Experimental analysis of asset markets, workplace tournaments and crime deterrence

A thesis submitted for the degree of

Doctor of Philosophy

Zhengyang Bao

BEcon (Honors) (University of Queensland, Australia)

Department of Economics

Monash University

December 3, 2020

Abstract

This thesis uses the experimental methodology to understand three important topics in economics and finance. In the first chapter, we explore the trading of assets without positive fundamental values and find that these assets are prone to larger bubbles than assets with positive fundamental values. In the second chapter, we study a complementary safeguard which guarantees a higher minimum wage. Contrary to our conjecture, such a safeguard does not decrease the gender gap in wages but increases it. In the final chapter, we develop a novel mechanism to utilize insider information to apprehend the most guilty criminal in the crime network and enhance deterrence. The experimental data confirm that the new mechanism is more effective compared to the random audit mechanism, but is less effective than the Nash equilibrium predicts.

Declaration Statement

This thesis is an original work derived from my research. It contains no material that has been accepted for the award of any other degree or diploma at any university or equivalent institution. To the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

....

Zhengyang Bao, Aug 2020

Acknowledgements

This thess could not have been completed without the help and support of my supervisors, friends and families, whom I wish to thank wholeheartedly. First and foremost, I would like to sincerely thank Professor Andreas Leibbrandt, Professor Lata Gangadharan and Dr. Matthew Leister. Thank you all for being incredible supervisors and my role models. I learned a lot from you.

Thank you to my thesis panel members Professor Klaus Abbink, Professor Nicholas Feltovich, Professor Erte Xiao and thesis external examiners Professor Daniel Houser and Dr. Redzo Mujcic for all the constructive comments. I would also like to thank the professional staff members of the department, especially Sue Ball, who has been kind enough to advise and help as the department manager.

Thank you to my officemates Ratul Das Chaudhury, Veasna Kheng, Cun Liu, Justin McKinley, Olamide Ayoade-Alabi, Kushneel Prakash, Chau Tran and Md Main Uddin. The past few years have not been an easy ride but your company makes it enjoyable. Although it would be impossible to list every name, there is a huge community of others who supported me deeply and who helped me flounder my way toward a solid set of chapters.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis uses novel laboratory experiments to explore new mechanisms to understand the trading of assets with no fundamental values, decrease the gender-wage gap in the labour markets, and deter criminal activities.

Chapter 1 investigates the trading of assets that lack positive fundamental values. Groups of traders are randomized into asset markets where fundamental values are either positive, zero, or negative. Our findings provide first evidence for the presence of much larger bubbles in asset markets that lack positive fundamental values than in asset markets with positive fundamental values. Further, we show that our findings are consistent with trader expectations but not with loss aversion and theories of complexity. This study provides experimental evidence that supports the need for particular scrutiny of markets with assets that have no positive fundamental value.

Chapter 2 explores a new mechanism that might be useful to generate gender neutral outcomes in workplace tournaments. Workplace tournaments are one likely contributor to gender differences in labor market outcomes. Relative to men, women are often less eager to compete and thrive less under competitive pressure. We investigate a competitive workplace environment that may produce more gender-neutral outcomes: tournaments with safeguards. In our experiment, participants take part in a tournament with a real-effort task and choose whether they want to have a complementary safeguard that guarantees higher wages for the low ranked. As expected, we find that women are more likely than men to choose such a safeguard. However, obtaining a safeguard comes at a cost. On average, the safeguard causes lower performance, creates a gender wage gap, and over-proportionally disadvantages women. Thus, we provide novel evidence that easing women into tournaments can backfire.

In chapter 3, we study a mechanism for crime deterrence that utilizes insider information contained in social networks. After a crime is committed the regulator randomly identifies one criminal in the network and this person is audited and punished for the crime, unless she can incriminate another person from the network who is deemed to be more responsible. The regulator audits the criminal activities of both and obtains two noisy signals about their actions. The one with the higher signal is punished and the other goes free. We show theoretically that, for a given probability and magnitude of the penalty, crime levels are lower with this mechanism than with a traditional random audit-punishment scheme. In equilibrium, crime levels depend on the given criminal's position in the network and on the network structure. Our experiment confirms that this mechanism effectively deters crime, but the magnitude is above the Nash equilibrium predictions and is less sensitive to changes in the network structure than theory predicts. Level-k reasoning helps to explain these patterns.

Economic experiments are a methodology that facilitates the identification of causal relationships and thus complements existing empirical data. Lab experiments provide insights regarding human behaviors, especially when real-world data are difficult to obtain or confounded with many other factors. However, making decisions in the lab can be very different from making decisions in the real world, and sometimes, the results from lab experiments may not apply outside. Cautions should be exerted when making inferences from results based on lab experiments.

Chapter 2

Experimental Evidence on Bubbles in Markets for Assets that lack positive Fundamental Value

[1]

2.1. Introduction

Assets that do not deliver any positive cash flow and even sometimes occur holding costs have significantly gained in popularity. For example, old cars and buildings can incur more maintenance costs than the value they generate. The storage and shipping costs of crude oil were likely to exceed the corresponding value during the outbreak of the COVID pandemic as the price of oil dropped below zero in many markets. In addition, some government bonds were offered at negative (real) interest rates.[2]

However, whether assets that lack a positive fundamental value induce more speculation

[1]This chapter reports research conducted jointly with Andreas Leibbrandt.

[2]For example, the interest rate of the 10-year Germany Government Bonds had a yield of -0.53%, and Switzerland, Denmark and Japan witnessed negative deposit interest rates in the year 2019.

than other traditional assets and, thus, should be treated differently remains an open question. In this experimental study, we provide a first glimpse into the extent and nature of mispricing of assets that have zero and negative fundamental value (FV, henceforth).[3] The data shows clear evidence of high levels of overpricing in the presence of assets that lack positive FV. The data show that zero FV shares are traded at a price that is much higher than zero, and the size of overpricing is larger for negative FV shares. Compared to markets with positive FV shares, we observe that the overpricing is approximately 3–4 times larger when shares have zero or negative FV, respectively. We also find that the variation in the cashflow increases the size of overpricing rather than decreases it.

A potential explanation for the observed bubble patterns is based on price expectations. More precisely, some traders might believe that other traders are willing to pay a premium above the FV and that this premium is systematically linked to the FV and dividend variance. We test this explanation in a guessing game where we incentivize inexperienced subjects to accurately predict the pricing patterns in some of the proposed treatments. We find that guessers successfully anticipate the existence of bubbles in zero FV markets and correctly guess relative bubble sizes in markets with different FVs. These findings suggest that markets with zero and negative FV assets have a natural tendency toward larger bubbles. This study substantially expands the literature on asset market bubbles. Since Smith et al. (1988), ample evidence shows considerable mispricing in traditional asset markets with positive FV (Ball & Holt 1998; Bostian, Goeree & Holt 2005; Dufwenberg, Lindqvist, & Moore 2005; Eckel & Fullbrunn 2015; Haruvy, Lahav, & Noussair 2007; Hussam, Porter, & Smith 2008; Lei, Noussair, & Plott 2001; Noussair & Powell 2010; Smith, Suchanek, & Williams 1988; Sutter et al. 2012; Kirchler 2008; Cason & Samek 2015; Cheung, Hedegaard, & Palan 2014; Breaban & Noussair 2015; Bao, Hommes, & Makarewicz 2017). In contrast to the existing literature, we investigate markets with non-positive FV assets.

Importantly, we find that applying the insights from existing studies on assets with positive FV to assets with non-positive FVs may be misleading. For example, the evidence from positive FV asset markets shows that bubbles are smaller in less complex

[3] An asset's fundamental value at time "t" is defined as the sum of discounted expected cash flows (those paid in the future). Throughout this manuscript, and due to the short horizon nature of our experiment setup, the discount rate is assumed to be zero.

	Negative FV	Zero FV	Positive FV
Risk-free dividends	No literature	No literature[a]	Literature available[b]
Risky dividends	No literature	No literature	Literature available[c]

Notes: This table summarizes the proposed 2×3 experimental design and related papers. Two treatment dimensions are considered: the type of FV and the type of dividend. We allow FV to be positive, zero, or negative and dividends to be fixed or variable.
[a] To our best knowledge, there is no SSW-type asset market literature fits in this cell. The "A1" markets in Smith et al. (2000) are related, as there is no dividend payment at the end each period, but the FV in their markets are still positive because the cumulative dividends are paid after the terminal period. There is also a strand of literature in fiat money flowing McCabe (1989), but the currency is always served as a medium of exchange.
[b] e.g., Porter & Smith (1995); Ball & Holt (1998).
[c] e.g., Bostian, Goeree, & Holt (2005); Dufwenberg, Lindqvist, & Moore (2005); Eckel & Fullbrunn (2015); Haruvy, Lahav, & Noussair (2007); Hussam, Porter, & Smith (2008); Lei, Noussair, & Plott (2001); Noussair & Powell (2010); Smith, Suchanek, & Williams (1988); Cason & Samek (2015).

Table 2.1: Experimental design in the context of existing literature

market environments (e.g., Huber & Kirchler 2012; Kirchler, Huber, & Stockl 2012), thus suggesting that bubbles are small in simple environments (such as in markets where FVs are always zero), which we do not observe. Further evidence shows the importance of loss aversion (Kahneman & Tversky 1979; Tversky & Kahneman 1991; Burgstahler & Dichev 1997; Haigh & List 2005; DellaVigna 2009, among the others) suggesting that bubbles are smaller when traders can make losses, which is at odds with our findings. In addition, contrary to Porter & Smith (1995) and Childs & Mestelman (2004), who find that an increase in the variance of the dividends has limited impact on the size of bubbles, we find that bubbles are larger when variations in the dividends are observed, and they can be negative.

2.2. Experimental Design

Table 2.1 summarizes how the proposed 3×2 experimental design expands the experimental asset market literature. Two treatment dimensions are considered: the type of FV and the type of dividend. More precisely, we allow the FV to be positive, zero, or negative and dividends to be fixed or variable.

Experimental asset markets with zero and negative FV

A total of 288 subjects (six groups of eight traders in each of the six treatments), inexperienced in asset market experiments, participated in the study. The experiment was programmed in z-Tree (Fischbacher 2007), and the subjects were recruited from SONA. Participants earned, on average, AUD 35, and the experiment lasted for approximately two hours. Before the beginning of the experiment, we asked traders to read an information sheet and sign a consent form. We then read the instructions out loud. Afterward, traders were given sufficient time to read the instructions on their own and ask questions. We then implemented two practice rounds and asked traders to answer a set of quiz questions.[4] After everyone correctly answered these questions, the asset market experiment started. When all traders completed the asset market experiment, we administered a short demographic questionnaire. When all traders completed the asset market experiment, we administered a short demographic questionnaire.

2.2.1. Market Environment

In the proposed asset market environment, eight traders form a trading group to trade shares with each other for 15 trading periods.[5] At the beginning of the first trading period, all traders have an initial endowment of tokens and shares to trade.[6] Traders have 100 seconds of trading time during which they can buy and sell shares in each of the 15 trading periods. At the end of each trading period, each share pays a dividend (and/or incurs a holding cost as a negative dividend, depending on the treatments). Individual inventories of shares and tokens owned by a trader carry over from one period to the next. After period 15, we convert tokens to Australian Dollar at the rate of 50 tokens=$1.

The double auction mechanism is similar to Smith, Suchanek, & Williams (1988), where traders can buy and sell as many times as they wish in each trading period, as long as

[4] Each practice round corresponds to a real trading period, except that we hide the realization of the dividend and earnings at the end of the practice round. We also inform the participants that the practice rounds are for them to get familiar with the experimental software, and they are not paid for these rounds.

[5] A trading group of eight traders is common in the experimental asset market literature. Many recent studies use trading groups with seven to ten individuals (e.g., Haruvy, Lahav, & Noussair 2007; Haruvy & Noussair 2006; Kirchler, Huber, & Stockl 2012; Lei, Noussair, & Plott 2001; Stockl, Huber, & Kirchler 2015; Bao et al. 2020).

[6] Providing participants with heterogeneous endowments and different cash/asset ratios is common (Palan 2013). We describe the endowment for each trader in Appendix A2.

they have enough tokens or shares. Shares are only traded in whole units, while the prices are quoted to two decimal places. To buy a share in the experiment, a trader can accept the ask offer at the minimum ask price. Alternatively, the trader can make a bid offer to buy cheaper than the minimum ask price; if the bid is accepted, then the trader buys the corresponding shares, and receives nothing otherwise. Analogously, to sell a share in the experiment, a trader can accept the bid offer with the highest bid price. Alternatively, the trader can make an ask offer that is higher than the maximum bid price; if this bid is accepted, then the trader sells the corresponding shares and sells nothing otherwise. It is important to note that ask and bid prices can take negative values. Outgoing tokens and shares in outstanding offers are frozen during the trading period in which the offers are made. All offers are canceled at the end of the trading period, and the corresponding frozen shares and tokens are released. The program is set up such that purchases occur at the minimum ask price and sales at the maximum bid price. Only one type of share exists, and no tax, brokerage, short-selling, or margin buying is considered. We report the experiment instructions and describe the treatments in Appendix A1.

The proposed experiment also involved a repetition (List 2003; Hussam, Porter, & Smith 2008; King et al. 1993; Van Boening, Williams, & LaMaster 1993): each group of traders took part in a second trading block of this experiment, identical to the first trading block. Importantly, traders were unaware of the second trading block before it took place. However, they were aware that another task was scheduled after the first trading block. They also knew that at the end of the experiment, the outcomes of one of the two tasks (the first or second trading block) were to be randomly paid out. For brevity and because they are qualitatively similar, we mainly focus on the first trading block and report the complete findings from both blocks in Appendix A3.

2.2.2. Treatments

Table 2.2 provides further detail of the comprehensive 3 $\times$ 2 experimental design. As mentioned above, we allow variations in two dimensions: three risk-free fundamental values (negative, zero, and positive FVs) and the dividend process (fixed vs. variable). The first dimension explores assets with zero and negative FV generating fixed (risk-free) cash flows. The zero and negative FV assets mimic the case where the benefit is smaller than the costs. The second treatment dimension allows us to investigate the trading behavior of assets that generate variable (risky) dividends and incur holding costs (which are larger, equal to, or smaller than the expected dividends in the negative, zero, and positive FV treatments). The comparison across the two treatment dimensions, holding the FV constant, allows us to test the proposed models, which predict opposite pricing patterns, as elaborated in Section 3.

	Positive FV	Zero FV
Fixed dividends	(2.5,2.5)	(0,0)
Variable dividends	(-5,10)	(-10,10)

Notes: This table provides details on the proposed 3×2 experimental design. The first treatment dimension varies the level of FV, and the second dimension varies the dividend processes. The numbers in the brackets are the net cash flows generated by a share in the high and low states, respectively. Each treatment is assigned six independent groups of eight traders.

Table 2.2: Experimental design

To investigate the behavior of assets with zero and negative fundamental values, share prices are allowed to take negative values in our experiment. To render subject bankruptcy unlikely, each participant receives a starting endowment of 1000 tokens (the equivalent of $20) in addition to a $15 show-up fee.[7] Further, to minimize the likelihood of bankruptcy, participants with insufficient funds are not allowed to sell shares at negative prices or buy shares at positive prices.

We follow the convention to frame negative dividends as holding costs (e.g., Noussair, Robin, & Ruffieux 2001; Noussair & Powell 2010; Stockl, Huber, & Kirchler 2015) and take additional measures to make losses (and gains) clear and salient. In particular, we include in the instructions a neutral sentence informing traders of the net loss (or profit) of holding one share in each period. Further, we deliberately use an example with a loss in the quiz (administered before the first trading period) for treatments that the shares may incur losses and reinforce that each share can incur a net loss (or benefit) in each period.

[7] To completely rule out bankruptcy in all the proposed treatments, a $108 show-up fee would be required. However, paying such a large show-up fee is problematic for various reasons, including the potential to distort incentives (Azrieli, Chambers, & Healy 2018), and causes a large house money effect in multi-period financial settings (Ackert et al. 2006). The proposed measures proved successful in reducing the risk of bankruptcy. Only three participants' payoffs were below $10, the minimum payment of a typical experiment conducted in the laboratory, in the first trading block.

Besides, we label the price axis in the price chat symmetric around zero, implying that negative prices are possible. Last, we inform each trader of the net cash flow per share and the net cash flow from all shares he/she holds at the end of each period and display the net benefit/loss added/subtracted from the account. These measures proved to be successful to imply the possibility of negative prices. For example, 5 out of 6 markets in the (-2.5,-2.5) treatment show non-positive prices in the first trading block. We also find the (-2.5,-2.5) treatment has a significantly lower price than (0,0) and (2.5,2.5) treatments. If participants are confused with the instructions, we would not expect such changes.

Treatments with shares that pay fixed dividends

In treatments (-2.5,-2.5), (0,0), and (2.5,2.5), the dividend in trading period $t \in \{1,2,3,\ldots,15\}$ is fixed. The shares in the zero FV treatment (0,0) do not generate any cash flow in any trading period. We denote this treatment as (0,0) because shares pay 0 tokens as a dividend in both high and low states.[8] We investigate the price patterns of shares with negative FVs by addressing treatment (-2.5,-2.5), where shares always charge 2.5 tokens as holding costs. We compare trading in these two markets to the standard assets addressed by the literature, which generate a positive FV. In treatment (2.5,2.5), the shares always pay 2.5 tokens, regardless of whether the state is high or low. Table 2.3 describes the FVs for each asset in each period under different treatments.

[8] The literature typically uses a two-point uniform distribution (e.g., Lei et al. 2001; Lei & Vesely 2009; Childs 2009) to capture randomness in the dividends. Lei et al. (2001) and van Boening et al. (1993) also suggest that using either a two-, four-, or a five-point symmetric dividend does not have a significant impact on the pricing patterns.

	Trading Period														
Treatments	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Positive	37.5	35	32.5	30	27.5	25	22.5	20	17.5	15	12.5	10	7.5	5	2.5
Zero	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Negative	-37.5	-35	-32.5	-30	-27.5	-25	-22.5	-20	-17.5	-15	-12.5	-10	-7.5	-5	-2.5

Notes: This table lists the risk-free FV for all treatments in each period. The value in period t (FV_t) is calculated as $FV_t = \Sigma_{i=t}^{15} E\left(d_t\right)$, where d_t is the (expected) dividend in period t.

Table 2.3: Risk-free fundamental values for positive, zero, and negative treatments

Treatments with shares that pay variable dividends

In treatments (-5,10), (-10,10), and (-15,10), we allow variations in the dividend process and implement a holding cost but keep the FV constant to the corresponding treatments with fixed dividends. The positive, zero, and negative FVs mimic the cases where the expected dividends are more than, equal to, and less than the holding costs. The dividends in this set of treatments always follow two-point uniform distributions depending on the dividend state at the end of each trading period. The states are pre-drawn with the help of a fair coin, and they are fixed across all treatments in this set. We describe the realizations of the state in Appendix A2. In the treatment with positive FV, the shares generate 10 tokens when the state is high and charge five tokens when the state is low. We denote this treatment as (-5,10). In the treatment with zero FV, the shares generate 10 tokens when the state is high and charge 10 tokens when the state is low. We denote this treatment as (-10,10). In the treatment with negative FV, the shares generate 10 tokens in the high state and charge 15 tokens in the low state. We denote this treatment as (-15,10). The risk-free FV of each treatment is the same as its counterpart with a fixed dividend. We describe the dividend for each asset in each period in Appendix A2.

2.3. Hypotheses

While standard theories (e.g., Samuelson 1965; Fama 1970) predict no systematic mispricing in the studied asset markets, a rich body of experimental evidence suggests otherwise. The two primary insights from previous asset market experiments are that (i) bubbles are common (most previous asset market experiments find price bubbles), and (ii) the bubble size increases with the complexity (i.e., number of mathematical manipulations required to calculate the FV) of the cash flow process (Huber & Kirchler 2012; Kirchler, Huber, & Stockl 2012; Galanis 2018). In the proposed experimental design, we systematically vary complexity across treatments. First, the zero fundamental value treatment (0,0) arguably presents the least complex environment as there is no cash flow—no mathematical manipulation is required to calculate the FV; analogously, (-10,10) is less complex than the other treatments with variable dividends as the only mathematical step required is to take the expectation. Second, the cash flow is less complex in the treatments with constant FV (-2.5,-2.5), (0,0), (2.5,2.5) as compared to the treatments with variable FV (-10,5), (-10,10), (-10,15) since the latter involves taking an extra step— taking the expectation. Thus, we derive the following prediction based on the complexity argument:

Hypothesis (complexity): *(i) Bubbles occur in (0,0); (ii) among the fixed dividend treatments, (0,0) is characterized by the smallest bubbles, and among the variable dividend treatments, (-10,10) is characterized by the smallest bubbles; (iii) for each level of FV, the treatment with fixed dividend is characterized by smaller bubbles than its variable dividend counterpart.*

However, substantial evidence indicates that people are loss averse (Kahneman & Tversky 1979). In the context of the proposed experiment, the theory suggests that a cash outflow generates a more significant impact than a cash inflow of the same amount.[9] A novel feature of the proposed experimental design is that we can test the role of loss aversion in asset markets as traders can have gains and losses in each trading period in the treatments with variable dividends.[10] While our research hypothesis based on the previous experimental asset market literature predicts that bubbles are larger when dividends are variable, loss aversion generates the opposite prediction.

Hypothesis (loss aversion): *(i) Among fixed dividend treatments, (-2.5,-2.5) has the smallest bubbles; (ii) for each level of non-negative FV, the treatment with fixed dividend is characterized by larger bubbles than its variable dividend counterpart.*

[9] Loss aversion depends on the assumption of the reference point. We follow Kahneman & Tversky (1979); Benartzi & Thaler (1995); Odean (1998) and assume that zero net cash flow (the status quo) serves as the neutral reference point. To make the gain and losses salient, we inform participants of the net cash flow generated by the dividends and holding costs after each trading period.

[10] Breaban and Noussair (2015) find that a greater loss aversion measure of the trader cohort correlates to with a lower volume transacted, but they did not provide a clear interpretation of how loss aversion affects prices.

2.4. Experimental Findings

In this section, we present the data obtained from the six treatments that took place in the Monash Laboratory of Experimental Economics (MonLEE) between July 2018 and March 2019.[11] We first report mispricing in the markets with zero and negative FVs (4.1) and then compare the results to mispricing in markets with positive FVs (4.2). Finally, this section compares mispricing in markets with fixed and variable dividends (4.3).

2.4.1. Bubbles in markets with zero and negative fundamental values

Figure 2.1 provides a first overview of the trading patterns for the three assets (0,0), (-2.5,-2.5), and (2.5,2.5) by illustrating average prices over trading periods across all groups.[12] We observe that price trends are similar across the three assets. As in the standard asset with positive fundamental value, prices slowly decline over trading periods in (0,0) and (-2.5,-2.5). More importantly, we observe that average prices are positive in all the 15 trading periods for the asset with zero FV and in the first 10 periods for the asset with negative FV. The market with zero FV has an average price of 18.6 (tokens), significantly higher than zero, thus providing evidence of the presence of significant bubbles ($p<.001$).[13] The average price in market (2.5,2.5) is 26 (the average FV is 20), which is not statistically higher than in (0,0) ($p=.2$). The average price in (-2.5,-2.5) is 4.9, smaller than in the other two asset markets ($p<.05$ in both cases) but still significantly larger than zero ($p<.01$).

Finding 1: *Bubbles occur in asset markets with zero and negative FVs.*

[11] 98% of the participants are students from Monash University studying different majors. 43% of the participants are females and their average age is 22 years old. The demographics are balanced across all treatments.

[12] We report trading volumes in Appendix A4. There are no significant differences in trading volumes across the different treatments.

[13] Without further specification, all p-values reported in this paper are from two-tail Mann-Whitney tests on variables aggregated at the group level (so that all observations are independent).

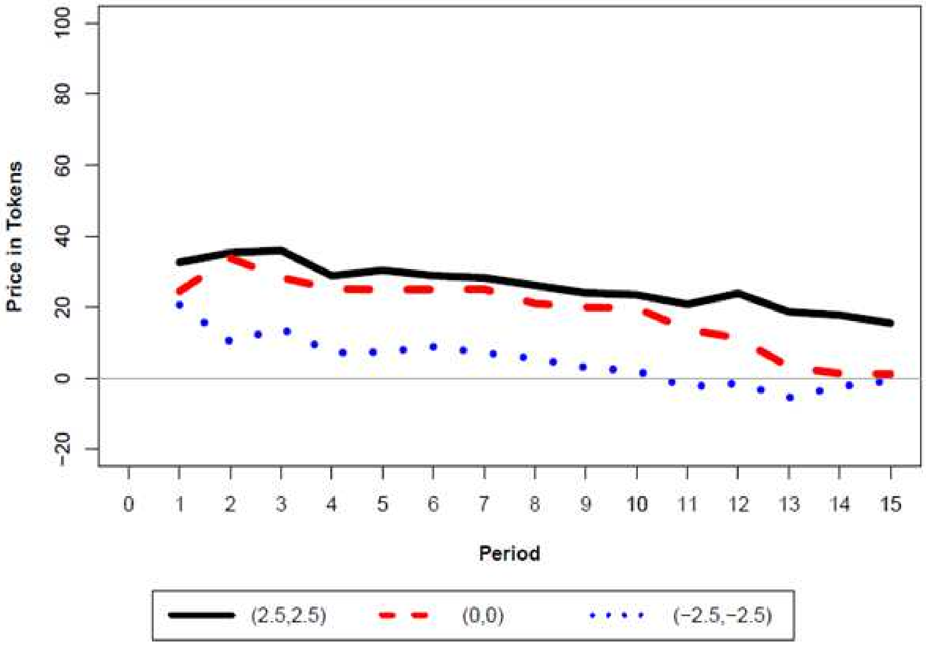

Notes: This graph shows the average price of six markets for each treatment with fixed dividend over the 15 trading periods in trading block 1.

