0.014)		(0.014)	(0.010)
$N = 10$		-0.038**		-0.038**	0.029*
		(0.012)		(0.012)	(0.011)
$r = 2$			0.017***	0.017***	
			(0.003)	(0.003)	
$r = 3$			-0.014***	-0.014***	-0.084***
			(0.004)	(0.004)	(0.009)
$r = 4$			0.023***	0.023***	0.028***
			(0.004)	(0.004)	(0.006)
$r = 5$			0.005	0.005	-0.020**
			(0.004)	(0.004)	(0.007)
Observations	18355	18355	18355	18355	980
Adjusted R^2	-	0.052	0.054	0.106	0.225
F-statistic	-	2.607	31.289	19.288	27.003

Results of OLS regressions on specification 1.13 are reported. In model (5), the dependent variable is the change in excess market power, $\Delta L^x_{i,t}$. Below estimates, robust standard errors that are clustered at the market level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

during the second and fourth rounds, and lower during the third round relative to the first round. In model (5) where we truncate the data, the coefficient of rounds 3 and 5 are negative and that of round 4 is positive. Furthermore, we reject the hypothesis $H_0 : \alpha + \delta_s + \gamma_e = 0$ at a confidence level of 0.05 (i.e., $\alpha + \delta_{\{N=3,4,10\}} + \gamma_5 = 0$). We can thus say that immediately after a negative (the large positive) shock, the excess market power increases (decreases). Third, coefficients of group size specific indicator variables are negative and significant, although at marginal level for $N = 6$ (p-value= 0.071) in model (2). Here, we also reject the hypothesis $H_0 : \alpha + \delta_s = 0$ for all s (p-value< 0.01).

This suggests that tacit collusion is present in all markets but its magnitude is smaller when $N > 2$. The sign of these coefficients in model (5) suggests that markets larger than size 3 increase their market power in response to the first negative shock. Lastly, we generally reject the hypothesis $H_0 : \alpha + \delta_s + \gamma_e = 0$ in the most unrestricted model (4) (11 times out of 15 tests at p-value< 0.05). The overall interpretation of these tests provide the basis of our second result:

Result 2: *Excess market power (i.e., tacit collusion) is not invariant to shock direction and group size. It is persistently higher in duopolies, and in larger-sized markets it rises following negative cost shocks.*

1.5.3 Deviations from Best-Response

Finally, we assess the deviation of subjects' prices from the profit-maximizing best-response action that is computed by using the submitted expectations/guesses (i.e., $p_{i,t} - p_{i,t}^{BR|\mathbb{E}}$). Conditional on the subjects accurately reporting their expectations, they risk lower profits if they fail to best-respond to these reported expectations. Consequently, the observed deviations can be attributed either to error or alternatively to strategic motives (i.e. attempts to collude). To argue that the deviations we observe in our experiments are not entirely due to erroneous behavior, we compare the magnitudes and directions of such deviations to the average magnitude of expectation errors (i.e., $\mathbb{E}_{i,t-1}[p_{-i,t}] - p_{-i,t}$) and the average of absolute expectation errors.[24] If subjects deviate from their best-response action in a way that is different from their expectation errors, this suggests that deviations reflect intentional behavior.

Figure 1.4(a) depicts the average value of these deviations over time. The average

[24]We label these latter two as "errors" rather than as "deviations" as there is no strategic benefit to knowingly submitting inaccurate guesses/expectations in our experiment.

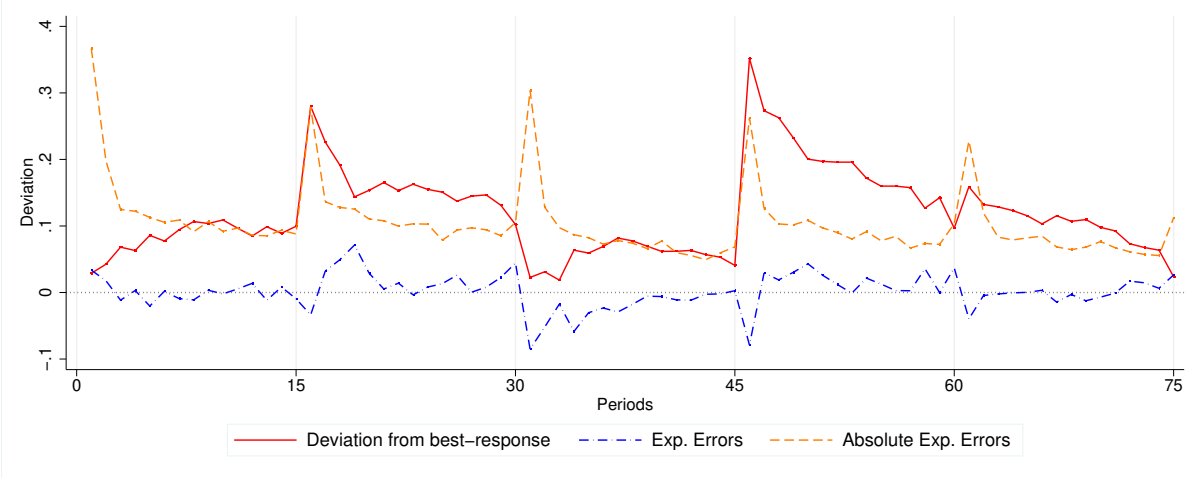

(a) Average of different deviations across all markets

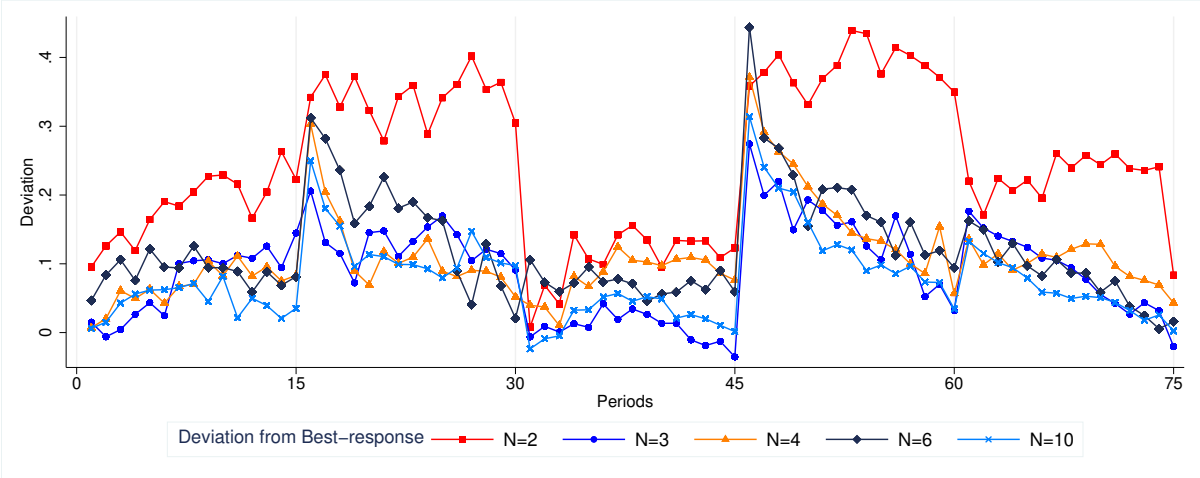

(b) The deviations from best-response action by group size

Figure 1.4: Deviations from best-response action and errors in expectations.

expectation errors are remarkably close to zero, with no obvious trend across periods. Although this suggests that beliefs are on average correct, it does not imply the complete absence of errors: the average measure of absolute expectation errors lies well above zero throughout the experiment. The latter peaks following both positive and negative cost shocks but subsequently trends downward. The patterns of high initial absolute expectation errors and slow convergence are consistent with those of prior experiments in which prices are strategic complements (e.g., Hommes et al. 2005[49]; Cooper et al.,

2021[50]). However, deviations from the best-response action reveal a different and rather interesting pattern: they peak sharply following negative shocks and remain high during these rounds, but do not peak similarly following positive shocks. The second graph in Figure 1.4 depicts the average of deviations from best-response action by group size. The same pattern can be traced across our treatment groups.

We perform OLS regressions to study deviations from best-response. Consider the following specification:

$$p_{i,t} - p_{i,t}^{BR|\,\mathbb{E}} = \alpha + \sum_{s \neq 2} \delta_s \cdot [N = s] + \sum_{e \neq 1} \gamma_e \cdot [r = e] + \epsilon_{i,t} \tag{1.14}$$

where the deviation of subject i's price from the best-response action conditional on the submitted guess is modeled as a function of group size- and round-specific indicator variables. Our null hypothesis concerns the overall significance of this model.

Table 1.4 reports the estimates in a step-wise manner. In model (5), we truncate the data to the periods where shocks shift marginal cost the same way as we did in the previous section, and replace the dependent variable with the change in deviation from best-response action following a cost shock.

We reject the null hypothesis in all specifications. The hypotheses $H_0 : \alpha + \gamma_s = 0$ in model (2) and $H_0 : \alpha + \delta_s = 0$ in model (3) can be rejected at a significance of $p < 0.01$. This points out to the following two results: (i) Sellers deviate more (less) from the associated best-response action following negative (positive) shocks and (ii) deviations are lower when $N > 2$. We see that the sign of group size indicator coefficients in models (4) and (5) are flipped. In model (4), they reflect the fact that groups of size 3 and larger deviate less, on average, relative to duopolies. In model (5), they correspond to the immediate reaction of these groups to the first negative shock. These deviations rise further when a large negative shock shifts the marginal cost down ($\hat{\gamma}_4 = 0.131$) while they

Table 1.4: Deviations from best-response

	(1)	(2)	(3)	(4)	(5)
Constant	0.120***	0.083***	0.248***	0.212***	0.044
	(0.013)	(0.012)	(0.041)	(0.040)	(0.028)
$N = 3$			-0.159**	-0.159**	0.122*
			(0.049)	(0.049)	(0.055)
$N = 4$			-0.138**	-0.138**	0.163**
			(0.050)	(0.050)	(0.051)
$N = 6$			-0.128*	-0.128*	0.210***
			(0.051)	(0.051)	(0.041)
$N = 10$			-0.170***	-0.170***	0.144**
			(0.044)	(0.044)	(0.046)
$r = 2$		0.080***		0.080***	
		(0.010)		(0.010)	
$r = 3$		-0.028**		-0.028**	-0.260***
		(0.010)		(0.010)	(0.031)
$r = 4$		0.112***		0.112***	0.131***
		(0.015)		(0.015)	(0.025)
$r = 5$		0.018		0.018	-0.118***
		(0.013)		(0.013)	(0.024)
Observations	18355	18355	18355	18355	980
Adjusted R^2	-	0.044	0.051	0.095	0.170
F-statistic	-	37.840	4.013	19.546	40.442

Results of OLS regressions on specification 1.14 are reported. In model (5), the dependent variable is the change in deviation from the best-response action following a cost shock, $\Delta(p_{i,t} - p_{i,t}^{BR|\mathbb{E}})$. Robust standard errors are clustered at the market level and are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

drop significantly in response to the large positive shock ($\hat{\gamma}_3 = -0.260$). In consequence, we reach to the following results:

Result 3.1: *Sellers deviate on average above their best-response action.*

Result 3.2: *Deviations from the best-response action grow (shrink) following negative (positive) shocks.*

1.6 Discussion

Our results point to the co-appearance of asymmetric price transmission and tacit collusion. The latter seems to be the result of strategic behavior, as our analysis of deviations from best-response action reveals. These findings are consistent with theories that cast tacit collusion as having a significant role in the emergence of APT, such as the trigger price model in Borenstein et al. (1997)[6]. Most of the other theoretical explanations of APT in the literature cannot account for the pricing behavior observed in our results. We can reasonably exclude, for example, the influence of explicit collusion (i.e., involving direct communication), capacity constraints, inventory limitations, (a)symmetric menu costs, consumer loss aversion, (a)symmetric search costs, contexts of alternating price moves and price lockup periods, and so forth, as being necessary conditions for APT, since these features are excluded by our design.

We cannot, however, claim a monotonic relation between the magnitude of APT and tacit collusion: the pricing behavior of duopolies in our experiment is revealed to be consistent with 1) a negative APT response in the immediate aftermath of shocks, and 2) elevated and relatively stable pricing behavior between shocks. We propose an explanation for the exceptionality of the duopoly result follows, tackling the latter result first: in duopolies collusion is so strong that sellers are, by and large, able to maintain cooperative (tacitly collusive) pricing over a sustained period of time, with pricing showing no reversion to Nash. We therefore argue that APT requires significant, but imperfect, collusion.[25]

[25]The fact that our duopolies reached almost stable collusion, while larger markets did not, is consistent with the literature we review in Section 2.3. This can be attributable to a combination of two factors: first, coordination between market participants becomes increasingly difficult with each new seller, and three may well be the number from which the difficulties and costs involved in maintaining coordination start to exceed the marginal benefits; second, our duopolies are unique in that each participant can deduce the choices made by the other participant by observing aggregate market outcomes. In a triopoly or larger market, by contrast, it is not possible for sellers to detect whether an aggregate market outcome is due to the defection of a single competitor, or from a broader but shallower defection by multiple

The success of duopoly markets in maintaining average prices well above equilibrium levels between shocks, and well above average price levels of triopoly and larger markets, provides a plausible explanation for the negative APT result we observed with duopolies for large shocks: as is apparent in an examination of (1.9), the high price deviations above NE prices in the lead-up to shocks in duopoly markets implies a correspondingly greater incentive for firms to deviate downward following negative than positive shocks. Thus we have the interesting result that while the dynamics of our duopoly markets appears to have enabled more success in attempts to tacitly collude between shocks, this success (in the form of elevated prices) also made these markets relatively more prone to respond strongly to downward than upward cost shocks.[26]

If tacit collusion is indeed a significant causal force behind APT, then our work has important implications for antitrust enforcement policy against collusion and price-fixing. In particular, regulators may consider APT in a market as a signal for the presence of collusion between firms in that market. Since many real-world interactions between competitive firms are repeated indefinitely, such collusion may even be sustainable as a NE. Further research is needed to determine whether collusion is an important cause of APT behavior in field settings, and if so, whether suitable forms of regulatory intervention might exist to reduce such collusion without increases in inefficiency.

We propose that follow-on research may yield further insights into the mechanisms through which tacit collusion leads to APT, as well as potential policy responses that might diminish its frequency and magnitude. In particular, future experiments should

competitors.

[26]As we noted in the Results section, our dupoloy markets exhibit pricing behavior consistent with *negative* APT, while larger markets' behavior was consistent with *positive* APT. We see these apparently contradictory results as indicating a tension between two phenomena operating in opposite directions: on the one hand, collusive dynamics encourage competitors to respond more sharply to upward than to downward cost shocks, as we have argued throughout the paper; however, as prices rise above NE levels the differential incentives to deviate become stronger following negative shocks than positive shocks. At some point the success of tacit collusion may be so great as to reverse the direction of the resulting APT.

address the impact of different levels of information transparency. Most notably, testing the effects of providing feedback on individual prices and/or payoffs of rivals on APT may provide particularly helpful insights. The latter is shown to lead to more rivalistic outcomes in experimental oligopoly studies, as it initiates imitation dynamics (Fiala and Suetens, 2017), while the former can lead to more collusive levels.[38] Nevertheless, both may reduce the degree of asymmetry in price transmission. Another area of needed research is to explore the roles of market power and market concentration in shaping APT pricing behavior. In our experimental design we explicitly kept the incentives provided to sellers the same across markets of varying sizes to study the pure number effect similar to Hanaki and Masiliunas (2021)[15]. Finally, future studies may benefit from testing the robustness of our findings to alternative demand specifications. The parameters of demand in our experimental markets are consistent with goods such as retail gasoline, which are demanded inelastically but for which suppliers face elastic demand. Our results are also more relevant to markets in which there are close substitutes.

Chapter 2

Mutual Payoff Information and Strategic Uncertainty

2.1 Introduction

Game-theoretic models are typically motivated by the idea that players reason about the behavior of others and choose their strategies accordingly. This reasoning can be informed directly by observing the payoff structure of the game or indirectly by observing and learning from the actions of other players. The type of introspective reasoning supported by directly observing others' payoffs is often embedded in models of strategic decision making, such as higher-level reasoning in level-k models (Stahl and Wilson, 1994[51], Nagel, 1995[52]), which assume mutual payoff information. Additionally, experiments using eye-tracking have found that subjects devote a sizable amount of attention to the payoffs of other players (Knoepfle et al, 2009[53], Polonio and Coricelli, 2019[54]), and it has been documented that subjects engage in higher-level reasoning when other players' payoffs can be observed. In contrast, in providing an interpretation for his seminal equilibrium concept, Nash (1950) makes it explicit that, "it is unnecessary to assume

that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes."[55] Similarly, theoretical models of learning explore how equilibria can be reached and selected through processes of learning, adaptation, and/or imitation rather than introspection (Fudenberg and Levine, 2009[56]), and uncoupled learning models describe how equilibria can be reached in the absence of information about other players' incentives or even their existence. (Hart and Mas-Colell, 2006[57], Foster and Young, 2006[58], Young, 2009[59], and Babichenko, 2010[60])

It is not well understood, however, how mutual payoff information—in addition to observing others' actions—affects players' behavior in strategic interactions. On the one hand, knowledge of payoffs can reveal opportunities to coordinate on Pareto-dominant outcomes; on the other hand, being aware of an opportunity to cooperate can increase the strategic tension of a game if the cooperative outcomes are associated with actions that are dominated for at least one player. Furthermore, it is not clear through which channels such an effect would operate. Mutual payoff information may affect players' initial perceptions of a game, the process through which they learn and gain experience, or both.

We present an experiment designed to study how mutual payoff information affects play in the Prisoner's Dilemma (PD) and the Stag Hunt (SH). Subjects are asked to play these games repeatedly with random re-matching of opponents each round. In our partial-information treatment, subjects observe their own payoffs and the action of their opponent after each round, but they never observe the other's payoffs. Comparing this partial-information version to the full-information baseline treatment in which subjects additionally observe the payoffs of the other player allows us to detect differences in play that arise due to the presence or lack of mutual payoff information.

The appeal of contrasting the PD with the SH in our experiment is grounded in

our conjecture that mutual payoff information affects behavior differently in these two games. While the SH exhibits a tension between a mutually desirable outcome and avoiding personal risk, the PD exhibits a tension between a socially optimal outcome and personal gain. In the SH, knowledge of the other's payoffs arguably reduces the strategic tension of the stage game by revealing that a mutually beneficial outcome exists. In the PD, on the other hand, mutual payoff information may increase the strategic tension of the stage game. Without knowing the other's payoffs, there is little reason not to choose the payoff-dominant action. Indeed in this circumstance there is arguably no "dilemma," as subjects are unaware of the socially optimal, albeit dominated, cooperative action pair.[1] By introducing mutual payoff information, however, players observe a socially optimal outcome which can only be reached at personal expense. In the presence of social preferences, either player might prefer such an outcome but still wonder whether the other player shares these preferences and will reciprocate, thereby potentially making one's own attempt to cooperate futile. This uncertainty may increase the strategic tension of the PD game.

Our main result is that the presence or absence of mutual payoff information has a strong effect on play in both games. The fraction of subjects who initially cooperate in the PD or who coordinate on the payoff-dominant equilibrium in the SH is substantially higher under full-information than under partial-information. In the SH, this effect is remarkably persistent: throughout all rounds of the game, the vast majority of subjects choose the action consistent with the payoff-dominant equilibrium of the SH in the full-information treatment, while choosing the risk-dominant action under partial-information.[2] By contrast, in the PD play converges toward the unique NE of the game

[1]For consistency with prior literature, we use the words "cooperate" and "defect" to describe the action choices in the PD with acknowledgement that they are only well-defined from the player's perspective in the presence of mutual payoff information.

[2]Strictly speaking, following the canonical definition of risk dominance from Harsanyi and Selten (1988), an equilibrium cannot be "risk-dominant" for a player who does not have access to full payoff

under both information treatments.

To investigate the channels driving this information treatment effect, we estimate a belief-learning model that is a special case of Camerer and Ho's Experience-Weighted Attraction (EWA) model.[62] In this model, players choose the action with the higher expected payoff (higher attraction) in each round, given their beliefs about the actions of other players. Beliefs are formed as weighted averages of the observed history of play and a prior belief. We estimate a set of four parameters for each game and information treatment: two parameters for the initial attractions (one for each strategy), one parameter related to how subjects weight past observations when updating beliefs, and one parameter governing how precisely subjects respond to their estimated beliefs (i.e. to the updated attractions they hold for each action).

To examine the importance of these parameters, we conduct simulations where we flip the estimated parameters at the information treatment level to isolate the effect associated with initial play versus learning. We find that doing so results in simulated play that sharply diverges from what we actually observe, suggesting that our information treatment creates significant differences both in initial play and also in learning across rounds.

We propose that our results may be explained, at least in part, by the impact mutual payoff information has on strategic uncertainty and structural uncertainty and the relationship these have to the strategic tensions inherent in each game. While strategic uncertainty reflects a player's uncertainty about which strategies the other players will choose, structural uncertainty is the uncertainty about the parameters of a game.[63] Removing mutual payoff information mechanically introduces structural uncertainty to both games, which we argue increases strategic uncertainty in the PD while decreasing

information.[61] For the sake of clarity and consistency, we will refer to the "risk-dominant" and "payoff-dominant" actions for both the full- and partial-information treatments in the SH game.

it in the SH, and affects subject behavior in ways that are captured by changes in the parameters of our learning model.

First, for subjects who face greater strategic uncertainty, experimenting may be relatively more attractive. Indeed, we find that subjects respond less precisely in the SH partial-information treatment than in the full-information counterpart, while the opposite is true of the PD. Additionally, when faced with greater strategic uncertainty, subjects may place greater weight on recent observations as opposed to their initial perception of the game. This is reflected in our estimate of the belief updating parameter in the SH, which suggests that subjects place greater weight on recent observations under the partial-information treatment. While the effect is not significant in the PD, the direction of our estimate is consistent with this interpretation.

Our paper makes several contributions. First, we provide evidence that even in the absence of mutual payoff information most subjects eventually choose actions that correspond to Nash equilibria. Second, we show, to our knowledge for the first time, that mutual payoff information can affect equilibrium selection in the SH throughout all rounds and initial play in the PD. These differences in play are associated with greater coordination on the payoff-dominant equilibrium across all rounds of play for the SH and greater cooperation in the short-run for the PD. Finally, while a number of previous experiments (e.g., Mookherjee and Sopher, 1994[64], Cox et al., 2001[65], and Feltovich and Oda, 2014[66]) employ learning models in their analysis, the goal has been to evaluate such models in isolation or relative to one another. We take a different approach and instead interpret our learning model parameters as indicative of underlying changes in the way players perceive, update, and react to the availability of opponent payoff information. We propose that these changes reflect how strategic and structural uncertainty vary across our information treatments.

Related experiments. A closely related collection of experiments studies how subjects behave in incomplete information environments similar to ours. We are not aware of an experiment, however, that employs the same information treatments we use to the PD and the SH game. Mookherjee and Sopher (1994) find that in a matching pennies game with fixed partners, choice frequencies tend towards the unique mixed strategy equilibrium whether or not subjects are presented with opponent payoff information.[64] Oechssler and Schipper (2003) have subjects play incomplete information treatments in the SH and PD;[67] however, our experiments are difficult to compare as their design creates incentives for experimentation in the initial rounds by paying subjects more later and by providing rewards for correct answers regarding the opponents' payoffs. Perhaps most notable for our experiment, subjects' play converges towards Nash equilibrium despite the inability of subjects to fully perceive the game structure.

Most similar to our design is Feltovitch and Oda (2014) in which subjects play incomplete information versions of six games, including the SH and PD, with treatments for random re-matching and fixed matching.[66] Their results suggest that the matching mechanism does matter in the incomplete information environment, with fixed-pair often leading to increased coordination on pure strategy equilibria, higher payoffs, and faster convergence, but no full-information treatments are run for comparison. McKelvey and Palfrey (2001) run an ambitious number of games and information treatments, including versions of the PD and SH, but do not report choice frequencies for direct comparison to our results.[68][3]

[3]Additional studies have altered the information structure in ways that make them less comparable to our environment. For example, Cox et al. (2001) and Danz et al. (2012) inform subjects about a set of payoffs from which the actual opponent payoffs may be drawn,[65][69] Niscklisch (2011) introduces information asymmetries into the environment,[70] and Nikolaychuk (2012) matches subjects with a computer following a learning algorithm and let them observe their own earnings after each round, or the whole payoff matrix, in versions of the PD, SH, and Battle of the Sexes.[71] Andreoni et al. (2007) vary the information that bidders have about their rivals' valuation in first- and second-price auctions, and document that subjects' behavior in response to this information is consistent with theoretical predictions.[72]

Our paper also relates to a literature examining the impact of payoff information on the formation of cooperation. In both Friedman et al. (2015) and Huck et al. (2017), subjects play Cournot games that exhibit tension similar to our PD between competition and cooperation.[73][74] Subjects do not have access to their own or others' payoff function but are told payoff functions are symmetric and time invariant, and they receive feedback on their own and others' payoffs and actions. Huck et al. (2017) introduce a comparison treatment where subjects are shown the possible payoffs they could have received based on their partners last action. Contrary to our results, in these games payoff information tends to lead to play that is more competitive, suggesting that payoff information hinders cooperation. Nax et al. (2016) study cooperation in a voluntary contribution game under different information structures, and find that contribution rates are similar when players have full information about the game and when they only get to observe their own payoffs.[75]

Another strand of literature we contribute to addresses the question of which equilibrium play converges to in the SH. With one payoff-dominant and one risk-dominant equilibrium, the SH embodies a tradeoff between maximizing social efficiency and minimizing personal risk. Many experiments (e.g. Battalio et al., 2001[76], Schmidt et al, 2003[77], Dubois et al., 2012[78], and Kendall, 2020[79]) have been designed to better understand the conditions under which play converges to either equilibrium. A common feature of these studies is that subjects play SH games where payoffs are commonly known, and by varying these payoffs across games, diverse theoretical predictions can be disentangled. Our paper is related to this literature but takes a different approach; we keep payoffs constant across treatments but vary whether players get to observe the other's payoffs. Doing so allows us to identify mutual payoff information as an important factor for the payoff-dominant equilibrium to arise in the SH.

This paper proceeds as follows: Section 2.2 provides details of our experimental de-

Figure 2.1: The Games and Information Treatments From the Row Player's Perspective

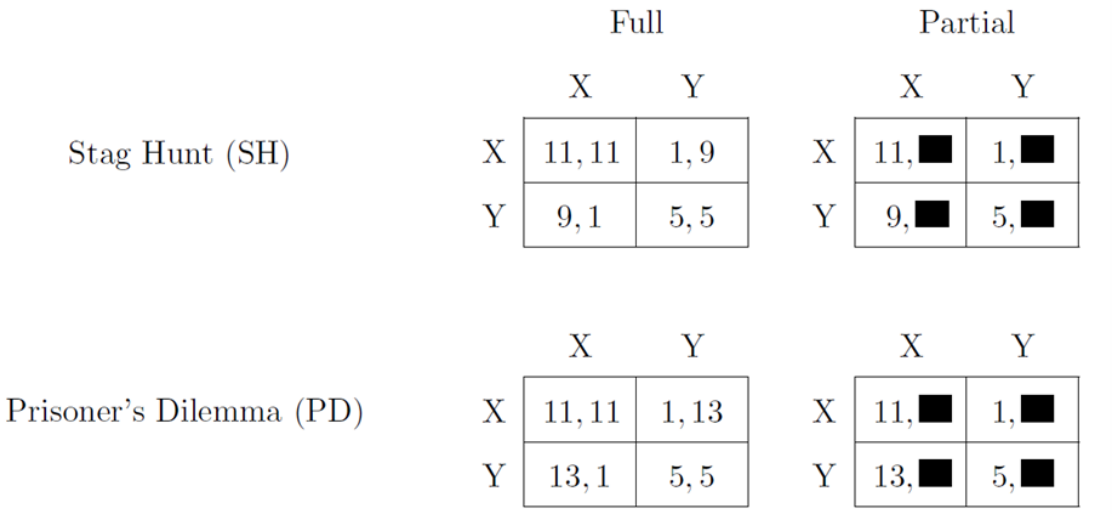

sign. Section 2.3 presents our main results in a descriptive manner. Section 2.4 introduces and estimates the belief learning model. Additionally, we present simulations to isolate the effect of initial play versus learning and to explore the robustness of our estimates. Section 2.5 provides a discussion and Section 2.6 concludes.