Figure 2.1: Prices for treatments with fixed dividends (trading block 1)

2.4.2. Relative bubble size

Figure 2.2 corresponds to Figure 2.1 but illustrates the average bubble size or mispricing (i.e., price – fundamental value).[14] Bubbles are common in asset markets with zero and negative FV. More precisely, we observe that bubbles occur in all trading periods of (0, 0) and (-2.5, -2.5). In (2.5,2.5), bubbles occur in 13 out of the 15 trading periods. Importantly, we observe much larger bubbles in (0,0) and (-2.5, -2.5) than in (2.5, 2.5) in the first 10 trading periods. The average bubble size is 18.6 (p=.03) in (0,0) and 24.9 (p<.01) in (-2.5, -2.5), significantly larger than in (2.5, 2.5), where the average bubble size is 6.0.

Table 2.4 corroborates the previous result using four Ordinary Least Squares (OLS) models of the following regression specification:

$$pd_{g,t} = c + \delta_1 I_{g,t}^{zero} + \delta_2 I_{g,t}^{negative} + \sum_{\tau=2}^{15} \alpha_\tau I_{g,t}^{\tau} + \gamma_1 female_{g,t} + \gamma_2 knowledge_{g,t} + \epsilon_{g,t}, \quad (2.1)$$

[14] Haruvy & Noussair (2006) introduced "price dispersion," which we call "mispricing" in this study. Powell (2016) provides a comprehensive survey on the measurements of price bubbles. However, most of these measures are not well-defined in the zero FV treatments addressed by this study. Mispricing allows us to compare different treatments in the proposed experiment.

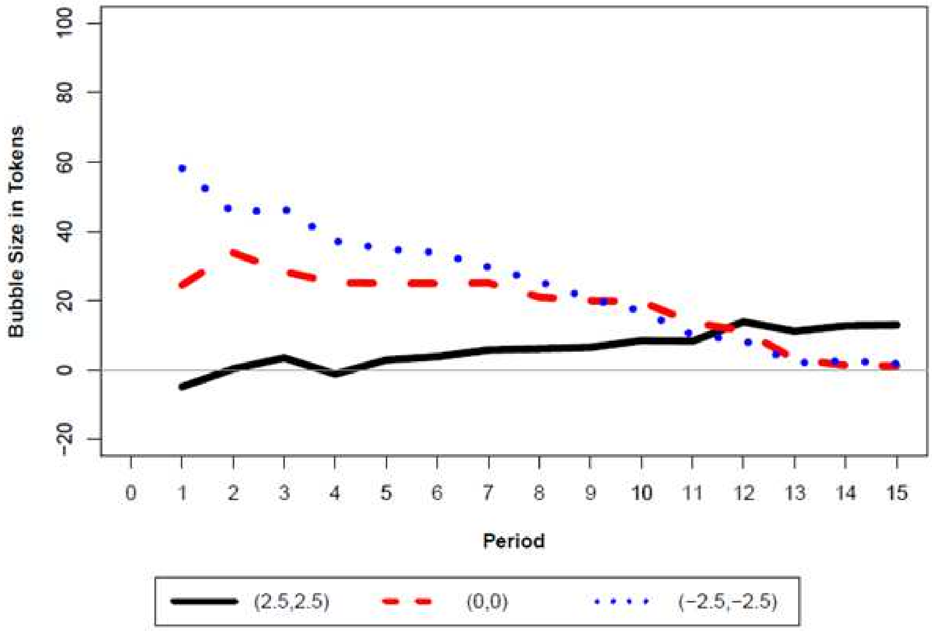

Notes: This graph shows the bubble size of six markets for each treatment over the 15 trading periods in trading block 1.

Figure 2.2: Bubble size of treatments with fixed dividends (trading block 1)

where $pd_{g,t}$ is the mispricing of group g in trading period t; $I^{zero}_{g,t}$, $I^{negative}_{g,t}$, and $I^{\tau}_{g,t}$ are indicator variables that take value 1 when the corresponding logic statement is true, and 0 otherwise. For example, $I^{zero}_{g,t}$ and $I^{negative}_{g,t}$ take value 1 when the FV is zero and negative, respectively; and $I^{\tau}_{g,t}$ takes value 1 when $t = \tau$; $female_{g,t}$ documents the number of female traders in the group, and $knowledge_{g,t}$ records the number of traders who have taken courses in asset pricing. $I^{zero}_{g,t}$ and $I^{negative}_{g,t}$ capture the differences in mispricing caused by changes in FV; $I^{\tau}_{g,t}$ controls for the heterogeneity across periods; $female_{g,t}$ and $knowledge_{g,t}$ control for some potentially important variations across groups.[15] We cluster the standard errors at the group level and control the variables reported in Table 2.4 to verify the robustness of our second finding.

In Model 1, we exclude the control variables for time fixed effects ($\sum_{\tau=2}^{15} \alpha_\tau I^{\tau}_{g,t}$), number of females in the group ($female_{g,t}$), and participants' knowledge on asset pricing ($knowledge_{g,t}$) and sequentially add these controls in Models 2–4. The point estimates of the treatment effect (δ_1 and δ_2) are economically large and significantly higher than zero in all regression

[15] Eckel & Fullbrunn (2015) and Wang et al. (2018) find that gender may play a role in the mispricing behvaior. Despite the main purpose of the paper is not about the economics of genders, we include the control variable female to capture how the gender composition of the group can affect the mispricing size. As the group size is fixed across sessions, the number of females in each group is perfectly collinear with the female proportion.

	Model 1	Model 2	Model 3	Model 4
Zero FV (δ_1)	12.60**	12.67**	14.86**	14.11**
	0.019	0.021	0.011	0.011
Negative FV (δ_2)	20.46***	20.01***	19.66***	20.94***
	0.000	0.000	0.000	0.000
Constant (c)	5.96*	15.05**	20.50***	21.92***
	0.063	0.017	0.010	0.006
Time fixed effects ($\Sigma\alpha_\tau I^t$)	No	Yes	Yes	Yes
# of females ($female$)	No	No	Yes	Yes
# have taken asset pricing classes ($knowledge$)	No	No	No	Yes

Notes: This table shows the estimation results for whether FVs have a significant impact on the bubble size. We use 262 observations for each regression, and all standard errors are clustered at the group level (we exclude eight trading periods where no trade takes place). We report the p-value under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2.4: Comparison of the bubble sizes in treatments with positive, zero, and negative FVs

specifications. We observe, for example, that the coefficient for negative (zero) FV in Model 1 is approximately four (two) times as large as the constant for the positive FV.

Finding 2: *Bubbles in asset markets with zero and negative FVs are larger than in markets with positive FVs.*

2.4.3. Bubbles and variances in the dividends

Figure 2.3 illustrates the trading patterns in the treatments with variable dividends (-5, 10), (-10, 10), and (-15, 20). We observe that prices are positive in all trading periods of treatments (-10, 10) and (-5,10), and in 14 out of 15 periods in (-15, 10). In line with the asset markets with fixed dividends, prices start very high and then decline over trading periods. The average price in (-10, 10) is 28 (tokens), which is not statistically different from 27.4 in (-15, 10) ($p=1$). The average price in (-5, 10) is 39.9, larger than in (-10, 10) and (-5,10) but statistically insignificant ($p=.34$ and $p=.15$, respectively). The average price in (-10, 10) and (-15, 10) is significantly higher than 0 ($p<.01$), and the price is insensitive to the drop of the FV from positive to zero and then negative when the dividend is variable. This result provided further evidence of the existence of large bubbles in zero and negative FV asset markets.

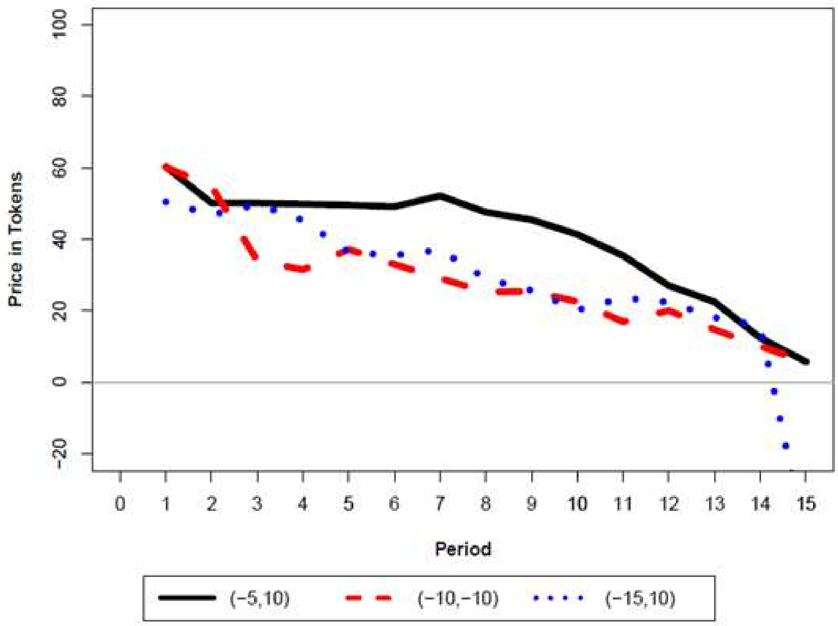

Notes: This graph shows the average price of six markets for each treatment with variable dividend over the 15 trading periods in trading block 1.

Figure 2.3: Prices for treatments with variable dividends (trading block 1)

Figure 2.4 corresponds to Figure 2.3 but plots the average bubble size. Bubbles are prevalent in all three treatments. We observe that bubbles occur in all trading periods in all treatments except the last period of (-15,10). In addition, (-15,10) is characterized by the largest bubbles in the first 14 periods, while (-5,10) is characterized by the smallest bubbles in the first seven periods and behaves similarly to (-10,10) afterward. The average bubble size of (-15,10) is 47.4, which is larger than 28 (p=.15) in (-10,10) and 19.9 (p=.055) in (-5,10).

Next, we investigate whether bubbles increase (complexity hypothesis) or decrease (loss aversion hypothesis) when we introduce the variance of dividends. Figure 2.5 provides a first illustration by pooling the three fixed and three variable dividend treatments and illustrates the average bubble size throughout the experiment. A clear pattern emerges. Bubbles are larger when dividends are variable, which is consistent with our complexity hypothesis. The average bubble size is 31.8, with variable dividends compared to only 16.5 when dividends are fixed (p<.05).

Figures 2.6-2.8 correspond to Figure 2.5 but compares bubbles across the different treatments holding FV constant. We observe that the bubble size is larger in most periods

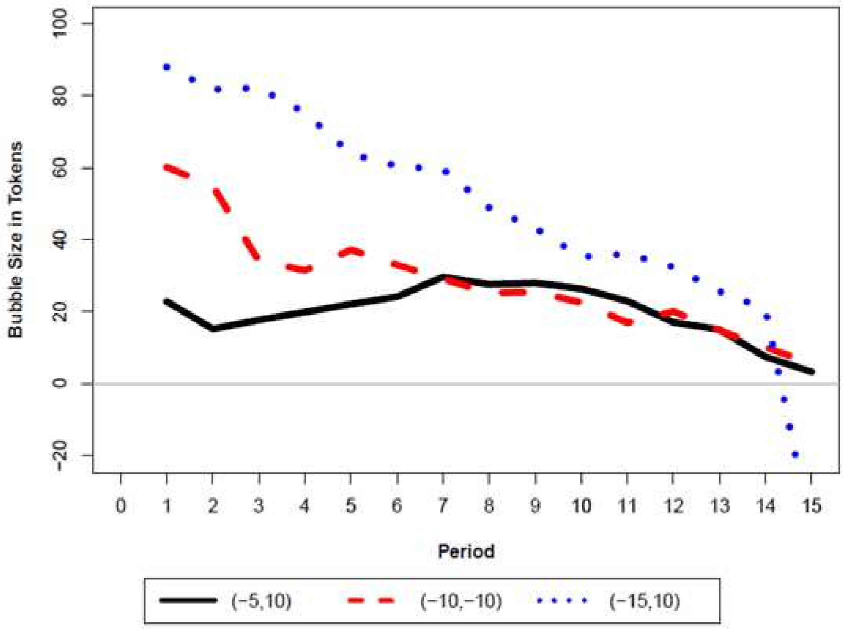

Notes: This graph shows the average bubble size of six markets for each treatment over the 15 trading periods in trading block 1.

Figure 2.4: Bubble size of treatments with variable dividends (trading block 1)

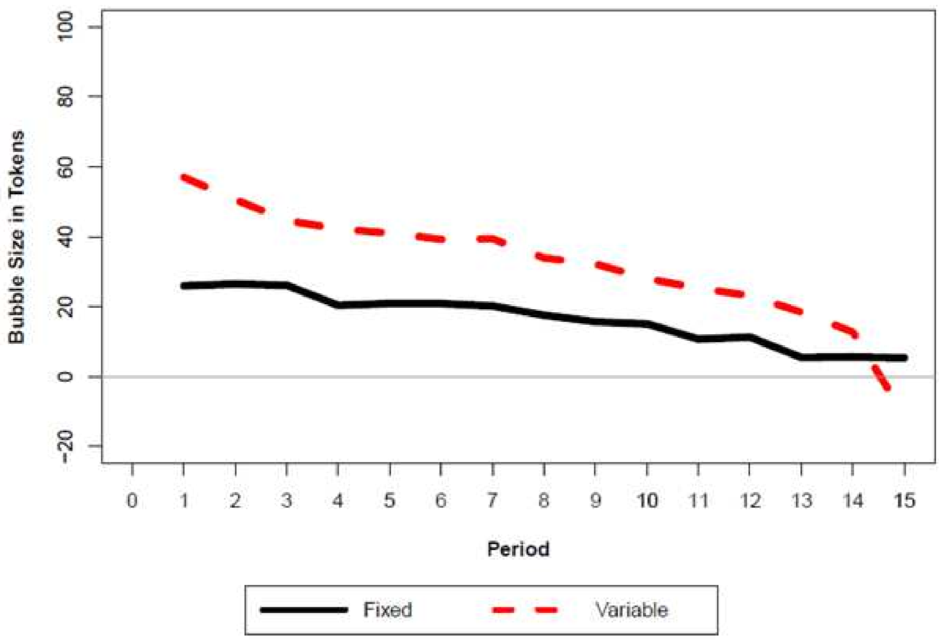

Notes: This graph shows the average bubble size of three treatments where dividends are variable and the average mispricing of three treatments where dividends are fixed over the 15 trading periods in trading block 1.

Figure 2.5: The average bubble size of three variable dividend treatments and three fixed dividend treatments

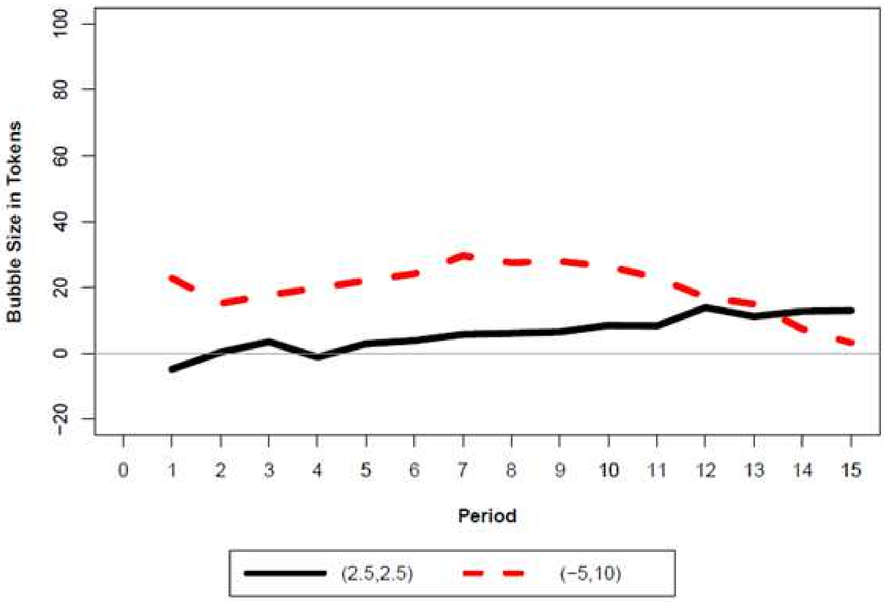

Notes: This graph shows the average bubble size of six markets for the (2.5,2.5) and (-5,10) treatments over the 15 trading periods in trading block 1.

Figure 2.6: The average bubble size of (2.5,2.5) and (-5,10)

	Positive FV	Zero FV	Negative FV
Fixed dividends	6.05	18.56	24.92
Variable dividends with loss possible	19.91	28.00	47.42
p-value	0.037	0.63	0.055

Notes: This table reports the bubble size (mispricing) for each treatment. The p-value in the bottom row compares the bubble size of the fixed dividend treatment with the variable one at the corresponding level of FV.

Table 2.5: Actual bubble size in different treatments

when the dividend is variable, regardless of the FV (negative, zero, or positive). In addition, (-10,10) has a larger bubble than (0,0) in all 15 periods, (-15,10) has a larger bubble than (-2.5,-2.5) in 13 out of 15 periods, and (-5,10) has a larger bubble than (2.5,2.5) in 14 out of 15 periods. In Table 2.5, we report the average bubble size for each treatment. The bubble size for (-10,10) is 28, 18.6 for (0,0) (p=.63), 19.9 for (-5,10), and 6.0 for (2.5,2.5) (p=.037). The bubble size is 47.4 for (-15,10) and 24.9 for (-2.5,-2.5) (p=.055).

We also run different OLS regressions to include all dependent observations across trading periods. As a result, we obtain a comprehensive view on the bubble pattern across

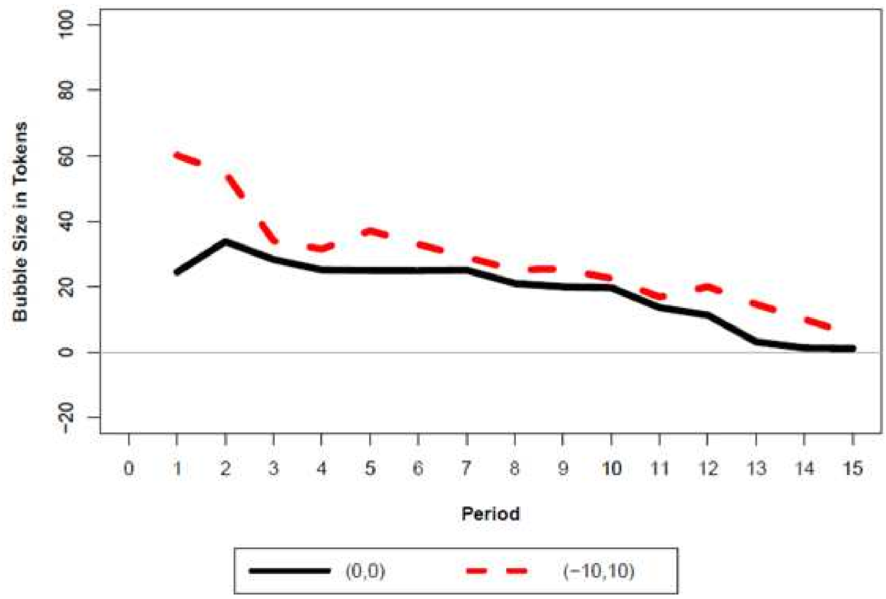

This graph shows the average bubble size of six markets for the (0,0) and (-10,10) treatments over the 15 trading periods in trading block 1.

Figure 2.7: The average bubble size of (0,0) and (-10,10)

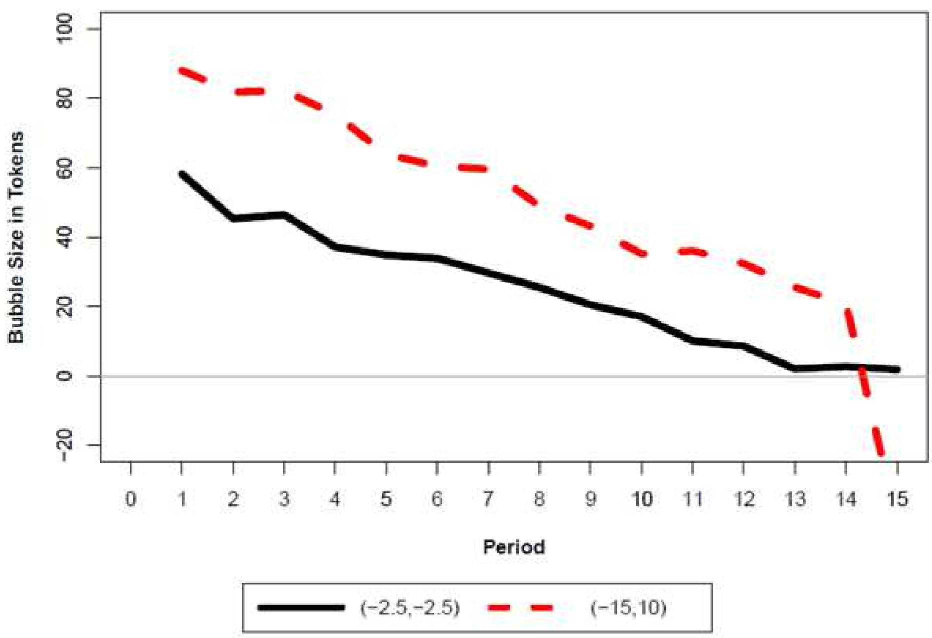

Notes: This graph shows the average bubble size of six markets for the (-2.5,-2.5) and (-15,10) treatments over the 15 trading periods in trading block 1.

Figure 2.8: The average bubble size of (-2.5,-2.5) and (-15,10)

	Model 1	Model 2	Model 3	Model 4
Positive Variable (β_1)	20.61*	20.69*	20.56*	21.54
	0.094	0.097	0.093	0.131
Zero fixed (β_2)	12.60**	12.68**	13.59**	13.24**
	0.012	0.013	0.020	0.020
Zero variable (β_3)	22.05**	22.13**	22.58**	21.87**
	0.022	0.024	0.025	0.040
Negative fixed (β_4)	20.46***	19.72***	19.58***	20.13***
	0.000	0.000	0.000	0.000
Negative variable (β_5)	41.47***	41.55***	41.42***	41.00***
	0.000	0.000	0.000	0.000
Constant (c)	5.96*	23.67***	25.91***	26.45***
	0.051	0.000	0.006	0.007
Time fixed effects ($\Sigma\alpha_\tau\tau^t$)	No	Yes	Yes	Yes
# of females ($female$)	No	No	Yes	Yes
# have taken asset pricing classes ($knowledge$)	No	No	No	Yes

Notes: This table shows the estimation results for whether the nature of the dividends has a significant impact on the bubble size. We use 532 observations for each regression, and all standard errors are clustered at the group level (we exclude eight trading periods where no trade takes place). The data refer to the first trading block. We report the p-value under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2.6: Comparison of the bubble sizes in treatments with fixed and variable dividends

treatments. We estimate the following model:

$$\begin{aligned} pd_{g,t} = c &+ \beta_1 I_{g,t}^{(-5,10)} + \beta_2 I_{g,t}^{(0,0)} + \beta_3 I_{g,t}^{(-10,10)} + \beta_4 I_{g,t}^{(-2.5,-2.5)} \\ &+ \beta_5 I_{g,t}^{(-15,10)} + \sum_{\tau=2}^{15} \alpha_\tau T_{g,t}^{\tau} + \gamma_1 female_{g,t} + \gamma_2 knowledge_{g,t} + \epsilon_{g,t}, \end{aligned} \tag{2.2}$$

where I is an indicator variable, which takes value 1 if group g participates in the treatment indicated by the superscript, and 0 otherwise. The set of indicator variables I compares the bubble size across treatments. Table 2.6 plots the estimation of the regression model with variations in the control variables.

We estimate Equation 2.2 and report the results in Table 2.6. In Model 1, we exclude the control variables for the time fixed effects ($\alpha_\tau T_{g,t}^{\tau}$), number of females in the group ($female_{g,t}$), and participants' knowledge on asset pricing ($knowledge_{g,t}$) and subsequentially include these controls in Models 2–4. The relative bubble size (point estimates of βs) is robust to the inclusion of different control variables. In Model 4, which includes most controls, the previous results hold: (i) holding the nature of the dividend (fixed or variable) constant, the treatment with the lower FV is always characterized by the larger bubble size, (ii) fixing the level of the FV, the price is always higher in the variable dividend treatment. With respect to bubbles sizes, we find the following ordering:

$$(2.5,2.5)<(0,0)<(-2.5,-2.5)<(-5,-10)<(-10,10)<(-15,10).$$

The point estimates from the four models also show that the treatment effects are meaningful and large. For example, in Model 4, the point estimate for the positive variable dividend treatment ($I_{g,t}^{(-5,10)}$) is 21.54, which implies that the bubble size of (-5,10) is more than 84% bigger than that of (2.5,2.5) in the first trading period.

Finding 3: *Bubbles are larger in asset markets with variable dividends than in the presence of fixed dividends. Bubbles are the largest in asset markets with negative FV and variable dividends and the smallest in asset markets with positive FV and fixed dividends.*

Bubbles occur in all treatments, but their size is hard to reconcile with loss aversion predictions based on Prospect Theory. For example, (-2.5,-2.5) markets are characterized by larger bubbles than (2.5,2.5) and (0,0) markets, and in treatments with variable dividends and potential losses, bubbles are larger than in their fixed dividend counterparts, in contrast with the predictions of loss aversion. A model based on complexity correctly predicts that variations in the dividends inflate the bubble, but it neither explains the large bubbles in the markets with zero FV nor predicts the relative bubble size within each set (fixed dividend or variable dividend) of treatments. One explanation of these patterns is that there are systematic differences in terms of understanding of the tasks across treatments, but this does not seem to be the case. First, from the post-experiment questionnaire, only 6% of participants find the instructions difficult to understand and there is no significant difference across treatments. Second, there is no significant difference in the speed of finishing quiz questions, nor in the demographic variables we collected across treatments ($p>0.1$ using F-tests). More importantly, five out of six groups in (-2.5,-2.5) treatment have negative price records, but this is not observed in (0,0) or (2.5,2.5) treatments. This implies that there is a treatment effect caused by the change in the FV and participants are able to understand and bid in negative prices. The next section provides and tests an alternative explanation based on trader expectations.

2.5. An Explanation Based on Price Expectations

Bubbles may be driven by trader expectations about the behavior of other traders in the different treatment scenarios (e.g., Cheung, Hedegaard, & Palan 2014; Holt, Porzio, & Song 2017). In particular, traders may believe that other traders are willing to pay high prices for assets with zero FV and higher prices for assets with variable dividends compared to fixed dividends. In this case, an optimal strategy is to pay high prices for these assets to then sell them for even higher prices. To test this alternative explanation, we conduct a second experimental study.

2.5.1. Study 2: The guessing game

We invited 83 participants, inexperienced in asset market experiments, to analyze how expectations help explain the observed price patterns. The proposed guessing game was programmed in z-Tree (Fischbacher 2007), and the participants were recruited from SONA. Participants earned, on average, AUD 17, and sessions lasted for approximately 45 minutes. Before the experiment, participants read an information sheet and signed a consent form. Participants read the instructions and were then asked to guess prices in different asset markets. The instructions and a short demographic questionnaire are reported in Appendix A5. The set-up of the guessing game is as follows. Each participant guesses the average price of each of the 15 periods in the first trading block of one asset market and then the average price for 15 periods in another asset market. We randomize the two sets of instructions. We do not provide participants with any information about their performance throughout the experiment. Guessers are monetarily incentivized using the binary scoring rule (Hossain & Okui 2013): each guess within five tokens of the actual price is awarded $2. To examine (i) whether guessers expect large bubbles in zero FV treatments, and (ii) whether they expect different bubble sizes among markets with different FV, the analysis considers the following two treatments.

2.5.2. Treatments (guessing game)

(0,0) vs. (-10,10)

We ask the first set of participants (n=54) to guess the average price across the studied six markets (to moderate some extreme values) in all the 15 trading periods of asset markets (0,0) and (-10,10). We denote this treatment as "(0,0) vs. (-10,10)". This approach allows us to investigate trader expectations about the bubbles in zero FV treatments.

(-15,10) vs. (-10,10)

We let the second set of participants (n=29) guess the average prices (across the six markets of all the 15 trading periods) of (-15,10) and (-10,10). We denote this treatment as "(-15,10) vs. (-10,10)". This treatment explores guessers' expectation about the relative bubble sizes between (-15,10) and (-10,10).

2.5.3. Results (guessing game)

Figure 2.9 plots the median bubble size based on all the prices guessed. Among the 83 guessers, only four guessed at least one zero price, and none of them guessed a single negative price. Participants predicted large bubbles and anticipated the relative bubble size between these three treatments. The average predicted bubble size is 41 (tokens) for (0,0), 51 for (-10,10), and 99 for (-15,10), which are all larger than the actual values ($p<.05$ for all three cases).

To compare these bubble sizes rigorously, we estimate the following model:

$$\bar{pd}_t = c + \beta_1 I_i^{(-15,10)} + \beta_2 I_i^{(-10,10)} + \gamma_1 I_i^{male} + \gamma_2 I_i^{knowledge} + \epsilon_i,$$

where $\bar{pd}_t$ is the average bubble size guessed by guesser i, I is an indicator variable that takes value 1 when the logic condition in its superscript is true. For example, $I_i^{(-15,10)}$

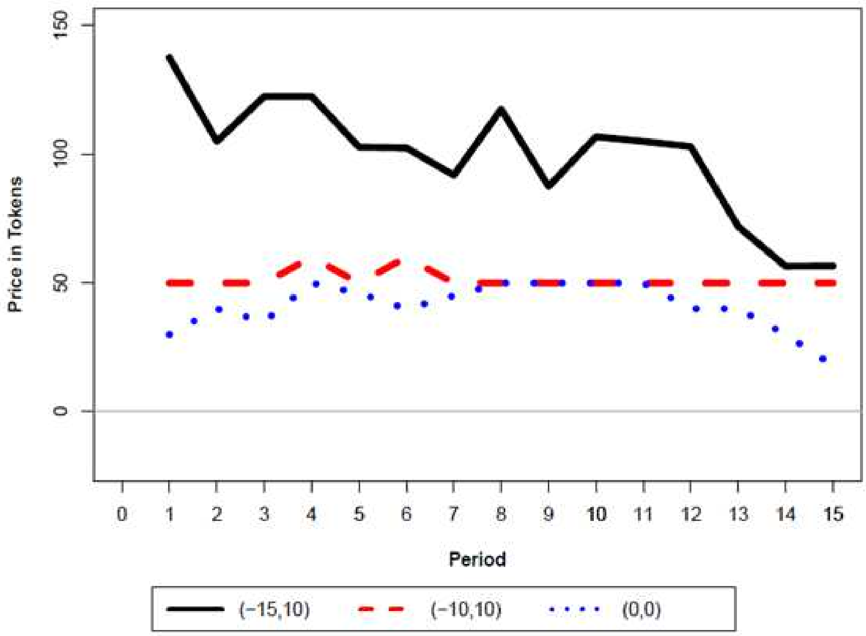

Notes: This graph shows the median guessed price of all 15 trading periods in (-15,10), (-10,10), and (0,0).

Figure 2.9: Prices predicted for (-15,10), (-10,10), and (0,0)

takes value 1 when the treatment guessed by i is (-15,10); I_i^{male} takes value 1 when i is a male; $I_i^{knowledge}$ takes value 1 when i had taken courses related to asset pricing.

Table 2.7 plots the estimation of the regression results considering variations in the control variables and estimation methods. Model 1 only includes treatment dummies and is estimated by OLS. Model 2 includes the same set of variables as Model 1 but is estimated using the quantile regression (QR) method to moderate the impact of outliers in the sample. Model 3 includes all the control variables and is estimated by OLS to check the robustness of this specification. Model 4 corresponds to Model 3 but uses QR. We cluster the errors at the individual level in all regressions. These models show that the relative bubble sizes observed in Figure 2.9 are robust to different model specifications. In Model 4, after controlling for individual characteristics and moderating the impact of outliers, we find that guessers predict a large bubble of 58 (tokens) in (0,0), which is significantly higher than zero (p<.01). Guessers also correctly predict that (-15,10) has the largest bubble, which is 77.56 higher than (0,0) and 56.36 higher than (-10,10) (p<.05 in both cases).

Finding 4: *Individuals predict (i) the positive price in the treatments with zero FV, (ii)*

	Model 1	Model 2	Model 3	Model 4
Negative variable (β_1)	99.36*	67.33**	98.23*	77.56**
	0.064	0.014	0.082	0.017
Zero variable (β_2)	20.13	17.86	19.73	21.2
	0.398	0.247	0.429	0.160
Gender (γ_1)	-	-	-70.79*	-11.33
			0.091	0.634
Knowledge (γ_2)	-	-	-8.14	-24.87
			0.884	0.296
Constant (c)	125.47***	50.00***	158.13***	58.00***
	0.000	0.000	0.000	0.007
Estimation method	OLS	QR	OLS	QR

Notes: This table shows the estimation results for whether the nature of the dividends has a significant impact on the bubble size. We use 166 observations for each regression, and all standard errors are clustered at the individual level. Quantile regressions are estimated for the median response (quantile=0.5). We report the p-value under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. The reason why OLS regressions have much larger constants is that some guessers indicated large numbers, and OLS estimation is sensitive to outliers.

Table 2.7: Comparison of the guessed bubble sizes in different treatments

that variations in the dividends increase the bubble size, and (iii) that the bubble size increases when the FV drops from zero to negative values. Expectations motivate the price patterns that complexity theory cannot explain.

2.6. Conclusion

Assets that lack any positive fundamental value are commonplace in modern financial markets. In this study, we provide a first glance into the price patterns in markets with zero and negative FV assets and their mechanisms using a controlled laboratory experiment.

We manipulate the risk-free FV of shares (positive, zero, or negative) and allow the nature of the dividends to vary in two sets of treatments. In the first set of treatments, the shares generate risk-free (fixed) cash-flows. In the second set of treatments, the shares pay risky (variable) cash flows. This represents a more realistic scenario where the value/costs have some variances like crude oil, old cars, or companies from sunset industries. We find several interesting patterns. First, overpricing occurs in all treatments, even in the simple zero FV case. Second, the overpricing inflates when the FV drops to zero and negative values, ceteris paribus. These findings suggest that zero and negative FV per se foster price overpricing.

Beyond exploring the relative bubble sizes across treatments, we also investigate the relationship between trader expectations and bubbles. We find that expectations match the actual relative bubble size. Our findings suggest that traders have a natural tendency toward trading assets with non-positive FV at positive prices, even without considering other features that may lead to price surges. This result also implies that the existing literature may underestimate the severity of price bubbles for these types of assets.

To what extent can we learn from lab experiments has been debated. In our experiment, the participants are mainly university students who are inexperienced with asset markets. In addition, the endowment they receive to trade with each other is given by the experimenters rather than their own money. However, ample evidence find lab experiments capture the essence in the real-word and can provide informative predictions (see, e.g., Levitt & List 2007; Fehr & Leibbrandt 2011; Alm et al. 2015). There is also a large literature showing that using student subjects pool (Smith et al. 1988; King et al. 1993) and house money (Ang et al. 2010; Corgnet et al. 2013) does not change the results from professional participants and earned money in the asset market experiment context.

2.7. References

Ackert, L. F., Charupat, N., Church, B. K., & Deaves, R., 2006, 'An experimental examination of the house money effect in a multi-period setting', Experimental Economics, 9(1), 5-16.

Alm, J., Bloomquist, K.M. & McKee, M., 2015, 'On the external validity of laboratory tax compliance experiments', Economic Inquiry, 53(2), pp.1170-1186.

Ang, J. S., Diavatopoulos, D. & Schwarz, T. V. 2010, 'The creation and control of speculative bubbles in a laboratory setting', in 'Handbook of Quantitative Finance and Risk Management', Springer, pp. 137-164.

Azrieli, Y., Chambers, C. P., & Healy, P. J., 2018, 'Incentives in experiments: A theoretical analysis', Journal of Political Economy, 126(4), 1472-1503.

Ball, S & Holt, C, 1998, 'Classroom games: Speculation and bubbles in an asset market', Journal of Economic Perspectives, vol. 12, no. 1, pp. 207-18.

Bao, T, Hommes, C, & Makarewicz, T 2017, 'Bubble Formation and (In)Efficient Markets in Learning-to-forecast and optimise Experiments', The Economic Journal, 127(605), pp. 581-609.

Bao, Z, Kalayci, K, Leibbrandt, A & Oyarzun, C 2019, 'Do regulations work? A Comprehensive Analysis of Price Limits and Trading Restrictions in Experimental Asset Markets with Deterministic and Stochastic Fundamental Values', working paper

Bartos, J 2015, 'Does Bitcoin follow the hypothesis of efficient market?', International Journal of Economic Sciences, Vol. 4(2), pp. 10-23

Benartzi, S, & Thaler, R.H, 1995, 'Myopic loss aversion and the equity premium puzzle', The Quarterly Journal of Economics, 110(1), pp.73-92.

Benz, M. and Meier, S., 2008, 'Do people behave in experiments as in the field?—evidence from donations', Experimental economics, 11(3), pp.268-281.

Biais, B., Bisiere, C., Bouvard, M. and Casamatta, C., 2019, 'The blockchain folk theorem', The Review of Financial Studies, 32(5), pp.1662-1715.

Bohme, R., Christin, N., Edelman, B. and Moore, T., 2015, 'Bitcoin: Economics, technology, and governance', Journal of Economic Perspectives, 29(2), pp.213-38.

Borri, N., 2019, 'Conditional tail-risk in cryptocurrency markets', Journal of Empirical Finance, 50, pp. 1-19.

Bostian, A, Goeree, J & Holt, CA 2005, 'Price bubbles in asset market experiments with a flat fundamental value', in Draft for the Experimental Finance Conference, Federal Reserve Bank of Atlanta September, vol. 23, p. 2005.

Breaban, A. and Noussair, C.N., 2015, 'Trader characteristics and fundamental value trajectories in an asset market experiment', Journal of Behavioral and Experimental Finance, 8, pp.1-17.

Burgstahler, D. and Dichev, I., 1997, 'Earnings management to avoid earnings decreases and losses', Journal of Accounting and Economics, 24(1), pp.99-126.

Cason, T.N. and Samek, A., 2015, 'Learning through passive participation in asset market bubbles', Journal of the Economic Science Association, 1(2), pp.170-181.

Charness, G. and Neugebauer, T., 2019, 'A Test of the Modigliani-Miller Invariance Theorem and Arbitrage in Experimental Asset Markets', The Journal of Finance, 74(1), pp.493-529.

Cheah, E.T. and Fry, J., 2015, 'Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin', Economics Letters, 130, pp.32-36.

Cheung, S.L., Hedegaard, M. and Palan, S., 2014, 'To see is to believe: Common expectations in experimental asset markets', European Economic Review, 66, pp.84-96.

Childs, J. 2007, 'Rate of return parity and currency crises in experimental asset markets', Journal of International Financial Markets, Institutions and Money, no. 19, pp.157–170.

Childs, J., & Mestelman, S. 2006, 'Rate-of-return parity in experimental asset markets', Review of International Economics, 14(3), 331-347.

Cipriani, M., Fostel, A. and Houser, D., 2018, 'Collateral constraints and the law of one price: an experiment', The Journal of Finance, 73(6), pp.2757-2786.

Corgnet, B., HernÃ¡n-GonzÃ¡lez, R., Kujal, P. and Porter, D., 2015, 'The effect of earned versus house money on price bubble formation in experimental asset markets', Review of Finance, 19(4), pp.1455-1488.

Cong, L.W. and He, Z., 2019, 'Blockchain disruption and smart contracts', The Review of Financial Studies, 32(5), pp.1754-1797.

Corbet, S., Lucey, B., Urquhart, A., & Yarovaya, L. 2019, 'Cryptocurrencies as a financial asset: A systematic analysis', International Review of Financial Analysis, no. 62, pp. 182-199.

DellaVigna, S., 2009, 'Psychology and economics: Evidence from the field', Journal of Economic literature, 47(2), pp.315-72.

Dufwenberg, M, Lindqvist, T & Moore, E 2005, 'Bubbles and experience: An experiment', American Economic Review, vol. 95, no. 5, pp. 1731-7.

Dyhrberg, A.H., 2016, 'Bitcoin, gold and the dollar–A GARCH volatility analysis', Finance Research Letters, 16, pp.85-92.

Eckel, CC & Fullbrunn, SC 2015, 'Thar she blows? Gender, competition, and bubbles in experimental asset markets', American Economic Review, vol. 105, no. 2, pp. 906-20.

Fehr, E. and Leibbrandt, A., 2011, 'A field study on cooperativeness and impatience in the tragedy of the commons', Journal of Public Economics, 95(9-10), pp.1144-1155.

Fischbacher, U 2007, 'z-Tree: Zurich toolbox for ready-made economic experiments', Experimental Economics, vol. 10, no. 2, pp. 171-8.

Foley, S., Karlsen, J.R. and PutniņÂš, T.J., 2019, 'Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? ', The Review of Financial Studies, 32(5), pp.1798-1853.

Fry, J. and Cheah, E.T., 2016, 'Negative bubbles and shocks in cryptocurrency markets', International Review of Financial Analysis, 47, pp.343-352.

Galanis, S 2018, ‘Financial complexity and trade’, Games and Economic Behavior, vol 112, pp. 219-230

Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M.C. and Siering, M., 2014, 'Bitcoin-asset or currency? Revealing users' hidden intentions', Revealing Users' Hidden Intentions (April 15, 2014). ECIS.

Haigh, M.S. and List, J.A., 2005, 'Do professional traders exhibit myopic loss aversion? An experimental analysis', The Journal of Finance, 60(1), pp.523-534.

Haruvy, E & Noussair, CN 2006, 'The effect of short selling on bubbles and crashes in experimental spot asset markets', The Journal of Finance, vol. 61, no. 3, pp. 1119-57.

Haruvy, E, Lahav, Y & Noussair, CN 2007, 'Traders' expectations in asset markets: Experimental evidence', American Economic Review, vol. 97, no. 5, pp. 1901-20.

Hileman, G. and Rauchs, M., 2017, 'Global cryptocurrency benchmarking study', Cambridge Centre for Alternative Finance, 33.

Holt, C.A., Porzio, M. and Song, M.Y., 2017, 'Price bubbles, gender, and expectations in experimental asset markets', European Economic Review, 100, pp.72-94.

Hossain, T. and Okui, R., 2013, 'The binarized scoring rule', Review of Economic Studies, 80(3), pp.984-1001.

Huber, J., & Kirchler, M., 2012, 'The impact of instructions and procedure on reducing confusion and bubbles in experimental asset markets', Experimental Economics, 15(1), 89-105.

Hussam, RN, Porter, D & Smith, VL 2008, 'Thar she blows: Can bubbles be rekindled with experienced subjects?', American Economic Review, vol. 98, no. 3, pp. 924-37.

Jakub, B., 2015, 'Does Bitcoin follow the hypothesis of efficient market? ', International Journal of Economic Sciences, 4(2), pp.10-23.

Kahneman, D. and Tversky, A., 1979, 'Prospect theory: An analysis of decision under risk', In Handbook of the fundamentals of financial decision making: Part I (pp. 99-127).

Kahneman, D., & Tversky, A., 1984, 'Choices, Values, and Frames' American Psychologist, Vol. 39, No. 4, pp. 341-350

Kahneman, D., Knetsch, J.L. and Thaler, R.H., 1991, 'The endowment effect, loss aversion, and status quo bias: Anomalies', Journal of Economic Perspectives, 5(1), pp.193-206.

King, RR, Smith, VL, Williams, AW & Van Boening, M 1993, 'The robustness of bubbles and crashes in experimental stock markets', Nonlinear Dynamics and Evolutionary Economics, pp. 183-200.

Kirchler, M., 2009, 'Underreaction to fundamental information and asymmetry in mispricing between bullish and bearish markets. An experimental study', Journal of Economic Dynamics and Control, 33(2), pp.491-506.

Kirchler, M., Huber, J., & Stockl, T. 2012, 'Thar she bursts: Reducing confusion reduces bubbles', American Economic Review, 102(2), 865-83.

Lei, V, Noussair, CN & Plott, CR 2001, 'Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality', Econometrica, vol. 69, no. 4, pp. 831-59.

Lei, V., & Vesely, F. 2009, 'Market efficiency: evidence from a no-bubble asset market experiment', Pacific Economic Review, 14(2), 246-258.

Levitt, SD., & List, JA. 2007, 'What do laboratory experiments measuring social preferences reveal about the real world?' Journal of Economic Perspectives, 21(2), 153-174.

List, JA. 2003, 'Does Market Experience Eliminate Market Anomalies?', Quarterly Journal of Economics, 118(1): 41-71.

Malkiel, B.G. and Fama, E.F., 1970, 'Efficient capital markets: A review of theory and empirical work', The Journal of Finance, 25(2), pp.383-417.

McCabe, K.A. 1989, 'Fiat Money as a Store of Value in an Experimental Market', Journal of Economic Behavior and Organization 12, 215-231.

Noussair, C, Robin, S, & Ruffieux, B, 2001, 'Price bubbles in laboratory asset markets with constant fundamental values', Experimental Economics 4(1), pp. 87-105.

Noussair, CN & Powell, O 2010, 'Peaks and valleys: Price discovery in experimental asset markets with non-monotonic fundamentals', Journal of Economic Studies, vol. 37, no. 2, pp. 152-80.

Odean, T., 1998, 'Are investors reluctant to realize their losses? ', The Journal of finance, 53(5), pp.1775-1798.

Palan, S 2013, 'A review of bubbles and crashes in experimental asset markets', Journal of Economic Surveys, vol. 27, no. 3, pp. 570-88.

Porter, D. P., & Smith, V. L. 1995, 'Futures Contracting and Dividend Uncertainty in Experimental Asset Markets', Journal of Business, 1(2), 111-128.

Powell, O. (2016). 'Numeraire independence and the measurement of mispricing in experimental asset markets', Journal of Behavioral and Experimental Finance, Vol.9, pp. 56-62.

Powell, O. and Shestakova, N., 2016, 'Experimental asset markets: A survey of recent developments', Journal of Behavioral and Experimental Finance, 12, pp.14-22.

Samuelson, P.A., 2016, 'Proof that properly anticipated prices fluctuate randomly', In The World Scientific Handbook of Futures Markets (pp. 25-38).

Smith, V. L., Van Boening, M. & Wellford, C. P. 2000, 'Dividend timing and behavior in laboratory asset markets', Economic Theory, 16(3), 567-583.

Smith, VL, Suchanek, GL & Williams, AW 1988, 'Bubbles, crashes, and endogenous expectations in experimental spot asset markets', Econometrica: Journal of the Econometric Society, pp. 1119-51.

Stockl, T, Huber, J & Kirchler, M 2015, 'Multi-period experimental asset markets with distinct fundamental value regimes', Experimental Economics, vol. 18, no. 2, pp. 314-34.

Sutter, M., Huber, J. and Kirchler, M., 2012, 'Bubbles and information: An experiment', Management Science, 58(2), pp.384-393.

Tversky, A. and Kahneman, D., 1991 'Loss aversion in riskless choice: A reference-dependent model', The Quarterly Journal of Economics, 106(4), pp.1039-1061.

Tversky, A., & Kahneman, D., 1981 'The framing of decisions and the psychology of choice', Science, 211(4481), pp. 453-458.

Urquhart, A., 2016, 'The inefficiency of Bitcoin', Economics Letters, 148, pp.80-82.

Van Boening, MV, Williams, AW & LaMaster, S 1993, 'Price bubbles and crashes in experimental call markets', Economics Letters, vol. 41, no. 2, pp. 179-85.

Wang, J., Houser, D. and Xu, H., 2018. Culture, gender and asset prices: Experimental evidence from the US and China. Journal of Economic Behavior & Organization, 155, pp.253-287.

Yermack, D., 2015, 'Is Bitcoin a real currency? An economic appraisal', In Handbook of digital currency (pp. 31-43). Academic Press.

2.8. Appendix

2.8.1. Instructions and quiz questions for the main study treatment (-10,10)

Here, we provide the instructions and quiz questions for the treatments with zero FV and variable dividends. The materials for the other treatments only differ in the value of the dividends and holding costs and available upon request.

Instructions

Thanks for your participation! There are 2 tasks in this experiment, and you will be paid for only 1 randomly chosen task in private and in cash at the end of the experiment. The instructions below describe task 1. The instructions of task 2 will be announced after task 1.

1. General Instructions (Task 1)

Please read this instruction carefully. The amount you will earn depends on the decisions you make, thus a good understanding of the instructions is crucial to your earnings. After reading the instructions, there will be a 2-period practice round to help you familiarize with the experimental software. The earnings in the practice round do not enter your final payment. After the practice round, there will be some questions to ensure that everyone understands the task. The main part will start after everyone has answered the questions correctly.

The game money in this task is called Token and the game commodity that can be traded is called Goods. Prices quoted in the task have two decimal places and Goods must be traded in whole units.

If you have any questions during the task please raise your hand and the experimenter will come to you. Please do not ask your questions out loud, or attempt to communicate with other participants, or look at other participants' computer screens at any time during the task. Please turn your phone to silent mode and place it on the floor.

2. Market Environment

Market setup

In this task, 8 participants (including you) form a market. Each participant can appear as a buyer and a seller at the same time to buy and sell Goods. This market will be open for 15 trading periods. At the beginning of the first period, participants are endowed with Tokens and Goods for trading. Tokens and Goods you hold at the end of each trading period will carry over to the subsequent period. In each trading period, the market will be open for 100 seconds, during which you may buy and sell Goods.

Dividends and holding costs

After each trading period, each Good pays a dividend and incurs a holding cost. The dividend can either be 20 Tokens or 0 Token with equal probability, while the holding cost is always 10 Tokens per Good.

At the end of period 15, each Good pays a dividend and incurs a holding cost (as in previous periods) and then become useless.

To sum up

Each Good earns a dividend minus a holding cost after each trading period ends. The dividend in each period can either be 20 or 0 Tokens per Good, but the holding cost is always 10 Tokens per Goods. So the net benefit of holding a Good is either 10 Tokens or -10 Tokens per period with equal probability. Goods become useless after the payment of the dividend and the charge of the holding cost in period 15.

3. Earnings

The amount of Tokens you will earn in task 1 is equal to:

Tokens you receive at the beginning of the task

+ Dividends you receive throughout the task

- Holding costs you paid throughout the task

+ Tokens received from sales of Goods

− Tokens spent on purchases of Goods

The Tokens will be converted to Australian Dollars with the exchange rate of 50 Tokens=$1. If task 1 is chosen for the payment, then your will receive the amount converted from the Tokens plus a $15 show-up fee.