2.2 Experimental Design

Overview. In our experiment, subjects played SH and PD games in randomly rematched pairs over many periods. We employed two information treatments per game: In full-information ("Full") treatments, subjects were shown the complete payoff matrix, including own and opponent payoffs. In partial-information ("Partial") treatments, on the other hand, we showed subjects only their own payoffs. In both treatments, players made choices simultaneously. At the end of each round we notified subjects of their opponent's action, reminded them of their own action, and displayed their resulting payoff, but not their counterpart's.

Games and information treatments. For each game and information treatment, Figure 2.1 depicts the payoff matrices, as well as the available payoff information from the row player's perspective. In the SH, there are two pure-strategy equilibria. One is payoff-dominant (X, X) and the other one is risk-dominant (Y, Y). In the PD, there is one equilibrium in strictly dominant strategies, (Y, Y). Note that the Full and Partial treatments of each game have the same payoffs (though they are partially hidden from subjects in the Partial treatment), thus keeping the best-response correspondences and equilibria constant across the two information treatments of a game.[4]

40 rounds with random re-matching. In each treatment, subjects played 40 rounds of a game and were randomly and anonymously re-matched with other subjects each round.[5] The information treatment was the same for all subjects within a session. That is, in the Full treatment, it was common knowledge that subjects were being re-matched with other subjects who could observe the whole payoff matrix. Similarly, in the Partial treatment, it was common knowledge that subjects were being re-matched with other subjects who could only observe their own payoffs and that any opponent they would be matched with had the same payoffs as their previous opponents.

To credibly implement the Partial treatment, it was crucial that subjects were not able to infer that they were playing a symmetric game.[6] We therefore employed a two-population matching mechanism, in which subjects were randomly divided into two groups (labeled A and B) at the beginning of the experiment and were exclusively

[4]This also means that the size of the basin of attraction for pure strategies X and Y—in our context the beliefs one would have to have that one's counterpart will select the opposite action in order for expected payoffs of one's own actions to be equal—are held constant between information treatments in the SH. Embrey et al., 2017 provide experimental evidence of a positive correlation between the size of the basin of attraction for a given strategy and the frequency with which subjects select this strategy.[80]

[5]Ghidoni, 2019 find that cooperation rates in a PD game with 10 rounds are very similar when subjects are randomly re-matched in groups of 6 or with a new opponent each round.[81]

[6]This would enable subjects to determine their opponent's payoffs, thereby undermining the information treatment.

matched with subjects of the opposite group throughout the session. Subjects were told that while they could not observe the payoffs of participants from the other group, all participants of the other group faced the same payoffs. Likewise, participants from the other group would always be matched with somebody of one's own group, facing the same payoffs as themselves. We used this two-population matching mechanism in both information treatments for consistency.

Feedback after each round. After selecting an action in each round, subjects learned what action their opponent had chosen in that round and their resulting payoff. Throughout the 40 rounds of the game, we showed subjects a table on the left side of the screen with a list of their choices, their counterparts' choices, and their own payoffs for the current and previous rounds of the game. Figure 2.2 contains example screenshots depicting the information shown to subjects, before, during, and after selection of an action in the PD Partial treatment.[7]

Comprehension. The experimenters handed out written instructions, which they also read aloud to subjects. To participate, subjects had to correctly answer a comprehension quiz on paper.[8] The instructions and the comprehension questions we used are provided in Appendices D.1 and D.2.

We explained to subjects how to read a payoff matrix based on an example of an asymmetric game with all payoffs visible. In sessions that included a Partial treatment, we then told subjects that in part of the experiment the other's payoffs would be covered. This was illustrated by using the same example game matrix, this time with gray squares

[7]See Appendix D.3 for corresponding screenshots from experiments conducted under the SH Full treatment.

[8]Only 5 of 196 subjects failed to answer all quiz questions correctly the first time. We pointed out to these subjects what they did wrong and provided them with a second quiz version with different matrix entries; each of these 5 subject correctly answered all of these questions on the second attempt.

Figure 2.2: Screenshots of Experimental Interface, PD Partial Treatment

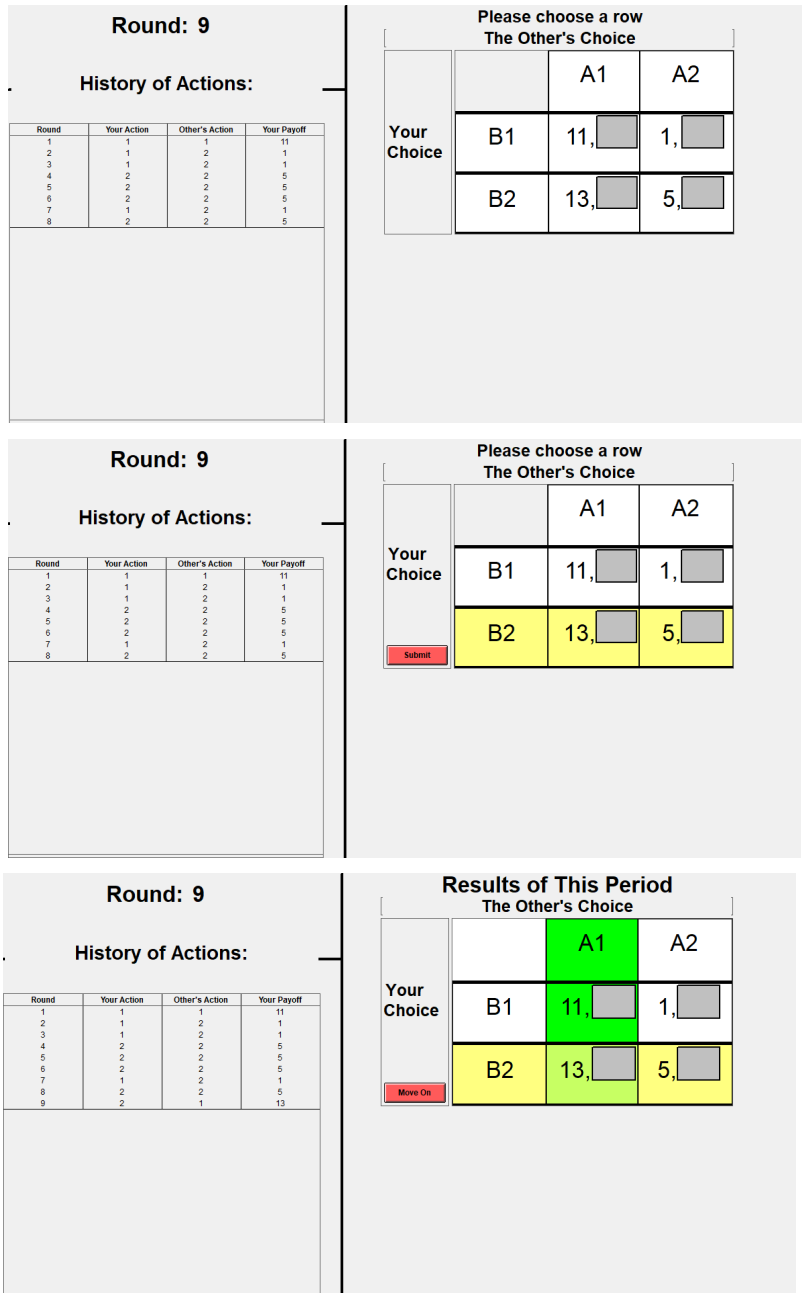

Note: This figure shows screenshots of the experimental interface of the PD Partial treatment for a subject that was assigned to Group B. The first panel displays the interface at the beginning of the ninth round; the subject sees the history of the first eight rounds but has not chosen their next action yet. The second panel depicts the same interface after the subject has selected—but not yet committed to—Action B2. Finally, the third panel shows the feedback the subject receives at the end of Round 9; they can see that the person from Group A they were matched with in Round 9 chose Action A1. Note in this panel that the History of Actions table has been updated with the results from Round 9.

covering the other's payoffs in each cell, hereby mimicking the interface used in the experiment. Since all subjects in a session were given the same instructions, it was implicitly communicated that all subjects would only see their own payoffs in the Partial treatment and thus their opponent (who would also only see their own payoffs) would not know the subject's payoffs.

Organization of treatments, between versus within analysis. Each experimental session consisted of two blocks of 40 rounds each, one with a Full treatment and the other with a Partial treatment, for a total of 80 rounds of play. Before each session began we administered the relevant instructions followed by a comprehension quiz. We did not inform subjects of details of the second game until the first game was completed.

In sessions 1-6 subjects first played a Partial treatment of a game (SH in sessions 1-3, and PD in sessions 4-6), followed by the Full treatment of the same game. This design feature thus enables a within-subjects analysis. In sessions 7-12 subjects first played a Full treatment of one game (SH in sessions 7-9 and PD in sessions 10-12), followed by a Partial treatment of the other game (PD and SH, respectively). This enables a between-subjects analysis and addresses potential order effects that could occur if subjects only played the Full treatment of a game after first playing its Partial counterpart.

The allocation of treatments for each part of the 12 experimental sessions we conducted and the number of subjects per session are depicted in Table 2.1. Note that subjects never got to play the Full version of a game before the Partial version of the same game to avoid them inferring that the payoffs in the second game were the same as in the first game.

Our organization of treatments enables us to conduct both between-subjects and within-subjects analyses, as summarized in Table 2.2. We center our analysis on the pooled data that uses all available data from both the first and second games of all

Sessions	Part 1	Part 2	# Subjects per Session
1-3	SH - Partial	SH - Full	16, 16, 20
4-6	PD - Partial	PD - Full	16, 16, 18
7-9	SH - Full	PD - Partial	16, 14, 14
10-12	PD - Full	SH - Partial	16, 18, 14

Table 2.1: Treatments by Session

Analysis	Game	Data
Between-subjects	SH	first part sessions 1-3, first part sessions 7-9
Between-subjects	PD	first part sessions 4-6, first part sessions 10-12
Within-subjects	SH	sessions 1-3 (first and second part)
Within-subjects	PD	sessions 4-6 (first and second part)

Table 2.2: Between- vs. Within-Subjects Analysis

sessions, either for the SH or the PD. In results not reported, we find that separately analyzing the between-subjects or within-subjects data yields qualitatively similar results.[9]

Experimental details. We programmed the experimental interface using Z-Tree (Fischbacher, 2007[83]), and conducted all sessions at the University of California, Santa Barbara's Experimental and Behavioral Economics Laboratory (EBEL) in April and September of 2018. We recruited subjects for the experiment from the EBEL subject pool, using the Online Recruitment System for Economic Experiments (ORSEE) tool (Greiner, 2015[84]). A total of 194 subjects participated in the twelve separate experimental sessions described in Table 2.1. Subjects were between 18 and 68 years old with a median age of 20, and 16% of them indicated Economics as their major or intended

[9]This is consistent with Duffy and Fehr, 2018, who find that the frequency of playing the action that is associated with the Pareto-efficient outcome in the PD or the SH does not depend on the order in which these two games are played.[82]

major.

Experimental sessions lasted between 45 and 55 minutes. Subjects were paid their payoff from a randomly selected round of the total of 80 rounds, plus an additional $7.00 show-up fee. The average total payment was $13.22, while the minimum payment was $8.00 and the maximum was $20.00.

2.3 Descriptive Results

In this section, we describe the main results of the experiment. We use the pooled sample of data as detailed in Section 2.2 but note that our results are qualitatively similar when using alternative samples.

We first look at the impact of the information treatment on choosing action X in aggregate terms. Recall that action X is associated with the Pareto-efficient outcome in both games. Figure 2.3 shows how the average rate of choosing action X evolves in each game and information treatment. For the SH, there is a large difference in play with Partial versus Full across all rounds. For the PD, there is initially a large difference in the rate of playing action X in the earlier rounds that deteriorates by the final rounds.

To test the significance of these results, we estimate regressions of the following form:

$$Y_{i,r,s} = \alpha + \beta PartialInfo + \gamma C_s + u_{i,r,s}, \tag{2.1}$$

where $Y_{i,r,s}$ is a binary indicator for subject i choosing action X in round r of session s. The vector C_s is a set of dummy variables that flexibly controls for session size.[10] The variable $PartialInfo$ is an indicator and our main variable of interest. It equals one if the action choice is made under the Partial treatment and zero otherwise. Thus,

[10]We do not use session fixed-effects in our main specification since the estimate $\hat{\beta}$ would only exploit the within-subjects data. Results are qualitatively similar when using the within-subjects data.

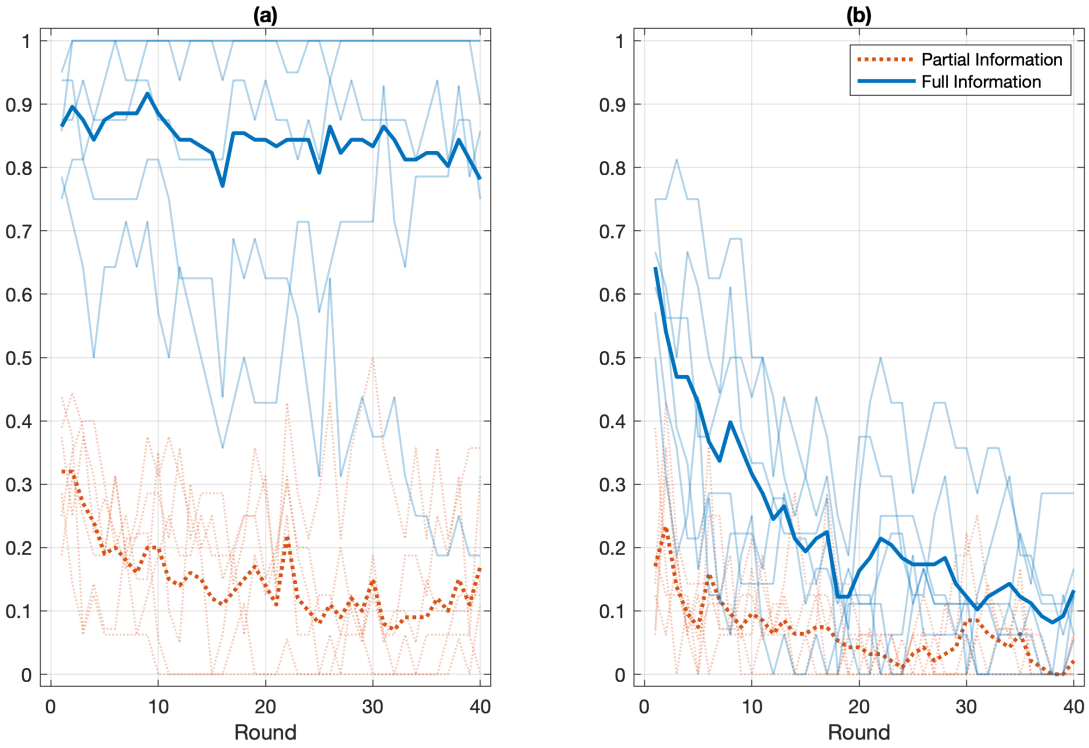

Figure 2.3: Mean Action Rates by Treatment

Note: Each line represents the mean rate of the action X grouped by treatment. Panel (a) shows rates for the SH, while panel (b) shows rates for the PD. Faded lines represent the mean rates for each session separately. Note that there are two pure strategy equilibria for the SH (corresponding to both players selecting action X with a probability of 0 or 1, respectively) and one mixed strategy equilibrium (corresponding to both players selecting action X with a probability of $\frac{2}{3}$). For the PD the sole equilibrium exists where both players select action X with a probability of 0.

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) Stag Hunt					
Treatment effect	-0.676	-0.669	-0.683	-0.677	-0.673
	(0.034)	(0.030)	(0.041)	(0.047)	(0.040)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Control mean	0.844	0.881	0.836	0.822	0.838
Number of clusters	144	144	144	144	144
N	7,840	1,960	1,960	1,960	1,960
b) Prisoner's Dilemma					
Treatment effect	-0.173	-0.304	-0.147	-0.156	-0.086
	(0.022)	(0.037)	(0.028)	(0.031)	(0.026)
Cluster p-value	0.000	0.000	0.000	0.000	0.001
Control mean	0.232	0.433	0.205	0.176	0.113
Number of clusters	142	142	142	142	142
N	7,680	1,920	1,920	1,920	1,920

Table 2.3: Treatment Effect for Selection of Action X

Note: The sample uses the pooled data. The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

the estimated coefficient $\hat{\beta}$ can be interpreted as the percentage point difference in the probability of choosing action X under the Partial treatment compared to the Full treatment. This regression equation is used to compare the session-size-adjusted outcome means across treatments, not to fully explain behavior of subjects, which will be investigated more carefully in Section 2.4. We employ an ordinary least squares regression to estimate equation 2.1 and cluster standard errors at the subject-session-level. For our main results, we present estimates of equation 2.1 in Table 2.3, with the results for the SH in panel a) and the PD in panel b).

Initial play. For Full, the fraction of subjects who initially choose action X is substantially larger than for Partial in both games. In column (2) of panel a) in Table 2.3, we can see the treatment effect in the first 10 rounds of play for the SH is a 66.9 percentage point (pp) reduction in the probability of playing action X. The control group selects action X about 88.1% of the time, so the effect in percent terms is an 80.1% reduction relative the control mean. In column (2) of panel b), we can see that for the PD there is a treatment effect of -30.4pp in the probability of selecting action X in the first 10 rounds of play. This result is about -70.2% relative to the control mean given the control group plays action X about 43.3% of the time.

The very first round of each game is special as subjects in the Partial treatment could not observe any previous actions of their opponents and thus have absolutely no information about their incentives. We find that the pattern emerging in the first 10 rounds is qualitatively similar to the very first round: in the SH-Full treatment, 86.5% of subjects chose action X in the first round, but only 32.0% chose that action in the SH-Partial treatment. For the PD, the corresponding values are 64.3% and 17.0%.

Result 1. *Mutual payoff information has a large effect on initial play in both games. Under full-information, the fraction of subjects choosing action X (the action supporting Pareto-superior outcomes) is substantially higher.*

Equilibrium selection and convergence. Next, we analyze how play evolves across the 40 rounds of a game. In the SH, the initial effect is remarkably persistent throughout the game. As can be seen from panel a) of Table 2.3, the vast majority of subjects choose action X in the SH-Full treatment and action Y in the SH-Partial treatment throughout all rounds of the experiment. In column (1) of panel a), the average treatment effect across all rounds is -67.6pp, which equates to a treatment effect of -80.1% relative to the control mean of 84.4%.

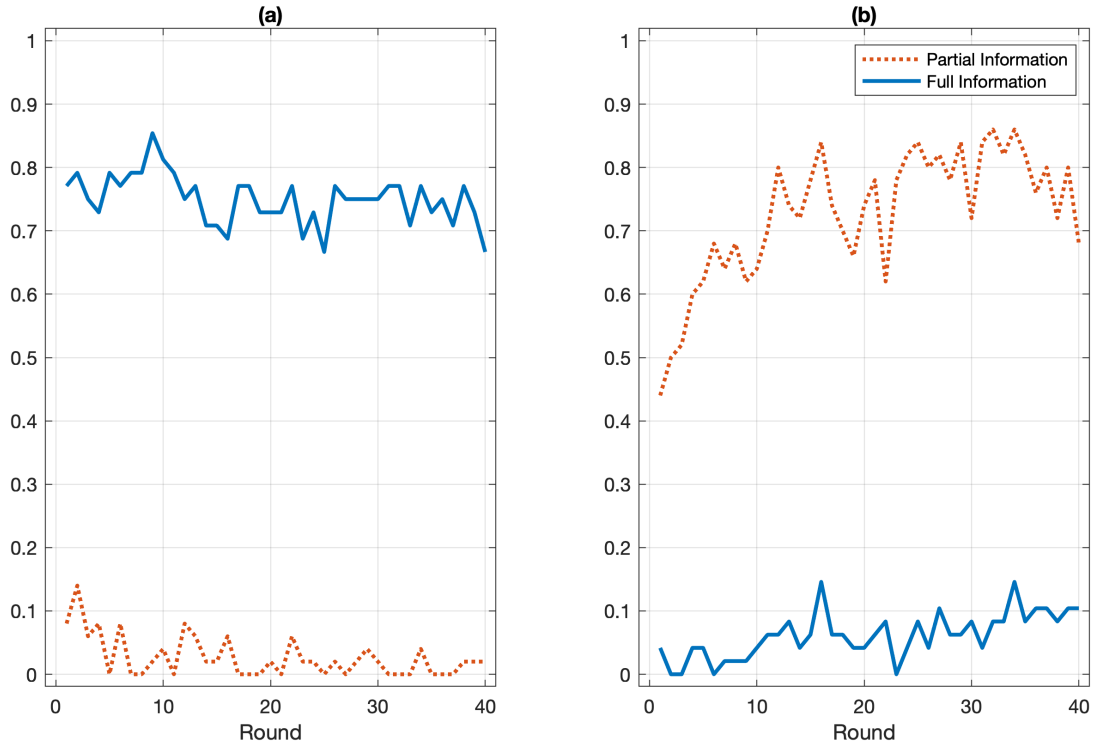

Figure 2.4: Equilibrium Rates for SH Sessions

Note: Panel (a) shows rates of the payoff-dominant equilibrium. Panel (b) shows rates of the risk-dominant equilibrium. Note that the sum of the solid blue lines adds up to less than one and the sum of the dotted red lines adds up to less than one. The remaining fraction is accounted for by the fact that a pure strategy Nash equilibrium is not always achieved in a given round.

These results directly impact equilibrium and efficiency in the respective treatments for the SH. While these results are partly driven by the random mechanism through which pairs of subjects were matched each round, we present them here to show the stark impacts mutual payoff information has for each game. For Partial, subjects tend to reach the risk-dominant equilibrium, while for Full, subjects tend to reach the payoff-dominant equilibrium, as can be seen in Figure 2.4.

We estimate equation 2.1 using a binary indicator for reaching a pure strategy Nash equilibrium, which for the SH is (X, X) or (Y, Y), and report the results in Table 2.4.[11]

[11] Our SH game also has a mixed strategy Nash equilibrium where subjects choose action X two-thirds of the time and choose action Y one third of the time. In Table E.1 of Appendix E we show the share of

It is clear from panel a) that there is only a small difference in the rate of reaching an equilibrium due to treatment. For the average across all 40 rounds of play in column (1), we can see that the control group reaches an equilibrium about 81.1% of the time and the treatment group only does so 10.0pp less often (cluster p-value of 0.001). However, since the payoff-dominant equilibrium is more efficient than the risk-dominant one, partial information leads to lower efficiency. We again estimate equation 2.1 but this time use an efficiency ratio as the outcome.[12] In panel a) of Table 2.3, the subjects under Partial experience lower efficiency compared to the subjects under Full. The average across all 40 rounds is a reduction of -0.396 of the ratio, or -45.8% relative to the control ratio of 0.864.

Result 2. *Throughout all rounds of play in the SH, for the Full treatment, the vast majority of subjects select action X, which corresponds to reaching the payoff-dominant Nash equilibrium, while in the Partial treatment, most subjects choose action Y, which corresponds to reaching the risk-dominant Nash equilibrium.*

In the PD, on the other hand, play converges toward the unique Nash equilibrium of the game under both information treatments. We define convergence as occurring at a given round of play at which at least 80% of subjects (averaged across all sessions) play the equilibrium action, Y, and at least 80% play action Y for the remaining rounds of the game. For the PD-Full, this occurs at round 24, and for the PD-Partial, this occurs at round 3. As evidenced by Figure 2.3, by the end of the game both treatments converge toward the deviating action Y. Panel b) of Table 2.3 quantifies how the choice of action X evolves across rounds in the PD. By the last 10 rounds, the treatment effect degrades

subjects in each of four 10-round periods whose mix of actions are within 10pp of $p_X = 0.667$. Subjects exhibiting such patterns of play are in the minority in all treatments and periods but are more common during earlier rounds and in the full information treatment.

[12]The efficiency ratio is the total payoffs of both subjects in a given round of play divided by the total payoffs of the efficient outcome. Naturally, the random re-matching of subjects will induce some variation here.

Table 2.4: Treatment Effect for Reaching an Equilibrium Outcome

	(1)	(2)	(3)	(4)	(5)
	Overall	1-10	11-20	21-30	31-40
a) Stag Hunt					
Treatment effect	-0.100	-0.174	-0.050	-0.076	-0.099
	(0.028)	(0.036)	(0.036)	(0.033)	(0.033)
Cluster p-value	0.001	0.000	0.167	0.023	0.004
Control mean	0.811	0.808	0.798	0.831	0.808
Number of clusters	144	144	144	144	144
N	7,840	1,960	1,960	1,960	1,960
b) Prisoner's Dilemma					
Treatment effect	0.266	0.405	0.241	0.258	0.158
	(0.019)	(0.032)	(0.030)	(0.033)	(0.027)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Control mean	0.618	0.347	0.643	0.696	0.788
Number of clusters	142	142	142	142	142
N	7,680	1,920	1,920	1,920	1,920

Note: The sample uses the pooled data. The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

Table 2.5: Treatment Effect for Efficiency Ratio

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) Stag Hunt					
Treatment effect	-0.396	-0.412	-0.386	-0.390	-0.394
	(0.018)	(0.015)	(0.021)	(0.026)	(0.021)
Cluster p-value	0.000	0.000	0.000	0.000	0.000
Control mean	0.864	0.883	0.856	0.857	0.859
Number of clusters	144	144	144	144	144
N	7,840	1,960	1,960	1,960	1,960
b) Prisoner's Dilemma					
Treatment effect	-0.078	-0.147	-0.063	-0.066	-0.034
	(0.011)	(0.020)	(0.014)	(0.015)	(0.015)
Cluster p-value	0.000	0.000	0.000	0.000	0.025
Control mean	0.554	0.650	0.539	0.527	0.498
Number of clusters	142	142	142	142	142
N	7,680	1,920	1,920	1,920	1,920

Note: The sample uses the pooled data. The regressions include controls for session size. Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. Clustering is at the session-subject level. Cluster p-value indicates the p-value from a two-sided t-test of the null hypothesis that the treatment effect is zero using the cluster-robust standard error.

to a difference of only -8.6pp between the Full and Partial treatments (cluster p-value of 0.001), and the control group also decline to only selecting action X 11.3% of the time. Though this difference is still statistically significant by the last 10 rounds, it diminishes greatly and almost monotonically across all rounds of play. As panel b) of Table 2.4 shows, by the last 10 rounds of play, the treatment group is only slightly more likely to reach an equilibrium relative to the control group, whereas the difference is much larger in the initial rounds. The efficiency implications of this convergence can be seen in panel b) of Table 2.5. In the first 10 rounds of play, the treatment group is much less efficient than the control group (-22.6%), but by the last 10 rounds, this treatment effect is greatly attenuated (-6.8%).

Result 3. *In the PD, play in both treatments converges toward the unique Nash equilibrium of the game.*

In results not reported, we vary the sample (within-subjects or between-subjects samples) and the controls used (no controls for session size, session-fixed effects), and find that Results (1)-(3) are qualitatively similar regardless of sample or specification.

2.4 Estimating a Learning Model and Simulations

While our descriptive results show that mutual payoff information has substantial effects on initial play in both games and the long-run outcome in SH, it is not clear if these effects are driven by initial play, learning, or both. For example, it could be that initial play is all that is impacted and any long-run differences are due solely to a history dependence. Alternatively, opponent payoff information may impact the dynamic learning process used by players, which has long-run implications for convergence and equilibrium selection.

In order to better understand the channels through which mutual payoff information

operates, we estimate a learning model and perform simulations. Our purpose here is not to test models to determine which more accurately matches the data, but rather to apply one that helps distinguish between the impacts of initial play and learning dynamics. As such, we apply a model of fictitious play, specified as a special case of the Experience-Weighted Attraction (EWA) learning model of Camerer and Ho (1999).[62][13] In introducing this model, the authors comment that they "consider the scientific problem of figuring out how people choose their initial strategies as being fundamentally different than explaining how they learn." The advantage of using the EWA model in our setting is that it allows us to investigate whether the effect of having mutual payoff information operates through initial play, ongoing learning, or both.

Model basics. Each round, players choose their actions (X or Y) based on the updated attractions of these two actions. Loosely speaking, attractions are players' expected payoffs conditional on their beliefs, explained below in more detail. The attractions for X and Y depend on past observations of other players' actions and a prior attraction that players bring into each particular game. We model subject behavior by a single representative agent, i.e., we assume that all subjects within a treatment take actions according to the same learning and decision-making mechanism, governed by the same parameter values. As different subjects may observe different histories of play, however, they may choose different actions as a result.

For each player i at the end of round t, action j has attraction $A_i^j(t)$, and these attractions determine play in round $t+1$ according to the following logistic choice probability

[13] We choose this particular model to economize on the number of parameters to be estimated. Belief-learning models perform favorably compared with alternatives; for example, see Nyarko and Schotter (2003).[85]

function:

$$P_i^j(t+1) = \frac{e^{\lambda \cdot A_i^j(t)}}{\sum_{k=1}^{2} e^{\lambda \cdot A_i^k(t)}}, \tag{2.2}$$

where $P_i^j(t+1)$ is the probability that action a_i^j is chosen by player i in round $t+1$, and λ is the response sensitivity parameter, explained below in more detail. Attraction $A_i^j(t)$ is given as

$$A_i^j(t;\phi) = \frac{\phi^t \cdot A_i^j(0) + \sum_{m=0}^{t-1} \phi^m \cdot \pi_i(a_i^j, a_{-i}(t-m))}{\sum_{n=0}^{t} \phi^n}, \tag{2.3}$$

and can be interpreted as subject i's expected payoff from action j after round t, given the subject's beliefs about the action chosen by other players, conditional on their own actions. That is, the underlying beliefs are defined over two states: the state that one's opponent plays X, conditional on oneself playing X, and the state that one's opponent plays X, conditional on oneself playing Y.[14] Each belief is a weighted average of the history of play that has been observed and a prior—i.e., $A_i^j(0)$ —the initial attraction the representative agent brings into the particular game (e.g., SH Full treatment, PD Partial treatment). The parameter ϕ is the weighting decay rate, explained below in more detail. The functions $a_i(t)$ and $a_{-i}(t)$ are the chosen actions in round t of player i and of the opponent that the player faced in that round. Finally, $\pi_i(a_i^j, a_{-i}(t))$ is player i's hypothetical payoff from choosing action j in round t, conditional on $a_{-i}(t)$, the actual actions of all other players in round t. Similarly, $\pi_i(a_i(t), a_{-i}(t))\pi_i(t)$ is the realized payoff for player i in round t.

[14] We do not require these probabilities to be the same. In other words, we allow for beliefs where the probability assigned to the other's action depends on a player's own action. We do this to allow for the possibility that subjects might perceive the occurrence of symmetric outcomes—i.e., (X, X) or (Y, Y)—as disproportionately likely.

Interpreting the parameters we estimate. For each treatment, we estimate a set of four parameters: two learning parameters (response sensitivity parameter λ and weighting decay rate ϕ) as well as the initial attractions ($A^X(0)$ and $A^Y(0)$).

The response-sensitivity parameter (also known as noise parameter) λ models players' sensitivity to differences between attractions (Camerer and Ho, 1999[62]). This parameter can be thought of as a measure of the noise with which players respond to updated attraction values in each round. λ can take on values from 0 to $+\infty$. At one extreme, if $\lambda = 0$, subjects select from their action set in a uniformly random manner. At the other extreme, if λ approaches ∞, subjects always best respond—i.e., they strictly choose the action with the largest attraction value in each round.

The weighting decay rate ϕ captures how much weight is put on the observations of previous rounds, relative to those of the most recent round. In particular, after having played t rounds, ϕ^r captures how much weight is put on the results of round $t-r$, relative to the results of round t. Parameter ϕ can take on values between 0 and $+\infty$.[15] While values of ϕ approaching 0 would indicate that subjects only take their last opponent's action into account when forming beliefs, $\phi = 1$ would indicate that all observed actions are weighted equally, and values of ϕ approaching ∞ would indicate that subjects do not engage in learning at all but instead base their beliefs solely based on their initial attractions—i.e., they do not update or "learn."

The initial attractions $A^X(0)$ and $A^Y(0)$ can be interpreted as subjects' expected payoffs from each respective action prior to beginning play in the first round, given their beliefs about the play of the other subject with whom they are matched, conditional on their own actions.[16] Note that attractions are defined in terms of own payoffs, which

[15] Camerer and Ho (1999) comment that values for this decay rate are "presumably between zero and one."[62] We do not limit the value of ϕ from exceeding 1, but we do note that our results are consistent with estimated values of ϕ that are somewhat less than 1 in all treatments.

[16] Note that for a given expected probability of one's counterpart playing X in the first round, $P^X(1)$, there is a one-to-one correspondence between the expected value of X and the expected value of Y. For

means that $A^j(0)$ is constrained to take a value between the lowest and highest possible payoffs that can be derived from choosing action j. This is done mainly to maintain consistency with attractions in later rounds and to ease interpretation of these attractions.

Initial play is captured by the initial attractions as well as λ. First and foremost, $A_i^X(0)$ and $A_i^Y(0)$ represent the expected payoff of each respective action, conditional on beliefs. In addition, λ can affect initial play as lower values of λ reduce the probability that the action with the higher attraction is chosen. Learning, on the other hand, is captured by the learning parameters ϕ and λ, as well as the continually updated, history-dependent attractions for each subject, $A_i^X(t)$ and $A_i^Y(t)$.

Estimation and simulation details. The estimation is performed numerically, using maximum likelihood techniques. We employ the bootstrap procedure for estimating parameter sampling distributions, with $B = 2,000$ bootstrap samples per estimation, from which we then conduct inference using the Bias Corrected-accelerated (BCa) confidence interval method of inference pioneered by Efron and Tibshirani (1993).[86][17]

Consider a treatment with N subjects and $T = 40$ rounds. Then, the likelihood of observing subject i's action history $\{a_i(1), a_i(2), ..., a_i(T-1), a_i(T)\}$, given $(A^x(0), A^y(0), \phi, \lambda)$ is

$$\Pi_{t=1}^{T} P_i^{a_i(t)}(t|A^x(0), A^y(0), \phi, \lambda), \tag{2.4}$$

example, if a player assesses in our SH game that $\mathbb{E}(X) = 8$ in the first round, this implies $P^X(1) = 0.7$ (note that $0.7 * 11 + (1 - 0.7) * 1 = 8$), which in turn implies that $\mathbb{E}(Y) = 0.7 * 9 + (1 - 0.7) * 5 = 7.8$. Thus, $\mathbb{E}(X) = 8 \Leftrightarrow \mathbb{E}(Y) = 7.8$. However, we allow both initial attractions to range independently of one another, between the minimum and maximum possible payoff for each respective action. As discussed above, this allows for "skewed beliefs," beliefs in which subjects assign too high (or too low) a probability on symmetric outcomes. For example, suppose subjects playing the SH Partial treatment believe that conditional on themselves playing X (Y), their opponent will play X with a probability of 0.7 (0.6) in the first round. Given this set of beliefs, $A^X(0)$ would be 8 ($= 0.7 * 11 + 0.3 * 1$), while $A^Y(0)$ would be 9.8 ($= 0.6 * 13 + 0.4 * 5$).

[17] We are unable to use more commonly used standard error techniques due to extreme levels of skew and kurtosis in the sample distribution of bootstrapped parameter estimates.

where $P_i^{a_i(t)}$ corresponds to the logistic probability function defined in equation (2.2). The joint likelihood function $\mathcal{L}(A^x(0), A^y(0), \phi, \lambda)$ of observing all subjects' action histories is given by

$$\mathcal{L}(A^x(0), A^y(0), \phi, \lambda) = \Pi_i^N \{\Pi_{t=1}^T P_i^{a_i(t)}(t|A^x(0), A^y(0), \phi, \lambda)\}.$$

To test whether our parameter estimates lead to predicted behavior that is consistent with actual observed behavior, we conduct simulations of $1,000$ sessions per treatment. Each one of these sessions has an even number of subjects between 14 and 20, chosen randomly.[18]

2.4.1 Results: Learning Model and Simulations

Parameter estimates. Table 2.6 reports estimated parameters of the learning model by game and information treatment. Estimates of response sensitivity parameter λ are higher for the SH Full treatment than for the SH Partial treatment, and lower for PD Full treatment than for the PD Partial treatment, indicating that mutual information increases the sensitivity of responses to attractions in the SH while decreasing it in the PD. As reported in Appendix E Table E.3, these differences are significant at a 95% confidence level.

Similarly, estimates of weighting decay rate ϕ are higher for the SH Full treatment than for the SH Partial treatment, and are lower for the PD Partial treatment than for the PD Full treatment. As reported in Table E.4, however, these differences are only marginally significant for the SH ($p = 0.062$) and are not significant for the PD ($p = 0.301$).[19]

[18]The simulation results are robust to changes in the number of subjects in each session.

[19]Because ϕ enters the learning model with an exponent equal to the number of rounds since an observation has been made, small differences in the value of ϕ can lead to significantly different modeled behavior. For example, our estimates of ϕ for the SH imply that an observation made 3 rounds previously in the SH Full treatment would have a weight of 90% relative to the current round's observation,

Estimates of the initial attraction for the coordinating action (in the SH) and cooperative action (in the PD) both increase under mutual information. These differences are significant for the SH at a 95% confidence level and at a 90% confidence level for the PD. It is worth pointing out that the estimates of $A^X(0)$ and $A^Y(0)$ in the Partial treatment are reasonably close to the expected payoffs associated with either action if one's prior belief that the other player chooses X versus Y is $50:50$.[20] Considering that subjects did not get to observe any actions before the first round, holding a prior of 0.5 seems very intuitive in the Partial treatment.

Simulations with flipped parameters, initial play versus learning. Next, we shed light on whether the treatment effect operates primarily through initial play or learning. To do so, we conduct simulations wherein we swap the value of each parameter with the corresponding parameter value of the other treatment in the same game. For example, to isolate the role that the initial attractions play for explaining behavior in the SH Full treatment, we perform simulations using the estimated parameters of the SH Full treatment, except we use the initial attractions parameters from SH Partial treatment. Note that while the parameter estimates would already allow us to hypothesize in which direction average behavior may change, conducting this simulation exercise with flipped parameters has the advantage of providing insights on how economically meaningful these changes are.

while a similarly aged observation in the SH partial treatment would have only 63% relative weighting. After 6 rounds, the weights drop to 81% and 40%, respectively. This suggests that learning under the SH Partial treatment is much more heavily weighted to recent observations than it is in the SH Full treatment, implying a learning process that is much more sensitive to potential changes in behavior of one's counterparts.

[20]For example, if a subject in the SH believes that her opponent is equally likely to choose action X and action Y, she expects to earn 6 $(= 0.5 * 11 + 0.5 * 1)$ by choosing X, and 7 $(= 0.5 * 9 + 0.5 * 5)$ by choosing Y. Similarly, a subject in the PD with the same prior belief would expect to earn 6 by choosing X and 9 by choosing Y.

Table 2.6: Learning Model Parameter Estimates and 95% Confidence Intervals

Parameter	SH−Full	SH−Partial	PD−Full	PD−Partial
λ	1.5831	0.6516	0.4297	0.7426
	(1.2159 − 1.8776)	(0.5184 − 0.7673)	(0.3347 − 0.5242)	(0.6615 − 0.9146)
ϕ	0.9649	0.8290	0.8011	0.8769
	(0.9011 − 1.0223)	(0.7112 − 0.9911)	(0.3009 − 0.9398)	(0.6019 − 85.013)
$A^X(0)$	8.5765	5.9218	9.4480	7.1258
	(7.8948 − 9.7541)	(5.1815 − 6.4317)	(8.4351 − 10.5367)	(6.9829 − 8.5193)
$A^Y(0)$	6.4400	7.0199	6.4995	8.5930
	(5.8775 − 6.6354)	(6.8786 − 7.2157)	(5.5125 − 7.3160)	(8.5141 − 12.9087)
$L(\lambda)$	−1227.06	−1542.83	−1934.22	−882.90
n	3840	4000	3920	3760
$P^X(1)$	0.9672	0.3284	0.7802	0.2517
(learning model)	(0.8002 − 0.9999)	(0.1636 − 0.5069)	(0.1991 − 0.9334)	(0.0614 − 0.3999)
$p_X(1)$	0.8646	0.3200	0.6429	0.1702
(binomial model)	(0.8038 − 0.9358)	(0.2537 − 0.4366)	(0.5641 − 0.7486)	(0.1231 − 0.2825)

Notes: Results of tests of significance of the information treatment effect on parameters estimates for λ and ϕ in SH and PD are reported in Appendix E Tables E.3 and E.4, respectively.
Differences in values for λ are significant for both SH and PD ($p < 0.05$), while difference in values of ϕ are not significant for PD, and are only weakly significant for SH.
For estimates of $p_X(1)$ we employ Agresti-Coull binomial confidence intervals; see [87] and [88] for mathematical definition and motivation for their use over the more commonly used Wald confidence interval approach.

Figure 2.5 shows model simulations with our estimated parameters (in solid lines) alongside simulations with flipped parameters (in dotted lines). This is done separately for the SH (left panels) and the PD (right panels). In panel (a) and (b), estimates for the initial attractions are flipped at the information treatment level. Likewise, in panel (c) and (d), estimates of λ are flipped. Finally, in panel (e) and (f), estimates of ϕ are flipped.

For the SH, switching the initial attraction estimates results in quite dramatic differences of the simulated data. For the PD, differences are initially notable but vanish by about round 15. In both games, flipping the initial attractions reduces the simulated fraction of subjects playing X in the Full treatment but increases that fraction in the Partial treatment. This echoes the estimates of $A^X(0)$ in Table 2.6, where the SH Full and PD Full treatments have bigger values of $A^X(0)$ than their Partial counterparts. When $A^X(0)$ ($A^Y(0)$) is higher, the simulated fraction playing X goes up (down), which is an intuitive result.

Next, we look at whether our estimated differences for λ have meaningful consequences on behavior. When simulating behavior in the SH Full treatment with the λ of the SH Partial treatment, λ is now smaller than in the original simulations, moving behavior towards the *smaller* attraction, which is $A^Y(t)$ for the SH Full treatment. Likewise, if we simulate behavior in the SH Partial treatment of the λ of the SH Full treatment, we face a higher λ, leading behavior to move towards the *bigger* attraction, which is $A^Y(t)$ for the SH Partial treatment. Consequently, as panel (c) of Figure 2.5 shows, flipping the estimated λ parameters results in a lower fraction of playing X in both treatments of the SH. The same logic applies to the PD. Note that for the PD Full treatment, the attraction is initially higher for X, but starting in about round 5, the attraction is higher for Y. This is why in panel (d) of Figure 2.5, the simulation of the PD Full treatment with a flipped ϕ is initially lying above the original simulation but

Figure 2.5: Mean Action X Rates in Simulated Data of Six Sessions per Treatment, versus Same Rates in Simulated Data Using Parameters from Other Information Treatment of Same Game

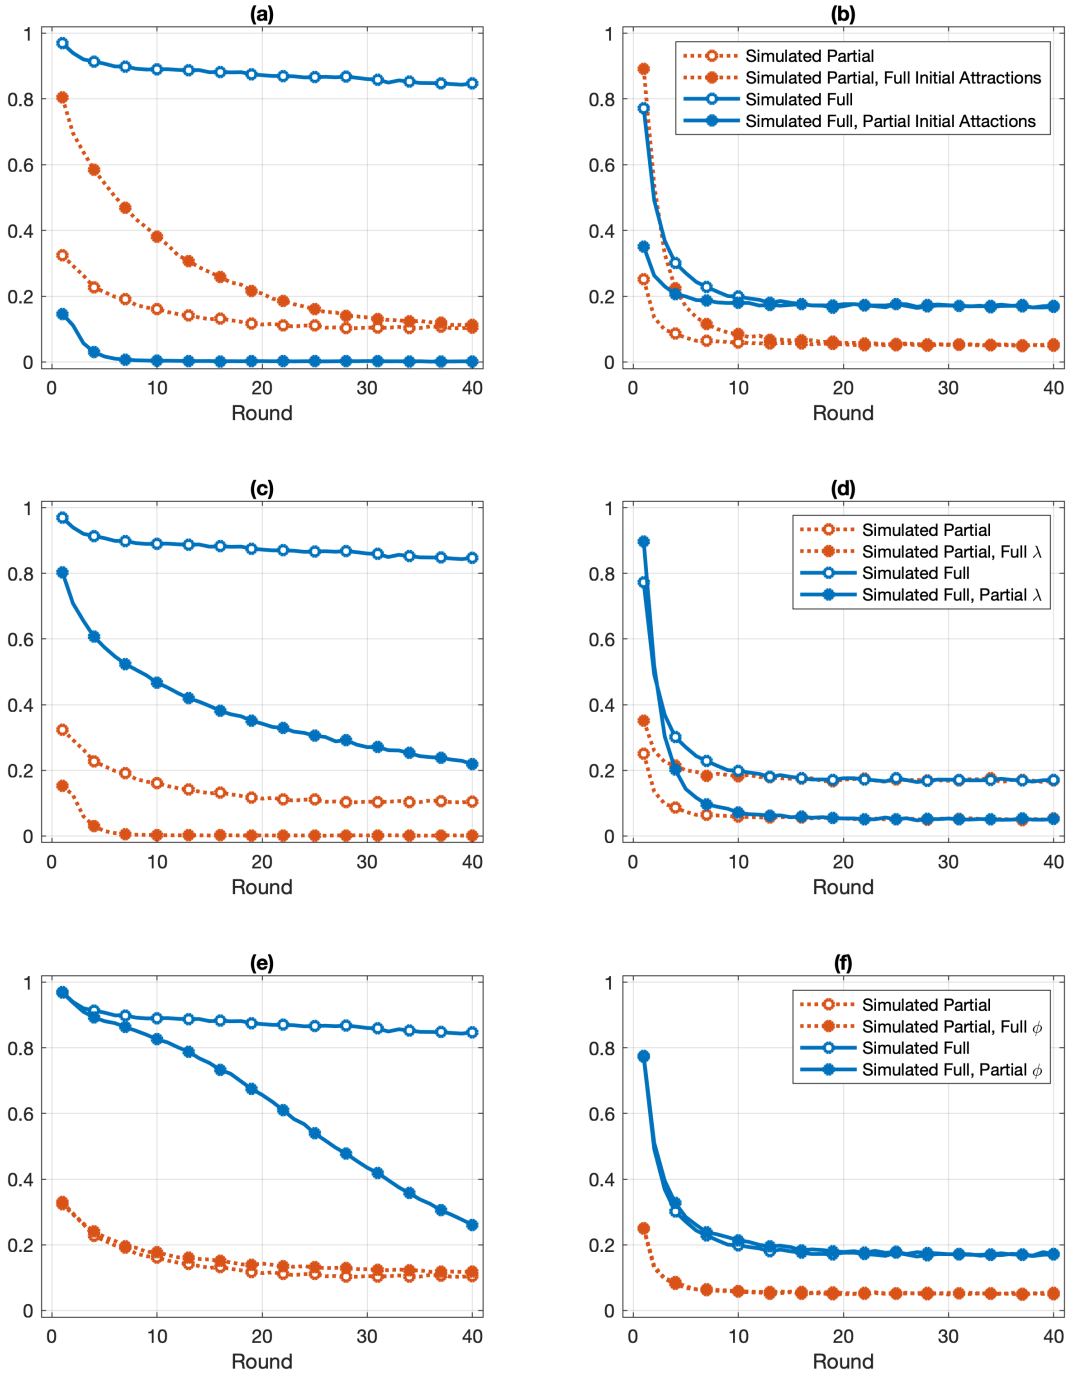

Note: Left hand Panels (a, c, e) show these rates for the SH, and right hand Panels (b, d, f) show them for the PD

then falls below it.

When flipping ϕ, note that our estimates of ϕ are not statistically distinguishable for the PD and are only marginally statistically distinguishable for the SH (see Table 2.6). While we find no effect for the SH Partial, PD Full, and PD Partial treatments, there is a substantial drop of the fraction playing X over time in the SH Full treatment when we flip the ϕ parameters. This suggests that coordinating on the payoff-dominant Nash equilibrium is highly sensitive even to small changes of the model parameters. Note that the lower ϕ is, the more rapidly the weight placed on observations of earlier rounds is depreciated. Therefore, low values of ϕ make play more sensitive to volatility in behavior, as players are placing higher weights on a smaller set of (recent) observations when forming beliefs.

Taken together, the simulations where we flip parameters at the treatment level indicate that each of the four parameters we estimate (with the exception of ϕ in the PD) play an important role in explaining the differences in behavior across treatments. In fact, for our SH game the payoff dominant equilibrium is only maintained through a combination of the estimated parameters for initial attractions, belief updating, and response sensitivity. However, in the PD the effect is less persistent. While subjects' initial perceptions of the game and their response sensitivity change under mutual information, the way they process observations does not appear to change. That is, the initial perception of the game retains similar importance for future play under both information treatments in the PD.

Result 4. *Simulations suggest that the information treatment effect can be attributed to differences in both initial play and learning in both games.*

Model validation. To test whether the model performs reasonably well in fitting the observed data, we conduct simulations of 1,000 sessions per treatment using the param-

eter values we estimate. Each simulated session has an even number of subjects between 14 and 20, chosen randomly with equal probabilities. Figure 2.6 compares the observed average behavior in our experiment with simulated data that we generate by using the estimated learning model parameters. For each game and information treatment, solid lines depict the average fraction of subjects that chose action X in a given round, and dashed lines depict the corresponding simulated data averaged across $1,000$ simulations. As Figure 2.6 shows, the simulated data is qualitatively very similar to the actual observed average behavior of subjects. Appendix Table E.2 provides regression analyses showing further support that the simulated data is on average statistically indistinguishable from the observed behavior.

2.5 Discussion

As discussed in Section 2.4, the information treatment effect operates through both initial play and learning. In this section, we discuss potential channels behind this treatment effect.

Recall that the appeal of contrasting the PD and the SH in our experiment is based on the idea that the structural uncertainty which is implied by the absence of mutual information affects strategic uncertainty differently in these games: Observing the other's payoffs arguably reduces strategic uncertainty in the SH, as it facilitates coordinating on the payoff-dominant Nash equilibrium. In the PD, however, mutual payoff information makes the cooperative action more attractive for players with social preferences. At the same time, players face uncertainty about the social preferences of others, and consequently mutual payoff information may increase strategic uncertainty.

We propose that differences in strategic uncertainty are captured by differences in the values of the learning model parameters. Conditional on the attractions, more prior

Figure 2.6: Mean Action X Rates in Pooled Data of Six Sessions per Treatment, versus Same Rates for 1000 Simulated Sessions

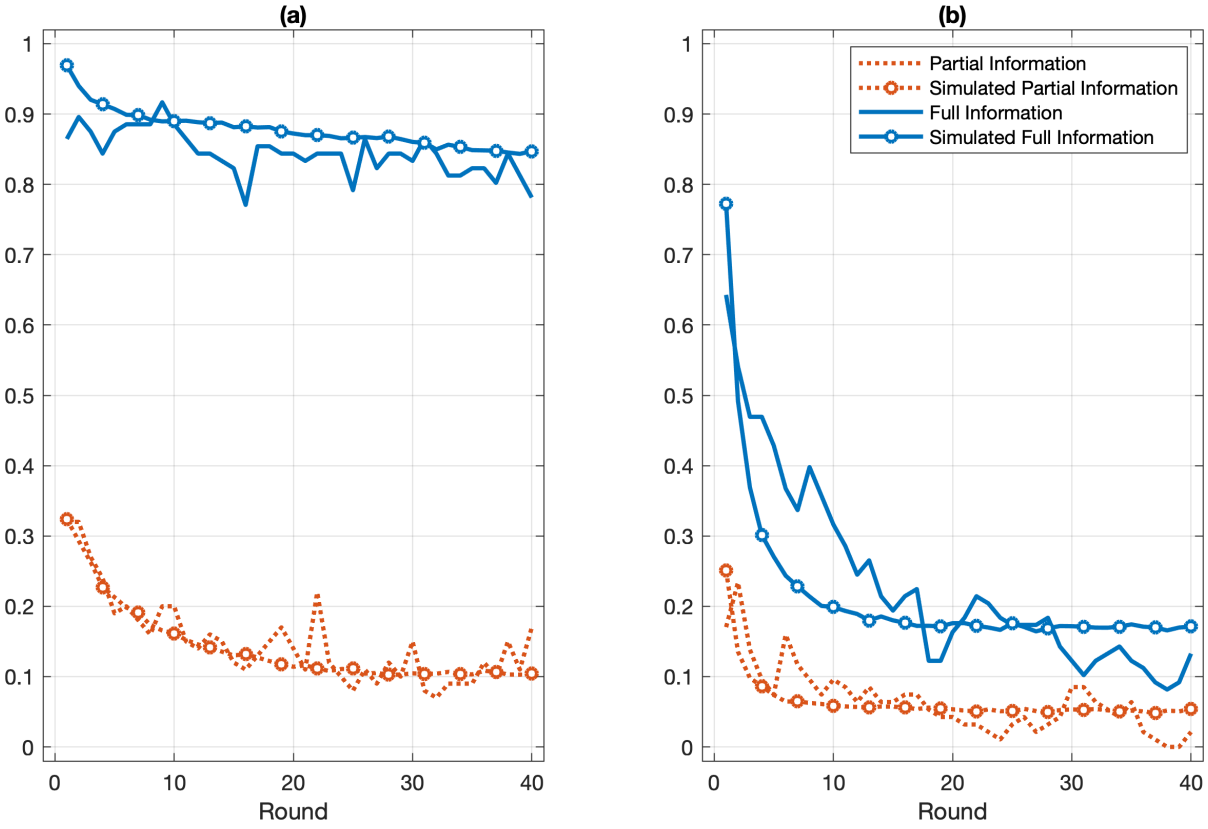

Note: Panel (a) shows these rates for the SH, and Panel (b) shows them for the PD.

uncertainty makes experimentation more beneficial, reflected by lower values of λ. Low values of λ are therefore consistent with more strategic uncertainty. While we cannot identify that low values of λ indeed imply more strategic uncertainty, one way to think of low values of λ would be that this is suggestive evidence that the game is perceived to be more strategically uncertain, all else equal. Similarly, we argue that higher levels of strategic uncertainty are consistent with lower values of ϕ: If subjects place more weight on results from the most recent rounds (relative to that of earlier rounds), their behavior is more sensitive to noise and to potential emerging trends in the patterns of play of their counterparts.

Indeed, our estimates of λ imply that there is more noise in play in the SH Partial treatment than in the SH Full treatment and in the PD Full treatment than in the PD Partial treatment, which is consistent with the idea that mutual payoff information increases strategic uncertainty in the PD while decreasing it in the SH. Furthermore, our estimates of the weighting depreciating parameter ϕ indicate that subjects place more weight on counterpart actions from recent rounds, and thus tend to update their attractions more rapidly in the SH Partial treatment than in the SH Full treatment and in the PD Full treatment than in the PD Partial treatment, which is again consistent with the idea that the strategic uncertainty fueled by mutual payoff information goes in different directions in these two games.

Aside from our learning model, what other indicators of strategic uncertainty could one use to test the idea that mutual payoff information affects strategic uncertainty differently in the SH and the PD? While the literature addressing strategic uncertainty describes the concept in largely qualitative terms, Calford and Oprea (2017) utilize the size of the Basin of Attraction (BOA) as a simple quantitative measure of strategic risk[89]. Applied in our context, the BOA index of a subject's least cooperative strategy (Y) is the probability they must assign to their counterpart playing the most cooperative

of their strategies (X) in order to be indifferent between playing one's own most and least cooperative strategies. The BOA index depends only on one's own payoffs and thus cannot explain the information treatment effects we examine. Moreover, this measure cannot provide useful insights into the strategic uncertainty that may exist within a given one-shot PD game: By convention, the BOA index is set to a value of 1.0, regardless of the degree to which one's "temptation" payoff exceeds their "reward" payoff or to which their "sucker" payoff falls below their "punishment" payoff, etc. The BOA index's invariant values for a one-shot PD do not reflect the varying levels of cooperation that are generally observed in this game.[21]

What other features of a game change as mutual payoff information is removed? Observing other players' payoffs may potentially also affect signaling and reputation building (as noted by Feltovich and Oda, 2014), focal points, and/or the expression of social preferences.[66] Given that our experimental design employs random re-matching, we do not think that signaling or reputation concerns notably contribute to the information treatment effect we document. As for focal points, Crawford et al. (2008) suggest that four features of a game can potentially generate them: labels, pre-play communication, precedents based on shared histories, and the payoff structure of a game.[91] Labels are kept constant across treatments in our design, and there is no pre-play communication in our experiment. While precedents of shared histories could be an issue if play in the first game of a session affects the second game, we do not find that the order in which games are presented to subjects in the experiment affects the behavior we observe. Finally, the payoff structure is a strategic element and as such not separate from the channel of strategic uncertainty. One could argue that non-strategic use of payoff structure, such as awareness that the game is symmetric, might affect play. However, this cannot predict

[21] See for example Charness et al. (2016), who show that increasing the reward payoff monotonically increases rates of cooperation.[90]

the direction of the treatment effect we observe.

While our experiment is not designed to study the effects of counterpart payoff information on social preferences, it is worth noting that our results for initial round play are nevertheless consistent with models that incorporate social preferences.[22] We observe higher rates of initial coordination on socially optimal action pair (X, X) in the Full treatments of both our SH and PD games.[23] This is consistent with a social preference model in which agents behave as follows: If the socially optimal outcome can be identified—e.g., through mutual payoff information—they select the action associated with this outcome. In absence of such information, they select the level-1 action—i.e., the expected payoff-maximizing action assuming a counterpart who uniformly randomizes from their action set. A robust exploration of the impacts of counterpart payoff information on the expression of social preferences is beyond the scope of this paper and is left to further research.