4. How to Use the Software

There are two ways to buy Goods. Firstly, you can buy from an existing seller by accepting the ask offer made by the seller. To do this, you need to enter the quantity you want to buy and then click the “Buy” button. The quantity you get is the minimum of the quantity offered by the seller and the quantity you entered, and the amount of Token you pay is the ask price times the quantity. You must buy from the lowest (cheapest) ask price.

Secondly, you can make a bid offer to be accepted by potential seller(s). To do this, you need to enter the price and quantity you want to buy and then click the “Bid to buy” button. The quantity you get is the amount accepted by the potential seller(s). If no one accepts your bid offer, then you buy nothing. You must submit a bid price lower than the lowest ask price, otherwise, you should use the “Buy” button.

Similarly, there are two ways to sell Goods. Firstly, you can sell to an existing buyer by accepting the bid offer made by the buyer. To do this, you need to enter the quantity you want to sell and then click the “Sell” button. The quantity you sell is the minimum of the quantity offered by the buyer and the quantity you entered, and the amount of Tokens you receive is the bid price times the quantity. You must sell to the highest (most lucrative) bid price.

Secondly, you can make an ask offer to be accepted by potential buyer(s). To do this, you need to enter the price and quantity you want to sell and then click the “Ask to sell” button. The quantity you sell is the amount accepted by the potential buyer(s). If no one accepts your ask offer, then you sell nothing. You must submit an ask price higher than the highest bid price, otherwise, you should use the “Sell” button.

Bid or ask offers cannot be withdrawn once made and the corresponding Tokens and Goods are frozen. This means that if you have non-transacted ask/bid offers outstanding, then the amount of Goods/Token you may further use to sell/buy is less than the balance that appears on your screen. Non-transacted offers will be cancelled at the end of each trading period and the frozen Tokens and Goods will defreeze. You cannot buy from or sell to yourself. Please turn to the next page to see the window you will see during the trading time.

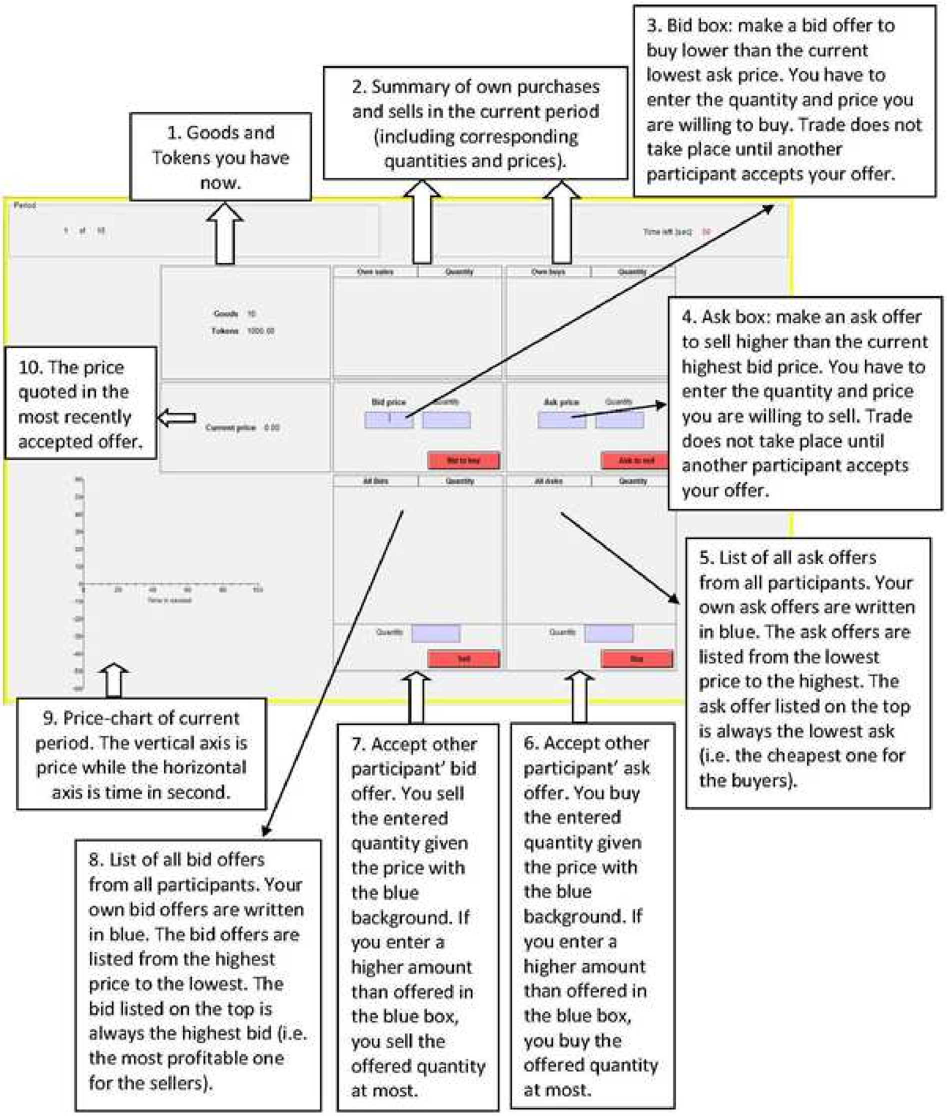
1. Goods and Tokens you have now.
2. Summary of own purchases and sells in the current period (including corresponding quantities and prices).
3. Bid box: make a bid offer to buy lower than the current lowest ask price. You have to enter the quantity and price you are willing to buy. Trade does not take place until another participant accepts your offer.
4. Ask box: make an ask offer to sell higher than the current highest bid price. You have to enter the quantity and price you are willing to sell. Trade does not take place until another participant accepts your offer.
5. List of all ask offers from all participants. Your own ask offers are written in blue. The ask offers are listed from the lowest price to the highest. The ask offer listed on the top is always the lowest ask (i.e. the cheapest one for the buyers).
6. Accept other participant' ask offer. You buy the entered quantity given the price with the blue background. If you enter a higher amount than offered in the blue box, you buy the offered quantity at most.
7. Accept other participant' bid offer. You sell the entered quantity given the price with the blue background. If you enter a higher amount than offered in the blue box, you sell the offered quantity at most.
8. List of all bid offers from all participants. Your own bid offers are written in blue. The bid offers are listed from the highest price to the lowest. The bid listed on the top is always the highest bid (i.e. the most profitable one for the sellers).
9. Price-chart of current period. The vertical axis is price while the horizontal axis is time in second.
10. The price quoted in the most recently accepted offer.

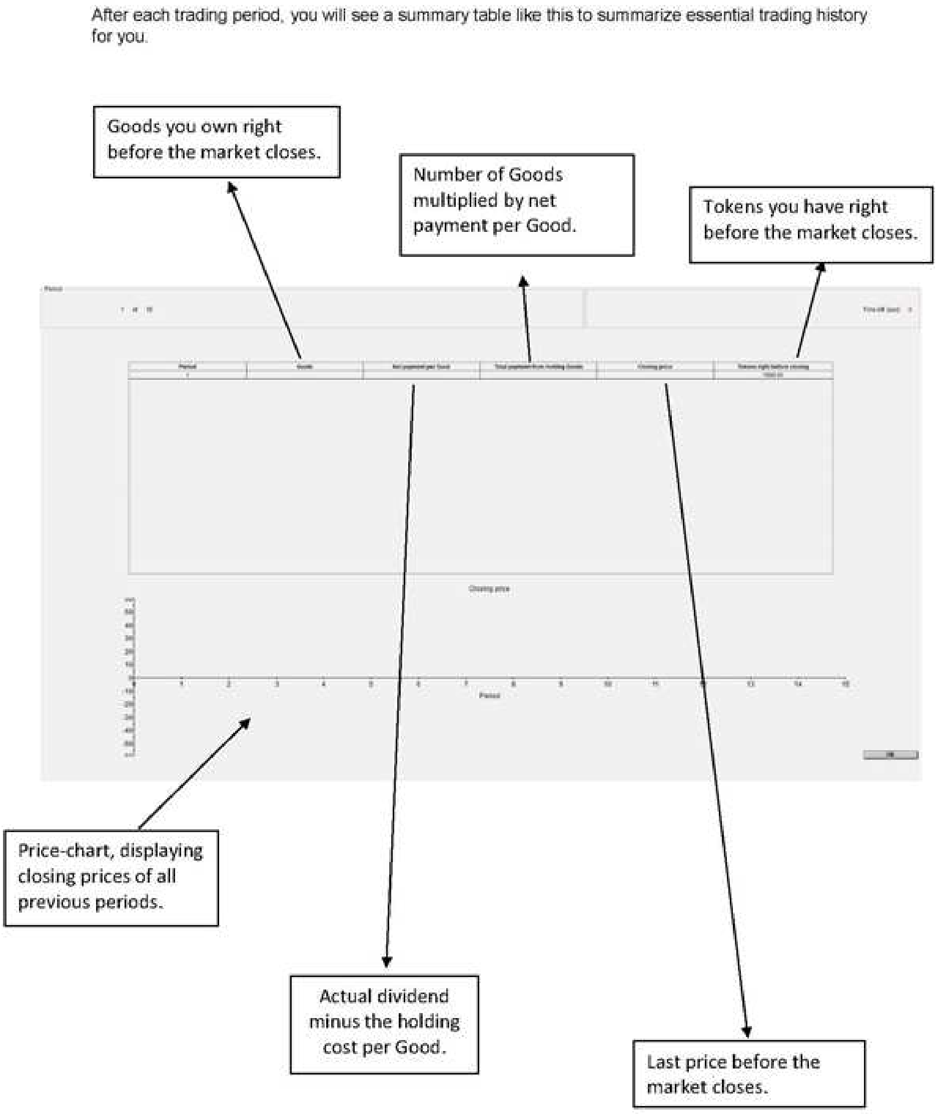

Quiz questions

Q1: If the dividend was 20 Tokens in the previous period, what is the probability that the dividend is 20 Tokens in the current period? Answer: 50%.

If the participant provides a wrong answer, then a message “Your answer is incorrect! Please refer to "Dividends" section in the instruction to find the right answer. Hint: Dividend state in the current period is independent of the dividend in the previous period.” pops up.

If the participant enters a correct answer, then the software displays “Your answer is correct! As mentioned in the "Dividend" section of the instructions, the dividend per Good in each period will be either 0 or 20 tokens with equal chance.”

Q2: If the dividend is 0 in period 15, what is the total amount of Tokens you can get from holding a Good at the end of this period? Answer: -10.

If the participant provides a wrong answer, then a message “Your answer is incorrect! Please refer to "To sum up" part of the instruction to find the right answer. Hint: the number of Tokens you get from holding one Good at the end of period 15 is the dividend at the period minus the holding cost.” pops up.

If the participant enters a correct answer, then the software displays “Your answer is correct! As mentioned in "To sum up" part of the instruction, one Good pays a dividend (0 Token) and incurs a holding cost (10 Tokens) and then becomes useless. So it costs you 10 Tokens from holding a Good at the end of this period.”

Q3: If you entered 3 units and click "Buy" to buy from an ask offer with price 10 Tokens per Good and quantity 1 unit, how much money will you pay for this action? Answer: 10.

If the participant provides a wrong answer, then a message “Your answer is incorrect! As mentioned in "How to use the software" section, if you enter a higher amount than offered, you buy the offered quantity at most. So you only buy 1 unit at the price 10 Tokens per Good.” pops up. If the participant enters a correct answer, then the software displays “Your answer is correct! As mentioned in "How to use the software" section, if you enter a higher amount than offered, you buy the offered quantity at most. So you only buy 1 unit at the price 10 Tokens per Good.”

Q4: Can you make a bid offer with a price higher than the current lowest ask offer? Answer: No.

If the participant provides a wrong answer, then a message “Your answer is incorrect!”

Endowment groups	Assets	Tokens	Number of traders
1	5	1000	4
2	10	1000	4

Notes: This table lists the exact endowment of all 8 participants at the beginning of each trading block.

Table 2.8: Endowments at the beginning of the first trading period

	Trading Period														
Trading block	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	H	L	L	H	H	L	H	L	H	H	H	L	L	H	L
2	L	L	H	H	L	L	L	H	L	L	L	L	H	H	H

Notes: This table lists our dividend states when the dividends are variable. All traders in markets with variable dividends treatments observe the same dividend realization. H and L indicate high and low dividend states, respectively.

Table 2.9: Realized dividend series for variable FV treatments

pops up.

If the participant enters a correct answer, then the software displays "Your answer is correct! You can simply accept the lowest ask offer to earn more!"

2.8.2. Details about endowments and dividends in the main study

In this section, we provide more details about endowments, and realization of the dividends in each trading periods.

2.8.3. Combined analysis for both trading blocks

In this section, we replicate the analysis presented in the main paper but include data from both trading blocks. Results using both trading blocks are qualitatively similar as compared to using only the first block.

Bubbles in markets with zero and negative fundamental values Relative price levels between (2.5,2.5), (0,0), and (-2.5,-2.5) do not change when including the second trading block. The market with zero fundamental value has an average price of 10.1 (Tokens), which is significantly higher than zero providing evidence for significant mispricing ($p<.01$). The average price in market (2.5,2.5) is 23.1 (the average fundamental value is 20), which is statistically higher than in (0,0) ($p=.016$). The average price in (-2.5,-2.5) is -2.89 (the average fundamental value is -20), which is smaller than in the other two asset markets ($p<0.05$ in both cases) but significantly bigger than its FV -20 ($p<.01$).

Finding A1: *There are bubbles in asset markets with zero and negative fundamental values.*

Relative bubble size Bubbles are omnipresent in asset markets with zero and negative fundamental value. More precisely, we observe that there are bubbles in all trading periods of (0, 0) and (-2.5, -2.5). In (2.5,2.5), there are bubbles in 14 out of the 15 trading periods. Importantly, we observe larger bubbles in (0,0) and (-2.5, -2.5) than in (2.5, 2.5) in the first 11 trading periods. The average bubble size is 10.1 in (0,0) and 17.1 in (-2.5, -2.5), which is significantly larger in than in (2.5, 2.5) where it is 3.1 ($p=.1$ for (0, 0) and $p=.078$ for (-2.5, -2.5)).

Table 2.10 correspond to Table 2.4 and compares bubbles in zero and negative FV markets using OLS models of the following regression specification:

$$pd_{g,t} = c + \delta_1 I_{g,t}^{zero} + \delta_2 I_{g,t}^{negative} + \sum_{\tau=2}^{15} \alpha_\tau I_{g,t}^{\tau} + \gamma_1 female_{g,t} + \gamma_2 knowledge_{g,t} + \epsilon_{g,t}, \quad (2.3)$$

We cluster the standard errors at the group level in our analyses and vary control variables in Table 2.10 to check for the robustness of our second finding.

We exclude the control terms for time fixed effects ($\Sigma_{\tau=2}^{15} I_{g,t}^{\tau}$), number of females in the group ($female_{g,t}$) and participants' knowledge on the asset pricing ($knowledge_{g,t}$) in Model 1 and sequentially add these controls from Model 2 to Model 4. The point estimates of treatment effects (δ_1 and δ_2) are all positive and economically large and δ_1 is significantly bigger than zero in all regression specifications. We observe, for example, that in model 1 bubbles are almost two (more than two and half) times as large in the negative (zero) FV markets as in markets with positive FV. In model 4, where time fixed

	Model 1	Model 2	Model 3	Model 4
Zero FV (δ_1)	6.33	6.07	7.05	5.77
	0.133	0.154	0.154	0.183
Negative FV (δ_2)	9.71*	9.64*	9.53*	11.59***
	0.059	0.059	0.061	0.005
Constant (c)	6.15*	16.61***	18.89**	21.79***
	0.069	0.005	0.011	0.003
Time fixed effects ($\Sigma\alpha_\tau I^t$)	No	Yes	Yes	Yes
# of females ($female$)	No	No	Yes	Yes
# have taken asset pricing classes ($knowledge$)	No	No	No	Yes

Notes: This table shows the estimation results for whether FVs have any impact on the bubble size. There are 509 observations for each regression and all standard errors are clustered at the group level (we exclude 31 trading periods where no trade took place in either trading block). We report the p-value under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2.10: Compare the bubble size between treatments with positive, zero, and negative FVs

effects and group heterogeneities are controlled for, the bubble size in zero and negative FV markets are still 26-53% higher than in the positive FV control, in the first trading period.

Finding A2: *Bubbles in asset markets with zero and negative fundamental values are larger than in markets with positive fundamental values.*

Bubbles and variances in the dividends As in the asset markets with fixed dividends, prices start very high and then decline over trading periods. The average price in (-10, 10) is 17.36 (Tokens), which is not statistically different from 26.42 in (-5, 10) (p=.2). The average price in (-15, 10) is 15.8 which is very similar to (-10, 10), but statistically smaller than (-5, 10) (p=.025). The average price in (-10, 10) and (-15, 10) is significantly bigger than 0 (p<.01) and the price is insensitive to the drop of the FV from 0 to -2.5 when the dividend is variable, providing further evidence for bubbles in zero and negative FV asset markets.

When we convert the price level to mispricing, we find bubbles are prevalent in all three treatments. We observe that there are bubbles in all trading periods among all three treatments but in the second period of (-5,10). (-15,10) has the biggest bubbles, followed

	Positive FV	Zero FV	Negative FV
Fixed dividends	3.06	10.12	17.11
Variable dividends with loss possible	6.42	17.36	35.80
p-value	0.1	0.34	0.01

Notes: This table reports the bubble size (price dispersion) of each treatment. The p-value in the bottom row compares the bubble size of the fixed dividend treatment with the variable one at the corresponding level of FV.

Table 2.11: Actual bubble size

by (-10,10) and (-5,10) the smallest bubbles. The average bubble size of (-15,10) is 35.8, which is larger than 17.36 in (-10,10) and 6.42 in (-5,10) (p=.1 for (-10,10) and p=.037 for (-5,10)).

Next, we study whether bubbles increase (hypothesis on complexity) or decrease (loss aversion) when we introduce variances in dividends. We pool the three fixed and three variable dividend treatments and find bubbles are larger when dividends are variable, which is consistent with our complexity hypothesis. The average bubble size is 19.86 with variable dividends compared to only 10.10 when dividends are fixed (p=.018). Table 2.11 show the pattern is qualitatively true regardless of fundamental value (negative, zero, and positive). The price dispersion for (-10,10) is 17.36 while it is 10.12 for (0,0) (p=.34); it is 6.42 for (-5,10) and 3.06 for (2.5,2.5) (p=.1); and it is 35.80 for (-15,10) and 17.11 for (-2.5,-2.5) (p=.01).

We also run OLS regressions to include all dependent observations across trading periods to have a panoramic view on the bubble pattern across treatments. In concrete, we estimate the following model:

$$\begin{aligned} pd_{g,t} = c &+ \beta_1 I_{g,t}^{(-5,10)} + \beta_2 I_{g,t}^{(0,0)} + \beta_3 I_{g,t}^{(-10,10)} + \beta_4 I_{g,t}^{(-2.5,-2.5)} + \beta_5 I_{g,t}^{(-15,10)} \\ &+ \sum_{\tau=2}^{15} \alpha_\tau T_{g,t}^{\tau} + \gamma_1 female_{g,t} + \gamma_2 knowledge_{g,t} + \epsilon_{g,t}. \end{aligned} \tag{2.4}$$

Table 2.12 correspond to table 2.6 and plots the estimation of the regression model with variations on the control variable.

We exclude the control terms for time fixed effects ($\alpha_\tau T_{g,t}^{\tau}$), number of females in the group ($female_{g,t}$) and participants' knowledge on the asset pricing ($knowledge_{g,t}$) in Model 1 and include these controls from Model 2 to Model 4. The relative bubble

	Model 1	Model 2	Model 3	Model 4
Positive Variable	9.17	9.08	8.92	11.82
	0.209	0.218	0.218	0.119
Zero fixed	6.33	5.91	6.92	5.85
	0.119	0.150	0.132	0.163
Zero variable	14.85**	14.88**	15.06**	13.31**
	0.016	0.017	0.017	0.037
Negative fixed	9.71**	9.58**	9.47*	11.05***
	0.048	0.049	0.053	0.007
Negative variable	33.40***	33.32***	33.49***	31.73***
	0.000	0.000	0.000	0.000
Constant (c)	6.15*	21.41***	23.79***	25.78***
	0.057	0.000	0.000	0.000
Time fixed effects ($\Sigma\alpha_\tau\tau^t$)	No	Yes	Yes	Yes
# of females ($female$)	No	No	Yes	Yes
# have taken asset pricing classes ($knowledge$)	No	No	No	Yes

Notes: This table shows the estimation results for whether the nature of the dividends have any impact on the bubble size. There are 1042 observations for each regression and all standard errors are clustered at the group level (we exclude 38 trading periods where no trade took place). We report the p-value under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2.12: Compare the bubble size between treatments with fixed and variable dividends

size (point estimates of βs) are robust to the different control variables. In Model 4, with the most comprehensive controls, our previous results hold: (i) holding the nature of dividend (fixed or variable) constant, the treatment with the lower FV always has the larger bubble size, (ii) fixing the level of FV, the price is always higher in the variable dividend treatment. In concrete, for bubbles sizes we find the following ordering: (2.5,2.5)<(0,0)<(-2.5,-2.5)<(-5,-10)<(-10,10)<(-15,10).

The point estimates from the four models also show that the treatment effects are meaningful and large. For example, in Model 4, the point estimate for positive variable dividend treatment ($I_{g,t}^{(-5,10)}$) is 11.82, implying that the bubble size of (-5,10) is more than 46% bigger than that of (2.5,2.5) in the first trading period.

Finding A3: *Bubbles are larger in asset markets with variable dividends than in asset markets with fixed dividends. Bubbles are largest in asset markets with negative fundamental value and variable dividends and smallest in asset markets positive fundamental value and fixed dividends.*

	Positive	Zero	Negative
Fixed	38.62	55.02	41.88
Variable	40.42	42.22	45.08

Notes: This table shows the average period trading volume in each treatment in trading block 1.

Table 2.13: Compare the trading volume across treatments (block 1)

	Positive	Zero	Negative
Fixed	33.2	42.63	37.55
Variable	35.87	36.39	38.1

Notes: This table shows the average period trading volume in each treatment in trading block 1 and 2.

Table 2.14: Compare the trading volume across treatments (block 1 & 2)

2.8.4. Transaction volume

Table 2.13 shows the average trading volume per period in all 6 treatments in trading block 1. On average, 43.87 shares were bought and sold in each trading period. There is no clear pattern across treatments and Mann-Whitney tests indicate that there is no treatment effect between any of these two treatments (p>.1 for all pairwise comparison).

Table 2.14 shows the average trading volume per period in all 6 treatments for the two trading blocks pooled. On average, 37.29 shares were bought and sold in each trading period. Same as the data from only block 1, there is no clear pattern across treatments and Mann-Whitney tests confirm this (p>.1 for all pairwise comparisons).

2.8.5. Instructions of the guessing game

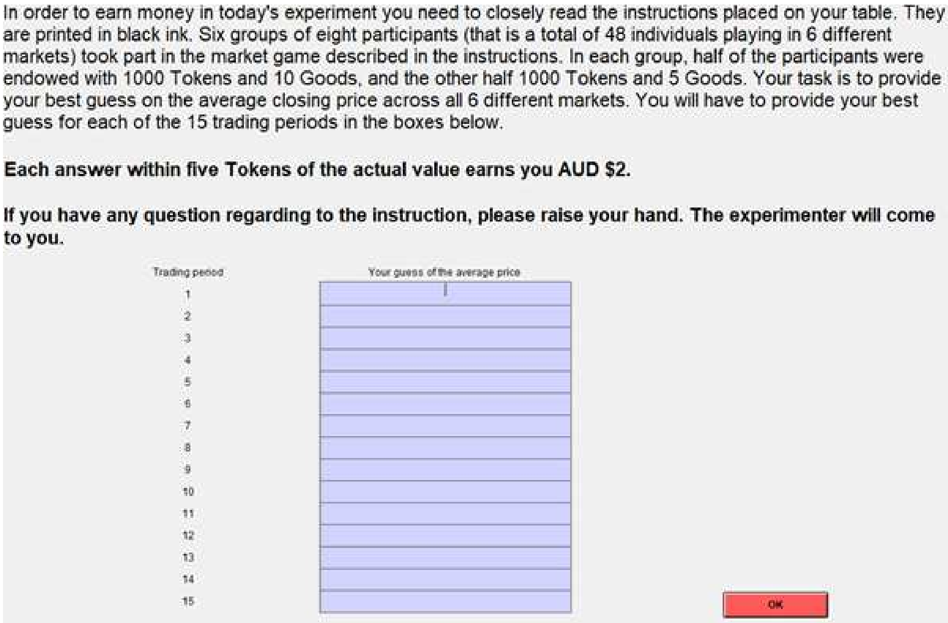
In order to earn money in today's experiment you need to closely read the instructions placed on your table. They are printed in black ink. Six groups of eight participants (that is a total of 48 individuals playing in 6 different markets) took part in the market game described in the instructions. In each group, half of the participants were endowed with 1000 Tokens and 10 Goods, and the other half 1000 Tokens and 5 Goods. Your task is to provide your best guess on the average closing price across all 6 different markets. You will have to provide your best guess for each of the 15 trading periods in the boxes below.

Each answer within five Tokens of the actual value earns you AUD $2.

If you have any question regarding to the instruction, please raise your hand. The experimenter will come to you.

Trading period	Your guess of the average price
1	
2	
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	
15	

OK

Chapter 3

Tournaments with Safeguards: A Blessing or a Curse for Women?