2.6 Conclusion

We conduct an experiment in which subjects play repeated versions of SH and PD games with random re-matching. In the full-information treatment subjects observe the entire payoff matrix, while in the partial-information treatment subjects observe only their own payoffs. In both treatments, subjects observe the chosen actions after each round.

We find that mutual payoff information has a strong effect on initial play in both

[22] We also compare the first-round play observed in our experiment with several other contemporary models of one-shot and initial game play (e.g. QRE, level-k, Cognitive Hierarchy, etc.). Of these, only models of social preference are able to rationalize results for both our SH and PD games.

[23] Notably, in our PD Full treatment we observe initial rates of coordination on socially-optimal action profile (X, X) of 64.3%, with a 95% confidence interval that excludes 50% (see the final row of Table 2.6). If we apply the QRE model to analyze the forces shaping initial moves, this fact implies that errors or randomizing alone cannot rationalize our high rates of initial cooperation in the PD Full treatment: The attraction of X must be higher than that of Y for the rate of selection of X to exceed 50%.

games. In the SH, the vast majority of subjects selects the action consistent with the payoff-dominant Nash equilibrium under full-information but the action consistent with the risk-dominant Nash equilibrium under partial-information. This effect persists through all 40 rounds of the game. In the PD, on the other hand, we initially observe a pronounced difference—subjects in the full-information treatment are much more likely to cooperate—but play converges toward the unique Nash equilibrium in later rounds of the game.

We estimate a belief-learning model to study our information treatment effect on both initial play and on the learning process. These results and related simulations suggest that mutual payoff information alters not only the way subjects initially perceive the game but also the way they update and respond to their beliefs. The values of our estimated learning parameters are important in explaining the observed behavior, suggesting that learning by observing is an important feature of behavior in the long run regardless of whether or not mutual payoff information is available.

We propose that these effects are evidence of the effect that mutual payoff information has on strategic uncertainty. In the SH strategic uncertainty decreases as players can reason about the mutually beneficial, Pareto-efficient outcome. However, in the PD mutual payoff information reveals the tension inherent in the game and thus increases strategic uncertainty. In both cases, social preferences may also be at play when opponent payoff information is revealed.

Chapter 3

Disentangling the Effects of Risk Aversion and Risk Perceptions

3.1 Introduction

Experimental and survey research spanning the last two decades concludes that people who are more risk-tolerant are more likely to engage in risky health activities such as smoking and heavy alcohol consumption, and are more likely to be obese (Anderson and Mellor, 2008[92]; Barsky, Kimball, Juster, and Shapiro, 1995[93]; Dave and Saffer, 2007[94]; Dohmen et al., 2011[95]). Separately, and independent of one's level of risk tolerance, subjective perceptions of the risk associated with different activities have also been found to be associated with health behaviors (Khwaja, Sloan, and Chung, 2007[96]; Lundborg and Lindgren, 2002[97], 2004[98]; Viscusi, 1990[99], 1991[100]; Viscusi et al., 2000[101]). While there are numerous studies that link risk perceptions with risky behavior, it is notable that none of these control for risk aversion. Similarly, studies that control for risk aversion fail to control for risk misperceptions.

Numerous studies have documented that health-risk misperceptions are common. For example, Viscusi (1990[99], 1991[100]) finds that smokers tend to overestimate the lung cancer and diminished life-expectancy risks from smoking. Riddel and Shaw (2006)[102] find that survey subjects believe that the public health risks related to transport of nu-

clear waste are many times what is proffered by nuclear scientists and engineers. Riddel and Hales (2018) find that survey subjects tend to be overly optimistic about the efficacy of cancer-prevention behaviors and also optimistic about the benefits of avoiding behaviors, such as smoking, that can increase cancer risk[103]. Given the apparent prevalence of risk misperceptions in the health domain and the dearth of studies that simultaneously control for optimism and risk aversion, it is clear that a study that merges these two strands of the research is needed.

One of the challenges of estimating subjective optimism is having a solid basis for objective risks associated with different risk-increasing or risk-reducing behaviors. One of the strengths of the current study is that our measures of cancer-risk optimism rest on comparing subjective assessments of the relative risk (RR) of different health-related behaviors to objective estimates contained in the Harvard Cancer Risk Index (HCRI) (Colditz et al., 2000[104]). The HCRI was established by a working group of "epidemiologists, clinical oncologists, and other Harvard faculty with quantitative expertise focused on cancer and risk assessment" at the Harvard Center for Cancer Prevention (Colditz et al., 2000). The HCRI establishes generally agreed on quantitative RR factors for demographic and behavioral characteristics that bear on the incidence of a given cancer.

The purpose of this article is to examine the relationship between risk misperceptions and choices to engage in behaviors that either increase or decrease the risk of cancer, controlling for risk aversion and other factors such as cognitive ability and demographic variables that may affect health choices. The data are collected using a national survey of 474 men and women, aged 18 and over, in the United States. We focus on four cancers: bladder and colon cancer in men and women as well as prostate cancer in men and breast cancer in women. We elicit a coefficient of life-duration risk aversion from each subject using a sequential price list auction (Holt and Laury, 2002[105]). On the risk-increasing side, we use probit models to explore how subjective optimism concerning the

risk of contracting one of these cancers correlates with the decision to engage in activities that increase cancer risk such as smoking, eating excessive amounts of red meat, heavy alcohol consumption, and being medically obese. On the cancer prevention side, we model engaging in regular physical activity, taking a daily aspirin, and taking a daily multivitamin.

Following Spinnewijn (2013), we construct our models based on two sources of optimism.[106] Baseline optimists believe their cancer risk is lower than an expert would assess their risk to be, given the subjects' current and past behavior and health history. Control optimists believe that efforts to reduce cancer risk, including engaging in prevention measures and avoiding risky activities, are more effective than the current state of scientific knowledge would indicate.[1] We measure baseline and control optimism by comparing subjective assessments of cancer risk and efficacy of prevention with objective estimates of these same risks taken from the HCRI.

We find that optimism – in particular baseline optimism – is highly correlated with people's decisions both to engage in cancer-prevention behavior, and to avoid activities that increase cancer risk.[2] In general, baseline optimists expend less effort to reduce their cancer risk than do their more pessimistic counterparts. Such persons are more likely to smoke, to eat excessive amounts of red meat, and to be medically obese. They are also less likely to engage in prevention behavior such as regular physical exercise and taking

[1]It is important to note that our definition of control optimism differs from that used in the psychology literature. Control optimism in the current context arises from two sources. First, a subject is control optimistic if he or she believes that prevention efforts are more effective in reducing cancer risk than what is reported in the HCRI. Second, a subject is control optimistic if he or she believes that avoiding risky behaviors such as smoking is more effective in reducing risk than what is stated in the HCRI. So, for example, if a subject overestimates the risk of smoking he or she is classified as a control optimistic with respect to smoking.

[2]The probit models we estimate determine correlation between optimism, risk aversion, and behavior, controlling for cognitive ability and demographic variables. There could well be unobservable variables that jointly influence both optimism and behavior, thus we need to take care not to falsely infer causation when only correlation is known. Of course, this is also true for studies such as Barsky et al. (1997)[93], Anderson and Mellor (2008)[92], Dohmen et al. (2011)[95], and Viscusi (1991)[100].

a daily aspirin or multivitamin.

The effects of control optimism on health behaviors appear to have less of a clear pattern. Control optimists engage in a generally higher level of prevention behavior and are less likely to smoke and drink excessively, but more likely to be obese. Control optimists are no more likely than their pessimist counterparts to engage in the specific prevention behaviors studied.

The model results also indicate that risk aversion is far less important in explaining risky behavior than are baseline optimism and control optimism. We show that while models that exclude optimism and cognitive ability frequently indicate a significant correlation between risky behavior and risk aversion, the relationship disappears when these other potentially important variables are controlled for. This result calls into question previous research linking risk aversion with risky health behaviors. We conclude that empirical models that fail to control for subjective beliefs about disease risk may erroneously conclude that risk aversion correlates with health choices, when instead it is optimism that is actually behind the result.

3.2 Literature Review

Two strands of literature have developed that seek to explain why people engage in risky health or prevention behaviors. Early work argued that people were frequently misinformed about health risks, so that potentially flawed risk perceptions, rather than objective assessments of risk, influenced behavior. The later strand of the literature correlates risk aversion with the choice to engage in risky and health behaviors. We discuss each strand of the literature in turn.

3.2.1 Risk Perceptions, Cognitive Ability, and Health Behaviors

Numerous studies explore the relationship between subjective risk perceptions and the decision to smoke. Viscusi (1990) finds that smokers and nonsmokers alike tend to overestimate smoking risks[99]. Moreover, the higher the magnitude of overestimation (i.e., the more control optimistic the subject is about lung cancer risk), the less likely the subject is to smoke. Using the same data set, Viscusi (1991) concludes that younger cohorts believe smoking risks are higher and are therefore less likely to smoke[100]. Viscusi et al. (2000) come to a similar conclusion about Spanish smokers. Lundborg and Lindgren (2004) use cross-sectional survey data of Swedish adolescents to analyze risk perceptions and smoking behavior[98]. They find that subjects overestimate the risk of smoking and that the probability a subject smokes is decreasing in their perceived lung cancer risk. Khwaja et al. (2007) use data from the Health and Retirement Study to compare subjective beliefs about expected longevity to objectively assessed probabilities[96]. They conclude that, on average, people are fairly accurate about their beliefs, but smokers tend to be relatively optimistic about their likelihood of living another 10 years.

Lundborg and Lindgren (2002) use a cross-sectional survey of young Swedes to study the link between risk perceptions of alcohol use and drinking behavior[97]. They find that the higher the perceived risk of alcohol consumption, the less likely the subject is to consume alcohol.

Carman and Kooreman (2014) study the link between subjective risk perceptions and engaging in prevention behaviors such as flu shots, mammograms, and taking aspirin to guard against heart disease[107]. They find that subjects do not have very accurate risk beliefs, but beliefs have a significant correlation with the choice to engage in prevention behaviors. Winter and Wuppermann (2014) study subjective perceptions of obesity-related health risks for a sample of obese adults aged 50–62 years in the United States,

based on data from the Health and Retirement Study[108]. They find that subjects tend to overestimate the risk of heart attack, stroke, and chronic lung disease but underestimate other health risks such as arthritis and high blood pressure. They conclude that subjects tend to underestimate the risk of small-probability diseases and overestimate risks from higher-probability diseases.

Two recent studies examine the relationship between risky behavior, cognitive ability, and optimism. Bijwaard, van Kippersluis, and Veenman (2015) investigate the influence of cognitive ability and education on survival of a Dutch cohort born between 1937 and 1941 with observed mortality between ages 55 and 75[109]. They find that survival is increasing in both variables, with much of the survival differences across education cohorts explained by cognitive ability. Riddel and Hales (2018) use data from a survey of adults in the United States to elicit individual perceptions of cancer risk and demand for a hypothetical cancer insurance policy[103]. They find that demand for cancer insurance does not vary with risk aversion and cognitive ability, but is decreasing in baseline optimism. Moreover, they show that risk-reducing effort is decreasing in baseline optimism, increasing in control optimism, but unrelated to risk aversion or cognitive ability.

3.2.2 Risk Aversion and Health Behaviors

Other studies look more broadly at perceptions and health behaviors. Weber, Blais, and Betz (2002) study the effects of both risk perceptions and risk attitudes on behavior[110]. They conclude that risk preference measures must carefully control for perceptions of both probabilities and outcomes. If these are neglected, then correlations between behavior and risk perceptions will likely be biased, as they depend not only on intrinsic risk preferences, but also on risk perception.

Several studies correlate risk aversion with a greater propensity to engage in risky health behaviors. An early study by Barsky et al. (1997) uses data from the Health and

Retirement Survey Study to calculate subject-specific coefficients of financial risk aversion based on choices over gambles of lifetime income[93]. They measure risk tolerance using responses to a series of questions from the Panel Study of Income Dynamics of the form: "Would you take a job with 50-50 chance it will double your income and a 50-50 chance it will cut you income by x%?" where x = 50, 33, and 20. They find that subjects who are more risk tolerant are more likely to smoke, drink alcohol, work in risky jobs, and have no health insurance. Dave and Saffer (2007) study the effect of risk tolerance on the demand for alcohol using the same risk preference elicitation approach[94]. They find that demand for alcohol is increasing in risk tolerance.

Anderson and Mellor (2008)[92] calculate subject-specific coefficients of financial risk aversion using Holt and Laury's (2002)[105] price list approach. Controlling for demographic characteristics, they find that the risk-averse subjects are less likely to engage in unhealthy behaviors such as smoking, heavy alcohol consumption, and being obese and more likely to wear seatbelts.

Dohmen et al. (2011) analyze subjects' willingness to take risks in general, as well as over five domains, including driving automobiles, financial matters, sports and leisure, career, and health, based on a large representative survey (n = 22,000) of the German population[95]. They measure risk tolerance using two approaches. The first approach asks the subjects to grade their willingness to take risks in each of the domains and then generally on a 10-point scale. The second risk preference measure is determined by choices over a real financial stakes lottery. They find that self-reported measures of willingness to take risks are highly correlated with risk aversion, derived from the real financial stakes lottery. They further show that smokers are likely to report a higher willingness to take health risks, and a higher willingness to take risks in general, than are non-smokers.

At least one study fails to find a significant association between risk tolerance and

risky behavior. Picone, Frank, and Taylor (2004) find that the probability of undergoing preventative medical tests such as breast self-exam, mammography, and pap smears are increasing in risk tolerance, but the associations are only marginally statistically significant[111].

In what follows, we discuss a survey designed to elicit risk aversion, optimism, and current and past health behaviors, as well as demographic characteristics of subjects. We then model the choice to engage in risky and prevention behaviors as a function of these variables.

3.3 The Survey

We conducted an online survey of 474 men and women aged 18 and over living in the United States using Amazon Mechanical Turk (AMT). AMT is a web-based tool for matching workers with employers who need small tasks performed that require human intelligence. Tasks range from completing marketing and academic surveys to interpreting photos for anomalies. Employers post descriptions of the tasks on the web platform along with compensation amounts. Employers are only required to pay for successfully completed tasks. However, employers who refuse to pay when tasks are adequately completed may well lose their ability to recruit subjects.

The ease of use and streamlined process for contacting subjects has made AMT increasingly popular over the past seven years with social science and business researchers. Buhrmester, Kwang, and Gosling (2011) conclude that AMT is a cost-effective source for high-quality data[112]. They analyze the demographic variation in AMT survey subjects, determining that participants tend to be more diverse than a typical Internet sample or a sample based on university students. They also found that the data quality is comparable to conventional Internet or telephone surveys based on random sampling.

We described the task to potential respondents as a survey that aimed to elicit behavior and beliefs related to cancer risks. Subjects were given between $2.50 and $4.50 to complete the survey. The questionnaire has six components, described below.

3.3.1 Risk Perception

In this section, we elicit subjective beliefs about the likelihood of contracting three cancers: breast, bladder, and colon for women and prostate, bladder, and colon for men. Subjects were asked to gauge their risk for each of the three cancers relative to someone their age and gender. Allowed responses were assigned both a numerical and textual description. There were eight categories of responses that ranged from 0 (there is no risk of me getting this cancer) to very much above average (risk is five times or more than the average) with six other categories in between.

Following that, subjects were asked to rate the increase (decrease) in their risk arising from engaging in risky (preventative) activities. The risk factors were taken directly from Colditz et al. (2000) HCRI and varied with the cancer[104]. For example, risk factors for bladder cancer included smoking and exposure to chemicals at work, whereas preventative activities related to colon cancer included taking a daily aspirin and regular physical activity. The RRs for each characteristic or behavior were presented as both ranges of RR and a qualitative descriptor as follows: no risk increase, small risk increase, risk is higher but less than double the average risk, moderate risk increase where the risk is two to four times the average risk, large risk increase of four to eight times the average risk, and very large risk increase where risk is more than eight times the average risk.

3.3.2 Risk Preference

We used a sequential price list auction to elicit a risk-aversion coefficient based on the constant relative risk aversion utility function defined over life-duration risk. Details

on the exact choice list can be found in Riddel and Kolstoe (2013)[113]. Briefly, the following text was presented to the subjects:

> *Hypothetical Health Risk: Assume you have been diagnosed with a disease that will certainly be fatal in a year without treatment. There are two treatments, but neither is effective 100% of the time. Assume the costs of the treatment are the same, and neither treatment has side effects.*

The subjects were subsequently presented with a sequence of paired lotteries and asked to choose the one they preferred. For example, the first gamble presented was:

> *Treatment A means a 30% chance of eight more years of life and a 70% chance of two more years. Treatment B gives a 90% chance of one more year (the treatment fails) and a 10% chance of 13.5 more years.*

In the subsequent gambles, the better outcome in treatment B is increased so that $\mathbb{E}[A] - \mathbb{E}[B]$ declines, eventually becoming negative. The researcher notes where the subject switches from choosing lottery A to lottery B. The switch defines an inequality over the prospects, with later switch points indicating relatively higher levels of risk aversion.

It should be noted that the sequential price list approach offers a general measure of risk aversion that does not allow one to decompose risk aversion into utility-based risk aversion and the probabilistic risk aversion associated with cumulative-prospect type preferences (Tversky and Kahneman, 1992)[114]. Still, our primary goal in this study is to understand the correlation between health behaviors and risk perceptions, controlling for other factors, such as risk aversion, that affect preferences. We leave it to future research to parse out the effects of different sources of risk aversion on behavior.[3]

[3] A handful of papers have looked at either risk aversion or probability weighting in the health domain.

Subjects were not compensated for their choices because it is clearly impossible to play out life-duration gambles. We are optimistic that the lack of incentives did not bias the estimates. In a review of 74 studies that analyzed the influence of compensation in experiments, Camerer and Hogarth (1999) concluded that "the effects of incentives are mixed and complicated."[118] They generally concluded that when financial gambles were considered, the mean response was unaffected, but the variance tended to be higher when incentives were not offered.

3.3.3 Health History and Objective RRs from HCRI

The survey elicited a detailed health history related to the three cancers in question. Questions included their family history of the cancers and risky or preventative behaviors such as smoking, exercise, vitamin use, dietary habits, chemical exposure, and alcohol use. In the current study, we are interested in a subset of the behaviors, including the following preventative behaviors: (1) taking a daily multivitamin, (2) taking a daily aspirin, and (3) regularly engaging in exercise, defined as 30 minutes of exercise most days, or three or more hours of exercise per week. Of the risky behaviors, we evaluate (1) being a current smoker, (2) being an ever smoker (described as having smoked at least 100 cigarettes in the subject's lifetime), (3) eating more than three servings of red meat per week, (4) having a BMI greater than 30, and (5) drinking excessive amounts of alcohol, defined as more than two drinks daily.

The HCRI is constructed using a series of multiplicative risk factors, such that when multiplied together one can objectively estimate one's overall risk of a given cancer, relative to those of the same age and gender. For example, an RR factor of 1.3 for a

An early paper by Wakker and Deneffe (1996) analyzes preferences over gambles in the life domain[115]. They find that greater observed risk aversion in the life-duration domain than the money domain can be attributed to deviations from expected utility. Bleichrodt and Pinto (2000) find evidence of probability weighting in the medical decision-making domain[116]. Attema, Brouwer, and l'Haridon (2013) estimate CP preferences in the domain of life duration[117]. They also find significant deviations from EU preferences.

given risky activity (or demographic factor) implies that those who regularly engage in this activity (or belong to the demographic) are 30% more likely to get the particular cancer, ceteris paribus, than the U.S. population average for those of the same gender and age. Similarly, an RR factor of 0.8 for a healthy behavior implies that those who regularly engage in this activity have a cancer risk that is 20% lower than other Americans of their same gender and age. By multiplying all of the RR risk behavioral factors that apply to each subject, one arrives at a comprehensive RR factor for the individual in question. This is described in more detail in Section 3.4.1.

3.3.4 Cognitive Ability

Subjects were given a series of seven questions commonly asked in the Wonderlic cognitive ability test. Subjects scored a point for each correct answer. Their total cognitive ability score was the sum of the individual scores.

3.3.5 Demographic Variables

Subjects were queried about their age, income, insurance status, education level, and marital status. Table 3.1 gives descriptive statistics for cognitive ability (mean score out of possible seven correct), risk aversion (mean CRRA coefficient), and the demographic variables. The mean risk aversion coefficient is 0.685, indicating that the average subject is modestly life-duration risk averse. This value is similar to that found in other studies (Riddel and Kolstoe, 2013)[113]. For the most part, the demographics compare reasonably well to the U.S. population. Approximately 45% of the subjects are male compared to the U.S. population percentage of 49%. The median income of the sample of $45,839 is somewhat less than the national median of $51,939. The survey subjects are relatively well educated, with over 43% having a bachelor's degree or higher compared to the U.S. percent of just under 29%. Sample subjects are less likely to be

Table 3.1: Summary Statistics for Demographics, Risk Aversion, and Cognitive Ability

	Sample		U.S. Census (2013)
Variable	Mean or Median	Std. Dev.	Mean or Median
Risk averse (mean CRRA coefficent)	0.685	0.320	—
Cognitive ability (mean out of 7)	4.814	1.227	—
Male (percent)	45.600	49.850	49.200
Age (median)	32.000	10.903	36.800
Income (median $thousand)	45.839	31.019	51.939
4-year college degree (percent)	43.700	49.700	28.800
Married (percent)	36.400	48.200	51.200
Uninsured (percent)	11.300	31.700	13.300

married: 36% married for the sample compared to 51% for the U.S. population. They are slightly more likely to have health insurance than the general U.S. population.

3.4 Indexes for Baseline and Control Optimism

Based on subjects' perceptions of their personal cancer risk and their beliefs about the risk and efficacy of risky behaviors and prevention efforts, we seek to calculate measures of baseline and control optimism, respectively.[4] Although there is no generally agreed on method for calculating optimism in these two dimensions, we maintain that a key criterion for such measures is that they clearly represent the magnitude of the deviation of actual risk from perceived risk. Consistent with this thinking, we developed the measures described below.

3.4.1 Measures of Baseline Optimism

Subjects are defined as baseline optimistic if they underestimate their true cancer risk relative to a person in the United States of their same age, gender, and health behaviors. Based on this, we calculate the variable Baseline Optimism for each subject by comparing his/her stated, subjective population-RR estimate of cancer incidence with the subject's

[4]Note that in this section we apply indexes for baseline and control optimism that we developed and described in earlier work (Riddel and Hales, 2018)[103].

"actual" population-RR factor (ARR).We calculate the ARR by applying each subject's responses to demographic, family history, and lifestyle questions in the survey to the risk estimates recorded for those behaviors in the HCRI. As such, we calculate the ARR of subject i's risk of incidence of cancer j as follows:

$$ARR_{ij} = \frac{1}{PD(\underline{x}_i)} \prod_{k=1}^{k_0(j)} RR_{jk}(\underline{x}_i), \tag{3.1}$$

where $k_0(j)$ is the number of RR factors for cancer j identified in the HCRI, $\underline{x}_i$ is a vector of subject i's demographic characteristics, family history, and lifestyle choices, $RR_{jk}(\underline{x}_i)$ is the HCRI-RR measure for subject i for factor k of cancer j, and $PD(\underline{x}_i)$ is a population denominator derived from the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) Program.[5] The resulting population-RR factor gives expert opinion-derived estimates of a given subject's risk of incidence of cancer j, relative to the U.S. population of persons the same age and gender (Colditz et al., 2000)[104]. An ARR_{ij} value of 1.0 implies that subject i has an average risk of cancer j incidence equal to the average of persons in the United States of the same age and gender; a value of 2.0 suggests cancer risk that is twice the average, etc.

As reported above, the survey asked respondents to report their perceived risk of contracting each of three cancers, again relative to people of their same age and gender. Following the methodology outlined in Colditz et al. (2000), we constructed allowed responses to range from "very much below average risk," corresponding to a RR value of 0.2, to "very much above average risk," corresponding to an RR value of 5.0.[104] We also allowed survey respondents to select "no risk." Since our risk measure is based on log differences, we cannot use a zero value for "no risk" because it would entail dividing

[5]Note that as we did not have access to the SEER population denominator for prostate cancer, we used an estimate of 1.107372, based on the average (non-normalized) relative risk factors of our sample of 218 men.

by zero. We calculated three different values for baseline optimism, based on values for "no risk" of 0.0001, 0.001, and 0.01, respectively. The models we present in the main body of the article are based on the assumption of a value of 0.0001 because this appears to fit the data the best. We report the results under the other two assumptions in a sensitivity analysis reported in Appendix Tables F.1 and F.2.

We define subject i's stated estimate of RR of cancer j as SRR_{ij}, and create a measure of subjects' baseline optimism as follows:

$$BaselineOptimism_{ij} = log_2\left(\frac{ARR_{ij}}{SRR_{ij}}\right) \tag{3.2}$$

As defined, a *BaselineOptimism* value of 0.0 indicates that a subject's own estimate of cancer risk is identical to the HCRI estimate. A Baseline Optimism value of 1.0 indicates that the subject's estimates of cancer incidence risk are half of the expert-derived value (making her risk perceptions relatively optimistic), and a Baseline Optimism value of –1.0 indicates the subject's cancer risk estimates are twice that of the expert value (making her risk perception relatively pessimistic). The logarithmic nature of our index means that each increase (decrease) of one point in our measure thus has the effect of doubling the amount by which expert risk assessments exceed (are exceeded by) the subjects' own risk estimates.[6]

Next, we calculate an overall estimate of each subject's tendency to exhibit baseline optimism by averaging the separate measures for each of the three cancers considered in our study:

$$BaselineOptimism_i = \frac{1}{3}\sum_{j=1}^{3} BaselineOptimism_{ij} = \frac{1}{3}\sum_{j=1}^{3} log_2\left(\frac{ARR_{ij}}{SRR_{ij}}\right) \tag{3.3}$$

[6]Note that we chose to use base-2 logarithms as the resulting measures have a more straightforward interpretation than would measures using a natural or base-10 logarithm.

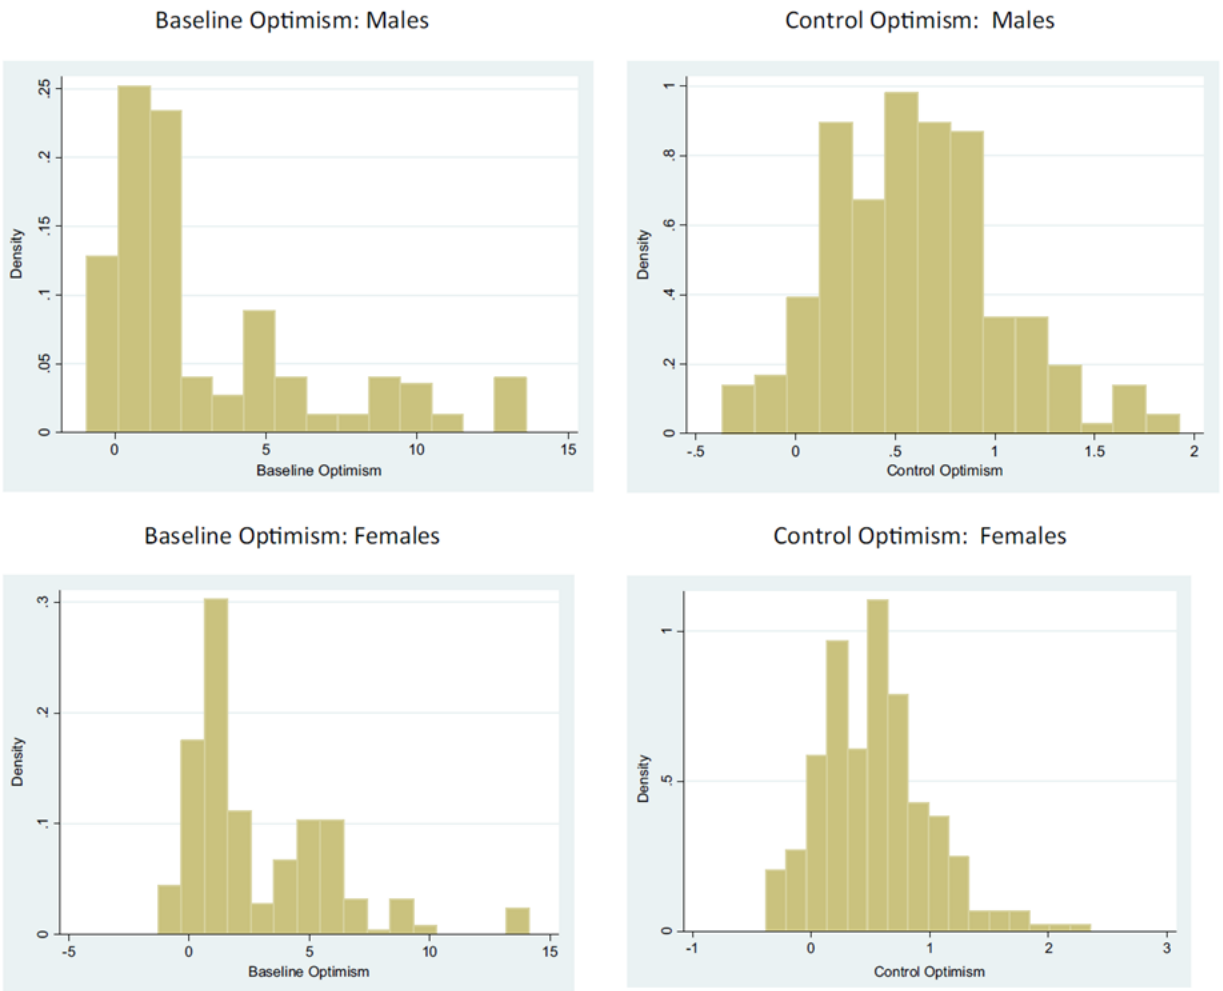

Figure 3.1: Histograms of baseline and control optimism for males and females

Fig 3.1 gives the distribution of *BaselineOptimism* for males and females. Roughly 90% of subjects of each gender are baseline optimistic. That is, they believe their risk of contracting one of the cancers is lower than that predicted by the HCRI. The mean and variance are somewhat higher for females (mean = 2.88, std. dev. = 3.04) than males (mean = 3.04, std. dev. = 3.61).