[1]

3.1. Introduction

There is little disagreement that gender equality is desirable in labor markets, yet there are still significant gender differences in key labor outcomes. For example, of the CEOs who lead the companies that compose the 2018 Fortune 500 list, less than 5% are women. Such low representation of women at the CEO level may be partly attributed to tournaments, because females tend to shy away from competitions (Niederle & Vesterlund 2007) and thrive less when they compete (Gneezy, Niederle, & Rustichini 2003; Gneezy & Rustichini 2004). A few mechanisms that are studied in the literature encourage women to join tournaments ranging from team competitions to individual competitions (Dargnies 2012; Flory, Leibbrandt, & List 2015; Healy & Pate 2011) to gender quotas (Leibbrandt, Wang, & Foo 2018; Niederle, Segal, & Vesterlund 2013) and making tournaments the default choice (Erkal, Gangadharan, & Xiao 2019). However, luring women into tournaments may

[1]This chapter reports research conducted jointly with Andreas Leibbrandt.

not be sufficient to narrow gender gaps. Once they have participated in a tournament, women need a competitive environment where they can thrive.

Many organizations employ tournaments, but they vary the risk exposure for low-ranked employees. At one extreme, some use up or out schemes while others use different types of safeguards to protect low-ranked employees. For example, organizations can provide tenure and significant base salaries so that the consequences for even the lowest-ranked workers are moderate.[2] However, it is difficult to identify whether such safeguards are useful in tackling gender gaps in labor outcomes as their implementation is typically highly correlated with organizational and industry characteristics.

In this experimental study, we explore a tournament environment that may help women without hurting men. We design a real-effort rank-order tournament with an elective safeguard, a device that mitigates the consequences of being low ranked (defined by wage rank). More precisely, workers are informed of the availability of this safeguard and decide whether to obtain it before the start of the tournament. The safeguard is complementary and increases the user's minimum wage if their relative performance falls into the lowest ranking category. Our conjecture is that the safeguard is particularly popular among women, improves women's outcomes relative to men's outcomes, and alleviates psychological competitive pressure to perform.

Our findings show that women are indeed more likely than men to select a complementary safeguard and that risk-aversion is an important predictor for the choice of safeguard. However, we also find that this safeguard increases the gender wage gap compared to more standard tournaments. Further, we observe that the safeguard reduces the performance of both women and men regardless of whether it was voluntarily selected or automatically implemented. Our survey findings suggest that a safeguard does not alleviate pressure more for women than for men, and the safeguard encourages genders to slack off.

Our study is closely related to the literature on gender differences in the selection of incentive schemes (Dohmen & Falk 2011; Eckel & Grossman 2008; Flory, Leibbrandt, & List 2015; Gneezy, Leonard, & List 2009). This literature provides evidence that

[2] For example, in many countries, government jobs are more secure than jobs that require similar skills in the private sector because the government sector usually has a lower dismissal rate.

women are more likely than men to be disadvantaged by competitive environments such as workplace tournaments (Gneezy, Niederle, & Rustichini 2003; Gneezy & Rustichini 2004; Niederle & Vesterlund 2007). We depart from this literature by studying an environment where competition cannot be avoided altogether. Such environments are common in hierarchical organizations where promotion to higher levels is often based on competition.

We also contribute to the literature on gender quotas and other policies aimed at improving female labor outcomes (Dargnies 2012; Healy & Pate 2011; Niederle, Segal, & Vesterlund 2013). While there is some evidence that gender quota and policies can address gender gaps (Erkal, Gangadharan, & Xiao 2019), there is also evidence that they might backfire (Gangadharan, Jain, Maitra & Vecci 2016; Leibbrandt, Wang, & Foo 2018; Leibbrandt & List 2018). Our intervention deviates from most affirmative action policies because it does not treat women differently than men but only provides an additional choice that is accessible to either gender. Nevertheless, we find that our instrument can also backfire because the difference in the selection process distorts the incentives disproportionately across genders. The results provide evidence that supports the need for caution when designing seemingly harmless policies that may attract one gender more than the other.

In addition, we contribute to the tournament literature.[3] A key subject in this literature is how individuals react to exogenously imposed variations in the prize structure (see e.g., Harbring & Irlenbusch 2003; Moldovanu & Sela 2001; Orrison, Schotter, & Weigelt 2004; Sheremeta 2011) and how this relates to behavioral aspects (e.g., Delfgaauw, Dur, Sol, & Verbeke 2013; Hong, Hossain, & List 2015; Sheremeta 2015). We complement this literature by investigating reactions to safeguards that allow for endogenous selection of the lowest prize and how this relates to gender differences in tournament outcomes. Finally, our work is related to the literature studying the link between incentive schemes and labor productivity (see e.g., Anderhub, Gachter & Konigstein 2002; Corgnet & Hernan-Gonzalez 2019; Corgneta, Fehr & Goette 2007; Gomez-Minambresbc & Hernan-Gonzalezd 2018). In a typical setup in this literature, workers cannot change or choose their incentive schemes. We fill this gap by giving workers the freedom to choose between two incentive schemes.

[3]For a survey on the tournament literature, the reader is referred to Dechenaux, Kovenock, & Sheremeta (2015).

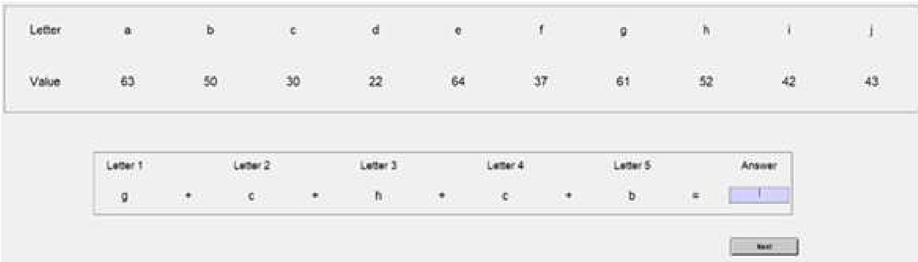

Notes: the figure provides a screenshot of an example puzzle that workers have to solve during the experiment. They first need to decipher the code in the lower box using the upper table and then sum up all the values.

Figure 3.1: Illustration of task

3.2. Experimental Design

3.2.1. The real-effort task

In this experiment, we used a real-effort task, which mimics tedious work assignments that require focus as well as mathematical and verbal skills (see, e.g., Carpenter, Matthews, & Schirm 2010; Erkal, Gangadharan, & Nikiforakis 2011; Gill & Prowse 2012; Leibbrandt, Wang, & Foo 2018;). In this task, participants were asked to solve as many puzzles as possible in a time span of 40 minutes. Fig. 3.1 provides one example of such a puzzle. Participants first deciphered the value of five letters and then summed up the corresponding values. In this example, letters 1 through 5 have values 61 (g), 30 (c), 52 (h), 30 (c) and 50 (b), respectively, and the correct answer is 223 (61+30+52+30+50). After typing their answers in the answer box and pressing the Next button, the computer displayed a different set of numbers until the participants ran out of time.[4] Participants were not allowed to use calculators but could write on scrap paper.

3.2.2. The rank-order tournament

All participants took part in a rank-order tournament and their performance (number of correctly solved puzzles) relative to other participants determined their monetary compensation. There were three compensation levels: (i) the top 10% of participants in

[4]To ensure each quiz had similar difficulty across workers and treatments, all values were two-digit integers, and all letters in the encode game were randomly drawn between a to j.

	Choice (n=160)		Baseline (n=141)	Compulsory (n=130)
	No safeguard	Safeguard		
Top 10%	\$60	\$60	\$60	\$60
Top 10%-50%	\$30	\$30	\$30	\$30
Bottom 50%	\$15	\$20	\$15	\$20

Notes: This table shows the tournament payment structure for each of the three treatments.

Table 3.1: Treatment overview

a given session receive \$60, (ii) the top 10 to 50% of participants receive \$30, and (iii) the remaining participants (bottom 50%) receive \$15 or \$20, depending on treatment and choice. Thus, the tournaments used a multiple-prize payment rule common to many workplaces (Cason, Masters, & Sheremeta 2010; Kalra & Shi 2001; Vandegrift, Yavas, & Brown 2007).[5]

3.2.3. The treatments

This study had three main treatments. The choice treatment allowed participants to select their own tournament incentives. To properly identify the impact of this choice, there were two other treatments: compulsory and baseline. All participants took part in one treatment only. The experiment instructions are in the Appendix.

Choice treatment

In this treatment, workers chose whether to have a complementary safeguard before the task started. The safeguard guaranteed a higher minimum payment if performance was in the lower half but did not affect payment if performance was in the upper half. More precisely, as can be seen in Table 3.1, if a participant chose the safeguard, they received \$20 if their performance was in the bottom 50% instead of only \$15 if they did not choose the safeguard. If their performance was in the top 50%, their payment was not affected.

[5] When the cut off between prize levels was not an integer, we rounded up the number of workers who received the better prize in favor of workers. For example, if there were 28 workers in a group, then the top three performers received \$60.

Compulsory treatment

In this treatment, the safeguard was already embedded in the incentive structure for all participants. That is, the minimum payment was $20, just as it was in the choice treatment if participants chose a safeguard. The safeguard guaranteed a minimum payment of $20 for all participants. We tested two versions of this treatment. In one frame, we made the safeguard explicit (i.e., we told participants that the minimum compensation was only $15 in some of the other sessions). In the other frame, the safeguard was implicit and unknown to the participants (i.e., we gave workers no reference concerning the other group's prize structures).

Baseline

In this treatment, there was no safeguard available, and the bottom 50% received $15. Participants were unaware that participants in other sessions had access to a safeguard.

3.2.4. Conjectures

The standard prediction for all three treatments was that all workers would be incentivized to provide effort and thus increase their likelihood of higher compensation. Thus, the first conjecture was that there are no treatment differences in effort across treatments. The second conjecture applied to the choice treatment, for which we conjectured that all participants chose the safeguard as it weakly dominated not choosing the safeguard.

These standard predictions did not take into account mental effort costs associated with fatigue and the temptation to rest. While tournaments provide a significant incentive to work hard (possible tripling to quadrupling compensation), it is possible that the cost of the mental effort for some participants is sufficiently high that it prevents them from providing (maximal) effort. Accordingly, the alternative prediction is that effort is lower for the compulsory treatment than it is for the baseline because the incentive is less pronounced. Further, it is possible that some participants rejected the safeguard if they believed that it might undermine their effort and opportunity to increase their payments. On the other hand, it is also possible that the safeguard relieves pressure and stress and thus affects performance. Whether less pressure and less stress increase or decrease performance is still an open question, although some progress has been made in this regard (e.g., Allen, Hitt, & Greer 1982; Compte & Postlewaite 2004; Hall & Lawler 1971; Harbring & Irlenbusch 2003; Van Dijk, Sonnemans, & Van Winden 2001).

So far, these predictions do not take gender into account. However, there is evidence that gender plays a crucial role in tournaments. In particular, there is evidence that women are more risk averse than men (Charness & Gneezy 2012; Croson & Gneezy 2009), less likely to enter tournaments (Flory, Leibbrandt, & List 2015; Gneezy, Leonard, & List 2009; Niederle & Vesterlund 2007), and underperform in competitions (Gneezy, Niederle, & Rustichini 2003; Gneezy & Rustichini 2004). For the choice treatment, we allowed for endogenous selection in two different payment rules, which permitted workers to limit their risk exposure to low compensation. Thus, our conjecture is that women are more likely than men to choose the safeguard. In turn, this may improve or harm their compensation relative to men depending on whether the safeguard increases the temptation to rest and has a gender-dependent impact on pressure and stress.

3.2.5. Experimental procedures

The experiment was programmed in z-Tree (Fischbacher 2007), and workers were recruited with the software SONA. In total, 431 workers (160 in choice, 130 in compulsory, and 141 in baseline) participated in the three treatments, and they earned, on average, $32 for an approximate 70-minute experimental session.[6] There was no mentioning of gender throughout the entire experiment to prevent potential experimenter demand effects.

Before the start of the experiment, workers read an information sheet and signed a consent form. Thereafter, we read the instructions aloud. After reading the instructions, workers had time to read the instructions on their own and ask questions. We then implemented a practice round that included identical questions to the real-effort task and quiz questions to ensure that workers understood the instructions and payment mechanism. After the

[6] The participants are mainly university students recruited across different disciplines. The demographic variables are balanced across treatments and are reported in the Appendix.

participants answered these questions correctly, they started the real-effort task. When all workers had completed the task, we administered a short post-experiment questionnaire to conclude. In the questionnaire, we collected the workers demographics, self-evaluation of psychological wellbeing during the experiment, and incentivized them to reveal their beliefs about their own performance, the group average performance.[7]

In addition, we elicited in this questionnaire risk preferences. We used the risk task proposed by Eckel & Grossman (2002). Each worker must choose one gamble out of six gambles with different risk exposure. The realization of the gamble prize is added to the worker's final payment. The choice is recorded as an ordinal categorical variable can be any integers between 1 to 6, where a larger number indicates less risk aversion.

3.3. Findings

3.3.1. Experimental findings

We observe high effort levels in our experiment. Fig. 3.2 illustrates the average number of attempted questions (effort) and correct questions (performance) across treatments. On average, 72.6 quizzes were attempted and 65.4 were solved correctly during the course of 40 minutes, and we observe few cases (7.9%) where participants attempted less than one quiz per minute. This suggests that the large majority of the participants made a substantial effort to do well in the experiment.

We found significant treatment differences in effort levels. In particular, we observed that workers attempted less quizzes in the compulsory treatment[8] than in the baseline (78.1 vs 66.7, $p<0.001$) and that the performance was also significantly lower (70.1 vs 60.7, $p<0.001$).[9] This finding is consistent with the alternative prediction based on fatigue and temptation but is inconsistent with the standard prediction of no treatment differences in

[7] We used the binarized scoring rule proposed by Hossain & Okui (2013) to incentivize participants to report their best guess about their own performance and the group average performance. Each answer within 5 quizzes around the true value is awarded with $2.

[8] For simplicity, we pool in the main analysis the two versions in the compulsory treatment. The findings in the two versions are qualitatively similar and discussed in Section 3.3.3.

[9] Without further specification, all p-values reported in this paper are from two-tail Mann-Whitney tests.

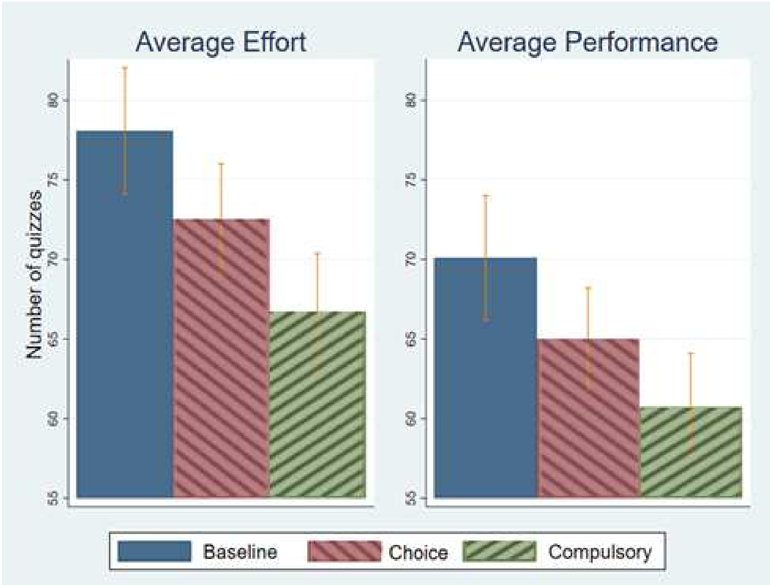

Notes: This figure shows the average attempts and performance depending on treatment. Participants had 40 minutes to attempt as many quizzes as possible. Error bars indicate 95% confidence intervals.

Figure 3.2: Effort and performance levels depending on treatment

effort levels across treatments. In addition, we observed that participants attempted less quizzes in the choice treatment than in the baseline (72.5 vs 78.1, p=0.037) and that the performance was also lower (65 vs. 70.1, p=0.062). Additionally, there were significant differences between the compulsory and choice treatments for both effort (66.7 vs 72.5, p=0.03) and performance (60.7 vs 65.0, p=0.095).

Table 3.2 reports the estimation results of an OLS regression with effort and performance as dependent variables. We observed the following. The compulsory safeguard significantly reduced the effort and performance by around 14.5% ($p<0.01$) and 13.2% ($p<0.01$), respectively, compared with the baseline. For workers in the choice treatment, those who chose to opt against the safeguard performed similarly to the baseline (p=0.61 for effort and p=0.37 for performance); those who chose to use the safeguard have a similar performance to those in the compulsory treatment(p=0.154) but exert more effort compared with those in the compulsory treatment (p=0.0799), and their performance was 8.1% to 8.3% worse than the baseline in terms of effort and performance ($p<0.05$ for both cases). Finally, the female dummy was significant at the 5% level and shows that men made more effort and had a higher performance ($p<0.05$ for both effort and performance).

	Performance	Effort
Compulsory	-8.985***	-10.962***
	(2.628)	(2.753)
Choice × not to use safeguard	-4.374	-2.627
	(4.837)	(5.136)
Choice × use safeguard	-5.480**	-6.283**
	(2.626)	(2.738)
Female	-4.961**	-5.259**
	(2.077)	(2.390)
Constant	72.555***	80.689***
	(2.264)	(2.309)

Notes: This table shows the OLS estimation results of the regression comparing the outcomes across treatments and the safeguard choice. Performance is defined by the number of correctly solved quizzes. Effort defines the number of attempted quizzes. The constant corresponds to the outcomes of male workers in the baseline, and all other independent variables are dummy variables. There are 431 observations for each regression. We report the estimate result of the regression models and include the robust standard error in parentheses under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 3.2: Effort and performance depending on treatment

Result 1: *The safeguard reduced effort regardless of whether it was implemented as a choice or as a default. Individuals who decided against using a safeguard provided more effort than those who decided to use a safeguard. Only workers that decided against using a safeguard performed similarly well as workers who did not have a safeguard.*

The safeguard was popular despite the negative impact on effort and performance (87% choose the safeguard). Importantly, and as conjectured, we found that women were more likely to choose the safeguard than men (93.1% vs. 81.8%, p=0.0368).

Table 3.3 compares the characteristics of the participants who chose the safeguard with the characteristics of those who did not choose the safeguard. We can see that the fraction of women is much lower in the sample of participants who did not choose the safeguard (23.8%) than in the sample of participants who chose the safeguard (48.2%). In addition, risk preference is also a significant predictor of safeguard choice. Participants make more risk averse choices in the risk task are more likely to be in the sample of those who chose the safeguard (p=0.02). As expected, we also observe gender differences in the willingness to take risks with females being more risk-averse (3.52 vs. 4.12, p<0.01). Characteristics other than gender and risk preference play less important roles in the safeguard choice.

Variable	Use safeguard (n=139)	Do not use safeguard (n=21)	Difference in absolute terms	p-value
Female (dummy)	48.2%	23.8%	24.4%	0.04
Risk	4.01	4.86	-0.85	0.02
Good at math (dummy)	64.7%	66.7%	-2.0%	0.86
Confidence $\left(\frac{\text{guessed own performance}}{\text{guessed group average}}\right)$	1.20	1.38	-0.18	0.30

Notes: This table compares key variables between the groups who choose to use the safeguard against those who opt out within the choice treatment. All these variables are collected by the post-experiment survey where we asked the participants questions on their gender, risk preference (Eckel & Grossman 2002), subjective belief about their mathematical abilities, and their estimate of their own performance relative to the group average. Female is a dummy variable that takes the value 1 if and only if the participant is a female; Risk is an ordinal categorical variable that can be any integers between 1 to 6, where a larger number indicates less risk aversion; Good at math is a dummy variable that takes the value 1 if and only if the participant's answer to this question is yes; Confidence is a continuous variable calculated as the ratio of the participant's guess about their own performance and their guess about the average performance in the respective experimental session. The numbers in the second and third columns in the table are the average values of the variables (Male, Risk > medium, Good at math, Confidence) of the group that choose to use the safeguard and the group that chooses not to use the safeguard, respectively. Each p-value reports the Mann-Whitney test result testing the null hypothesis that there is no difference between the group that uses the safeguard and the group that does not use the safeguard.

Table 3.3: Determinants of safeguard choice

	Choice of safeguard			
	Model 1	Model 2	Model 3	Model 4
Female (dummy)	0.112**		0.095**	0.097*
	(0.051)		(0.048)	(0.052)
Risk (categorical)		-0.035**	-0.030**	-0.027*
		(0.015)	(0.015)	(0.015)
Good at math (dummy)				0.019
				(0.058)
Confidence (guessed own performance / guessed group average)				-0.104
				(0.089)
Constant	0.818***	1.012***	0.949***	1.054***
	(0.030)	(0.059)	(0.068)	(0.131)

Notes: This table shows the OLS estimation of the linear probability model studying how individual characteristics affect the safeguard choice. Female is a dummy variable that takes the value 1 if and only if the participant is a female; Risk is an ordinal categorical variable that can be any integers between 1 to 6, where a larger number indicates less risk aversion. ; Good at math is a dummy variable that takes the value 1 if and only if the participant's answer to this question is yes; Confidence is a continuous variable calculated as the ratio of the participant's guess about their own performance and their guess about the average performance in the respective experimental session. There are 160 observations for the regression. We include the robust standard error in parentheses under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 3.4: Individual determinants of the safeguard choice

Task ability (math skills) and confidence (the ratio of guessed own performance and guessed group performance) are not significant drivers of safeguard choice.

In Table 3.4, we regress the choice of safeguard on these characteristics (gender, risk preferences, math skills, confidence) sequentially and simultaneously.[10] We observe that both gender and risk preference remain significant predictors in all models even if both are used simultaneously (p=0.03-0.06 for gender and p=0.03-0.07 for risk preference). In particular, in Model 1 and 2, we regress the choice of safeguard on gender and risk preference alone, respectively. We find females are 11.2 percentage points more likely to use the safeguard compared to males (p=0.03), and those choose the gamble with the risk level one number higher are 3.5 percentage points less likely to use the safeguard (p=0.03). These results are robust if add additional control variables, task ability and confidence, in Model 4 (despite gender and risk are only significant at the 10% level).

[10] We estimated the equation using a Linear Probability Model (LPM) method because it gives us the marginal effects directly. However, one drawback of LPM is that there are no probability bounds, so the estimated probability can be greater than 1 as in model 2 and 4.

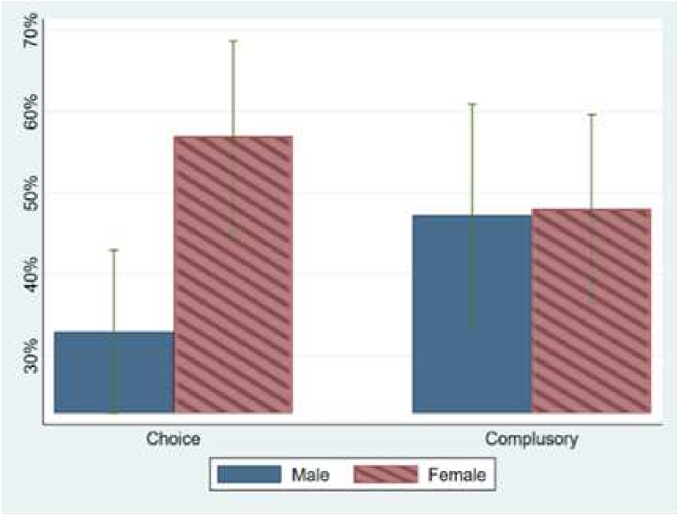

Notes: This figure illustrates the percentage of workers who received a higher minimum wage due to the materialization of the safeguard.

Figure 3.3: Materialization of safeguard depending on treatment and gender

Result 2: *Gender is an important predictor of safeguard choice. Women are more likely to choose a complementary safeguard in tournaments.*

The safeguard benefits women only at first sight. Fig. 3.3 shows how likely the safeguard is to materialize and increase the compensation in the choice treatment for each gender. The two bars on the left side show that in this treatment, 56.9% of the female workers and only 33% of the male workers receive the additional $5 from the materialization of the safeguard (p<0.01). In contrast, in the two bars on the right side for the compulsory treatment, we see that gender plays no role as an almost identical percentage of either gender receiving the safeguard payment of $20.

A closer look at the data reveals that women fare worse than men when there is a choice for safeguard. Fig. 3.4 compares the compensation for male and female workers across treatments. Compared to the baseline, we observed that compensation was higher for both men ($27.0 vs $30.3, p=0.013) and women ($24.9 vs $26.0, p=0.024) in the choice treatment, however, the treatment impact was gender specific. Men's average wages increased by $3.3 (12.2%) while women's wages only increased by $1.1 (4.4%). In addition, compared with the compulsory treatment, giving workers the choice of safeguard slightly increased men's payoff ($29.6 vs $30.3, p=0.60) but somewhat decreased women's payoff ($27.6 vs $26.0, p=0.138). Perhaps most importantly, we observed that the choice treatment created a significant gender wage gap of $4.2 (wage for women=$26, men=$30,

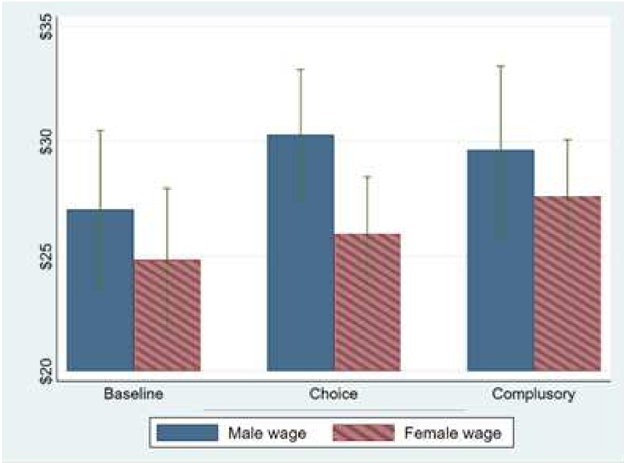

Notes: This figure shows the average wage for each gender across treatments. The error bars indicate 95% confidence intervals.