Table 3.2: Perceived Risk and Prevention Control Optimism

Variable	Mean Stated RR	Actual RR	Prevention Control Optimism = $\log_2$ (Actual Risk/Stated Risk)
Colon			
30 min/day of physical activity	0.529	0.600	0.18
Take multivitamin daily	0.684	0.500	−0.45
Take birth control pills 5+ years	0.923	0.700	−0.40
Take hormones 5+ years	0.921	0.800	−0.20
Take daily aspirin	0.835	0.700	−0.25
Recommended vitamin D dose	0.722	0.600	−0.27
Recent colon screening test	0.427	0.670	0.65
Maintain healthy weight	0.410	1.000	1.29
Bladder			
30 min/day of physical activity	0.566	1.000	0.82
Take multivitamin daily	0.681	1.000	0.55
Take birth control pills 5+ years	0.948	1.000	0.08
Take hormones 5+ years	0.943	1.000	0.08
Take daily aspirin	0.862	1.000	0.21
Maintain healthy weight	0.475	1.000	1.07
Breast			
30 min/day of physical activity	0.752	0.800	0.09
Take multivitamin daily	0.824	1.000	0.28
Take daily aspirin	0.937	1.000	0.09
Breastfeed 1+ years	0.807	0.800	−0.01
Two or more births	0.847	0.850	0.01
Take tamoxifen 5+ years	0.881	0.600	−0.55
First period 15+ years old	0.881	0.800	−0.14
Prostate			
30 min/day of physical activity	0.806	0.800	−0.01
Take multivitamin daily	0.867	1.000	0.21
Take daily aspirin	0.935	1.000	0.10
Eat tomatoes 3+ times per week	0.860	0.700	−0.30

3.4.2 Measures of Control Optimism

A subject is "Prevention Control Optimistic" if she believes that engaging in beneficial activities is more effective in reducing cancer risks than it actually is. Similarly, a subject is "Risk Control Optimistic" if she believes that engaging in a particular risky activity is more likely to lead to cancer than it actually is; we therefore infer that she overestimates her ability to *reduce* cancer risks by *avoiding or curtailing* the risky activity in question.

Our survey contained a set of questions for each cancer that elicited subjects' perceptions of the relative impact different activities have on cancer risk. For a given cancer j and beneficial activity k, subjects were asked to estimate risk-reducing factors between

no risk reduction effect ($RR = 1.0$) and a risk reduction of 10-fold ($RR = 0.1$).[7] Table 3.2 gives the mean subjective RR for each of the cancers for each activity. As Tables 3.2 and 3.2 indicate, some of the behaviors do not affect cancer risk. We included these in the survey to account for misperceptions about the efficacy of activities that do not bear on cancer risk. For example, some women may erroneously believe that taking a multivitamin reduces their breast cancer risk. If that misperception is excluded, then our estimates of control optimism would be biased. However, it could be that including an activity implied that a risk reduction/increase was certain. If this is the case, the design may have "nudged" subjects to give an optimistic response. We do, however, allow subjects to choose the categories of "no risk reduction" and "no risk increase" in the hope that this would mitigate any bias.

Comparing risk-reduction estimates with "actual" expert estimates for each cancer and preventative measure associated with each cancer, subject i's level of prevention control optimism is then:

$$PreventionControlOptimism_i = \frac{1}{\sum_{j=1}^{3} k_{prev}(j)} \sum_{j=1}^{3} \sum_{k=1}^{k_{prev}(j)} log_2 \left(\frac{APRR_{jk}}{SPRR_{jk}} \right), \qquad (3.4)$$

where $k_{prev}(j)$ is the number of preventative measures identified in the HCRI for cancer j, $APRR_{jk}$ is the HCRI-assessed "actual" post-preventative behavior k RR of cancer j, and $SPRR_{jk}$ is the subject's estimates of RR of incidence of cancer j, assuming behavior k (with possible responses coded with RR values ranging from 0.1 to 1.0).

By taking the base-2 logarithm of this ratio, and averaging over the total number of preventative measures identified for each of the three cancers in question, we arrive

[7] We chose the midpoint of the relative risk for each category as follows: not effective. Does not reduce cancer risk ($RR = 1$); somewhat effective. Reduces cancer risk by 10% to 30% of average ($RR = 0.8$); moderately effective. Reduces risk by 30% to 60% ($RR = 0.55$); very effective. Reduces risk by 60% to 80% of average ($RR = 0.3$); extremely effective. Reduces risk by more than 80% of average ($RR = 0.1$).

at a measure of optimism exhibited by subject i for a typical preventative measure. A $PreventionControlOptimism$ value of 0.0 suggests the subject's estimates of prevention effectiveness are, on average, equal to the "actual" expert estimates. A measure of 1.0 implies that, on average, the subject believes preventative measures are twice as effective as they actually are in reducing cancer risk; a measure of –1.0 implies that on average, the subject believes preventative measures are half as effective as they actually are.

Similarly, but with one crucial difference, we estimate risk control optimism as follows:

$$RiskControlOptimism_i = \frac{1}{\sum_{j=1}^{3} k_{risk}(j)} \sum_{j=1}^{3} \sum_{k=1}^{k_{risk}(j)} log_2 \left(\frac{SRRR_{jk}}{ARRR_{jk}} \right), \qquad (3.5)$$

where $k_{risk}(j)$ is the number of risky activities identified in the HCRI for cancer j, $ARRR_{jk}$ is the HCRI-assessed "actual" post-preventative behavior k RR of cancer j, and $SRRR_{jk}$ is the subject's estimates of RR of incidence of cancer j, assuming behavior k (with responses coded with RR values ranging from 1.0 to 10.0).[8] Note that to produce a consistent meaning, the ratio between stated and "actual" risk factors is inverted relative to preventative activities. When $RiskControlOptimism = 1.0$, the subject believes that avoiding risky activities is twice as effective in reducing cancer risk than is actually the case; a measure of –1.0 implies that the subject believes exerting such effort is half as effective in reducing cancer risk than it actually is.

Tables 3.2 and 3.3 give the mean stated RR, the actual RR, and the measure of control optimism for each of the prevention and risky behaviors, respectively. Care must be taken in interpreting these measures. For example, the true RR for regular physical activity with respect to colon cancer is 0.6, indicating that a subject who exercises regularly has

[8]We chose the midpoint of the relative risk for each category as follows: no risk increase ($RR = 1$); small risk increase. Risk is higher, but less than double the average risk ($RR = 1.5$); moderate risk increase. Risk is two to four times the average risk ($RR = 3$); large increase. Risk is four to eight times the average risk ($RR = 6$); very large risk increase. Risk is more than eight times the average risk ($RR = 10$).

Table 3.3: Perceived Risks and Risk Control Optimism

Variable	Mean Stated RR	Actual RR	Risk Control Optimism = $\log_2$ (Stated Risk/Actual Risk)
Colon			
3+ drinks/day	3.17	1.40	1.18
Excess red meat	3.64	1.20	1.60
BMI>30	5.21	1.50	1.80
Less than 3 servings dairy/day	1.82	1.30	0.49
Smoke 1–15 cigarettes/day	2.11	1.00	1.08
Smoke 15–25 cigarettes/day	2.71	1.00	1.44
Smoke 25 or more cigarettes/day	3.18	1.50	1.08
Bladder			
Smoke 1–15 cigarettes/day	3.68	1.30	1.50
Smoke 15–25 cigarettes/day	4.81	2.30	1.06
Smoke 25 or more cigarettes/day	5.53	3.00	0.88
Work in chemical industry 5–20 years	5.93	2.50	1.25
Work in chemical industry 20+ years	6.42	5.00	0.36
Exposed to arsenic 20+ years	5.87	2.00	1.55
Breast			
3+ drinks/day	2.80	1.30	1.11
Currently taking birth control pills	2.52	1.40	0.85
Gain 22–44 lbs since age 19	2.33	1.61	0.53
Gain 45+ lbs since age 18	3.42	1.99	0.78
Taking estrogen alone 5+ years	3.23	1.30	1.31
Taking estrogen +progesterone < 5 years	3.02	1.30	1.21
Taking estrogen +progesterone 5+ years	3.44	2.30	0.58
First birth over 35 years old	2.74	1.50	0.87
Prostate			
3+ drinks/day	1.874	1.000	0.906
Gain 22–44 lbs since age 19	1.520	1.000	0.604
Gain 45+ lbs since age 18	2.128	1.000	1.089
Smoke 1–15 cigarettes/day	2.892	1.000	1.532
Smoke 15–25 cigarettes/day	2.340	1.000	1.226
Smoke 25 or more cigarettes/day	1.911	1.000	0.934
Vasectomy	1.233	1.700	−0.463

a risk of colon cancer that is 60% of an individual of their age and gender who does not engage in exercise. The average subject in the sample believes the RR reduction is slightly more at 52.9%. Thus, the average subject overestimates the risk reduction from physical exercise and is therefore optimistic about the effectiveness of exercise in reducing colon cancer risk.

Table 3.2 indicates that there is a mix of optimism and pessimism for the effects of the different prevention behaviors on the four cancers in question, but optimism appears to dominate. Of the 26 combinations of behaviors and cancers, they are pessimistic in about 11 and optimistic in about the remaining 15. Subjects are more pessimistic about reducing colon cancer risk than bladder cancer, where they are universally optimistic.

Table 3.3 gives the individual measures of optimism for the risky behaviors for each cancer. Again, care must be taken when reviewing the results. A subject is *Risk Control Optimistic* if he or she overestimates the riskiness of the activity, thereby indicating that avoiding the activity has more risk-reducing power than what science might predict. The results indicate that the average subject is almost universally risk control optimistic. Of the 28 behavior and cancer combinations, subjects are only pessimistic on one: (1) having a vasectomy to prevent prostate cancer. For the remaining 27 behaviors, they believe these activities raise cancer risk more than they do. For example, the value of 1.5 for $RiskControlOptimism$ for the effect of smoking 1–15 cigarettes per day on the risk of colon cancer indicates that the average subject believes the risk is more than double the actual risk. For smoking 15–25 cigarettes per day, they believe the risk is slightly more than double the actual risk ($RiskControlOptimism$ equal to 1.06).

In our modeling effort, it is useful to have an aggregate measure of optimism that combines prevention and control optimism. Thus, we average the values of the two variables for each subject to arrive at a characteristic level of control optimism for each

subject:[9]

$$ControlOptimism_i = \frac{1}{2}(PreventionControlOptimism_i + RiskControlOptimism_i). \tag{3.6}$$

An overall Control Optimism measure of 0.0 indicates that the subject accurately assesses the efficacy of prevention efforts and danger of risky activities, in aggregate. A value of 1.0 indicates they believe control efforts are double what science would predict while a value of –1.0 indicates efforts are half as effective as the HCRI predicts.

The distribution of Control Optimism for males and females is given on the right-hand side of Fig 3.1. Roughly 93% of males and 90% of females are control optimistic, overestimating factors that increase cancer risk and overestimating the value of prevention efforts in the aggregate. Both the male and female Control Optimism distributions have significant right skew.

3.5 Results

The goal of this article is to understand how risk aversion, cognitive ability, and the different types of optimism correlate with risky and prevention behavior, controlling for demographic variables that may affect behavior. Toward that end, we estimate probit models first for the five risky behaviors in question followed by the three prevention behaviors. The table gives the marginal effects of a change in the independent variable on the probability the subject engages in the risky activity; p-values in parentheses are calculated using robust standard errors. Note that the marginal effects are calculated at

[9]We elected to weight values for prevention and risk control optimism equally in this estimate, rather than weighting by the total number of preventative and risk-related attributes for each cancer. We did this to avoid overweighting the influence of risk-related attributes, of which more are identified in the HCRI than are preventative-related attributes.

Table 3.4: Probit Models of Risky Behavior

	Current Smoker		Ever Smoker		Red Meat		High BMI		Heavy Drinker	
Variable	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)
Risk aversion	−0.008 (0.276)	−0.014 (0.050)	−0.017 (0.079)	−0.021 (0.020)	0.001 (0.887)	0.001 (0.87)	−0.010 (0.282)	−0.010 (0.276)	−0.008 (0.190)	−0.054 (0.053)
Cognitive ability	−0.046 (0.000)	—	−0.035 (0.020)	—	−0.007 (0.627)	—	−0.011 (0.543)	—	−0.018 (0.062)	—
Baseline optimism	0.018 (0.001)	—	0.020 (0.006)	—	0.015 (0.032)	—	0.017 (0.010)	—	0.006 (0.217)	—
Control optimism	−0.093 (0.024)	—	−0.074 (0.155)	—	−0.060 (0.241)	—	0.087 (0.095)	—	−0.087 (0.017)	—
Male	0.039 (0.315)	−0.026 (0.493)	0.069 (0.162)	0.022 (0.634)	0.083 (0.080)	0.078 (0.079)	0.247 (0.000)	0.248 (0.000)	0.018 (0.573)	−0.112 (0.419)
Age	0.003 (0.036)	3.03E-05 (0.975)	0.009 (0.000)	0.006 (0.000)	−0.0007 (0.688)	−0.001 (0.449)	0.002 (0.372)	0.002 (0.328)	−0.002 (0.079)	−0.018 (0.000)
Income ($000)	−0.002 (0.010)	−0.003 (0.000)	−0.002 (0.012)	−0.002 (0.002)	−0.0006 (0.431)	−0.001 (0.216)	−0.001 (0.150)	−0.001 (0.163)	−0.0009 (0.108)	−0.006 (0.018)
College	−0.096 (0.012)	−0.125 (0.001)	−0.168 (0.000)	−0.187 (0.000)	−0.048 (0.315)	−0.061 (0.189)	−0.054 (0.238)	−0.079 (0.076)	0.012 (0.704)	0.031 (0.832)
Married	0.030 (0.483)	0.024 (0.583)	−0.016 (0.768)	−0.008 (0.874)	0.006 (0.912)	0.017 (0.732)	0.081 (0.107)	0.093 (0.063)	0.040 (0.268)	0.118 (0.453)
Uninsured	−0.012 (0.836)	−0.035 (0.549)	−0.064 (0.387)	−0.071 (0.327)	−0.087 (0.222)	−0.068 (0.335)	−0.103 (0.126)	−0.010 (0.129)	−0.001 (0.981)	−0.099 (0.674)
Log likelihood	−226.260	−245.760	−297.960	−308.830	−318.180	−322.970	−276.830	−298.83	−181.230	−194.523

Note: Dependent Variable Equals 1 if the Subject Engages in that Behavior and 0 Otherwise for No Risk Value Equal to 0.0001

the means of the dependent variables.[10]

Table 3.4 gives the results of 10 probit models that relate optimism, risk aversion, cognitive ability, and demographic control variables to the probability of engaging in the risky activities. The risky behaviors evaluated are being a current smoker, being an ever smoker, excessively consuming red meat (eating red meat more than three times per week), being obese (having a body mass index [BMI] of 30 or higher), and excessively consuming alcohol (more than two drinks per day). For each behavior, we give a full model including all of the variables of interest and a reduced model that excludes cognitive ability and the optimism variables.

The first two columns address the probability the subject is a current smoker. The results indicate that older, lower-income subjects are significantly more likely to be current smokers than their younger, wealthier counterparts. College graduates are less likely to smoke currently than subjects with less education. Gender, marital status, and insurance status do not have a statistically significant correlation with the decision to smoke currently.

[10]Note that the models reported here are estimated under the "no risk" assumption of 0.0001 for baseline optimism. In Tables F.1 and F.2, we report model results assuming "no risk" values of 0.001 and 0.01.

Subjects who are relatively baseline optimistic about their risk of contracting one or more of the three cancers are significantly more likely to smoke (p-value = 0.001). Subjects who are control optimistic, that is, those who overestimate the efficacy of prevention efforts and avoiding risky activities, in aggregate, are less likely to smoke (p-value = 0.024). Cognitive ability is also statistically significant: subjects with higher cognitive abilities are less likely to smoke. Interestingly, the marginal effect of risk aversion in the reduced model is almost twice that of the full model that controls for optimism and cognitive ability. Moreover, the coefficient of risk aversion is significant in the reduced model but becomes insignificant in the full model. These results suggest that controlling for optimism significantly reduces the correlation between risk aversion on the decision to smoke.

As with current smokers, ever smokers tend to be older, lower income, and less likely to have a college degree. Marital status, insurance status, and gender are not statistically significant. This is true of both the full and reduced model.When subjects who have quit smoking are included in the smoker category, baseline optimism is still positively and significantly related to smoking. However, control optimism is only marginally significant (p-value = 0.155). As in the current smoker model, subjects with higher cognitive ability are less likely to have ever smoked in their lives. Risk-averse subjects are less likely to have ever smoked. Again, when optimism and cognitive ability are added to the model, the relationship between risk aversion and optimism is muted, with the marginal effect falling from –0.021 in the reduced model to –0.017 in the full model.

Excessive red meat consumption is increasing in baseline optimism, so that those who underestimate their personal risk of the three cancers are more likely to eat more than three servings of red meat per week. Risk aversion, cognitive ability, and control optimism are not statistically significant predictors of red meat consumption in the full model. Risk aversion remains insignificant in the reduced model. Of the demographic

variables, only gender is statistically significant.

The model results suggest that both types of optimism are positively correlated with having a high BMI. On one hand, subjects who are baseline optimistic and therefore underestimate their cancer risk are more likely to be obese than subjects with accurate baseline risk perceptions. On the other hand, subjects who believe that prevention activities and avoiding risky activities are a more effective strategy for reducing cancer risk than what experts would believe are also more likely to be obese. This latter correlation is counterintuitive, since we expect that control optimists would tend to overinvest in riskreducing effort. Males and married subjects are more likely to be obese, while age, income, education, and insurance status are not statistically significant predictors of obesity. Cognitive ability also does not correlate with obesity when controlling for optimism and demographic variables. Similarly, risk aversion is not statistically significant in either the full model or the reduced model.

As expected, heavy drinking is decreasing in control optimism: subjects who have optimistic views about the health benefits of avoiding risky activities and engaging in prevention measures are less likely to drink heavily. Subjects with higher cognitive ability are also less likely to consume detrimental amounts of alcohol. The probability of heavy drinking is decreasing in both age and income, but the marginal effect of income is only borderline significant in the full model (p-value = 0.108). None of the remaining demographic variables are significant. As with the model for current smoking, the coefficient of risk aversion is much smaller, in absolute value, in the full model than in the reduced model. Also note that the coefficient of risk aversion is not statistically significant in the full model that controls for optimism and cognitive ability, but becomes significant when these variables are excluded from the regression. Again, it appears that controlling for optimism and cognitive ability moderates the correlation between risk aversion and risky behavior.

Table 3.5: Probit Models of Healthy Behaviors

	Physically Active		Aspirin Daily		Multivitamin	
Variable	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)	*dy/dx* (*p*-value)
Risk aversion	−0.004 (0.678)	−0.005 (0.572)	−0.002 (0.235)	−0.0003 (0.828)	0.016 (0.070)	0.013 (0.137)
Cognitive ability	−0.004 (0.849)	—	−0.011 (0.015)	—	0.007 (0.706)	—
Baseline optimism	−0.014 (0.037)	—	−0.004 (0.042)	—	−0.002 (0.004)	—
Control optimism	0.046 (0.374)	—	0.0004 (0.962)	—	0.077 (0.127)	—
Male	0.120 (0.007)	0.113 (0.009)	−0.002 (0.806)	−0.001 (0.853)	−0.051 (0.266)	−0.045 (0.310)
Age	−0.003 (0.128)	−0.003 (0.131)	0.0005 (0.062)	0.001 (0.006)	0.004 (0.031)	0.003 (0.114)
Income ($000)	0.003 (0.681)	0.001 (0.504)	−0.0005 (0.030)	−0.001 (0.023)	0.0002 (0.717)	0.0004 (0.611)
College	0.119 (0.008)	0.123 (0.005)	−0.004 (0.650)	−0.004 (0.591)	0.105 (0.025)	0.110 (0.014)
Married	−0.035 (0.481)	−0.049 (0.314)	0.013 (0.239)	0.016 (0.228)	−0.031 (0.532)	−0.028 (0.562)
Uninsured	0.059 (0.380)	0.035 (0.606)	(omitted)	(omitted)	0.120 (0.112)	0.093 (0.191)
Log likelihood	−290.119	−294.47	−56.06	−50.52	−289.58	−298.83

Note: Dependent Variable Equals 1 if the Subject Engages in the Activity

Table 3.5 gives the marginal effects for a change in the probability of engaging in a cancer-risk-reducing behavior. We hypothesize that engaging in healthy behaviors will be increasing in risk aversion as more risk-averse subjects are more likely to take steps to reduce health risks. We expect that baseline optimists will be less likely to engage in these risk-reducing behaviors since they tend to underestimate their cancer risk and therefore under-invest in risk reduction. Finally, we expect the probability of engaging in the healthy behaviors to be increasing in control optimism since these subjects overestimate the efficacy of risk-reducing effort and would therefore engage in more of these activities than their more pessimistic counterparts.

The model results indicate that baseline optimists are significantly less likely to meet the exercise minimums, but risk aversion is not correlated with physical activity in either the full or reduced model. College-educated subjects are more likely than their less educated counterparts to achieve the minimum exercise requirement. Males are more likely to engage in regular physical activity than female subjects. Cognitive ability, control optimism, income, marital status, age, and insurance status are not statistically

significant correlates of exercise frequency.

Daily aspirin consumption is decreasing in baseline optimism, cognitive ability, and income and increasing in age. None of the other variables is a statistically significant predictor of the probability of taking a daily aspirin.[11]

More risk-averse subjects are more likely to take multivitamins. This is true in the full model, but the significance level of the variable declines in the reduced model. As expected, the probability of daily multivitamin use is decreasing in baseline optimism. Older subjects and college-educated subjects are also more likely to take multivitamins. Cognitive ability, control optimism, gender, income, marital status, and insurance status are not correlated with the choice to take a daily multivitamin.

3.6 Discussion

The models indicate that optimism is significantly correlated with people's decisions to engage in risky activities. On one hand, baseline optimists – those who underestimate their risk of contracting one of the three cancers – are more likely to smoke now or in the past, to eat excessive amounts of red meat, and to be medically obese. On the other hand, they are less likely to engage in the prevention behaviors such as regular physical activity or to take an aspirin and multivitamin daily. Thus, baseline optimistic subjects are more likely to engage in these risky behaviors, yet fail to undertake preventative health measures. However, while we have attempted to control for the effect of other factors that might have a causal effect on behavior, we cannot eliminate the possibility that we have included all that are relevant. Thus, while our results are suggestive of a causal relationship between baseline optimism and the probability of our subjects engaging in cancer-relevant behaviors, we cannot exclude the possibility that the observed correlations are caused by some other (unmodeled) relationship.

[11] Insurance status was dropped from the regression because it perfectly predicted daily aspirin use.

The relationship between control optimism and risky and prevention behavior is rather more complex. While control optimism is negatively correlated with smoking and heavy drinking, it is positively correlated with the probability that a subject is overweight. None of the other risky or prevention behaviors are correlated with control optimism. These findings support those of Viscusi's (1990) that subjects who tend to overestimate the specific risk of contracting lung cancer from smoking are less likely to smoke[99]. This is what one would predict from standard models of preferences: subjects with higher perceptions of the risk and therefore expected costs of engaging in an activity will be more likely, *ceteris paribus*, to avoid the costly activity.

The finding that obesity is increasing in control optimism is difficult to explain. We would expect that control optimists would engage in more prevention than those with accurate perceptions since they believe that the marginal return to effort is higher than it actually is. Presumably, one of the outcomes of increased effort would be a lower BMI and a decline in the likelihood that the subject is obese. Of course, body weight is an accumulation of past decisions about eating and exercise. It could well be that many subjects with a high BMI currently live healthy lifestyles, that is, engage in a significant amount of prevention activities, but continue to suffer from obesity. A more sophisticated analysis of current and past behavior may be able to give a deeper insight into control optimism, investing in prevention effort, and obesity.

One of the key goals of this article is to examine the complex interaction between risk aversion, baseline and control optimism, and cognitive ability and their joint influence on engaging in risky and prevention behaviors. Of the risky behaviors, risk aversion is statistically significant in the reduced models for current smoking, ever smoking, and heavy drinking, but the variable becomes insignificant when controlling for optimism and cognitive ability. Notably, the magnitude of the correlation is consistently smaller in the full models. Of the preventative behaviors, risk aversion is not statistically significant

in either the full or reduced models for aspirin use or physical exercise. Risk aversion is positive and significant in both of the models for multivitamin use, with the effect slightly smaller in the reduced model.

It is possible that a high correlation between optimism, risk aversion, and cognitive ability is inflating the standard errors of the respective coefficients and making the model estimates unstable. This is less likely to be of concern in a probit model because the inherent non-linearity between the dependent and independent variables helps identify the coefficients. Nevertheless, to test for multicollinearity, we calculated variance inflation factors (VIFs) for the risk aversion, cognitive ability, and the optimism variables for each of the models. The average VIF is 1.08, and the maximum VIF is 1.16. A frequently used rule of thumb suggests that multicollinearity is a problem if a variable or set of variables has a VIF of 10 or higher. Since the VIFs for the current set of variables are all well below 2, it is highly unlikely that multicollinearity is responsible for the high standard errors for the risk-aversion variable in the full regressions. Rather, we conclude that for the majority of our models, the correlation between behavior and risk aversion is attenuated when optimism and cognitive ability are considered.

This is at odds with Barsky et al.'s (1995)[93] and Anderson and Mellor's (2008)[92] findings that risk-averse subjects are less likely to smoke, less likely to drink excessively, and less likely to be obese. Of course, our results are not directly comparable to Barsky et al. (1995) and Anderson and Mellor (2008) because their measure of risk aversion is constructed over financial choices rather than the life-duration choices used in the current study. That said, previous research has found that financial and life-duration risk-aversion measures are highly correlated (Dohmen et al., 2011[95]; Einav, Finkelstein, Pascu, and Cullen, 2012)[119]. Still, our results indicate that at least in the domain of cancer-related health factors, optimism—particularly baseline optimism—is more correlated with health-risk behavior than is risk aversion.

3.7 Conclusions

In this article, we report the results of a survey of 474 men and women that analyzes variables that potentially correlate with cancer-prevention activities. In particular, we seek to sort out the effects of risk aversion, baseline and control optimism, and cognitive ability on the choice to engage in risky or prevention activities. Our results indicate that risk aversion is far less correlated with prevention and risk behavior than previously believed. In contrast, we find that choice patterns that were previously attributed to risk aversion more likely arise from optimism about one's cancer risk, paired with optimism about the efficacy of prevention and risk-avoidance behaviors. Of course, we caution the reader that the statistical models presented in the article offer correlations rather than allowing us to directly infer causation. Still, our finding that the correlation between risk aversion and behavior is generally attenuated when controlling for optimism and cognitive ability is compelling.

This finding is potentially important to public health experts. If choices to engage in risky behavior and prevention effort are primarily motivated by risk aversion, then there is not much room for information to change behavior since the shape of the preference function is rooted in individual preferences. However, if these decisions rest largely on misperceptions of the risk of contracting cancer, then interventions in the form of accurate information could be more useful. However, an interesting public health conundrum arises since better information would actually lead to less healthy behavior and higher cancer risks for the control optimists who comprise much of our sample. Finally, if a third variable is responsible for the ostensible link between misperceptions and behavior, then correctly aligning perceptions will not necessarily change behavior.

This is the first empirical study we know of that investigates the role of risk misperception in the domain of cancer-related health behaviors while simultaneously con-

trolling for cognitive ability and risk aversion. However, there are some limitations to our analysis. For one, we ask people about their perceptions of the risk and benefits of different behavior. Risk perceptions are difficult to elicit and problems with numeracy may well induce measurement error in these variables (Reyna, Nelson, Han, and Dieckman, 2009)[120]. We control for cognitive ability, which should offset bias introduced by numeracy issues. However, numeracy issues could still be a problem in the estimates. Also, we use self-reported measures of risky and prevention behavior. Some researchers question the validity of self-reported smoking and other risky behaviors, believing that subjects may understate how much risky behavior they engage in either because they underestimate the amount or intentionally provide a lower number because they do not want to be associated with an undesirable behavior. Patrick et al. (1994) found that under-reporting was less likely to happen in self-administered surveys such as the current survey[121]. Nevertheless, significant under-reporting of risky behavior or over-reporting of prevention behavior could potentially cloud our results.

Another limitation is that we implicitly assume that subjects are expected utility maximizers. A substantial body of work indicates that decisions under risk are often better characterized by cumulative prospect preferences that account for non-linear weighting of probabilities as well as loss aversion. Our primary interest is to evaluate the correlation between optimism and behavior while controlling for other sources of concavity in the preference function, which we believe we have accomplished. Of course, this modest approach is not capable of further disaggregating the influence of utility risk aversion, probability weighting, and loss aversion on behavior. We will leave this for future research.

Appendix A

Quadratic Utility and Linear Demand

Here we derive a linear demand system from an assumption of quasi-linear quadratic utility.[1]. Following the notation and proof in Amir et al. (2017) we have:

$$U(x, y) = a'x - \frac{1}{2}x'Bx + y,$$

where a is a strictly positive vector of size n, B is a positive-definite $n \times n$ matrix, x is an n-vector representing quantities of goods, and y is the quantity of the numeraire good with price $p_y = 1$.

Being that matrix B is positive definite B^{-1} exists and is also positive definite. Then, imposing the standard budget restriction $p'x + y \leq m$ with exogenous price vector p and budget m; assuming interiority condition $B^{-1}(a - p) > 0$ and feasibility condition

[1]We thank an anonymous reviewer for comments that inspired the approach followed in this appendix

$p'B^{-1}(a-p) \leq m$, we arrive at a system of linear demand functions:

$$q(p) = B^{-1}(a-p). \tag{A.1}$$

Now, to match the market environment we model in our experiment, we impose restrictions on a and B: we assume $a_i = a$, $b_{ii} = b$ and $b_{ij} = d$ $\forall i, j \in [1, n], i \neq j$ for some strictly positive constants a, b, and d, with $b > d$ to ensure B is positive definite. These assumptions are equivalent to assuming that the utility derived from consumption of each good x is symmetric both in terms of own- and cross-product parameter values. Intuitively and as we will see, this leads to symmetric (linear) demand functions for each good x.[2]

To explore this further and apply to our specific demand specification, we make the several definitions and impose the following additional restrictions on a, b, and d:

- Define parameters δ and $\gamma : \delta, \gamma \in \mathbb{R}_{++}$
- Define p_i as the ith element of price vector p and $p_{-i} \equiv \frac{1}{n-1}\sum_{j \neq i}^{n} p_j$ as the average of the other $n-1$ elements in the price vector
- Assume $n \in \mathbb{Z}, n \geq 2$
- Restrict $a = \delta n d$
- Restrict $b = d + \frac{n-1}{n\gamma}$
- For compactness of notation and clarity, define $\Delta \equiv \frac{n-1}{n\gamma} \cdot \frac{1}{d} \Rightarrow b = d(1+\Delta)$

[2] As Amir et al. (2017) point out, our use of the term "linear demand function" is a slight abuse of terminology. More correctly we have an affine function whenever the implied result is positive and zero otherwise.

Next, we impose these restrictions on (A.1) in order to show that:

$$\lim_{d\to\infty} q_i(p; d, n, \delta, \gamma) = \delta - \gamma(p_i - p_{-i})$$

Next, we must rationalize B^{-1}. Noting that $b = d(1 + \Delta)$ we can rewrite B as $d \cdot \mathbb{B}$, where the diagonal elements of $\mathbb{B}$ are all $1 + \Delta$ and the off-diagonal elements are all 1. For example, to illustrate with $n = 4$:

$$\mathbb{B} = \begin{bmatrix} 1+\Delta & 1 & 1 & 1 \\ 1 & 1+\Delta & 1 & 1 \\ 1 & 1 & 1+\Delta & 1 \\ 1 & 1 & 1 & 1+\Delta \end{bmatrix}$$

This leads to a specification of $B^{-1} = d^{-1} \cdot \mathbb{B}^{-1}$, with:

$$\mathbb{B}^{-1} = \frac{1}{\Delta(\Delta + n)} \begin{bmatrix} n-1+\Delta & -1 & -1 & -1 \\ -1 & n-1+\Delta & -1 & -1 \\ -1 & -1 & n-1+\Delta & -1 \\ -1 & -1 & -1 & n-1+\Delta \end{bmatrix}.$$

We then apply these restrictions to (A.1) and have:

$$q(p) = B^{-1}(a - p) = \left(\frac{1}{d}\right) \mathbb{B}^{-1}(a - p).$$

Then, the demand for arbitrary good i can be represented as:

$$q_i(p; d, n, \delta, \gamma) = \left(\frac{1}{d}\right) \cdot \frac{1}{\Delta(\Delta + n)} \left[(n - 1 + \Delta)(a - p_i) - (-1)\sum_{j\neq i}^{n}(a - p_j)\right]$$

$$= \frac{1}{d\Delta(\Delta+n)}\left[\Delta(a-p_i)+(n-1)(a-p_i)+(n-1)(a-p_{-i})\right]$$

$$= \frac{a-p_i}{d(\Delta+n)} - \frac{(n-1)(p_i-p_{-i})}{d\Delta(\Delta+n)}$$

Substituting in $a = \delta nd$ and $\Delta = \frac{n-1}{n\gamma d}$, we have:

$$q_i(p;d,n,\delta,\gamma) = \frac{\delta nd - p_i}{d\left(\frac{n-1}{n\gamma d}+n\right)} - \frac{(n-1)(p_i-p_{-i})}{d\left(\frac{n-1}{n\gamma d}\right)\left(\frac{n-1}{n\gamma d}+n\right)}$$

$$\Rightarrow q_i(p;d,n,\delta,\gamma) = \frac{\delta - p_i/nd - \gamma(p_i-p_{-i})}{\left(\frac{n-1}{n^2\gamma d}+1\right)} \tag{A.2}$$

Now, evaluating this expression as $d \to \infty$ we see that $\lim_{d\to\infty}\left(\frac{n-1}{n^2\gamma d}+1\right) = 1$ and $\lim_{d\to\infty}(p_i/nd) = 0$, which implies:

$$\Rightarrow \lim_{d\to\infty} q_i(p;d,n,\delta,\gamma) = \delta - \gamma(p_i-p_{-i}), \qquad \text{QED.} \tag{A.3}$$

We thus see that our model of linear demand, with own- and cross-price parameters equal in magnitude, is consistent with an assumption of quadratic, quasi-linear utility, albeit a special case taken in the limit as parameters a, b and d tend towards infinity in the pathway described above.

Appendix B

Experimental Material – Imperfect Tacit Collusion and Asymmetric Price Transmission

B.1 Instructions

Good morning, and thank you for agreeing to participate in this economics experiment.

Earnings

As compensation, you will be paid a show-up fee of $5. In addition to the show-up fee, you will have the opportunity to earn additional money. We anticipate that this experiment will run around 90-100 minutes. The experiment consists of 5 rounds of 15 periods each, a total of 75 rounds. The computer will randomly select a number between 1 and 5, corresponding to one of the 5 rounds of the experiment. At the end of the experiment, your additional fee will be equal to the average payout during that particular round. You will then be paid the added total of the show-up fee plus the

additional fee. You are free to leave at any time; if you do so you will still receive the $5 show-up fee, but if you leave before the experiment is complete you will not receive the additional fee. In all cases, your earnings will be paid individually and anonymously.

Market Setup

In this experiment we will simulate markets, in which you and the other participants each play the role of the CEO of a company that produces and sells a single product in your particular market. You will be randomly grouped with X $(n-1)$ other companies (participants), and together the Y (n) of you will form this market. You will stay matched with the same participants in your market for the duration of the entire experiment. Each company is largely identical, faces the same identical costs to produce each unit, and has the same profit function. The only thing that differs between companies is the price they set for the product. Demand for your product will be simulated by the computer, according to a formula shown on the payoff sheet we have left at your workstation. The higher your price, the fewer units you will sell. The higher the average price of the other participants in your market, the more units you will sell.

During each period of the experiment, all Y (n) of you will be asked to set a price at which you will each sell the product. You will not know anything about the price the other participants set, until after you have set your own price. We will then ask you to guess the average price the other participants set during that same round. Finally, after all Y participants have set their own prices, we will show you the average price set by the other X participants, and calculate and show you your payoff for that particular round. You will be able to see the history of your prices, the average prices of the other participants, and your resulting payoff for each of the previous periods within each round, to help you make future pricing decisions.

How to Set Your Price and Predict Your Payout each Period

In the first round of 15 periods, you will face an input cost of $0.90 per unit produced.

"Input cost" is shorthand for the total costs of raw materials, labor, etc., required to produce one unit of product. The payoff sheet at your workstation corresponds to this particular input cost. Based on what you guess the average of others' costs to be, shown along the columns, you can see how much you will earn for each potential price you would set, shown along the rows. For example, if the average price other participants set is \$1.50, and you set your price at \$2.00, your payoff would be \$4.35. As another example, if the average other participants set is \$2.90 and yours is \$2.30, your payoff would be \$17.01. Note that the maximum price you can set for your product is \$3.00, because consumers in this market are not willing to pay more than \$3.00 for the product. We only show prices that are multiples of \$0.10, because otherwise the payoff sheet would be too large to print on a single piece of paper. The price you set can be anywhere between the input price and \$3.00, in increments of \$0.01. Finally, note that negative numbers are shown in the payoff sheet in parentheses, so for example (\$1.00) means minus one dollar.

Changes in Input Cost

In each round, there will be a change to the input costs each company faces. Costs will increase or decrease by either \$0.40 or \$0.80. You will not know the size or direction of the change until it is announced at the beginning of each round. At that time, we will hand out a new payoff sheet that corresponds to the new input cost. Make sure you use only the payoff sheet that corresponds to the current input cost. Again, input costs will remain the same for all 15 periods of each round, but will change at the beginning of each new round.

Please raise your hand if you have a question, and one of us will come to you at your workstation. Please do not talk or discuss the experiment or anything else with your neighbors, until after the experiment is complete.

B.2 Comprehension Questions

Before we begin the experiment, please complete the following questions. Just write your answer on this sheet of paper. We will walk by each workstation and look at your answers, to ensure you understand. If you have questions please ask us when we come by, but DO NOT discuss or ask questions of your neighbors. If there is anything needing clarification, we will announce it to everyone in the group at the same time.

1) True/False: I will be rematched with different participants at the beginning of each round:

2) If I keep my price the same from one period to the next, but the average price others in my market set falls, my payout in that period will: (hint – look at the payout chart and see what happens as you go from right to left on any given row)

a. Rise

b. Fall

c. Stay the Same

d. May Rise or Fall, or Stay the Same

3) If the average price of others in my market stays the same between periods, but if I increase my price, my payout that period will: (hint – look at the payout chart, but this time see what happens as you go from top to bottom on any given column)

a. Rise

b. Fall

c. Stay the Same

d. May Rise or Fall, or Stay the Same

4) Use the Payoff Chart to answer the following questions:

a. If the average price others set in my market is $2.80, what is the price I could set that would maximize my own payout that period?

$

b. What would be the amount of that payout?

$

B.3 Payoff Table and Experimental Interface

PAYOFF TABLE - FOR USE ONLY WHEN INPUT COSTS = $ 0.90

Demand for your product = 8.50 - 7.275 * (Your Price - Average of Others' Prices)
Your Payoff = (Your Price - Input Cost) * Demand - $1.00

Others' Average Price ====>

		$ 1.30	$ 1.40	$ 1.50	$ 1.60	$ 1.70	$ 1.80	$ 1.90	$ 2.00	$ 2.10	$ 2.20	$ 2.30	$ 2.40	$ 2.50	$ 2.60	$ 2.70	$ 2.80	$ 2.90	$ 3.00
Your	$1.30	2.40	2.69	2.98	3.27	3.56	3.86	4.15	4.44	4.73	5.02	5.31	5.60	5.89	6.18	6.47	6.77	7.06	7.35
Price	$1.40	2.89	3.25	3.61	3.98	4.34	4.71	5.07	5.43	5.80	6.16	6.52	6.89	7.25	7.62	7.98	8.34	8.71	9.07
||	$1.50	3.23	3.66	4.10	4.54	4.97	5.41	5.85	6.28	6.72	7.16	7.59	8.03	8.47	8.90	9.34	9.77	10.21	10.65
||	$1.60	3.42	3.93	4.44	4.95	5.46	5.97	6.48	6.99	7.50	8.01	8.51	9.02	9.53	10.04	10.55	11.06	11.57	12.08
||	$1.70	3.47	4.05	4.64	5.22	5.80	6.38	6.96	7.55	8.13	8.71	9.29	9.87	10.46	11.04	11.62	12.20	12.78	13.37
||	$1.80	3.38	4.03	4.69	5.34	6.00	6.65	7.30	7.96	8.61	9.27	9.92	10.58	11.23	11.89	12.54	13.20	13.85	14.51
V	$1.90	3.14	3.86	4.59	5.32	6.05	6.77	7.50	8.23	8.96	9.68	10.41	11.14	11.87	12.59	13.32	14.05	14.78	15.50
	$2.00	2.75	3.55	4.35	5.15	5.95	6.75	7.55	8.35	9.15	9.95	10.75	11.55	12.35	13.15	13.95	14.75	15.55	16.35
	$2.10	2.22	3.09	3.96	4.84	5.71	6.58	7.45	8.33	9.20	10.07	10.95	11.82	12.69	13.57	14.44	15.31	16.18	17.06
	$2.20	1.54	2.48	3.43	4.38	5.32	6.27	7.21	8.16	9.10	10.05	11.00	11.94	12.89	13.83	14.78	15.72	16.67	17.62
	$2.30	0.72	1.73	2.75	3.77	4.79	5.81	6.83	7.84	8.86	9.88	10.90	11.92	12.94	13.96	14.97	15.99	17.01	18.03
	$2.40	(0.25)	0.84	1.93	3.02	4.11	5.20	6.29	7.39	8.48	9.57	10.66	11.75	12.84	13.93	15.02	16.12	17.21	18.30
	$2.50	(1.00)	(0.20)	0.96	2.12	3.29	4.45	5.62	6.78	7.94	9.11	10.27	11.44	12.60	13.76	14.93	16.09	17.26	18.42
	$2.60	(1.00)	(1.00)	(0.15)	1.08	2.32	3.56	4.79	6.03	7.27	8.50	9.74	10.98	12.21	13.45	14.69	15.92	17.16	18.40
	$2.70	(1.00)	(1.00)	(1.00)	(0.10)	1.21	2.51	3.82	5.13	6.44	7.75	9.06	10.37	11.68	12.99	14.30	15.61	16.92	18.23
	$2.80	(1.00)	(1.00)	(1.00)	(1.00)	(0.05)	1.33	2.71	4.09	5.47	6.86	8.24	9.62	11.00	12.39	13.77	15.15	16.53	17.91
	$2.90	(1.00)	(1.00)	(1.00)	(1.00)	(1.00)	(0.00)	1.45	2.91	4.36	5.82	7.27	8.73	10.18	11.64	13.09	14.55	16.00	17.46
	$3.00	(1.00)	(1.00)	(1.00)	(1.00)	(1.00)	(1.00)	0.04	1.57	3.10	4.63	6.16	7.68	9.21	10.74	12.27	13.79	15.32	16.85

Note: numbers in parentheses () indicate negative numbers. For example, (1.00) means -1

Figure B.1: Exemplary payoff matrix/table. Provided to subjects when marginal cost is equal to $0.90.

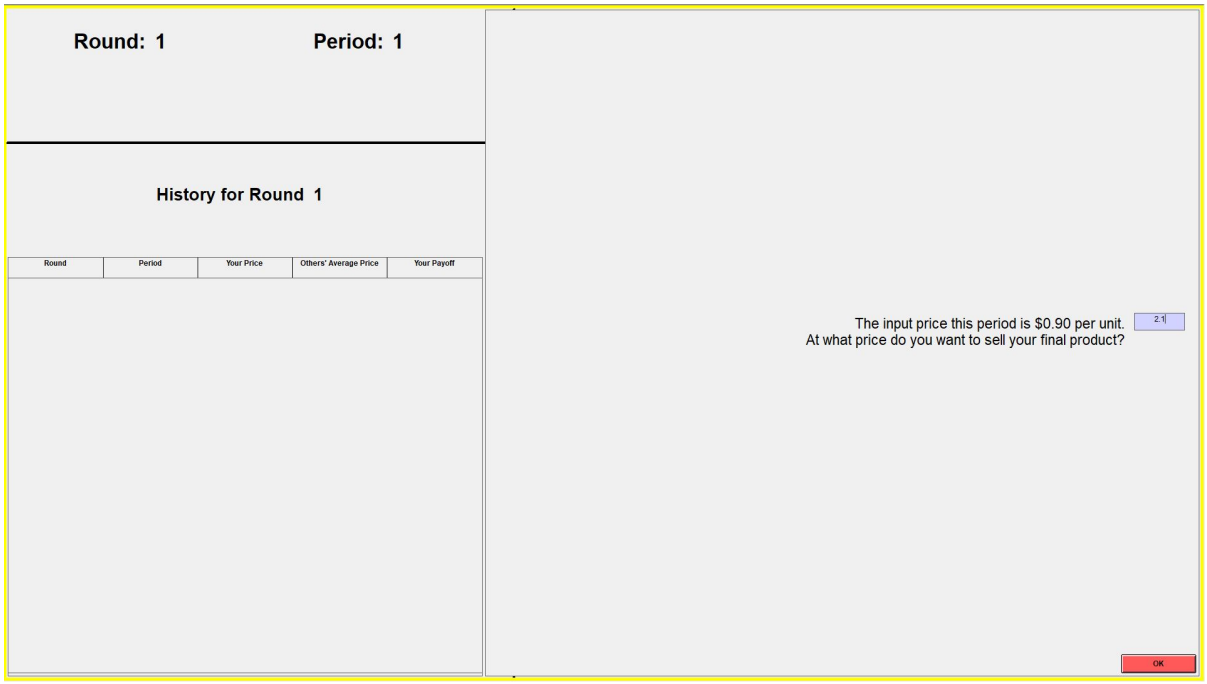

Figure B.2: An example for price setting screen.

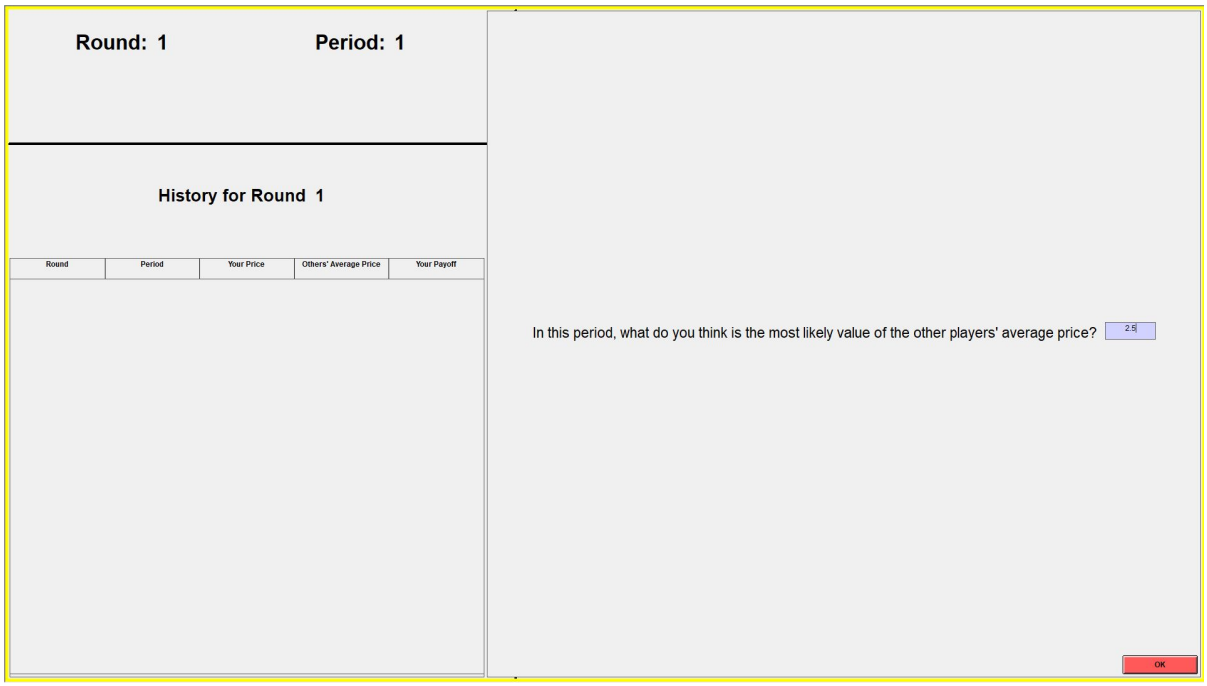

Figure B.3: An example for guess decision screen.

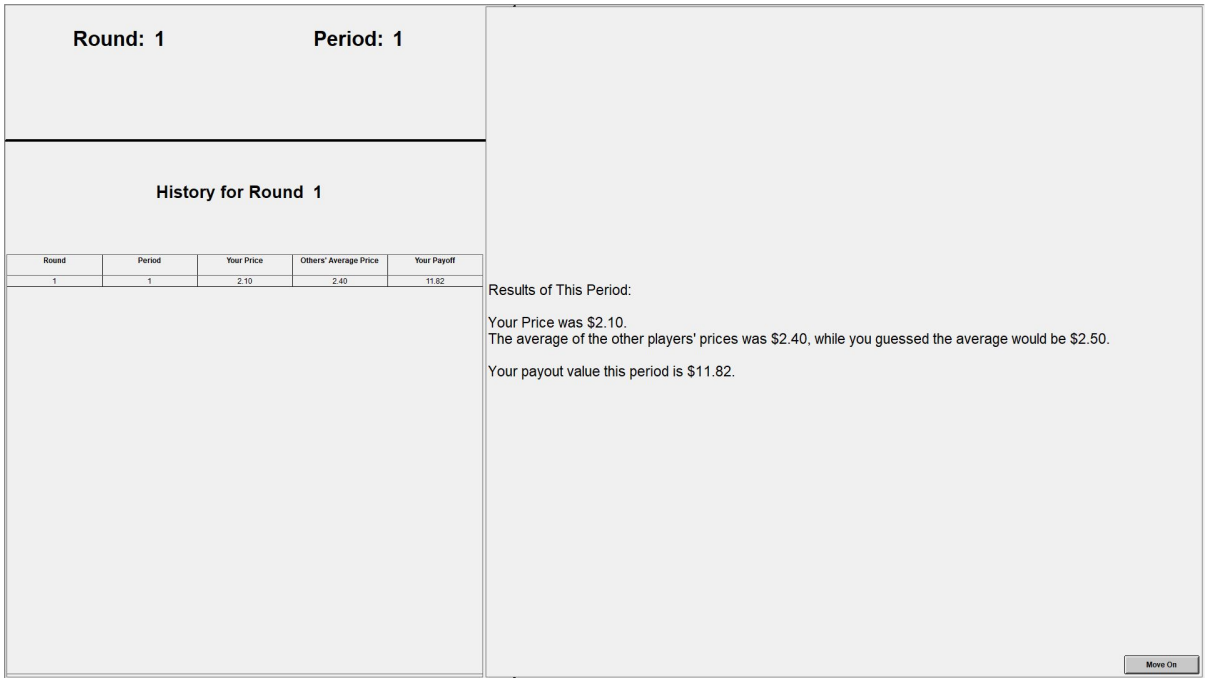

Figure B.4: An example for feedback screen.

Appendix C

Additional Analysis – Imperfect Tacit Collusion and Asymmetric Price Transmission

C.1 Regression Results of Asymmetry

The regression estimates used in Figure 2 of the main text are reported on Table C.1 under column (1). The remaining columns introduce control variables in a step-wise manner. In all models, the dependent variable is the change in output price (i.e., $\Delta p_{i,t} = p_{i,t} - p_{i,t-1}$). The robust standard errors that are clustered at market level are reported in parentheses, below estimates.

C.2 Pass-through Rates for $\tau = 14$

Table C.2 reports the pass-through rates when $\tau = 14$ for different aggregation levels of data.

C.3 Non-parametric Test on Excess Market Power

Table C.3 reports average of excess market power broken out by round and by group size. The rows 7 and 8 report the p-values resulting from Wilcoxon signed-rank test on the average of excess market power for small and large shocks, respectively. The null hypothesis is that the average of excess market power during rounds 2 and 5 (or 3 and 4) are equal.

Table C.1: Estimation of asymmetry

	(1)	(2)	(3)	(4)	(5)
Δmc_t	0.497***	0.487***	0.227**	0.224**	0.769***
	(0.050)	(0.051)	(0.066)	(0.066)	(0.070)
Δmc_{t-1}	-0.540***	-0.520***	-0.169	-0.165	-1.265***
	(0.076)	(0.076)	(0.088)	(0.088)	(0.125)
Δmc_{t-2}	0.566***	0.527***	0.362***	0.352**	1.257***
	(0.099)	(0.099)	(0.104)	(0.104)	(0.134)
Δmc_{t-3}	-0.295***	-0.265***	-0.172*	-0.164*	-0.582***
	(0.063)	(0.063)	(0.065)	(0.064)	(0.077)
Δmc_{t-4}	0.063***	0.053**	0.030	0.028	0.102***
	(0.016)	(0.016)	(0.016)	(0.016)	(0.018)
$k = 0$ (c_t)	0.412***	0.432***	0.583***	0.043	0.478***
	(0.106)	(0.109)	(0.140)	(0.144)	(0.118)
$k = 1$ (c_{t-1})	0.152	0.113	-0.026	-0.036	0.004
	(0.098)	(0.100)	(0.124)	(0.123)	(0.118)
$k = 2$ (c_{t-2})	-0.230	-0.152	-0.134	-0.113	-0.084
	(0.122)	(0.125)	(0.141)	(0.140)	(0.140)
$k = 3$ (c_{t-3})	0.161*	0.102	0.082	0.067	0.067
	(0.077)	(0.078)	(0.082)	(0.081)	(0.087)
$k = 4$ (c_{t-4})	-0.051*	-0.031	-0.024	-0.018	-0.008
	(0.021)	(0.021)	(0.020)	(0.020)	(0.022)
$N = 3$		-0.007***	-0.005***	-0.005	-0.010***
		(0.002)	(0.001)	(0.003)	(0.003)
$N = 4$		-0.007***	-0.005***	-0.008***	-0.009***
		(0.001)	(0.001)	(0.002)	(0.002)
$N = 6$		-0.009***	-0.007***	-0.012***	-0.013***
		(0.002)	(0.001)	(0.002)	(0.003)
$N = 10$		-0.008***	-0.006***	-0.007**	-0.011***
		(0.002)	(0.001)	(0.002)	(0.003)
$\Delta p_{-i,t-1}$			0.287***	0.287***	0.347***
			(0.022)	(0.021)	(0.024)
Three-way interaction terms with immediate asymmetry					
$(N = 3) \times c_t$				0.492*	
				(0.192)	
$(N = 4) \times c_t$				0.680***	
				(0.142)	
$(N = 6) \times c_t$				0.828***	
				(0.154)	
$(N = 10) \times c_t$				0.561***	
				(0.143)	
AR(4) included?	No	No	No	No	Yes
N	17130	17130	17130	17130	17130
adj. R^2	0.132	0.133	0.165	0.178	0.301
F-statistic	92.240	70.604	160.112	136.818	135.597

Table C.2: Asymmetry in the pass-through rates after 14 periods

	All	$N > 2$	$N = 2$	$N = 3$	$N = 4$	$N = 6$	$N = 10$
Small shocks							
β_{14}^{-}	0.881	0.957	0.483	1.004	1.002	1.088	0.749
	(0.0543)	(0.0518)	(0.192)	(0.145)	(0.0764)	(0.114)	(0.0816)
β_{14}^{+}	0.599	0.729	-0.0806	0.783	0.815	0.479	0.836
	(0.0518)	(0.0432)	(0.197)	(0.0809)	(0.0818)	(0.110)	(0.0564)
p-value	0.000	0.001	0.108	0.145	0.066	0.000	0.475
Large shocks							
β_{14}^{-}	0.821	0.902	0.398	0.811	0.961	0.878	0.936
	(0.0251)	(0.0205)	(0.0849)	(0.0598)	(0.0393)	(0.0417)	(0.0205)
β_{14}^{+}	0.802	0.866	0.465	0.637	0.978	1.002	0.796
	(0.0297)	(0.0270)	(0.105)	(0.0765)	(0.0450)	(0.0430)	(0.0399)
p-value	0.064	0.042	0.955	0.021	0.331	0.156	0.000
Observations	245	209	36	39	52	48	70

The averages of pass-through rates by differing group sizes are reported. Below averages, standard errors are reported in parentheses. p-values correspond to the result of the Wilcoxon signed-rank test on the equality of pass-through rates for small or large shocks (i.e. $H_0 : \beta_{14}^{+} = \beta_{14}^{-}$).