Figure 3.4: Average wage by gender in each treatment

p=0.02), which was insignificant in the other two treatments (p>0.38).

Fig. 3.5 provides insights as to why the safeguard choice backfires for women. This figure illustrates the ranking and prizes in the different treatments depending on gender. The middle panel shows that women were much more likely to score in the bottom 50% in the choice treatment than men, and that this was not paralleled in the compulsory and baseline treatment. In fact, such a gender difference was significant for the choice treatment (p=0.027) but not for the baseline (p=0.28) and compulsory treatment (p=1.0).

Fig. 3.6 provides a more fine-grained illustration of effort levels depending on gender and treatment. In this figure, we compared the effort level of each quantile in the no baseline against the same quantile in the choice treatment. While this quantile plot does not show the change in the effort level caused by our intervention at each quantile as we use a between subject design, it allows for causal interpretation at the aggregate level. For men, the choice of safeguard shifted the overall distribution of correct answers lower than the baseline, except for the quantiles near the two prize cut-offs. Women, in contrast, appeared to respond differently to the safeguard choice. All quantiles in the choice treatment were higher in the baseline, except for the few that were between the cut-offs. Thus, the choice of safeguard appeared to cause most men and women to exert less effort compared to the baseline, but this reduction was more pronounced for women whose performance was close and moved them up the ranks.

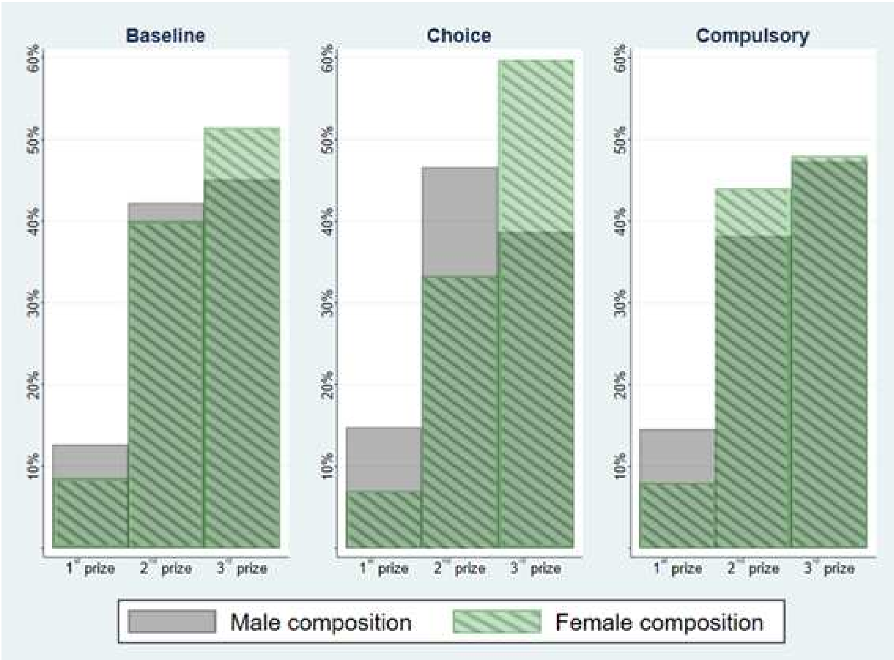

Notes: This figure plots the gender composition of each prize level in all treatments. The first price is awarded to the top 10% performers, the second is awarded to the top 10% to 50% performers, and the third is awarded to the bottom 50%.

Figure 3.5: The gender composition of each prize level for each treatment

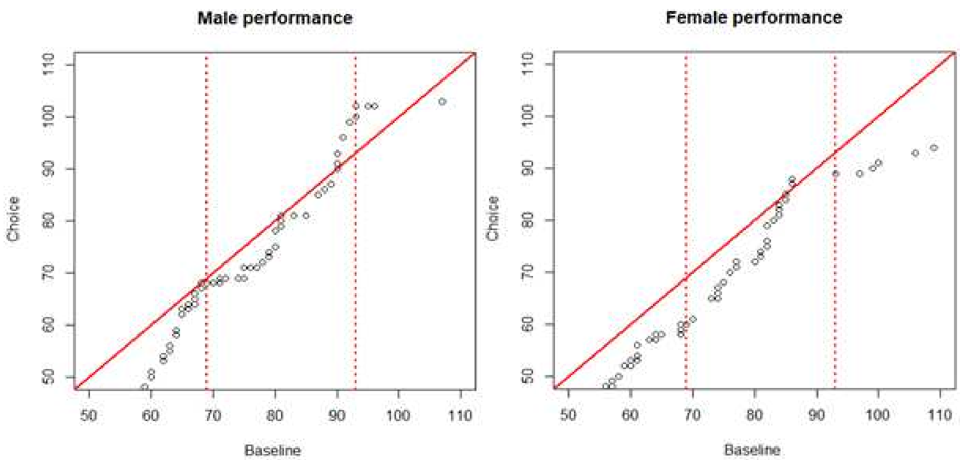

Notes: This figure compares the performance in the baseline with the choice treatment for both genders. Each quantile of the baseline is matched with that of the choice treatment and is plotted against the 45-degree line. Any point above the line indicates that the choice treatment has a higher performance for that particular quantile and vice versa. The two dashed vertical lines indicate the 50% and 10% cut-offs for the different prizes in the baseline.

Figure 3.6: Performance differences between choice and baseline treatment depending on gender

Result 3: *The availability of a complementary safeguard disadvantaged women more than men and created a gender wage gap.*

3.3.2. Survey findings

The experimental findings are consistent with our alternative conjecture based on the assumption that the safeguard increases temptation to rest and decreases stress. To provide insights on these potential underlying mechanisms and their relationships with gender, we conducted a survey with the participants after the tournament but before revealing information on their performance. In this survey, we asked participants to self report their stress levels and temptation to rest during the real-effort task and how a removal of the safeguard would change these factors.[11] The elicitation of stress level and temptation to rest was not incentivized. Participants had to choose between five stress levels ("Not at all", "A bit", "Somewhat", "Stressful", "Very stressful) and three temptation levels ("Not at all", "Somewhat", "A lot") and indicate how would a removal of safeguard change these levels from five choices ("Decrease by a lot", "decrease by a bit", "No Change", "Increase by a bit", "Increase by a lot"). A typical worker experience a bit stress during the task and somewhat temptation to rest.

Fig. 3.7 illustrates several corresponding survey insights. First, while we found that only a minority of workers experienced high levels of stress and temptation, many reported experiencing some stress and the temptation to rest suggesting the presence of mental costs. Second, we observed that a significant proportion of participants expected to experience higher stress level (45%) and less temptation to rest (35%) when there was no safeguard. Third, and perhaps most importantly, the psychological impact of the safeguard appeared to be gender independent. The reported stress and temptation levels were similar for men and women in the presence and absence of a safeguard ($p=0.91$ for change in stress and $p=0.465$ for change in temptations), suggesting that mental costs cannot explain the observed gender differences in safeguard choice, performance, and wage levels.

[11] We only asked participants in the compulsory treatment because all of them had experienced the safeguard. Experiencing the safeguard may facilitate the participants to understand the question and give more reliable answers.

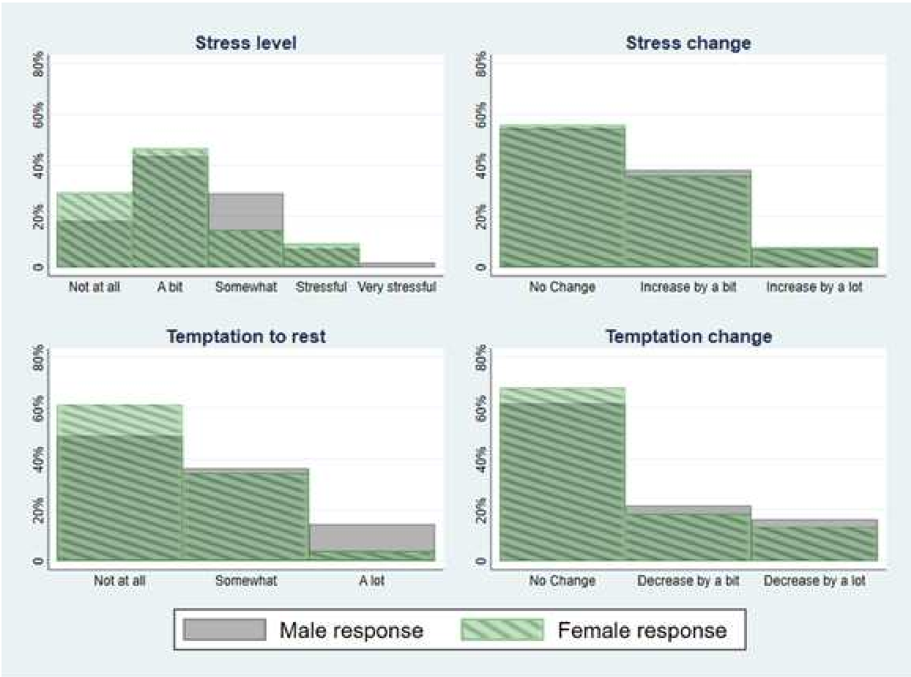

Notes: This figure shows the survey response regarding the stress level and temptation to rest during the tournament and how participants expect a removal of the safeguard to change these between male and female workers. All survey questions are implemented in the form of multiple-choice questions, and the possible choices map to the labels in the subfigures. The first row of subfigures plots the reported stress level and the expected change in the stress level from using to not using the safeguard. The second row shows the reported temptation to take a rest during the tournament and the expected change in the temptation from opting for the safeguard to opting against it. We collect data from the compulsory treatment to ensure that all participants have experienced the safeguard.

Figure 3.7: Treatment impact on stress and temptation depending on gender

3.4. Discussion

There is substantial evidence that uncertainty in labor relations is more detrimental to women than men (e.g., Frederiksen 2008; Garcia 2017; Hirsch & Schnabel 2012). There is also evidence that giving workers choice over their employment conditions can improve labor outcomes (e.g., Beckmann, Cornelissen, & Krakel 2017; Bloom, Garicano, Sadun, & Van Renee 2014; Leslie, Manchester, Park, & Mehng 2012). We investigate a workplace tournament that reduces uncertainty by providing workers with a safeguard choice that increases the minimum wage. Our findings suggest that giving workers the autonomy to select the incentive scheme can disadvantage women. This is because women are more tempted to choose safeguards, although they weaken the incentive to exert effort. Our findings provide novel evidence on the limitations of tournaments in creating gender neutral outcomes.

Our experiment also provides insights for the literature on incentive contracts beyond the economics of genders. First, providing safeguards to low-performing workers appears to be costly but also counter-productive as they lower effort and performance. At the same time, such safeguards can even have a detrimental impact on high-performing workers and, thus, decrease effort and performance throughout the whole distribution of workers. More generally, our study contributes to the discussion of optimal incentives for policies that target low-performing individuals (e.g., Heckman 2006; Mario, Nicole, Ute, Lisa, & Benjamin 2020; Rosen 1986). For example, poverty alleviation programs are more likely to affect the wages of low-skilled workers, and it is of key importance to understand whether and under what circumstances they cause less effort and thus might lead to larger, not smaller, wage gaps.[12]

Our study is a novel attempt to study the role of endogenous choice and safeguards in tournaments. We envision several extensions for future research. First, it seems important to further investigate in which environments workers benefit from having a choice in their compensation scheme and whether it is a general property that giving this choice is less beneficial for women than men. Second, our data suggest that opting against the

[12] However, Dariel et al. 2020 find that many of the gender gaps in competitiveness and other economic preferences are predominantly found inside the laboratory but not in real-world organizations. Cautions should Caution should be exerted to apply the results from the lab to other contexts.

safeguard serves as a self-control mechanism. It may be interesting to study ways to make individuals "burn the boat" to achieve greater success.

3.5. References

Anderhub, V., Gachter, S., & Konigstein, M. (2002). Efficient contracting and fair play in a simple principal-agent experiment. Experimental Economics, 5(1), pp.5-27.

Allen, R. D., Hitt, M. A., & Greer, C. R. (1982). Occupational stress and perceived organizational effectiveness in formal groups: An examination of stress level and stress type. Personnel Psychology, 35(2), 359–370.

Beckmann, M., Cornelissen, T., & Krakel, M. (2017). Self-managed working time and employee effort: Theory and evidence. Journal of Economic Behavior & Organization, 133, 285–302.

Bloom, N., Garicano, L., Sadun, R., & Van Reneen, J. (2014). The distinct effects of information technology and communication technology on firm organization. Management Science, 60(12): 2859–2885.

Carpenter, J., Matthews, P. H., & Schirm, J. (2010). Tournaments and office politics: Evidence from a real effort experiment. American Economic Review, 100(1), 504–517.

Cason, T. N., Masters, W. A., & Sheremeta, R. M. (2010). Entry into winner-take-all and proportional-prize contests: An experimental study. Journal of Public Economics, 94(9–10), 604–611.

Charness, G., & Gneezy, U. (2012). Strong evidence for gender differences in risk taking. Journal of Economic Behavior & Organization, 83(1), 50–58.

Compte, O., & Postlewaite, A. (2004). Confidence-enhanced performance. American Economic Review, 94(5), 1536–1557.

Corgnet, B., & Hernan-Gonzalez, R. (2019). Revisiting the trade-off between risk and incentives: The shocking effect of random shocks? Management Science, 65(3), pp.1096-1114.

Corgnet, B., Gomez-Minambres, J., & Hernan-Gonzalez, R. (2018). Goal setting in the principal–agent model: Weak incentives for strong performance. Games and Economic Behavior, 109, pp.311-326.

Croson, R., & Gneezy, U. (2009). Gender differences in preferences. Journal of Economic Literature, 47(2), 448-74.

Dargnies, M. P. (2012). Men too sometimes shy away from competition: The case of team competition. Management Science, 58(11), 1982–2000.

Dariel, A., Nikiforakis, N. and Stoop, J., (2020). Does selection bias cause us to overestimate gender differences in competitiveness? (No. 20200046).

Dechenaux, E., Kovenock, D., & Sheremeta, R. M. (2015). A survey of experimental research on contests, all-pay auctions and tournaments. Experimental Economics, 18(4), 609–669.

Delfgaauw, J., Dur, R., Sol, J., & Verbeke, W. (2013). Tournament incentives in the field: Gender differences in the workplace. Journal of Labor Economics, 31(2), 305–326.

Dohmen, T., & Falk, A. (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. American Economic Review, 101(2), 556–90.

Eckel, C. C., & Grossman, P. J. (2002). Sex differences and statistical stereotyping in attitudes toward financial risk. Evolution and Human Behavior, 23(4), 281–295.

Eckel, C. C., & Grossman, P. J. (2008). Differences in the economic decisions of men and women: Experimental evidence. Handbook of Experimental Economics Results, 1, 509–519.

Erkal, N., Gangadharan, L., & Nikiforakis, N. (2011). Relative wage and giving in a real-effort experiment. American Economic Review, 101(7), 3330–3348.

Erkal, N., Gangadharan, L., & Xiao, E. (2019). Leadership selection: Can changing the default break the glass ceiling? Available at SSRN 3022386.

Fehr, E., & Goette, L. (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. American Economic Review, 97(1), pp.298-317.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics, 10(2), 171–178.

Flory, J. A., Leibbrandt, A., & List, J. A. (2015). Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions. The Review of Economic Studies, 82(1), 122–155.

Frederiksen, A. (2008). Gender differences in job separation rates and employment stability: New evidence from employer-employee data. Labour Economics, 15(5), 915–937.

Gangadharan, L., Jain, T., Maitra, P. and Vecci, J. (2016). Social identity and governance: The behavioral response to female leaders. European Economic Review, 90, 302-325.

Garcia, A. N. (2017). Gender differences in unemployment dynamics and initial wages over the business cycle. Journal of Labor Research, 38(2), 228–260.

Gill, D., & Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. American Economic Review, 102(1), 469–503.

Gneezy, U., & Rustichini, A. (2004). Gender and competition at a young age. American Economic Review, 94(2), 377-381.

Gneezy, U., Leonard, K. L., & List, J. A. (2009). Gender differences in competition: Evidence from a matrilineal and a patriarchal society. Econometrica, 77(5), 1637–1664.

Gneezy, U., Niederle, M., & Rustichini, A. (2003). Performance in competitive environments: Gender differences. The Quarterly Journal of Economics, 118(3), 1049–1074.

Hall, D. T., & Lawler, E. E. (1971). Job pressures and research performance: Contrary to popular opinion, job pressures are not necessarily undesirable; in fact, certain pressures seem to enhance researchers' job attitudes and performance. American Scientist, 59(1), 64–73.

Harbring, C., & Irlenbusch, B. (2003). An experimental study on tournament design.

Labour Economics, 10(4), 443–464.

Healy, A., & Pate, J. (2011). Can teams help to close the gender competition gap? The Economic Journal, 121(555), 1192–1204.

Heckman, J. J. (2006). Skill formation and the economics of investing in disadvantaged children. Science, 312(5782), 1900–1902.

Hirsch, B., & Schnabel, C. (2012). Women move differently: Job separations and gender. Journal of Labor Research, 33(4), 417–442.

Hong, F., Hossain, T., & List, J. A. (2015). Framing manipulations in contests: A natural field experiment. Journal of Economic Behavior & Organization, 118, 372–382.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. Review of Economic Studies, 80(3), 984–1001.

Kalra, A., & Shi, M. (2001). Designing optimal sales contests: A theoretical perspective. Marketing Science, 20(2), 170–193.

Leibbrandt, A., & List, J. A. (2018). Do equal employment opportunity statements backfire? Evidence from a natural field experiment on job-entry decisions. National Bureau of Economic Research.

Leibbrandt, A., Wang, L. C., & Foo, C. (2018). Gender quotas, competitions, and peer review: Experimental evidence on the backlash against women. Management Science, 64(8), 3501–3516.

Leslie, L. M., Manchester, C. F., Park, T. Y., & Mehng, S. A. (2012). Flexible work practices: A source of career premiums or penalties? Academy of Management Journal, 55(6), 1407–1428.

Mario, B., Nicole, G., Ute, B., Lisa, F., & Benjamin, L. (2020). The German minimum wage: Effects on productivity, profitability, and investments. Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik), 240(2–3), 321–350.

Moldovanu, B., & Sela, A. (2001). The optimal allocation of prizes in contests. American Economic Review, 91(3), 542–558.

Niederle, M. & Vesterlund, L., 2007. Do women shy away from competition? Do men compete too much? The Quarterly Journal of Economics, 122(3), 1067-1101.

Niederle, M., Segal, C., & Vesterlund, L. (2013). How costly is diversity? Affirmative action in light of gender differences in competitiveness. Management Science, 59(1), 1–16.

Orrison, A., Schotter, A., & Weigelt, K. (2004). Multiperson tournaments: An experimental examination. Management Science, 50(2), 268–279.

Rosen, S. (1986). The theory of equalizing differences. Handbook of Labor Economics, 1, 641–692.

Sheremeta, R. M. (2011). Contest design: An experimental investigation. Economic Inquiry, 49(2), 573–590.

Sheremeta, R. M. (2015). Behavioral dimensions of contests. In R. D. Congleton & A. L. Hillman (Ed.), Companion to the political economy of rent seeking (pp. 150–164). Cheltenham: Edward Elgar Publishing.

Van Dijk, F., Sonnemans, J., & Van Winden, F. (2001). Incentive systems in a real effort experiment. European Economic Review, 45(2), 187–214.

Vandegrift, D., Yavas, A., & Brown, P. M. (2007). Incentive effects and overcrowding in tournaments: An experimental analysis. Experimental Economics, 10(4), 345–368.

3.6. Appendix

3.6.1. Instructions

Here, we provide the instructions for the choice treatment. The instructions for the other treatments only differ in the availability of the safeguard and are available upon request.

1. General Instructions

Thank you for participating in this experiment. Please read the following instructions carefully. After reading the instructions, there will be some questions to check that all participants have understood the experiment. Thereafter, the main experiment will start. Note that you will be paid in private and in cash at the end of the experiment.

If you have any questions during the experiment please raise your hand, and we will come to you. Please do not ask your questions out loud, attempt to communicate with other participants, or look at other participants' computer screens at any time during the experiment. Please turn your phone to silent mode and put it in your bag. Please do not use any calculator during the experiment including the one available with the Windows system.

2. The Task

You will have 40 minutes (2,400 seconds) to solve puzzles similar to the one in the graph below.

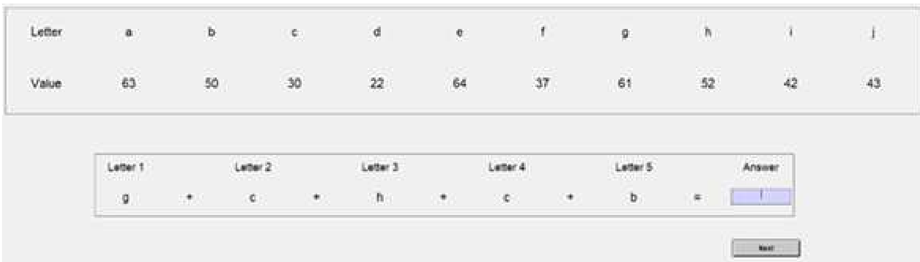

To solve this puzzle, you first need to decipher the code in the box below. For example, letter 1 (g) has a value of 61, letter 2 (c) 30, letter 3 (h) 52, letter 4 (c) 30 and letter 5 (b) 50. The correct answer is the summation of all five values: correct answer =

61+30+52+30+50 = 223. After typing your answer in the answer box, you have to press the Next button below to access the next puzzle.

3. Payment

Your payment depends on your performance relative to the other participants in this session and whether you choose to have a safeguard.

If you do not choose the safeguard, only your relative performance determines your wage:

- If your total number of correct answers is among the top 10%, then you will receive $60.

- If your total number of correct answers is among the top 10% to 50%, then you will receive $30.

- If your total number of correct answers is among the bottom 50%, then you will receive $15.

If you choose to have the safeguard, then you earn at least $20:

- If your number of correct answers is among the top 10%, then you will receive $60.

- If your number is among the top 10%-50%, then you will receive $30.

- If your number is among the bottom 50%, then you will receive $20.

3.6.2. Framing of the safeguard

We investigate the role of the framing of the safeguard here. We use two frames in the compulsory treatment, which vary depending on whether participants are made aware of the presence of a safeguard. More precisely, in the first frame, we only mention the payment rule without any mentioning of a safeguard. In contrast, in the second frame, we present participants with the concept of the safeguard and inform them that their minimum payment ($20) is higher than in some of the other groups ($15) because of the presence of a safeguard. We observe no differences between frames in terms of quizzes attempted and quizzes answered correctly (p=0.8 for attempt p=0.9 for correct). This also holds true if we analyze males and females separately. For males, the p-value testing for no difference caused by the framing is 0.37 for attempted questions and 0.30 for questions answered correctly; for females, the numbers are 0.62 and 0.38, respectively.

3.6.3. Demographic variables

We report the demographic variables across treatment in Table 3.5. In general, the demographic variables are balanced across treatments from F-tests.

	Gender (male=1)	Age	Student	Business school
Baseline	0.50	22.70	0.99	0.19
Choice	0.55	23.39	0.99	0.22
Compulsory	0.44	22.65	0.98	0.25
F_stat	2.3	1.2	0.74	0.59
p_value	0.11	0.29	0.49	0.55

Notes: This table presents the average demographic characteristics for each treatment. We also report the F-statistic and p-value testing the null hypothesis that there is no difference in the variable across all treatments.

Table 3.5: Balance check

Chapter 4

Deterrence Using Peer Information

[1]

4.1. Introduction

Drug dealers in neighboring areas usually observe each other's activities such as the size of stocks and sales. Sometimes, the police may randomly check car boots or patrol the area and capture one of them. The regulatory authority is often unaware of such insider information and lacks tools to invoke and leverage this valuable information upon finding the suspect. However, such information can be captured by criminal networks (Lindquist & Zenou 2019). We use this insight to propose a new peer-informed audit mechanism, such that the suspect is incentivized to reveal information regarding her neighbors' criminal activities.

We show the effectiveness of this novel mechanism that leverages individuals' information to apprehend the most severe offenders in the network and reduces overall criminal activities. The mechanism introduces a peer-informed audit stage that provides the suspect an opportunity to redirect the penalty. Consider a group of criminals linked by a network who are involved in similar kinds of criminal activities, and linked members

[1] This chapter reports research conducted jointly with Lata Gangadharan and C. Matthew Leister.

can observe each others' criminal activities with some noise. The regulator, possessing no information regarding the network, happens to capture a criminal in the network. The criminal can either pay a fine or challenge another linked criminal in the network. If the audited criminal does challenge another criminal, then the regulator compares the noisy signals of their crime levels. The regulator fines the criminal with the higher signal (i.e. with greater evidence of criminality) and lets the other one go free.