Table C.3: Non-parametric test on excess market power

Rounds	All Markets	$N > 2$	$N = 2$	$N = 3$	$N = 4$	$N = 6$	$N = 10$
1	0.025	0.021	0.051	0.019	0.020	0.029	0.017
2	0.043	0.037	0.077	0.034	0.033	0.046	0.034
3	0.011	0.008	0.029	-0.012	0.019	0.019	0.004
4	0.048	0.042	0.083	0.038	0.044	0.052	0.038
5	0.030	0.025	0.061	0.026	0.029	0.027	0.020
All Rounds	0.032	0.027	0.060	0.021	0.029	0.034	0.023
Small η (2vs5)	0.000	0.000	0.499	0.115	0.109	0.000	0.000
Large η (3vs4)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Observations	245	209	36	39	52	48	70

Appendix D

Material Presented to Subjects before the *Knowing Me, Knowing You* Experiment

D.1 Instructions

Welcome!

You are about to participate in an experiment on decision-making. In this experiment, you can earn a considerable amount of **money**, which will be paid to you in cash, privately, at the end of the experiment. How much you earn will depend on your decisions, the decisions of other participants, and chance.

Please **do not communicate** with the other participants at any point during the experiment. Make sure that your phone is turned off now.

To make sure that everybody understands the tasks in this experiment, we will begin with some basic instructions. **If you have any questions**, or need assistance of any kind, raise your hand and the experimenter will come and help you.

Basic Instructions

At the beginning of the experiment, the computer will randomly assign each participant to one of **two groups: *Group A* and *Group B***. Once assigned, participants will remain in the same group throughout the entire experiment. There will be the same number of participants in each group.

Participants assigned to the **same group** will each see the **same payoff table and other information** on their computer screen at the beginning of the experiment.

The experiment consists of **2 parts**. For now, we will explain to you what is happening in the first part of the experiment. Once the first part is completed, we will explain to you what is happening in the second part.

The first part of the experiment consists of **40 rounds**. In each round, you will be **randomly paired** with a participant from the other group. You will not know who of the other participants is assigned to which group, and you will also not know with whom

you are randomly paired in any given round. In each round, it is equally likely that you will be paired with any of the participants from the other group. You will never be paired with somebody from your own group.

In each of the 40 rounds, you and the person you are currently paired with will be asked to make a decision on the computer. In what follows, we will explain to you how you can make these decisions.

The Decision Tasks

In each of the 40 rounds, you will be able to **choose one of two actions**. The participant you are paired with will also be able to choose one out of two actions. In each round, everybody will have to **choose an action before seeing the action that the other participant has chosen**.

Below, we show you an example of how a decision task could look like on the computer. In the experiment, you will see a similar table on your computer screen, but with different numbers.

Example of a Payoff Table

The **table below shows the payoffs** associated with each combination of your choice and the choice of the participant you are paired with. This is an example of how a decision task could look like on the computer; please note that the actual numbers you will see in the experiment will be different from those shown in this example. We will now explain to you how you can interpret the numbers in the table.

The **first entry** in each cell (i.e. the number before the comma) represents **your payoff**. The **second entry** in each cell (i.e. the number after the comma) represents the **payoff of the person you are paired with**.

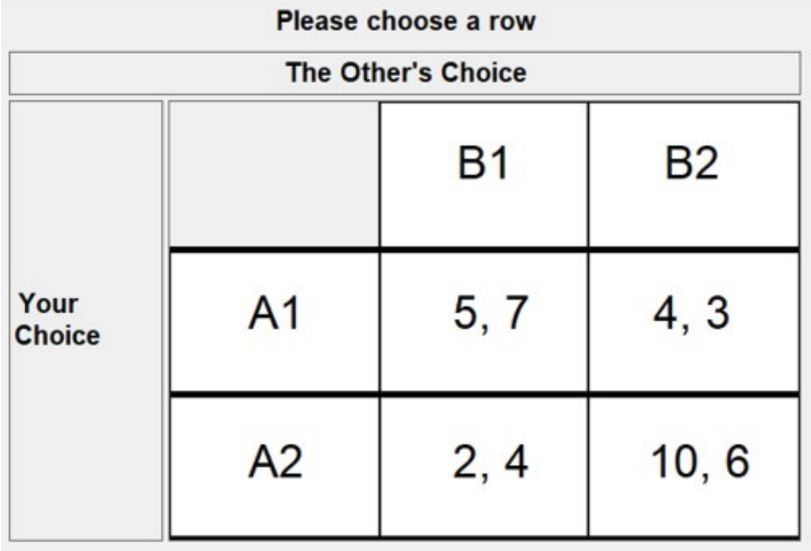

All **cell entries** of the table show the **payoffs that are associated with each combination of your choice and the other participant's choice**:

- For example, if you select "A1" and the other participant selects "B1", you earn 5 Dollars and the other participant earns 7 Dollars.
- As another example, if you select "A2" and the other participant selects "B1", you earn 2 Dollars and the other participant earns 4 Dollars.
- Another example: if you select "A1" and the other participant selects "B2", you earn 4 Dollars and the other participant earns 3 Dollars.
- Another example: if you select "A2" and the other participant selects "B2", you earn 10 Dollars and the other participant earns 6 Dollars.

How to Make Decisions

Suppose that the computer assigned you to *Group A*. In this example, you will be asked to choose either "A1" or "A2". Remember that if you are in *Group A*, then in each round, you will be paired with somebody from *Group B*, and they will be asked to choose

either "B1" or "B2". If you should get assigned to *Group B*, however, then **you** will be asked to choose between "B1" and "B2", and the **other participant** will be asked to choose between "A1" and "A2", as they would be assigned to *Group A*.

In the experiment, you will see a table similar to the example above on your computer screen. **To make a choice**, you will **click on one of the rows in the table**.

Once you select a row, it will change color and a red *SUBMIT* button will appear. Your choice will be finalized once you click on the *SUBMIT* button. After submitting your choice, you will need to wait until the other participant you are paired with has also made their choice. Once you and the participant you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear. **Remember that you will only see the choice of the other participant once you have submitted your own choice.** After each of the 40 rounds, you will see an overview of the choice you made, the choice the other participant made, and your payoffs of the round.

Example of How a Payoff Table Will Look Like in the First Part of the Experiment

In the example above, we showed you an example of a decision task. In each cell of the table above, you could see both your own payoff and the other's payoff for each combination of your and the other participant's choices. In the actual experiment, however, **you will only see your own payoffs**. The payoffs of the other person will be covered.

Here is an example of what a table in the experiment could actually look like:

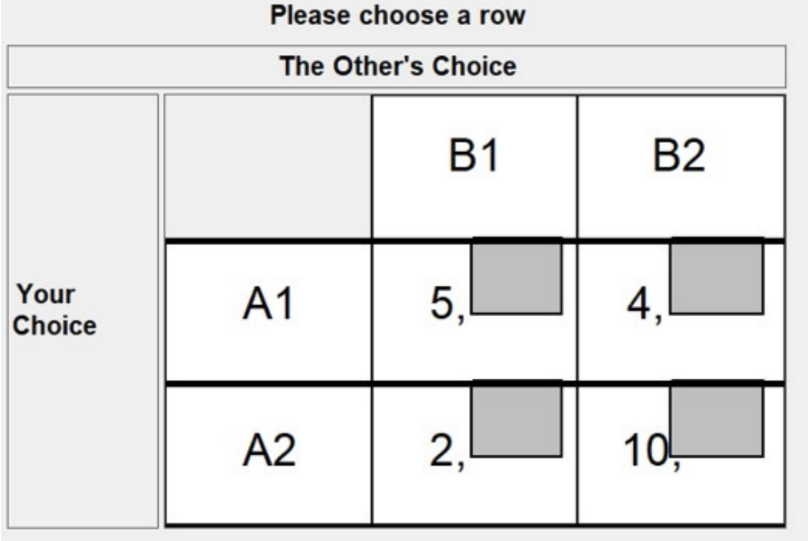

This **table shows your payoffs** associated with each combination of your choice and the choice of the participant you are paired with. It does not show, however, the payoffs of the person you are paired with. The **first entry in each cell** (i.e. the number before the comma sign) **represents your payoff**. For the first part of the experiment, you will never know how much the person you are paired with earns in each combination of your and their choice.

- For example, if you select "A1" and the other participant selects "B1", you earn 5 Dollars, but you don't know how much the other participant earns.
- As another example, if you select "A2" and the other participant selects "B1", you earn 2 Dollars, but you don't know how much the other participant earns.
- Another example: if you select "A1" and the other participant selects "B2", you earn 4 Dollars, but you don't know how much the other participant earns.
- Another example: if you select "A2" and the other participant selects "B2", you earn 10 Dollars, but you don't know how much the other participant earns.

Summary: What It Means to Not See the Other's Payoffs

1. You only **know your own payoffs**, and you know that everybody in your group has the same payoffs.

2. You **do not know** the payoffs of the person you are paired with.

3. This means that you do not know what payoffs participants in the other group are getting. You know, however, that every participant in the other group is getting the same payoffs.

How much will you get paid in the end?

At the end of the experiment, for each participant the computer will randomly select a number between 1 and 80, corresponding to each of the rounds of the experiment. Every participant will get paid, in US Dollars, the amount of their payoff in that particular round, PLUS the show-up fee of 7 dollars. Before that, a short questionnaire will appear on your screen.

Summary

- There are a total of 80 rounds in the experiment, divided into two parts of 40 rounds each.

- You will make a decision in each of these 80 rounds.

- **At the beginning of the experiment**, half of all participants will be randomly assigned to *Group A*, and the other half will be assigned to *Group B*.

- Participants stay assigned to the same group throughout the experiment.

- All *Group A* participants have the same payoffs, and all *Group B* participants have the same payoffs. These payoffs remain the same throughout all 40 rounds of that part of the experiment.
- **In each round**, you will be randomly paired with someone from the other group.
- This means that before each decision round, a new random pair will be formed.

D.2 Comprehension Quiz

Partial Information - Version 1

Comprehension Quiz

To make sure that you understand the instructions of this experiment, please answer the questions below.

Below, you see an example of a decision task, similar to the one that you might encounter in the experiment. In this example, you got assigned to Group A, and the person you are randomly paired with got assigned to Group B.

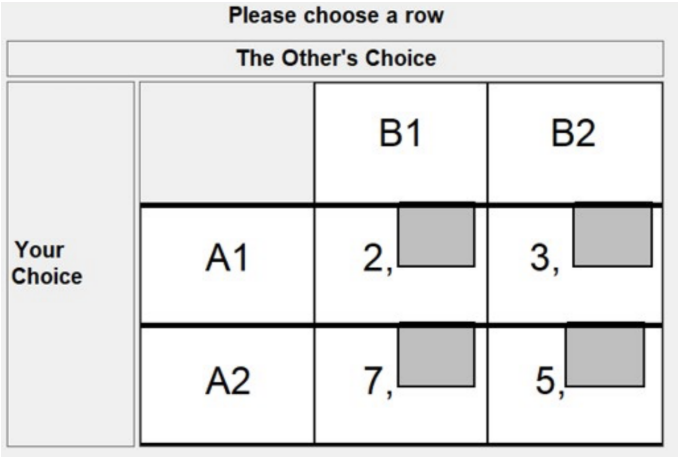

Please answer the following questions. If you don't know the answer for sure, please insert a question mark (?) into the blank space.

1. If you choose "A2" and the other chooses "B2", what is the other's payoff?

2. If you choose "A2" and the other chooses "B1", what is your payoff?

3. If you choose "A1" and the other chooses "B2", what is your payoff?

4. If you choose "A2" and the other chooses "B1", what is the other's payoff?

5. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 5?

6. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 7?

Full Information - Version 1

Comprehension Quiz

To make sure that you understand the instructions of this experiment, please answer the questions below.

Below, you see an example of a decision task, similar to the one that you might encounter in the experiment. In this example, you got assigned to Group A, and the person you are randomly paired with got assigned to Group B.

Please choose a row

	The Other's Choice		
		B1	B2
Your Choice	A1	2, 0	3, 3
	A2	7, 11	5, 3

Please answer the following questions. If you don't know the answer for sure, please insert a question mark (?) into the blank space.

1. If you choose "A2" and the other chooses "B2", what is the other's payoff?

2. If you choose "A2" and the other chooses "B1", what is your payoff?

3. If you choose "A1" and the other chooses "B2", what is your payoff?

4. If you choose "A2" and the other chooses "B1", what is the other's payoff?

5. In this example, which combination of your action and the other's action needs to happen so that you get a payoff of 5?

6. In this example, which combination of your action and the other's action needs to happen so that the other gets a payoff of 3?

D.3 Experimental Interface

In Figure 2.2 we presented screenshots of the experimental interface before, during, and after subjects selected their actions in the PD-Partial treatment. For completion's sake, here we present the corresponding screenshots for the SH-Full treatment:

Figure D.1: Screenshots of Experimental Interface, SH-Full Treatment

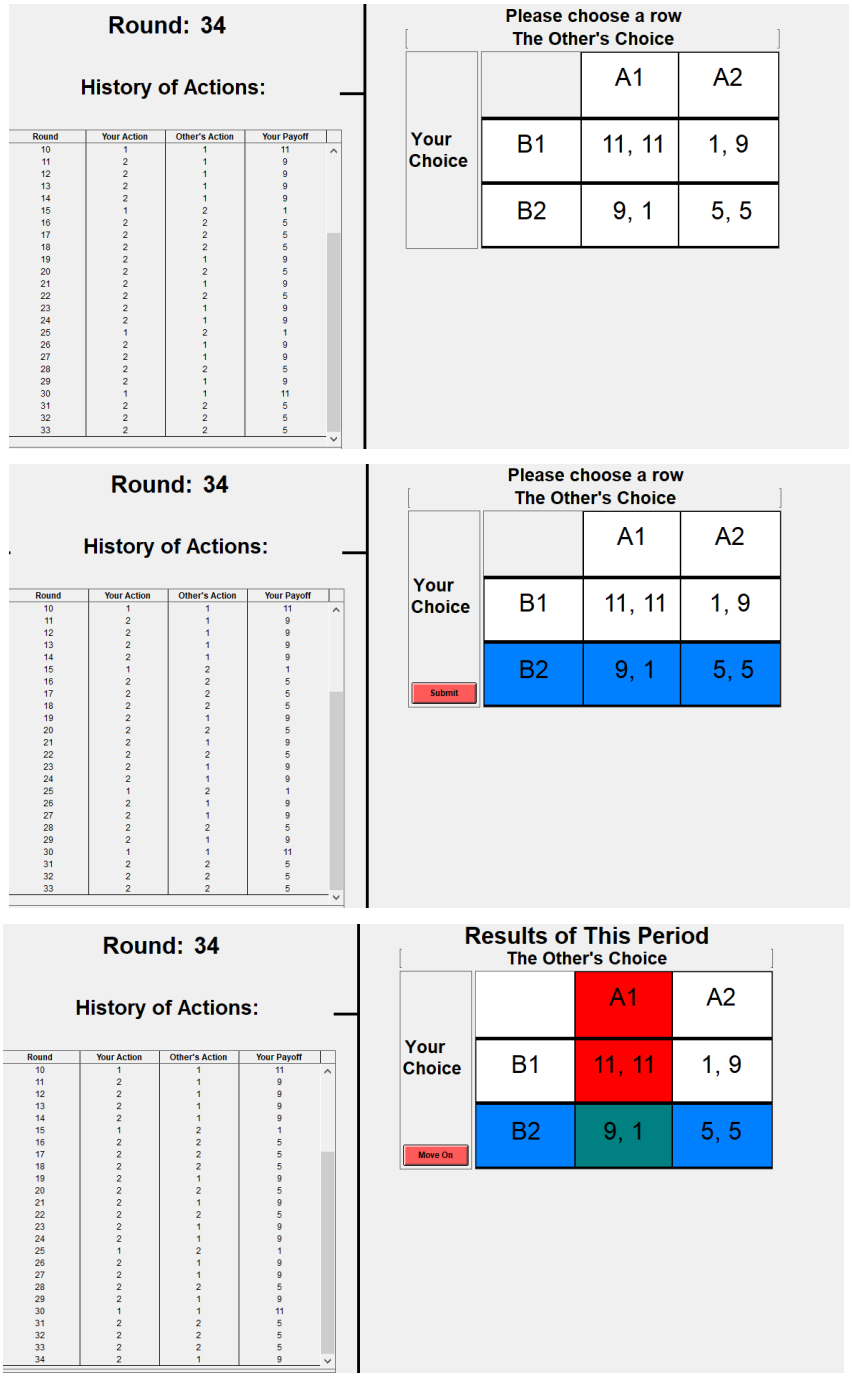

Note: We used a different color scheme for each of the two games in each experiment. We did this to help subjects understand and recall, when they are playing the second game of the experiment, that the current game they are playing is distinct from the first game they have already completed.

Appendix E

Additional Tables, *Knowing Me, Knowing You* Experiment

Table E.1: Share of subjects in the Stag Hunt choosing a mix of actions within 10 percentage points of the mixed strategy Nash equilibrium

	Rounds			
	1-10	11-20	21-30	31-40
a) Partial-information				
Fraction playing within .10 of MSNE	0.27	0.16	0.14	0.07
b) Full-information				
Fraction playing within .10 of MSNE	0.38	0.20	0.18	0.08

Note: Given that the mixed strategy Nash equilibrium is to play action X 66.67% of the time, the fraction of subjects playing within 10 percentage points of the mixed strategy Nash equilibrium (MSNE) equals the total number of subjects playing action X between 56.67% and 76.67% of the time divided by the total number of subjects exposed to that information treatment.

Table E.2: Comparison of Observed and Simulated Data

	(1) Overall	(2) 1-10	(3) 11-20	(4) 21-30	(5) 31-40
a) Stag Hunt, Partial					
Simulated	-0.007	-0.007	-0.010	-0.008	-0.002
	(0.028)	(0.021)	(0.040)	(0.039)	(0.037)
Cluster p-value	0.805	0.755	0.798	0.832	0.946
Bootstrap p-value	0.820	0.778	0.810	0.854	0.961
Outcome mean	0.149	0.228	0.142	0.120	0.107
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	686,400	171,600	171,600	171,600	171,600
b) Stag Hunt, Full					
Simulated	0.032	0.031	0.045	0.028	0.024
	(0.073)	(0.045)	(0.079)	(0.083)	(0.108)
Cluster p-value	0.663	0.496	0.574	0.735	0.825
Bootstrap p-value	0.714	0.573	0.597	0.789	0.916
Outcome mean	0.844	0.881	0.838	0.836	0.822
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	683,120	170,780	170,780	170,780	170,780
c) Prisoner's Dilemma, Partial					
Simulated	-0.002	-0.029	-0.010	0.017	0.016
	(0.006)	(0.011)	(0.012)	(0.010)	(0.010)
Cluster p-value	0.814	0.006	0.432	0.083	0.108
Bootstrap p-value	0.820	0.069	0.588	0.146	0.194
Outcome mean	0.065	0.126	0.065	0.035	0.036
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	687,360	171,840	171,840	171,840	171,840
d) Prisoner's Dilemma, Full					
Simulated	-0.018	-0.102	-0.025	-0.003	0.058
	(0.028)	(0.064)	(0.036)	(0.040)	(0.034)
Cluster p-value	0.522	0.114	0.496	0.931	0.084
Bootstrap p-value	0.577	0.225	0.539	0.960	0.148
Outcome mean	0.232	0.433	0.205	0.176	0.113
Number of clusters	1,006	1,006	1,006	1,006	1,006
N	686,080	171,520	171,520	171,520	171,520

Note: Standard errors presented in parentheses are calculated using the cluster-robust method allowing for correlation between observations within a cluster. The level of clustering is at the session. p-values indicate the p-value from the empirical sampling distribution found with the bootstrapping method.

Table E.3: Comparison of Estimates of λ by Information Treatment

Estimate	$\lambda-$Stag Hunt	$\lambda-$Prisoner's Dilemma
Full-Information Treatment	1.5831 $(1.2159 - 1.8776)$	0.4297 $(0.3347 - 0.5242)$
Partial-Information Treatment	0.6516 $(0.5184 - 0.7673)$	0.7426 $(0.6615 - 0.9146)$
$\underline{H_0 : \lambda_{partial} \geq \lambda_{full}}$ BCa interval test p-val	0.0014	
$\underline{H_0 : \lambda_{full} \geq \lambda_{partial}}$ BCa interval test p-val		0.0368
Mann-Whitney U test	0.0018	0.0018
Kolmogorov-Smirnov test	0.0000	0.0000

Note: One-sided BCa interval tests conducted using $B = 10,000$ bootstrap iterations, with bootstrapping being performed separately for full- and partial-information observations

Table E.4: Comparison of Estimates of ϕ by Information Treatment

Estimate	$\phi-$Stag Hunt	$\phi-$Prisoner's Dilemma
Full-Information Treatment	0.9649 $(0.9011 - 1.0223)$	0.8011 $(0.3009 - 0.9398)$
Partial-Information Treatment	0.8290 $(0.7112 - 0.9911)$	0.8769 $(0.6019 - 85.013)$
$\underline{H_0 : \phi_{partial} \geq \phi_{full}}$ BCa interval test p-val	0.0624	
$\underline{H_0 : \phi_{full} \geq \phi_{partial}}$ BCa interval test p-val		0.3005
Mann-Whitney U test	0.0017	0.0014
Kolmogorov-Smirnov test	0.0000	0.0000

Note: One-sided BCa interval tests conducted using $B = 10,000$ bootstrap iterations, with bootstrapping being performed separately for full- and partial-information observations

Appendix F

Sensitivity Analysis for Study on Risk Aversion and Perception

The survey queried subjects about their perceptions of contracting different cancers, allowing responses ranging from 0 (there is no risk of me getting this cancer) to very much above average (risk is 5 times or more than the average). The responses are used to construct the baseline optimism index using a logarithmic functional form. As such, the zero response is not valid and some number close to zero must be used instead for the "no risk" response. The models reported in the paper use a "no risk" value of 0.0001. We constructed two alternative baseline optimism indexes: one with the zero replaced with 0.01 and the other assuming a value of 0.001.

Tables F.1 and F.2 give the results of the risky behavior models for the baseline optimism indexes set to 0.01 and 0.001. Tables F.3 and F.4 report the models for prevention behavior. The coefficients and statistical significance for the 0.001 models are very close to that of the 0.0001 models. The most notable changes are that the marginal effects of baseline optimism on behavior are smaller, in absolute value, in most of the 0.0001 models than in the 0.01 models. This is not surprising since the lower value increases the

Table F.1: Sensitivity Analysis of Risky Behavior (baseline optimism index = 0.01)

	Current Smoker		Ever Smoker		Red Meat		High BMI		Heavy Drinker	
Variable	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)
Risk Aversion	-0.007 (0.346)	-0.014 (0.05)	-0.015 (0.111)	-0.021 (0.020)	0.002 (0.818)	0.001 (0.87)	-0.009 (0.348)	-0.010 (0.276)	-0.008 (0.206)	-0.054 (0.053)
Cognitive Ability	-0.043 (0.000)	----	-0.032 (0.086)	----	-0.006 (0.707)	----	-0.007 (0.680)	----	-0.018 (0.075)	----
Baseline Optimism	0.027 (0.036)	----	0.015 (0.364)	----	0.029 (0.076)	----	0.034 (0.024)	----	-0.0003 (0.976)	----
Control Optimism	-0.080 (0.053)	----	-0.063 (0.226)	----	-0.049 (0.332)	----	0.100 (0.054)	----	-0.084 (0.020)	----
Male	0.042 (0.288)	-0.026 (0.498)	0.073 (0.137)	0.022 (0.634)	0.085 (0.071)	0.078 (0.079)	0.249 (0.000)	0.248 (0.000)	0.019 (0.556)	-0.112 (0.419)
Age	0.003 (0.020)	3.00E-05 (0.975)	0.009 (0.000)	0.006 (0.000)	-3.99E-04 (0.826)	-0.001 (0.449)	0.002 (0.248)	0.002 (0.328)	-0.002 (0.113)	-0.018 (0.000)
Income ($000)	-0.002 (0.008)	-0.003 (0.000)	-0.002 (0.011)	-0.002 (0.002)	-0.001 (0.356)	-0.001 (0.216)	-0.001 (0.130)	-0.001 (0.163)	-0.0009 (0.120)	-0.006 (0.018)
College	-0.103 (0.007)	-0.125 (0.001)	-0.174 (0.000)	-0.187 (0.000)	-0.052 (0.276)	-0.061 (0.189)	-0.059 (0.201)	-0.079 (0.076)	0.009 (0.779)	0.031 (0.832)
Married	0.035 (0.414)	0.024 (0.583)	-0.007 (0.893)	-0.008 (0.874)	0.009 (0.860)	0.017 (0.732)	0.085 (0.092)	0.093 (0.068)	0.043 (0.239)	0.118 (0.452)
Uninsured	0.002 (0.971)	-0.035 (0.548)	-0.046 (0.537)	-0.071 (0.327)	-0.075 (0.294)	-0.068 (0.335)	-0.089 (0.188)	-0.010 (0.129)	0.003 (0.950)	-0.099 (0.674)
Log Likelihood	-228.540	-245.760	-300.840	-308.830	-317.770	-322.970	-277.630	-298.83	-181.984	-194.523

Note: probit Models of Risky Behavior with Dependent Variable Equal to 1 if the subject engages in the behavior and 0 otherwise. No Risk Value = 0.01.

average of the index and therefore the model coefficients. The only model that changes substantially in terms of the significance of the marginal effects is the Ever Smoker model. When the larger value of 0.01 is used to create the index, the marginal effect of baseline optimism becomes insignificant, whereas it is significant when the smaller value is used. Taken together, it appears that using the larger value exaggerates the marginal effects of baseline optimism on behavior. Still, the observation that baseline optimism significantly correlates with behavior is robust to specification of the values underlying the index.