As compared to the baseline mechanism in which punishment for the first suspect is certain, standard theory predicts that our mechanism deters crime more effectively, keeping the number and the magnitude of the penalties fixed. In equilibrium, an individual's crime level depends on the given criminal's position in the network and on the network structure. This is because if individuals observe the signals of others' crimes, then an individual audited by the authority in the peer-informed audit mechanism will optimally play out the strategy of challenging the linked criminal with the highest observed crime level. By backward induction, under this simple peer-challenge strategy, all connected criminals have the incentive to reduce their crime levels in order to avoid neighbors' challenges as well as successfully challenge neighbors when required. The likelihood of a successful challenge and being the target of another's challenge depends on the number of links the criminal has and the overall network structure.

We examine this mechanism using three commonly used network structures (star, circular, and complete networks), which also summarize many criminal network structures found in the literature. For example, Venkatesh & Levitt (2000) find that star networks closely represent many criminal gangs in Chicago, where there is a "manager" responsible for a syndicate of drug dealers; circular networks resemble neighborhood watches where observation is constrained by geographical proximity; complete networks resemble cartels where all members know and observe each other (see, e.g., Belleflamme & Bloch 2004). Theory predicts that the mechanism is most effective in the complete network, and least effective in the star network; within the star network, the center player would reduce her crime level more than the periphery players. Intuitively, the more links a player has, the more effective the mechanism is in deterring her crime level. This is because each link in the network represents a piece of insider information and, thus, more links imply that there is more insider information to be utilized by means of transmitting the information to the authority. Consequently, the theory shows that the peer-informed audit mechanism lowers the crime level in all three network treatments as compared to that in the baseline.

Next, we design an experiment to provide the first test of this mechanism. Experimental methods enable researchers to observe actions accurately and help identify a causal relationship between different enforcement mechanisms and actions.[2] Our experimental evidence provides strong support for the deterrence effect of the peer-informed audit mechanism. However, crime in the treatments with the peer-informed audit mechanism is above the levels predicted by the subgame perfect Nash equilibrium and is less sensitive to changes in network structures as compared to theoretical predictions. We also find that participants with higher cognitive ability commit less crime in more connected network treatments, thereby approaching optimal behavior given others' crime levels. In order to explore this further, we examine participants' responses in the post-experiment questionnaire. Certain participant responses are suggestive of a failure to discern that the level of crime affected their probability of being fined, while others indicate that they attempted to "produce a bit less than connected players to avoid paying the fine." An explanation for our findings, suggested by these responses, is that individuals are boundedly rational and have differential reasoning abilities. To test this explanation, we explore a level-k reasoning model where level-0 participants' decisions are based on the incorrect belief that punishment is independent of the crime level, level-1 best respond to level-0 participants, level-2 best respond to level-1 participants, and so forth. We find that this boundedly rational model of crime predicts that the reasoning capacity of individuals affects the overall crime level.

This paper is linked to several prominent strands of the economics literature. First, it is closely related to the literature that studies deterrence mechanisms. Several of these mechanisms rely on the authority's knowledge about potential offenders and randomly inspect a subset of them (e.g., Becker 1968; Polinsky & Shavell 1979; Alm & McKee 2004; Coricelli et al. 2010; Gilpatric et al. 2011; Friesen 2012; Cason et al. 2016).[3] However, it

[2] Our goal is to provide a clean first test of the mechanism guided by clear theoretical principles. In order to do this, we abstract from implementational issues—such as distrust of the mechanism, corruption, targeted retaliation, or other factors—that may complicate the interpretation of treatment differences. As we describe in Section 6, in future research it would be useful to examine these aspects more explicitly.

[3] While the standard economics-of-crime approach focuses on the monetary incentives for reducing criminal behaviour, a few of the papers listed above (e.g., Coricelli et al. 2010; Cason et al. 2016) also

is likely that in certain cases the authority does not know the potential offenders but can only apprehend one of them by chance. Mechanisms that focus on random or targeted inspections are not suitable in such cases. We fill this gap by introducing a new mechanism that enables the authority to utilize potential insider information.

Our research also connects to mechanisms that encourage whistleblowing by applying leniency on the whistleblowers in cases where the authorities are unaware of the illegal actions in the first place (see e.g., Miller 2009; Bigoni et al. 2012, 2015; Chen & Ray 2013; Abbink et al. 2014; Feltovich & Hamaguchi 2018). We diverge from the approach suggested in this literature along several dimensions. First, whistleblower leniency relies only on an endogenous reporter, while our peer-informed audit mechanism requires identifying one criminal from the entire network. More importantly, in our mechanism, the penalty depends purely on the relative crime levels, which solves several potential drawbacks of the leniency policy. For example, under leniency, firms may "form sacrificial cartels and apply for leniency in less valuable products to reduce convictions in more valuable products" (Marx et al. 2015). Our mechanism successfully eliminates this incentive, as the member with the higher observed crime level is always liable for the fine. Moroever, the whistleblower tends to be discriminated against (Carr & Lewis 2010), which weakens the incentive to report other criminals. Our mechanism preserves the anonymity of the peer informant in several important ways. Instead of punishing all members, our mechanism only punishes one, which reduces the potential for retribution. We also relax the implicit assumption in leniency schemes that every member of the cartel is aware of the other members involved (i.e., the complete network) by incorporating incomplete and irregular network structures, such as the circular and the star networks.

Second, this project is related to the nascent literature on social networks (Ballester et al. 2006), which finds that networks are crucial for education (Calvo-Armengol et al. 2009; Hahn et al. 2019), detection of crime (Calvo-Armengol & Zenou 2004; Lindquist & Zenou 2019), and other economic activities (Kaiser 2007; Nakajima et al. 2010; Charness et al. 2014; Helsley & Zenou 2014; Leibbrandt et al. 2015; van Leeuwen et al. 2015, 2019; Jackson et al. 2017; Bloch & Olckers 2019; Konig et al. 2019). Building on this literature, we focus on how regulators can use the information in networks to reduce

introduce additional considerations, such as emotional response or social stigma caused by engaging in illegal activities.

criminal activity and improve social welfare. In particular, we create incentives for the suspect to reveal information regarding the criminal activities of others in the network; in doing so, we bring together the research on leniency programs and social networks. We show theoretically that crime levels depend on the given criminal's position in the network and the network structure. Further, we provide experimental tests of the theory. While all peer-informed audit networks deter crime as compared to the baseline, in contrast to theoretical predictions, there are no statistically significant differences across network treatments. We show that an alternative model based on level-k reasoning explains these findings. This provides evidence that more sophisticated criminals will internalize the implications of peer-informed auditing over the network and, thus, reduce their criminal activity.

Finally, our mechanism is related to the level-k reasoning literature (Nagel 1995; Agranov et al. 2012; Gill & Prowse 2016; Jin 2018). To the best of our knowledge, we are the first to build a level-k model of interaction within a network structure. This model predicts that criminal activity decreases monotonically in k toward the Nash prediction (i.e., level-∞) for all network types we consider. Interestingly, we find that the complete network is the least effective network structure to deter crime (compared with star and circular networks) when the sophistication level is very low, but becomes the most effective one when the level is high. In other words, reasoning abilities can interact with the network structure and influence the incentives to reduce crime. Consistent with this theory, we find that experimental participants with higher cognitive abilities engage in crime in a manner that is consistent with the best response to prevailing criminal levels. We also find most participants range from level-0 to level-2, which matches the findings from level-k games without network structures. Our findings indicate the importance of studying such boundedly rational models of crime to understand the effectiveness of enforcement mechanisms.

The remainder of this paper is organized in the following manner. We present the theoretical framework and experiment in sections 2 and 3, respectively. We report our results in sections 4 and put forward a boundedly rational model in section 5 that further explains our findings. Section 6 discusses the implications of our findings and offers some ideas for future research.

4.2. Theoretical model

4.2.1. Model setup

We illustrate the peer-informed audit mechanism with an example of sales of illicit substances like drugs. The mechanism can be extended to other types of crimes (such as tax evation, financial misreporting, smuggling etc.) provided that the severity of criminal activities can be measured and compared.

Prior to the beginning of the game, $n \geq 2$ drug dealers are connected by an exogenously determined undirected graph represented by adjacency matrix g. In other words, if drug dealers i and j are connected, then $g_{ij} = g_{ji} = 1$, and $g_{ij} = g_{ji} = 0$ otherwise. The crime network structure defined by g is common knowledge among all drug dealers but is unknown to the regulator. We denote i's neighborhood $N(i) = \{j \in N : g_{ij} = 1\}$, which is assumed to be non-empty for each i.

The game has two stages. In the first stage, each drug dealer $i \in \{1, 2, ..., n\}$ selects a crime level e_i. This action can be interpreted as choosing how much drug to stock and sell. Higher e_i increases i's profit $\pi(e_i)$. We assume that the profit function $\pi(\cdot)$ is concave and assumes the form $\pi(e_i) = e_i^b$ for $0 < b < 1$.[4] After all drug dealers choose their crime levels, each drug dealer observes noisy signals informing them of the crime levels of connected drug dealers. Precisely, drug dealer i observes a noisy signal of e_j if and only if $g_{ij} = 1$. The signal observed by drug dealer i is given by $\hat{e}_j = e_j + \epsilon_j$, where ϵ_j follows cumulative distribution F.[5] We assume ϵ_j and ϵ_ℓ are independent for all $j \neq \ell$. Each drug dealer knows its own crime level by construction, but does not observe the signal informing others of its crime level.

In the second stage, with probability θ the regulator detects the crime. The regulator then audits one randomly-drawn drug dealer j (i.e., each drug dealer is audited by the regulator with probability θ/n) and observes the noisy signal of that drug dealer's crime

[4] The condition $\pi(0) = 0$ is without loss, as each drug dealer's realized profit may involve an additional constant capturing "legal" profit with $\pi(\cdot)$ yielding "excess" profit.

[5] The noise in the signal (ϵ) is required to define a well-behaved Nash equilibrium. Otherwise, players always have an incentive to undercut each other.

$\hat{e}_j$.[6] The payoff function for drug dealer i is

$$u_i(e_i, e_{-i}) = \pi\,(e_i) - f(\hat{e}_i) \cdot \mathbb{I}_i\,(\text{firm } i \text{ is fined})\,,$$

where $\mathbb{I}$ is an indicator function that takes value 1 if and only if the logical statement in its argument is true, and 0 otherwise. The non-negative payment $f(\hat{e}_i) \geq 0$ defines the size of the fine when drug dealer i is liable and the regulator observes signal $\hat{e}_i$.[7] In the next section, we introduce the mechanisms that determine which drug dealer is liable for the fine.

4.2.2. Enforcement mechanisms

We consider two enforcement mechanisms. In the baseline audit mechanism, the randomly audited drug dealer j pays the fine. In the peer-informed audit mechanism, the randomly audited drug dealer j can challenge one of her peers k such that $g_{jk} = 1$. The regulator compares the results of the audit of drug dealer j ($\hat{e}_j$) and drug dealer k ($\hat{e}_k$), and fines the one that has a higher observed crime (i.e., $\hat{e}_j > \hat{e}_k$ implies $\mathbb{I}_j$ (liable for the fine) $= 1$).

The peer-informed audit mechanism can be applied to all network structures. For simplicity, we focus on the following network structures: 1) a star network, in which drug dealer i is only connected to a unique center drug dealer if she is a periphery drug dealer, and is connected to all other (periphery) drug dealers if she is the center drug dealer; 2) a circular network in which each drug dealer i is only connected to two neighbors; and 3) a complete network in which each drug dealer i is connected to all other drug dealers in the group.

[6] We use this simple information structure where the drug dealers and the regulator observe the same signal for each drug dealer to reduce the complexity faced by experimental participants. The assumption that only one signal is generated for each drug dealer is inconsequential; this simplification yields qualitatively equivalent equilibrium challenge strategies and crime levels as the model in which players observe private signals with i.i.d. noise.

[7] In the event that the player liable for the fine has a negative signal, the fine imposed on this player is set to zero. However, it is possible for a player who selects zero crime level to be liable for the fine if the player obtains a positive signal. In the experiment, very few participants selected zero crime level. For example, only 4 out of 100 players in the one-shot game treatments selected the zero crime level and none of them were liable for the fine.

4.2.3. Equilibrium crime and profit-fine trade-offs

Given the model's assumptions, a Nash equilibrium in crime levels in the first stage solves the following system of equations:

$$\frac{\partial \mathbb{E}\left[u_i\left(e_i^*, e_{-i}^*\right)\right]}{\partial e_i} = 0, \tag{4.1}$$

for all $i \in \{1, 2, ..., n\}$. The expectation $\mathbb{E}[\cdot]$ is taken over all noise variables and whether the regulator detects crime. In the baseline audit mechanism, (4.1) becomes:

$$\frac{\partial \pi\left(e_i^*\right)}{\partial e_i} = \frac{\theta}{n} \frac{\partial \mathbb{E}\left[f(\hat{e}_i)|e_i^*\right]}{\partial e_i} \tag{4.2}$$

In the peer-informed audit mechanism, it is a weakly dominant strategy for any audited drug dealer to challenge the neighbor with the highest signal. When the drug dealers follow this strategy in the second stage, (4.1) becomes:

$$\frac{\partial \pi\left(e_i^*\right)}{\partial e_i} = \frac{\partial \mathbb{E}\left[f(\hat{e}_i) \cdot \mathbb{I}_i\left(\text{firm } i \text{ is fined}\right) \middle| e_i^*, e_{-i}^*\right]}{\partial e_i} \tag{4.3}$$

Appendix A1 provides closed-form calculations for the marginal expected fine. Here, we describe the basic trade-offs faced by drug dealers under this mechanism. Provided any audited drug dealer i always challenges their neighbor j with the highest $\hat{e}_{ij}$, (4.1) becomes:

$$\frac{\partial \pi\left(e_i^*\right)}{\partial e_i} = \left(\begin{array}{c} \frac{\theta}{n} \cdot \frac{\partial}{\partial e_i} \underbrace{\mathbb{E}\left[f\left(\hat{e}_i\right) \middle| L_i^{\rightarrow}, e_i^*, e_{-i}^*\right] \cdot \mathbb{P}\left(L_i^{\rightarrow} \middle| e_i, e_{-i}^*\right)}_{\text{Expected fine from being audited}} \\ + \frac{\theta|N(i)|}{n} \cdot \frac{\partial}{\partial e_i} \underbrace{\mathbb{E}\left[f\left(\hat{e}_i\right) \middle| L_i^{\leftarrow}, e_i^*, e_{-i}^*\right] \cdot \mathbb{P}\left(L_i^{\leftarrow} \middle| e_i^*, e_{-i}^*\right)}_{\text{Expected fine from being challenged by a neighbor}} \end{array} \right) \tag{4.4}$$

where $L_i^{\rightarrow}$ denotes the event $\{\hat{e}_i > \hat{e}_j, \hat{e}_{ij} = \max_{k \in N(i)}\{\hat{e}_{ik}\}\}$ that i is audited and loses her challenge of a neighbor j, and $L_i^{\leftarrow}$ denotes the event $\{\hat{e}_i > \hat{e}_k, \hat{e}_{jk} = \max_{k \in N(j)}\{\hat{e}_{jk}\}\}$ that i is challenged by a neighbor j and loses that challenge. Note that drug dealer i is liable for the fine under both events $L_i^{\rightarrow}$ and $L_i^{\leftarrow}$.

Under either mechanism, drug dealer i equates the marginal profit from additional crime with the marginal expected fine. In the baseline audit mechanism, the probability of being fined is fixed at θ/n, while in the peer-informed audit mechanism player i further accounts for the marginal effect of e_i on the probability of being fined. We see that, when calculating this probability, drug dealer i considers the possibility of being fined to originate from one of two channels: (1) i is audited and i's challenge of a neighbor j fails in avoiding the fine and (2) one of i's neighbors j is audited and successfully challenges i.

Given that each drug dealer's noisy signal is an estimate of the drug dealer's crime level, the network structure and the drug dealer's position in the network affect the likelihood of the drug dealer being liable for a fine. Generally speaking, the fewer the neighbors the drug dealer has, the fewer opportunities for $L_i^{\rightarrow}$ to occur, as drug dealer i has few accomplices to pass the blame to. Conversely, the more neighbors the drug dealer has, or the fewer the neighbors the drug dealer's neighbors have, the more opportunities for $L_i^{\leftarrow}$ to occur, as more neighbors are more likely to challenge drug dealer i.

Given the experimental parameters in Table 4.1, the benefits from adding a link (having more accomplices to pass the blame to) outweigh the corresponding costs (being more likely to be targeted by accomplices). Consequently, theory predicts that the mechanism is most effective in the complete network and least effective in the star network. In Appendix A1, we elaborate on these basic trade-offs in equilibrium within the star, circle, and complete network structures.

4.3. Experimental design

The experiment follows theory closely. We chose a simple parametrization that is consistent with the model. The parameters we use in the experiment are presented in Table 4.1. Participants are randomly and anonymously assigned into groups of five and play a two-stage game. In the first stage, each participant decides on their crime levels.[8] We set

[8]We framed the explanation of the experiment in neutral language in order to ensure maximum control. We refer to the penalty as a fine to introduce context and subtly indicate that production is illegal,

the range of e_i for each subject i to be between 0 and 40 units.[5] The upper threshold of 40 units is selected such that all decisions are incentive-compatible and the fine for the maximum crime is still covered by the endowment, thereby ensuring that participants bear the full consequence of their actions. Indeed, the more a participant produces, the more profit (in experimental currency— Tokens) they receive; we assume $\pi(e_i) = e_i^{0.7}$ (i.e. $b = 0.7$).

Notation/Functional form	Definition	Parameters
n	Number of participants	$n = 5$
e_i	Crime level selected by participant i	$e_i \in [0, 40]$
$\pi(e_i)$	Benefit from crime level e_i	$\pi(e_i) = e_i^{0.7}$
θ	Probability of detection	$\theta = 1$
$\epsilon_j \sim F$	The noise associated with observation	$F = \text{Unif}[-10, 10]$
$\hat{e}_i$	The signal of player i observed by others	$\hat{e}_i = e_i + \epsilon_i$
$f(\hat{e}_i) = f \cdot \hat{e}_i$	Magnitude of the fine	$f = 1.2$

Notes: This table presents the parametric setup of our experiment.

Table 4.1: Experiment parameters

In the second stage, the computer randomly draws a number from the uniform distribution with support $[-10, 10]$ for each participant and computes a signal that is equal to the actual crime level plus the random number drawn for the player.[9] As described earlier, each participant can observe the crime level selected by other connected participants with some noise. The signal mimics this noisy observation. Participants are informed that the random number has an equal likelihood of being any number between and including -10 and 10, with a mean of zero. This implies that, in expectation, the signal is equal to the actual crime level. For simplicity, we make the observational noise associated with each participant the same for the remaining participants in the same group. For example, participants 2–5 see the same signal for participant 1.[10]

without explicitly introducing a framing that leads to a potential demand effect (Alm 1991; Coricelli et al. 2010).

[9] In our experiment, the production level (e_i) has two decimal places and each random number (ϵ_i) has fourteen decimal places to closely proxy the continuity in the production and noise space, as assumed in the theoretical model. The discontinuity in the values does not change the predictions from the theory.

[10] To keep this feature of observability constant across all treatments, we allow participants to see the signals associated with the remainder of the group in all treatments rather than only observe the connected players. Standard theory predicts that the observations of signals of distant participants do not affect the equilibrium, as crime levels are selected before signals are observed. However, there are at least two reasons to keep this feature constant across treatments in the experiment. First, recent evidence suggests that varying such features could change the complexity perceived by

Further, participants do not fully observe the criminal activities of others in their group, only the signals associated with those choices. This mirrors the fact that drug dealers may not have perfect information regarding the illegal production levels of other drug dealers but rather possess rough estimates. Each participant also knows her own decision by construction, but is not aware of the random number (and thus the signal) that others receive that informs them of her selected crime level. Thus, an audited participant can construct an estimate of the likelihood that the regulator will render any challenged neighbor more criminally responsible.

After everyone observes the signals of others in the group, the process of determining which player is liable for a fine begins. First, the computer randomly identifies one participant to audit.[11] Depending on the treatment, the audited participant is either fined (in the baseline audit treatment) or is given the opportunity to challenge one neighbor (in the peer-informed audit treatments). In the peer-informed audit treatment, the computer compares the signals of the two participants, if the one initially audited decides to challenge another. The one with the higher signal is fined. The calculation of the size of the fine (in Tokens) is the same across treatments: the fine is 1.2×the signal of the liable participant (i.e. the payoff of the liable participant is reduced by 1.2×his/her signal). Tokens are converted to Australian Dollars at the exchange rate of 50 Tokens=$1. Table 4.1 summarizes the parameters used in the experiment.

4.3.1. Treatments

We implement a between-subjects design where subjects participate in only one treatment. Our treatments include the baseline audit mechanism, peer-informed audit mechanism with the star network, peer-informed audit mechanism with the circular network, and peer-informed audit mechanism with the complete network. These are described below.

participants and, thus, impact their decisions (e.g., Oprea 2020). Second, this may change group norm formation (e.g., Andreoni & Bernheim 2009) and would be important to control for in cases where the game is played more than once.

[11] Our theoretical model allows the possibility that the crime is not discovered and none of the participants is identified and punished. However, this does not change the predictions of the relative crime level across treatments because it affects all treatments equally. In the experiment, we set this probability to zero for simplicity and also to collect more observations on the challenging behavior of participants.

Baseline audit mechanism: Participants are informed that the participant randomly audited by the computer is fined for her observed crime level. This serves as a natural benchmark in our experiment. If not audited by the computer, the participant does not know who is audited and, thus, who is liable for the fine. If audited, the participant knows for certain she is liable for the fine and sees the amount deducted from her earnings.

Peer-informed audit mechanism: The mechanism allows the participant randomly audited by the computer to challenge another participant connected to her in the group in order to avoid paying the fine herself. If she chooses to challenge another, then the computer compares the noisy signals from the two participants and fines the one with the higher signal; otherwise, she is liable for the fine. Those not audited or challenged are not provided with information regarding who is audited or who is challenged. The one who challenges another knows the participant ID of the challenged by design, but the one being challenged does not know which participant challenged her, except for a challenged periphery player in the star treatment who infers that the challenge came from the center player. As in the baseline audit mechanism, the participant liable for the fine has the amount subtracted from her earnings. We test this mechanism in the following three network structures.

Peer-informed audit mechanism in the star network: In the star network, if the participant randomly audited by the computer is in the center of the network (i.e., participant 3), then she can challenge any of the other four participants in the group; if she is a periphery participant (i.e., participant 1, 2, 4, or 5), then she can only challenge the center participant. If the initially audited participant decides to challenge another participant, then the computer compares the noisy signals from the two participants and the one with the larger signal is liable for the fine. If the initially audited participant chooses to not challenge, then she remains liable for the fine.

Peer-informed audit mechanism in the circular network: The participant audited by the computer must decide whether to challenge one of the two participants whose IDs are adjacent to hers or pay the fine herself. For example, participant 5 can only challenge participants 1 or 4. If she chooses to challenge another participant, then the one who has the larger signal is liable for the fine. If she chooses not to challenge, then as in the treatment above, she is liable for the fine.

Peer-informed audit mechanism in the complete network: The participant audited by the computer must decide whether to challenge one of the other four participants in their group. If she chooses to challenge another participant, then the one who gets the larger signal is liable for the fine. If she chooses to not challenge, then, similar to the other treatments, she is liable for the fine.

In these treatments, our main focus is on the results using a one-shot game. We also conducted four treatments to examine the robustness of our peer-informed audit mechanism in a setting where participants interact with each other repeatedly for multiple rounds. We briefly describe the results from these repeated interaction treatments in the paper, with more details discussed in Appendix A3.

4.3.2. Equilibrium predictions and hypotheses

Figure 4.1 plots the equilibrium crime levels for each treatment using the Nash solution concept described in Section 2.3 and the parameters employed in the experiment.

In the baseline audit mechanism, the optimum crime level produced by each drug dealer is 35.5 experimental units. The average crime levels drop to 5.6, 3.4, and 1.2 units for our mechanism in the star, circular, and complete network structures, respectively. Thus, the peer-information audit mechanism leads to a striking reduction in the equilibrium crime level. In addition, we also observe that the network structure is crucial for the effectiveness of the mechanism. The circular and complete networks further reduce the crime level by approximately 40% and 80% in equilibrium as compared to the star network. Similarly, the complete network is able to reduce crime by approximately 65% compared to the circular network.[12]

[12] The effectiveness of the mechanism is partially driven by insider information to identify a second suspect who can be compared with the first one. If the authority can randomly find two suspects and punish the one with the higher signal, we show (in the Appendix) that the crime level drops to 6.58 units in Nash equilibriuim.

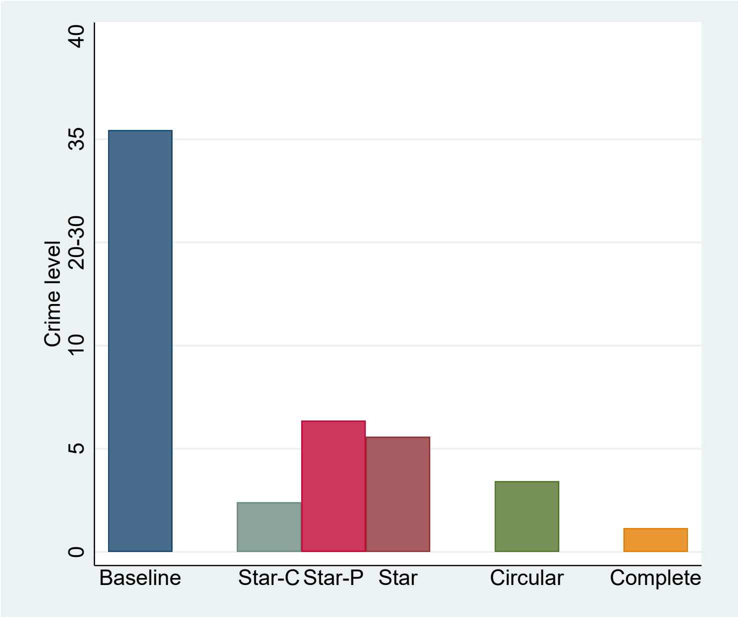

Notes: this figure plots the average crime level at the Nash Equilibrium for each treatment. Star-C indicates the Nash prediction of the center players in the star treatment, and Star-P indicates the Nash of the periphery players in the treatment. Star is the Nash equilibrium for the group average.