Table F.2: Sensitivity Analysis of Risky Behavior (baseline optimism index = 0.001)

	Current Smoker		Ever Smoker		Red Meat		High BMI		Heavy Drinker	
Variable	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)
Risk Aversion	-0.008 (0.269)	-0.014 (0.05)	-0.017 (0.079)	-0.021 (0.020)	0.001 (0.897)	0.001 (0.87)	-0.010 (0.283)	-0.010 (0.276)	-0.008 (0.192)	-0.054 (0.053)
Cognitive Ability	-0.046 (0.000)	-----	-0.036 (0.019)	-----	-0.008 (0.611)	-----	-0.010 (0.571)	-----	-0.018 (0.062)	-----
Baseline Optimism	0.024 (0.001)	-----	0.026 (0.006)	-----	0.021 (0.022)	-----	0.023 (0.006)	-----	0.006 (0.283)	-----
Control Optimism	-0.094 (0.023)	-----	-0.075 (0.153)	-----	-0.061 (0.233)	-----	0.089 (0.092)	-----	-0.086 (0.017)	-----
Male	0.040 (0.313)	-0.026 (0.493)	0.069 (0.160)	0.022 (0.634)	0.083 (0.079)	0.078 (0.079)	0.247 (0.000)	0.248 (0.000)	0.018 (0.567)	-0.112 (0.419)
Age	0.003 (0.043)	3.00E-05 (0.975)	0.008 (0.000)	0.006 (0.000)	-0.0008 (0.649)	-0.001 (0.449)	0.002 (0.368)	0.002 (0.328)	-0.002 (0.081)	-0.018 (0.000)
Income ($000)	-0.002 (0.010)	-0.003 (0.000)	-0.002 (0.012)	-0.002 (0.002)	-0.0006 (0.418)	-0.001 (0.216)	-0.001 (0.148)	-0.001 (0.163)	-0.0009 (0.108)	-0.006 (0.018)
College	-0.095 (0.013)	-0.125 (0.001)	-0.167 (0.001)	-0.187 (0.000)	-0.047 (0.327)	-0.061 (0.189)	-0.053 (0.250)	-0.079 (0.076)	0.012 (0.708)	0.031 (0.832)
Married	0.029 (0.491)	0.024 (0.583)	-0.016 (0.764)	-0.008 (0.874)	0.005 (0.926)	0.017 (0.732)	0.080 (0.112)	0.093 (0.063)	0.041 (0.266)	0.118 (0.453)
Uninsured	-0.012 (0.835)	-0.035 (0.549)	-0.064 (0.389)	-0.071 (0.327)	-0.0788 (0.216)	-0.068 (0.335)	-0.104 (0.121)	-0.010 (0.129)	0.049 (0.989)	-0.099 (0.674)
Log Likelihood	-225.950	-245.760	-300.840	-308.830	-317.770	-322.970	-276.410	-298.83	-181.42	-194.523

Note: probit Models of Risky Behavior with Dependent Variable Equal to 1 if the subject engages in the behavior and 0 otherwise. No Risk Value = 0.001.

Table F.3: Sensitivity Analysis of Prevention Behavior (baseline optimism index = 0.01)

	Physically Active		Aspirin Daily		Multivitamin	
Variable	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)
Risk Aversion	-0.004 (0.657)	-0.005 (0.572)	-0.003 (0.178)	-0.0003 (0.828)	0.016 (0.070)	0.013 (0.137)
Cognitive Ability	-0.011 (0.541)	-----	-0.013 (0.007)	-----	0.0008 (0.967)	-----
Baseline Optimism	-0.077 (0.000)	-----	-0.010 (0.033)	-----	-0.010 (0.000)	-----
Control Optimism	0.029 (0.572)	-----	-0.006 (0.612)	-----	0.047 (0.364)	-----
Male	0.118 (0.008)	0.113 (0.009)	-0.005 (0.659)	-0.001 (0.853)	-0.057 (0.207)	-0.045 (0.310)
Age	-0.003 (0.094)	-0.003 (0.131)	0.001 (0.137)	0.001 (0.006)	0.004 (0.046)	0.003 (0.114)
Income ($000)	0.001 (0.461)	0.001 (0.504)	-0.001 (0.052)	-0.001 (0.023)	0.0006 (0.384)	0.0004 (0.611)
College	0.118 (0.008)	0.123 (0.005)	-0.008 (0.502)	-0.004 (0.591)	0.010 (0.034)	0.110 (0.014)
Married	-0.033 (0.510)	-0.049 (0.314)	0.016 (0.248)	0.016 (0.228)	-0.018 (0.713)	-0.028 (0.562)
Uninsured	0.058 (0.392)	0.035 (0.606)	(omitted)	(omitted)	0.119 (0.110)	0.093 (0.191)
Log Likelihood	-280.54	-294.47	-44.02	-50.52	-278.22	-298.83

Note: probit Models of Prevention Behavior with Dependent Variable Equal to 1 if the subject engages in the behavior and 0 otherwise. No Risk Value = 0.01.

Table F.4: Sensitivity Analysis of Prevention Behavior (baseline optimism index = 0.001)

	Physically Active		Aspirin Daily		Multivitamin	
Variable	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)	dy/dx (p-value)
Risk Aversion	-0.004 (0.689)	-0.005 (0.572)	-0.003 (0.229)	-0.0003 (0.828)	0.016 (0.068)	0.013 (0.137)
Cognitive Ability	-0.004 (0.822)	-----	-0.012 (0.013)	-----	0.006 (0.745)	-----
Baseline Optimism	-0.022 (0.011)	-----	-0.006 (0.031)	-----	-0.035 (0.001)	-----
Control Optimism	0.047 (0.362)	-----	0.0002 (0.998)	-----	0.077 (0.127)	-----
Male	0.120 (0.007)	0.113 (0.009)	-0.003 (0.743)	-0.001 (0.853)	-0.052 (0.263)	-0.045 (0.310)
Age	-0.003 (0.133)	-0.003 (0.131)	0.001 (0.052)	0.001 (0.006)	0.005 (0.029)	0.003 (0.114)
Income ($000)	0.0003 (0.420)	0.001 (0.504)	-0.001 (0.034)	-0.001 (0.023)	0.0003 (0.695)	0.0004 (0.611)
College	0.117 (0.009)	0.123 (0.005)	-0.006 (0.580)	-0.004 (0.591)	0.103 (0.028)	0.110 (0.014)
Married	-0.033 (0.506)	-0.049 (0.314)	0.014 (0.262)	0.016 (0.228)	-0.027 (0.581)	-0.028 (0.562)
Uninsured	0.062 (0.351)	0.035 (0.606)	(omitted)	(omitted)	0.126 (0.097)	0.093 (0.191)
Log Likelihood	-289.13	-294.47	-55.55	-50.52	-287.32	-298.83

Note: probit Models of Prevention Behavior with Dependent Variable Equal to 1 if the subject engages in the behavior and 0 otherwise. No Risk Value = 0.001.

Appendix G

Scope Tests for Model Variables: Risk Aversion and Perception Study

Due to the stated preference nature of the survey, it is important to gauge whether subjects' responses are consistent with rational choices. A number of rationality tests have been proposed for stated preference surveys, but the scope test is likely the most popular (31). The current survey contained a single-bounded contingent valuation (CV) question addressing willingness to pay for insurance that covered all monetary costs associated with a cancer diagnosis. A companion paper (Riddel and Hales, 2018)[103] uses the results of the contingent valuation question to investigate adverse and advantageous selection in the market for cancer insurance. However, the CV data also provides for a test of scope. Economic theory indicates demand for insurance, hence willingness to pay for the full coverage should be 1) increasing in income if cancer insurance is a normal good, and 2) increasing in the insured risk. Moreover, the coefficient of the offered premium should be negative, indicating that the higher the cost of the insurance, the less likely the subject will purchase the coverage.

The results of probit models relating demand for coverage as a function of the offered

premium, perceived cancer risks, and income by gender are reported in Table 3.4. The dependent variable is equal to one if the subject agreed to pay the offered premium and zero otherwise. Note that the ratio of the coefficient of the perceived risk and income variables to the coefficient of the premium amount is equal to the marginal willingness to pay for that variable. Due to high correlation between the individual perceived cancer risks, the equations are estimated separately.

The coefficient of the offered premium is negative and statistically significant in all ten models, indicating that the higher the premium offered to the subject, the less likely they are to agree to purchase the hypothetical cancer insurance coverage. The coefficients of the perceived cancer risk variables are all positive and all but one are statistically significant for a p-value of 0.05 or better. Thus, willingness to pay for full coverage is increasing in the perceived risk of contracting the different cancers. Finally, willingness to pay is also increasing in income, supporting the hypothesis that cancer insurance is a normal good. Taken together, the scope tests indicate that people respond consistently to the survey.

Table G.1: Scope Tests for the Premium, Individual Perceived Cancer Risks, and Income

Variable	Females				
Premium	-0.004 (0.000)	-0.011 (0.000)	-0.008 (0.000)	-0.007 (0.000)	-0.009 (0.000)
Perceived Bladder Risk		0.633 (0.000)			
Perceived Breast Risk			0.162 (0.001)		
Perceived Colon Risk				0.212 (0.001)	
Income					0.009 (0.000)
Variable	Males				
Premium	-0.006*(0.000)	-0.010 (0.000)	-0.009 (0.000)	-0.008 (0.000)	-0.010 (0.000)
Perceived Bladder Risk		0.360 (0.014)			
Perceived Prostate Risk			0.246 (0.007)		
Perceived Colon Risk				0.090 (0.237)	
Income					0.008 (0.005)

Note: probit model where dependent variable equals one if the subject agreed to pay the offered premium. p-values in parentheses.

Bibliography

[1] M. Bulutay, D. Hales, P. Julius, and W. Tasch, *Imperfect tacit collusion and asymmetric price transmission*, *Journal of Economic Behavior & Organization* **192** (2021) 584–599.

[2] H. Alshaikhmubarak, D. Hales, M. Kogelnik, M. Schwarz, and K. Strauss, *Knowing me, knowing you: An experiment on mutual payoff information and strategic uncertainty*, *Available at SSRN 3915018* (2021).

[3] M. Riddel and D. Hales, *Predicting cancer-prevention behavior: Disentangling the effects of risk aversion and risk perceptions*, *Risk Analysis* **38** (2018), no. 10 2161–2177.

[4] J. Meyer and S. von Cramon-Taubadel, *Asymmetric price transmission: a survey*, *Journal of Agricultural Economics* **55** (2004), no. 3 581–611.

[5] G. Frey and M. Manera, *Econometric models of asymmetric price transmission*, *Journal of Economic Surveys* **21** (2007), no. 2 349–415.

[6] S. Borenstein, A. C. Cameron, and R. Gilbert, *Do gasoline prices respond asymmetrically to crude oil price changes?*, *Quarterly Journal of Economics* **112** (1997), no. 1 305–339.

[7] E. J. Green and R. H. Porter, *Noncooperative collusion under imperfect price information*, *Econometrica* **52** (1984), no. 1 87–100.

[8] J.-P. Loy, C. R. Weiss, and T. Glauben, *Asymmetric cost pass-through? empirical evidence on the role of market power, search and menu costs*, *Journal of Economic Behavior & Organization* **123** (2016) 184–192.

[9] S. Peltzman, *Prices rise faster than they fall*, *Journal of Political Economy* **108** (2000), no. 3 466–502.

[10] S. Huck, H.-T. Normann, and J. Oechssler, *Two are few and four are many: number effects in experimental oligopolies*, *Journal of Economic Behavior & Organization* **53** (2004), no. 4 435–446.

[11] M. Dufwenberg and U. Gneezy, *Price competition and market concentration: an experimental study*, *International Journal of Industrial Organization* **18** (2000), no. 1 7–22.

[12] M. A. Fonseca and H.-T. Normann, *Explicit vs. tacit collusion—the impact of communication in oligopoly experiments*, *European Economic Review* **56** (2012), no. 8 1759–1772.

[13] R.-C. Bayer and C. Ke, *What causes rockets and feathers? an experimental investigation*, *Journal of Economic Behavior & Organization* **153** (2018) 223–237.

[14] R. M. Isaac and J. M. Walker, *Group size effects in public goods provision: The voluntary contributions mechanism*, *The Quarterly Journal of Economics* **103** (1988), no. 1 179–199.

[15] N. Hanaki and A. Masiliūnas, *Market concentration and incentives to collude in cournot oligopoly experiments*, *ISER DP* (2021), no. 1131.

[16] R. W. Bacon, *Rockets and feathers: the asymmetric speed of adjustment of uk retail gasoline prices to cost changes*, *Energy Economics* **13** (1991), no. 3 211–218.

[17] J. Perdiguero-García, *Symmetric or asymmetric oil prices? a meta-analysis approach*, *Energy policy* **57** (2013) 389–397.

[18] E. Maskin and J. Tirole, *A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles*, *Econometrica* **56** (1988), no. 3 571–599.

[19] D. P. Byrne and N. De Roos, *Learning to coordinate: A study in retail gasoline*, *American Economic Review* **109** (2019), no. 2 591–619.

[20] S. P. Brown and M. K. Yucel, *Gasoline and crude oil prices: why the asymmetry?*, *Economic & Financial Review* (2000), no. Third Quarter 23–29.

[21] R. Benabou and R. Gertner, *Search with learning from prices: does increased inflationary uncertainty lead to higher markups?*, *Review of Economic Studies* **60** (1993), no. 1 69–93.

[22] H. Yang and L. Ye, *Search with learning: understanding asymmetric price adjustments*, *RAND Journal of Economics* **39** (2008), no. 2 547–564.

[23] P. B. Reagan and M. L. Weitzman, *Asymmetries in price and quantity adjustments by the competitive firm*, *Journal of Economic Theory* **27** (1982), no. 2 410–420.

[24] S. Borenstein and A. Shepard, *Sticky prices, inventories, and market power in wholesale gasoline markets*, tech. rep., National Bureau of Economic Research, 1996.

[25] L. Ball and N. G. Mankiw, *Asymmetric price adjustment and economic fluctuations*, *Economic Journal* **104** (1994), no. 423 247–261.

[26] S. Ahrens, I. Pirschel, and D. J. Snower, *A theory of price adjustment under loss aversion*, *Journal of Economic Behavior & Organization* **134** (2017) 78–95.

[27] J. Duffy, *Macroeconomics: a survey of laboratory research*, *Handbook of Experimental Economics* **2** (2016) 1–90.

[28] C. Cornand and F. Heinemann, *Experiments in macroeconomics: methods and applications*, in *Handbook of Research Methods and Applications in Experimental Economics*. Edward Elgar Publishing, 2019.

[29] C. A. Deck and B. J. Wilson, *Experimental gasoline markets*, *Journal of Economic Behavior & Organization* **67** (2008), no. 1 134–149.

[30] T. N. Cason and D. Friedman, *A laboratory study of customer markets*, *BE Journal of Economic Analysis & Policy* **2** (2002), no. 1 1–43.

[31] E. Fehr and J.-R. Tyran, *Does money illusion matter?*, *American Economic Review* **91** (2001), no. 5 1239–1262.

[32] P. Duersch and T. A. Eife, *Price competition in an inflationary environment*, *Journal of Monetary Economics* **104** (2019) 48–66.

[33] J. Morgan, H. Orzen, and M. Sefton, *An experimental study of price dispersion*, *Games and Economic Behavior* **54** (2006), no. 1 134–158.

[34] K. Abbink and J. Brandts, *Pricing in bertrand competition with increasing marginal costs*, *Games and Economic Behavior* **63** (2008), no. 1 1–31.

[35] H. Orzen, *Counterintuitive number effects in experimental oligopolies*, *Experimental Economics* **11** (2008), no. 4 390–401.

[36] D. Davis, *Pure numbers effects, market power, and tacit collusion in posted offer markets*, *Journal of Economic Behavior & Organization* **72** (2009), no. 1 475–488.

[37] N. Horstmann, J. Krämer, and D. Schnurr, *Number effects and tacit collusion in experimental oligopolies*, *Journal of Industrial Economics* **66** (2018), no. 3 650–700.

[38] L. Fiala and S. Suetens, *Transparency and cooperation in repeated dilemma games: a meta study*, *Experimental Economics* **20** (2017), no. 4 755–771.

[39] A. Deaton and J. Muellbauer, *An almost ideal demand system*, *American Economic Review* **70** (1980), no. 3 312–326.

[40] E. Zhelobodko, S. Kokovin, M. Parenti, and J.-F. Thisse, *Monopolistic competition: beyond the constant elasticity of substitution*, *Econometrica* **80** (2012), no. 6 2765–2784.

[41] Y. Chen and M. Riordan, *Price and variety in the spokes model*, *Economic Journal* **117** (July, 2007) 897–921.

[42] I. Bos and D. Vermeulen, *On the microfoundation of linear oligopoly demand*, *The B.E. Journal of Theoretical Economics* (2020) 20190114.

[43] S. C. Salop, *Monopolistic competition with outside goods*, *Bell Journal of Economics* **10** (1979), no. 1 141–156.

[44] G. W. Evans and S. Honkapohja, *Learning and expectations in macroeconomics*. Princeton University Press, 2001.

[45] U. Fischbacher, *z-tree: Zurich toolbox for ready-made economic experiments*, *Experimental economics* **10** (2007), no. 2 171–178.

[46] B. Greiner, *Subject pool recruitment procedures: organizing experiments with orsee*, *Journal of the Economic Science Association* **1** (2015), no. 1 114–125.

[47] J. F. Muth, *Rational expectations and the theory of price movements*, *Econometrica* (1961) 315–335.

[48] A. P. Lerner, *The concept of monopoly and the measurement of monopoly power*, *Review of Economic Studies* **1** (1934), no. 3 157–175.

[49] C. Hommes, J. Sonnemans, J. Tuinstra, and H. Van de Velden, *Coordination of expectations in asset pricing experiments*, *Review of Financial Studies* **18** (2005), no. 3 955–980.

[50] K. Cooper, H. Schneider, and M. Waldman, *Limited rationality and the strategic environment: Further evidence from a pricing game*, *Journal of Behavioral and Experimental Economics* **90** (2021) 101632.

[51] D. Stahl and P. Wilson, *Experimental evidence on players? Models of other players*, *Journal of Economic Behavior & Organization* **25** (1994), no. 3 309–327.

[52] R. Nagel, *Unraveling in guessing games: An experimental study*, *American Economic Review* **85** (1995), no. 5 1313–1326.

[53] D. T. Knoepfle, J. T.-y. Wang, and C. F. Camerer, *Studying learning in games using eye-tracking*, *Journal of the European Economic Association* **7** (2009), no. 2-3 388–398.

[54] L. Polonio and G. Coricelli, *Testing the level of consistency between choices and beliefs in games using eye-tracking*, *Games and Economic Behavior* **113** (2019) 566–586.

[55] J. Nash, *Non-cooperative games*, *Ph.D. Dissertation, Princeton University* (1950).

[56] D. Fudenberg and D. K. Levine, *Self-confirming equilibrium and the Lucas critique*, *Journal of Economic Theory* **144** (2009), no. 6 2354–2371.

[57] S. Hart and A. Mas-Colell, *Stochastic uncoupled dynamics and Nash Equilibrium*, *Games and Economic Behavior* **57** (2006), no. 2 286–303.

[58] D. Foster and H. Young, *Regret testing leads to Nash Equilibrium*, *Theoretical Economics* (2006).

[59] H. P. Young, *Learning by trial and error*, *Games and Economic Behavior* **65** (2009), no. 2 626–643.

[60] Y. Babichenko, *Uncoupled automata and pure Nash equilibria*, *International Journal of Game Theory* **39** (2010), no. 3 483–502.

[61] J. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.

[62] C. Camerer and T.-H. Ho, *Experience-weighted attraction learning in normal form games*, *Econometrica* **67** (1999), no. 4 827–874.

[63] A. Brandenburger, *Strategic and structural uncertainty in games*, in *Wise Choices: Decisions, Games, and Negotiations* (R. Zeckhauser, R. Keeney, and J. Sebenius, eds.), ch. 13, pp. 221–232. Harvard Business School Press, Boston, 1996.

[64] D. Mookherjee and B. Sopher, *Learning behavior in an experimental Matching Pennies game*, *Games and Economic Behavior* **7** (1994), no. 1 62–91.

[65] J. C. Cox, J. Shachat, and M. Walker, *An experiment to evaluate Bayesian learning of Nash Equilibrium play*, *Games and Economic Behavior* **33** (2001) 11–33.

[66] N. Feltovich and S. H. Oda, *Effect of matching mechanism on learning in games played under limited information*, *Pacific Economic Review* **3** (2014) 260–277.

[67] J. Oechssler and B. Schipper, *Can you guess the game you are playing?*, *Games and Economic Behavior* **43** (2003) 137–152.

[68] R. D. McKelvey and T. R. Palfrey, "Playing in the dark: Information, learning, and coordination in repeated games." Working paper, 2001.

[69] D. N. Danz, D. Fehr, and D. Kübler, *Information and beliefs in a repeated normal-form game*, *Experimental Economics* **15** (2012), no. 4 622–640.

[70] A. Nicklisch, *Learning strategic environments: an experimental study of strategy formation and transfer*, *Theory and Decision* **71** (2011), no. 4 539–558.

[71] O. Nikolaychuk, "Does it pay to know more in games of incomplete information?." Working paper, 2012.

[72] J. Andreoni, Y.-K. Che, and J. Kim, *Asymmetric information about rivals' types in standard auctions: An experiment*, *Games and Economic Behavior* **59** (2007), no. 2 240–259.

[73] D. Friedman, S. Huck, R. Oprea, and S. Weidenholzer, *From imitation to collusion: Long-run learning in a low-information environment*, *Journal of Economic Theory* **155** (2015) 185–205.

[74] S. Huck, J. Leutgeb, and R. Oprea, *Payoff information hampers the evolution of cooperation*, *Nature Communications* **8** (2017) 1–5.

[75] H. H. Nax, M. N. Burton-Chellew, S. A. West, and H. P. Young, *Learning in a black box*, *Journal of Economic Behavior & Organization* **127** (2016) 1–15.

[76] R. Battalio, L. Samuelson, and J. Van Huyck, *Optimization incentives and coordination failure in laboratory Stag Hunt games*, *Econometrica* **69** (2001), no. 3 749–764.

[77] D. Schmidt, R. Shupp, J. M. Walker, and E. Ostrom, *Playing safe in coordination games: the roles of risk dominance, payoff dominance, and history of play*, *Games and Economic Behavior* **42** (2003), no. 2 281–299.

[78] D. Dubois, M. Willinger, and P. Van Nguyen, *Optimization incentive and relative riskiness in experimental Stag-Hunt games*, *International Journal of Game Theory* **41** (2012), no. 2 369–380.

[79] R. Kendall, *Decomposing coordination failure in Stag Hunt games*, . Working paper.

[80] M. Embrey, G. R. Fréchette, and S. Yuksel, *Cooperation in the finitely repeated Prisoner's Dilemma*, *The Quarterly Journal of Economics* **133** (2017), no. 1 509–551.

[81] R. Ghidoni, B. L. Cleave, and S. Suetens, *Perfect and imperfect strangers in social dilemmas*, *European Economic Review* **116** (2019) 148–159.

[82] J. Duffy and D. Fehr, *Equilibrium selection in similar repeated games: experimental evidence on the role of precedents*, *Experimental Economics* **21** (2018), no. 3 573–600.

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

[84] B. Greiner, *Subject pool recruitment procedures: Organizing experiments with ORSEE*, *Journal of the Economic Science Association* **1** (2015) 114–125.

[85] Y. Nyarko and A. Schotter, *An experimental study of belief learning using elicited beliefs*, *Econometrica* **70** (December, 2003) 971–1005.

[86] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap*. Chapman and Hall, 1993.

[87] A. Agresti and B. A. Coull, *Approximate is better than "exact" for interval estimation of binomial proportions*, *The American Statistician* **52** (1998), no. 2 119–126.

[88] L. D. Brown, T. Cai, and A. DasGupta, *Interval estimation for a binomial proportion*, *Institute of Mathematical Statistics* **16** (2001), no. 2.

[89] E. Calford and R. Oprea, *Continuity, inertia, and strategic uncertainty: A test of the theory of continuous time games*, *Econometrica* **85** (2017), no. 3.

[90] G. Charness, L. Rigotti, and A. Rustichini, *Social surplus determines cooperation rates in the one-shot Prisoner's Dilemma*, *Games and Economic Behavior* **100** (2016) 113–124.

[91] V. P. Crawford, U. Gneezy, and Y. Rottenstreich, *The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures*, *American Economic Review* **98** (2008), no. 4 1443–58.

[92] L. R. Anderson and J. M. Mellor, *Predicting health behaviors with an experimental measure of risk preference*, *Journal of health economics* **27** (2008), no. 5 1260–1274.

[93] R. B. Barsky, F. T. Juster, M. S. Kimball, and M. D. Shapiro, *Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study*, *The Quarterly Journal of Economics* **112** (1997), no. 2 537–579.

[94] D. M. Dave and H. Saffer, *Risk tolerance and alcohol demand among adults and older adults*, 2007.

[95] T. Dohmen, A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner, *Individual risk attitudes: Measurement, determinants, and behavioral consequences*, *Journal of the european economic association* **9** (2011), no. 3 522–550.

[96] A. Khwaja, F. Sloan, and S. Chung, *The relationship between individual expectations and behaviors: Mortality expectations and smoking decisions*, *Journal of Risk and Uncertainty* **35** (2007), no. 2 179–201.

[97] P. Lundborg and B. Lindgren, *Risk perceptions and alcohol consumption among young people*, *Journal of Risk and Uncertainty* **25** (2002), no. 2 165–183.

[98] P. Lundborg and B. Lindgren, *Do they know what they are doing? risk perceptions and smoking behaviour among swedish teenagers*, *Journal of Risk and Uncertainty* **28** (2004), no. 3 261–286.

[99] W. K. Viscusi, *Do smokers underestimate risks?*, *Journal of political Economy* **98** (1990), no. 6 1253–1269.

[100] W. K. Viscusi, *Age variations in risk perceptions and smoking decisions*, *The Review of Economics and Statistics* (1991) 577–588.

[101] W. K. Viscusi, I. Carvalho, F. Antoñanzas, J. Rovira, F. J. Braña, and F. Portillo, *Smoking risks in spain: Part iii—determinants of smoking behavior*, *Journal of Risk and Uncertainty* **21** (2000), no. 2 213–234.

[102] M. Riddel and W. D. Shaw, *A theoretically-consistent empirical model of non-expected utility: An application to nuclear-waste transport*, *Journal of Risk and Uncertainty* **32** (2006), no. 2 131–150.

[103] M. Riddel and D. Hales, *Risk misperceptions and selection in insurance markets: An application to demand for cancer insurance*, *Journal of Risk and Insurance* **85** (2018), no. 3 749–785.

[104] G. A. Colditz, K. Atwood, K. Emmons, R. Monson, W. Willett, D. Trichopoulos, and D. Hunter, *Harvard report on cancer prevention volume 4: Harvard cancer risk index*, *Cancer causes & control* **11** (2000), no. 6 477–488.

[105] C. A. Holt and S. K. Laury, *Risk aversion and incentive effects*, *American economic review* **92** (2002), no. 5 1644–1655.

[106] J. Spinnewijn, *Insurance and perceptions: How to screen optimists and pessimists*, *The Economic Journal* **123** (2013), no. 569 606–633.

[107] K. G. Carman and P. Kooreman, *Probability perceptions and preventive health care*, *Journal of Risk and Uncertainty* **49** (2014), no. 1 43–71.

[108] J. Winter and A. Wuppermann, *Do they know what is at risk? health risk perception among the obese*, *Health Economics* **23** (2014), no. 5 564–585.

[109] G. E. Bijwaard, H. van Kippersluis, and J. Veenman, *Education and health: the role of cognitive ability*, *Journal of Health Economics* **42** (2015) 29–43.

[110] E. U. Weber, A.-R. Blais, and N. E. Betz, *A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors*, *Journal of behavioral decision making* **15** (2002), no. 4 263–290.

[111] G. Picone, F. Sloan, and D. Taylor, *Effects of risk and time preference and expected longevity on demand for medical tests*, *Journal of Risk and Uncertainty* **28** (2004), no. 1 39–53.

[112] M. Buhrmester, T. Kwang, and S. D. Gosling, *Amazon's mechanical turk: A new source of inexpensive, yet high-quality data?*, .

[113] M. Riddel and S. Kolstoe, *Heterogeneity in life-duration preferences: Are risky recreationists really more risk loving?*, *Journal of risk and uncertainty* **46** (2013), no. 2 191–213.

[114] A. Tversky and D. Kahneman, *Advances in prospect theory: Cumulative representation of uncertainty*, *Journal of Risk and uncertainty* **5** (1992), no. 4 297–323.

[115] P. Wakker and D. Deneffe, *Eliciting von neumann-morgenstern utilities when probabilities are distorted or unknown*, *Management science* **42** (1996), no. 8 1131–1150.

[116] H. Bleichrodt and J. L. Pinto, *A parameter-free elicitation of the probability weighting function in medical decision analysis*, *Management science* **46** (2000), no. 11 1485–1496.

[117] A. E. Attema, W. B. Brouwer, and O. l'Haridon, *Prospect theory in the health domain: A quantitative assessment*, *Journal of health economics* **32** (2013), no. 6 1057–1065.

[118] C. F. Camerer and R. M. Hogarth, *The effects of financial incentives in experiments: A review and capital-labor-production framework*, *Journal of risk and uncertainty* **19** (1999), no. 1 7–42.

[119] L. Einav, A. Finkelstein, I. Pascu, and M. R. Cullen, *How general are risk preferences? choices under uncertainty in different domains*, *American Economic Review* **102** (2012), no. 6 2606–38.

[120] V. F. Reyna, W. L. Nelson, P. K. Han, and N. F. Dieckmann, *How numeracy influences risk comprehension and medical decision making.*, *Psychological bulletin* **135** (2009), no. 6 943.

[121] D. L. Patrick, A. Cheadle, D. C. Thompson, P. Diehr, T. Koepsell, and S. Kinne, *The validity of self-reported smoking: a review and meta-analysis.*, *American journal of public health* **84** (1994), no. 7 1086–1093.