Figure 4.1: Prediction based on the Nash equilibrium

Comparing across network structures for the peer-information audit mechanism, we see that more neighbors imply lower levels of crime. This is observed both across networks and within the star network.[13] Thus, facing the additional risk of being challenged carries a stronger negative impact on crime than having additional opportunities to challenge and pass the blame to neighbors.

Overall, the mechanism leads to lower crime due to the following reasons. First, note that for a given level of crime under the peer-informed audit mechanism, the possibility of either challenging or being challenged increases the conditional expectation of the fine amount when deemed liable. More precisely, being liable for the fine implies that the drug dealer's signal is greater than that of some other drug dealers and, thus, the posterior expected fine is above that under the baseline audit mechanism. This channel unambiguously works toward decreasing overall crime.

Second, and also yielding an unambiguous decrease in crime, peer challenges increase the

[13]Note that the average number of connections in the star network is 1.6, which is below the average of 2 in the circular network.

marginal probability of being liable as the crime level of the drug dealer increases. In other words, an increase in one criminal's drug dealer's level deems challenges to others less effective and challenges from others more effective. Opportunities for peer-informed auditing are plentiful in the complete network, where all players can challenge any one of the other players. Peer-informed auditing opportunities are fewer in the circular network, as each player is constrained from challenging all but two players. In the star network, peer-informed auditing opportunities are fewest for the periphery players, who can only challenge and be challenged by the center player. The periphery players should always challenge the center, which provides the only option to these drug dealers to avoid a fine when audited. Thus, the center player is the one to always be challenged by any other player selected by the computer in this treatment. Consequently, the center player faces strong incentives to reduce her criminal activity.[14]

Based on the theoretical predictions of the model, we can construct the following three testable hypotheses.

Hypothesis 1: The peer-informed audit mechanism deters crime in all network structures as compared to the baseline audit mechanism.

Hypothesis 2: Of the three treatments with the peer-informed audit mechanism, the crime level is higher in the star network as compared to the circular network, which in turn has higher crime than the complete network.

Hypothesis 3: Within the star network, the center player has a lower crime level than the corresponding periphery players.

[14]As depicted in Figure 4.1, the lack of challenging opportunities faced by the periphery dominates the overall equilibrium incentives to lower crime levels, thereby resulting in high average crime levels in the star network. In addition, with the periphery selecting high levels of crime, the center partakes in more crime than members of the complete network in equilibrium.

4.3.3. Procedures

The experiment was conducted at the Monash Laboratory for Experimental Economics (MonLEE), with a total of 300 participants. In the first set of treatments, we recruited 100 participants to examine the effectiveness of the peer-informed audit mechanism using a one-shot game. In the second set of treatments, we recruited another 200 subjects to stress-test the peer-informed audit mechanism in a setting where there is repeated interaction. The procedures followed for the repeated interaction treatments are identical to the one-shot treatments, except that the same group of participants interact repeatedly for 20 rounds and they are paid for only one randomly selected round.

For the one-shot treatments, each subject constitutes an independent observation. We created 5 groups with 5 subjects in each group, and we obtained a total of 25 independent observations for each treatment. In the repeated game treatments, each subject observes others' decisions after the first round; therefore, each group is considered as an independent observation. We collected data from 10 groups with 5 subjects in each group, and we obtained 10 independent observations for each treatment.

For both one-shot and repeated sessions, we conducted only one treatment in each session, and the instructions explicitly state the different stages of the game and the rule for determining who is liable for the fine. These were read aloud and were, thus, common information in the session. Participants were also provided sufficient time to read the instructions on their own and ask questions. In order to ensure comprehension, we asked them to complete a set of incentivized quiz questions. Participants received AUD $2 for each correct answer. Incorrect responses were explained to subjects. At the end of the experiment, we administered a short demographic questionnaire. Instructions and quiz questions for the complete network treatment are provided in Appendix A2.

Subjects were university students recruited across different disciplines and had not previously participated in experiments involving audits and penalties.[15] The experiment was programmed in z-Tree (Fischbacher 2007), and the subjects were recruited using SONA. On average, each participant earned $24.9 with a minimum of $12 and a maximum of $34, and sessions lasted approximately 35 minutes (one-shot game) to one hour (repeated

[15] We provide more information regarding participants' demographic variables in the Appendix.

interaction game).

4.4. Results

We begin by discussing the results from the one-shot game treatments. Our main focus is on the crime levels selected by participants. In terms of whether participants challenge others in treatments where they had an opportunity to do so, our main finding is that, consistent with predictions, all randomly selected participants chose the dominant strategy to challenge the connected players who provided them the best opportunity to avoid the fine.

4.4.1. Crime levels

Figure 4.2 plots the average crime levels for the baseline and treatments with the peer-informed audit mechanism. Each participant selects their crime level before observing the activities of group members and, thus, each decision is treated as an independent observation. We find that the mechanism successfully reduces the crime level from 29.70 units in the baseline to 19.88 units in the peer-informed audit treatments— a 33% drop ($p<0.001$).[16]

[16] Without further specification, all p-values reported in the paper were obtained from two-sided non-parametric Mann-Whitney tests.

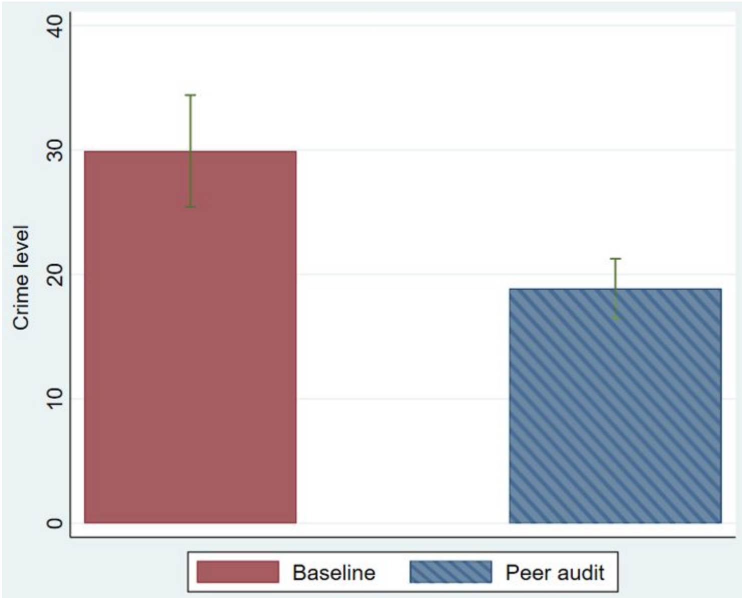

Notes: This figure plots the average crime levels for the baseline and three peer-informed audit treatments combined. The error bars indicate 95% confidence intervals.

Figure 4.2: Overview of crime levels

Figure 4.3 provides a comprehensive view of the crime levels in all treatments, allowing us to compare them against the corresponding predictions based on Nash. The new mechanism significantly decreases the crime level under all network structures we consider. Compared to the baseline, where the average crime level is 29.70 units, the mechanism reduces the levels to 21.68 ($p = 0.013$), 19.05 ($p = 0.001$), and 18.91 ($p < 0.001$) units under star, circular, and complete networks, respectively.[17] Consistent with theoretical predictions, the magnitude of crime in the star network is higher than the level in the circular network and similarly the circular network has higher crime levels than the complete network. However, there are no significant treatment differences between peer-informed audit treatments with different networks ($p > 0.1$ in all pairwise comparisons). In addition, there is also no significant difference between the center and periphery players in the star network. The average crime level of the center player is 23.6 units and the level of the periphery player is 21.2 units. This is in the opposite direction as predicted by Nash

[17] We also performed Kolmogorov–Smirnov tests to examine if the distribution of crime is different across treatments. Consistent with the Mann-Whitney test statistics on the crime levels, Kolmogorov–Smirnov tests show that the baseline has a different distribution than either treatment with the peer-informed audit mechanism ($p = 0.037$, 0.006 and 0.001 for comparisons with the Star, Circular and Complete treatment, respectively). There is, however, no significant difference between the treatments with peer-informed audit mechanism ($p = 0.699$, 0.281, 0.906 for comparisons between star v.s circular, star v.s. complete and circular v.s. complete).

but not statistically significant ($p > 0.1$).[18] The crime levels in all three peer-informed audit treatments are also significantly larger than that predicted by the Nash equilibrium ($p < 0.001$ in all three networks).

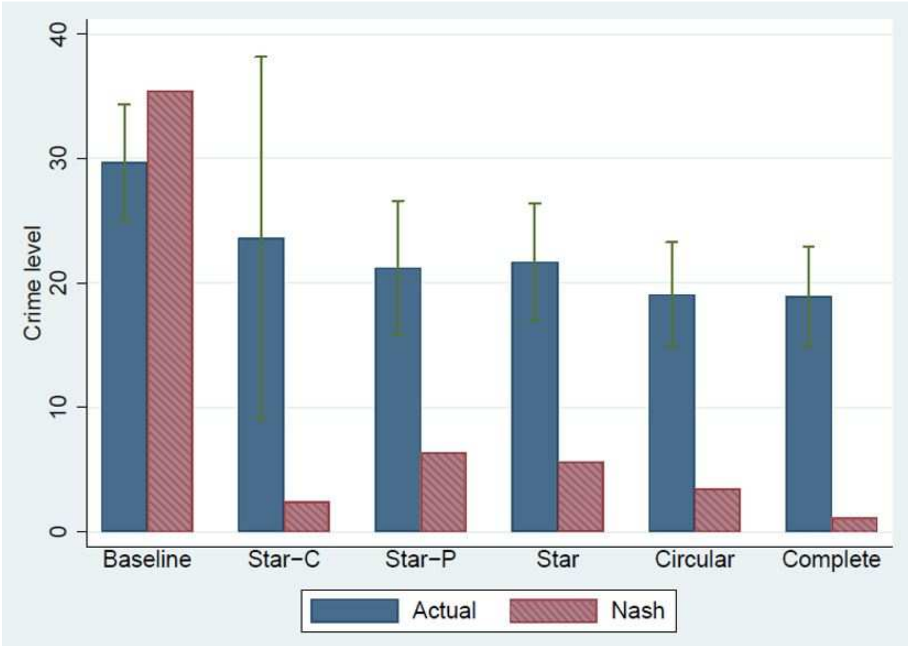

Notes: this figure plots the average crime levels for all treatments. The error bars indicare 95% confidence intervals. Star-C represents the numbers for the center players in the star treatment while Star-P stands for the periphery players in the treatment. Star is the aggregated measure for the group average.

Figure 4.3: Crime levels in all treatments

Table 4.2 corroborates the previous result using four ordinary least squares (OLS) models with the following regression specification:

$$\begin{aligned}\text{Crime}_i = c &+ \alpha_1 \mathbb{I}_i^{\{\text{treatment=star center}\}} + \alpha_2 \mathbb{I}_i^{\{\text{treatment=star periphery}\}} + \alpha_3 \mathbb{I}_i^{\{\text{treatment=circular}\}} \\ &+ \alpha_4 \mathbb{I}_i^{\{\text{treatment=complete}\}} + \beta_1 \mathbb{I}_i^{\{\text{gender=male}\}} + \beta_2 \text{Understand}_i \\ &+ \beta_3 \text{GPA-RA}_i + \beta_4 \text{GPA-PIA}_i + \epsilon_i,\end{aligned}$$

where $\mathbb{I}_i$ is an indicator variable which takes the value of 1 if and only if the logical

[18] One difference between the repeated treatments and one-shot treatments comes from the comparison between the center and periphery players in the star network. In the star network with repeated interaction treatment, center players started at a higher crime level than the periphery players as in the one-shot treatment. But these center players reduced the level as they gain experience and the overall level for the center players in repeated treatment is lower than the periphery players, which is as predicted by Nash.

statement in the superscript is true for player i, and 0 otherwise. Understand_i is a categorical variable that reveals player i's self-reported understanding of the experiment. It can have five possible values: 1 being very clear about the experiment, while 5 implying that the participant does not understand the experiment at all.

GPA-RA_i captures player i's deviation from the average grade point average (GPA) at the university if i participated in the baseline, and 0 otherwise; similarly, GPA-PIA_i measures participant i's deviation from the average GPA if i participated in one of the three peer-informed audit treatments, and 0 otherwise. GPA takes a value from 1 to 5: the higher the number, the better the academic performance. More specifically, 1 indicates that the participant's average GPA is between 50%–60%, 2 indicates that the GPA is between 61%–70%, and so on. In previous research, college GPA has been observed to be highly correlated with cognitive ability (see e.g., Thomson & Oppenheimer 2016). Based on this finding, we expect that participants with higher cognitive ability would make decisions closer to the equilibrium level. As our initial analysis suggests that the crime level is less than that predicted by Nash in the baseline but higher than the equilibrium level in the network treatments (see, Figure 4.3), we allow the GPA measure to have a differential effect on the peer-informed mechanism versus the baseline mechanism. To do this, we introduce interaction terms (GPA measure interacted with the enforcement mechanisms).[19]

[19] There are 14 participants who are either in their first year or do not study at the university; thus, the GPA measure is missing for them. We assign these participants the average score. As a robustness check, we re-estimate Model 4 without these observations and find that the results remain the same.

	Model 1	Model 2	Model 3	Model 4
Star center (α_1)	-6.10	-5.84	-5.51	-8.80
	(5.23)	(5.27)	(5.29)	(5.56)
Star periphery (α_2)	-8.50***	-8.69***	-8.76***	-9.33***
	(3.20)	(3.23)	(3.24)	(3.27)
Circular (α_3)	-10.65***	-10.86***	-11.26***	-12.17***
	(3.02)	(3.05)	(3.09)	(3.15)
Complete (α_4)	-10.79***	-10.79***	-11.08***	-10.98***
	(3.02)	(3.03)	(3.06)	(3.06)
Gender (β_1)		-1.28	-1.54	-1.72
		(2.19)	(2.22)	(2.19)
Understand (β_2)			-0.90	-0.97
			(1.09)	(1.07)
GPA-RA (β_3)				-1.24
				(1.98)
GPA-PIA (β_4)				-3.72**
				(1.88)
Constant (c)	29.70***	30.47***	32.67***	32.62***
	(2.14)	(2.51)	(3.67)	(3.65)
$H_0: \frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2 = \alpha_3$	p=0.21	p=0.28	p=0.20	p=0.19
$H_0: \frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2 = \alpha_4$	p=0.22	p=0.29	p=0.23	p=0.24
$H_0: \alpha_3 = \alpha_4$	p=0.98	p=0.97	p=0.94	p=0.62
$H_0: \alpha_1 = \alpha_2$	p=0.39	p=0.46	p=0.38	p=0.50

Notes: This table reports OLS estimates of crime levels under different network structures. There are 100 observations for each regression. The standard errors are reported in parentheses below each coefficient. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Pairwise t-tests of treatment differences caused by network structures and different positions in the star network are reported in the bottom panel.

Table 4.2: Regressions comparing crime levels

In Model 1, we exclude all control variables and sequentially add them in Models 2–4. Consistent with Figure 4.3, our mechanism significantly decreases the crime level in both a statistical sense and an economic sense. For example, the coefficient for the complete network (α_4) in Model 1 (-10.79 ($p<0.001$)) indicates that this network reduces by over 36% of the crime level compared to the baseline where the crime level is 29.70 units.[20]

However, there is no discernible treatment effect between network treatments and all pairwise t-tests fail to reject the null hypothesis of equality of coefficients across treatments. (i.e., there is no significant difference in pairwise comparison of $\alpha_s := \left(\frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2\right)$, α_3

[20]The coefficient for the star center is negative as predicted but is estimated with much less precision, which is partially due to the specific star network structure that would have a smaller number of observations for the center.

and α_4). In addition, the coefficients attached to all peer-informed audit treatments (α_1-α_4) indicate that the effectiveness of the mechanism in deterring crime is significantly less than that predicted by the Nash equilibrium ($p<0.001$ in all cases using two-tailed t-tests). Therefore, our results provide strong support for Hypothesis 1, but do not support Hypotheses 2 and 3.

Turning briefly towards the data from the repeated game, the results are qualitatively the same as those from the one-shot game. We find that the peer-informed audit mechanism successfully reduces crime from 26.66 units in the baseline to 18.73 units in the network treatments, which is a reduction of 30% ($p < 0.001$). There are also no significant differences in the crime level between the network structures ($p > 0.1$) in all the three pairwise comparisons). In addition, the center player selects a higher level of crime than the peripheral players in the star network in the first round, even though it is not statistically significant. In Appendix A3, we replicate the analysis of the one-shot game to the repeated game data and discuss the evolution of crime level in all treatments.

While the peer-informed audit mechanism deters crime overall, there are two patterns in the data that we wish to explore more in order to improve our understanding of how this mechanism works. The first pattern concerns the divergence of the data from the Nash equilibrium and the second relates to lack of treatment differences among the networks.[21] One likely explanation for these patterns is bounded rationality and belief-driven behavior, which is brought to our attention by two pieces of evidence in our data. First, those with higher cognitive ability (GPA) select significantly lower crime levels in our peer-informed audit treatments (i.e., β_4 is significantly smaller than 0 in Table 4.2). Second, the responses from the post-experiment questionnaire reveal differential strategies, which are perhaps based on the reasoning abilities of subjects. For example, we find that a few (approximately 5%) participants in the peer-informed audit treatments believed that the probability of being fined is purely random— and, thus, they failed to perform the peer-audit induction at all; on the other hand, approximately 35% of the participants

[21] These patterns are unlikely to be driven by participants not understanding the experiment. Evidence from the incentivized pre-experiment quiz reveals that in the peer-informed audit treatments, over 80% of the participants correctly solved at least 5 out of 7 or 8 questions (depending on the treatment). Responses from the post-experiment questionnaire indicate that only a small proportion (approximately 10%) of participants found the experiment to be not fully clear. Moreover, importantly, all the participants audited by the computer correctly challenged the player that gave them the best opportunity to avoid the fine in the peer-informed audit treatments.

reported that they chose to produce less than what they believed the other players chose in order to avoid the fine. We elaborate on this explanation in the next section using a level-k reasoning model.

4.5. Level-k model with network structures

We consider a simple model of bounded rationality to help explain the difference in crime levels predicted in Nash equilibrium with those observed in the data. Rather than assuming that all players have the same level of sophistication, we apply a level-k thinking model which assumes that any given player believes that all other players have lower reasoning ability than her own and best responds to such beliefs. As is convention, we assume that each player believes that all of her peers have reasoning levels that one-less than her own. We also assume that the criminal decisions made by any level-0 subject is equivalent to the Nash crime prediction in the baseline treatment (i.e., level-0 subjects fail to see the relationship between their crime level and the probability of being liable for a fine).[22]

If we assume that everyone plays the weakly dominant strategy in the second stage (which is supported by the data) and behaves as a level-k player in the first stage, then the equilibrium condition that predicts player i with reasoning level k to play e_i^k against everyone else playing e_{-i}^k is expressed in the following manner:

$$\frac{\partial \mathbb{E}\left[u_i\left(e_i^k, e_{-i}^{k-1}\right)\right]}{\partial e_i^k} = 0$$

for all $i \in \{1, 2, ..., n\}$, where the superscripts of e_i^k and e_{-i}^{k-1} denotes the sophistication

[22] We use the Nash prediction in the baseline as the level-0 crime level for several reasons. First, it is documented in the post-experiment questionnaire that a few participants in peer-informed audit treatments consider the fine to be random and independent of the crime levels. Second, it is independent of the choice of the action space. The action space (i.e., participants must choose a crime level between 0 and 40 units) is arbitrarily determined. If we were to assume that people make random choices along the action space, as assumed in the beauty contest literature, then the level-0 benchmark depends on the choice of action space. This is problematic in our case, as the reasoning level of the participants can be shifted by using different action spaces. Third, the choice of the level-0 benchmark does not affect the conclusion that the effectiveness of the mechanism is positively related to the sophistication levels of the participants, as long as this choice is larger than the corresponding Nash equilibrium.

level of the players.

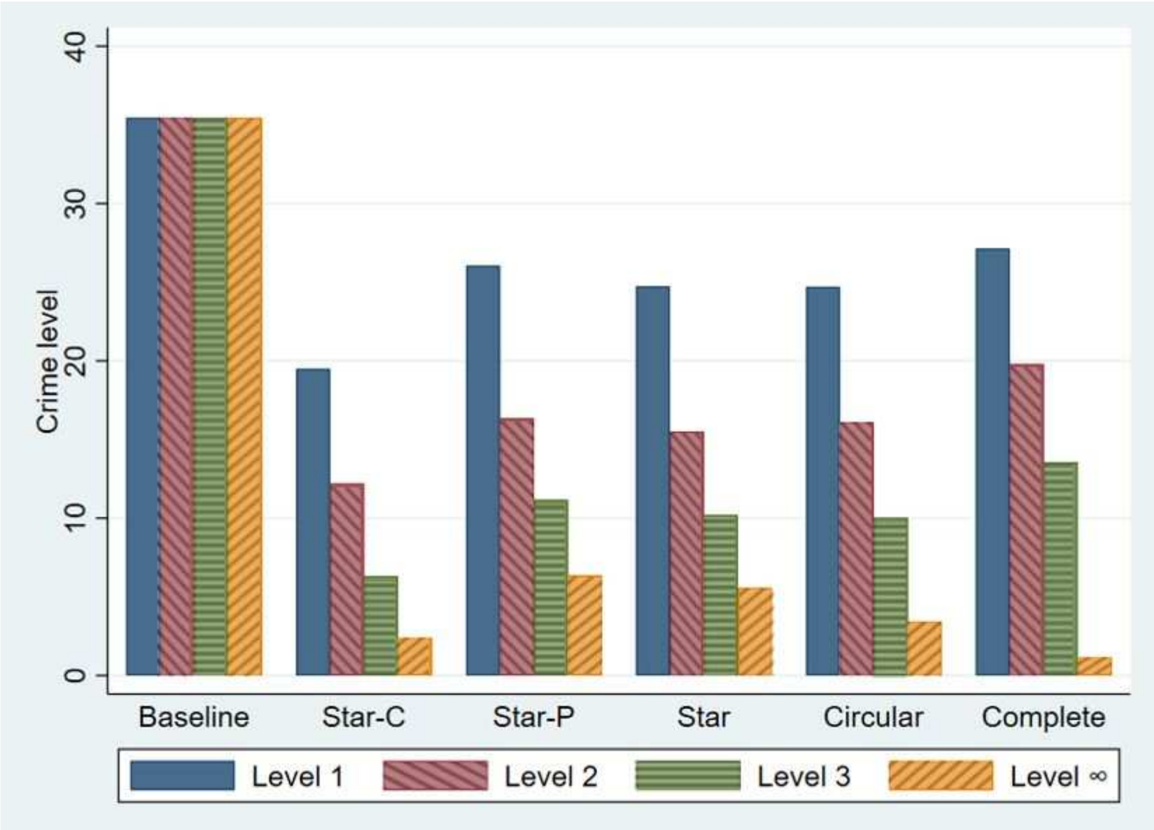

Notes: This figure plots the average crime level at different levels of sophistication for each treatment. We assume that level-0 drug dealers do not take our mechanism into account and produce the Nash equilibrium of the baseline audit scheme.

Figure 4.4: Theoretical prediction based on level-k reasoning

Figure 4.4 presents the predicted decisions of players with level-k sophistication.[23] The predicted crime level is lower for level-3 players as compared to level-2 players, and similarly level-2 players would engage in less crime as compared to level-1 players, and so on. It must be noted that the crime levels in all networks and network positions decrease monotonically toward the Nash prediction (denoted by $k = \infty$) under the peer-informed audit mechanism. Under the predictions of the level-k model, our mechanism still reduces the crime level; however, importantly, the magnitude depends on the beliefs and sophistication levels of participants.

Interestingly, the level-k model predicts that the complete network has a higher crime

[23] The derivation of the numerical solutions involves deriving and applying the drug dealer's best response functions, as done in the Nash equilibrium. However, rather than requiring that all drug dealers best respond to each other, the level-k model calculates best responses to all peers, taking into consideration strategies that are one level below one's own in terms of sophistication.

level than the circular network and the star network when the level of sophistication is close to zero. This is the reverse of the corresponding Nash prediction because there is an extra incentive to "exploit" the less sophisticated players in one's neighborhood in the level-k model and such an incentive is stronger when there are more connected players to be exploited. This exploitation opportunity weakens with an increase in the sophistication level (i.e., the predicted action gap between level-k and level-k+1 converges to zero as k increases in all networks and the exploitation opportunity vanishes) and the level-k predictions approach the Nash predictions. This observation may explain why there are no significant differences between treatments with the peer-informed audit mechanism if participants have a spectrum of sophistication from high to low levels.

We plot the histogram from all four treatments against the level-k predictions in Figure 4.5. In the baseline, more than half of the participants selected crime levels that are within 5 units of the Nash equilibrium, and the highest bar indicates that crime levels between 35 to 40 units were the most popular strategies among participants. However, the patterns for peer-informed audit treatments are different from that of the baseline. Less than 30% of the participants in the star, circular, and complete treatments selected crime levels within five units of the Nash prediction. In addition, the graph for each of the peer-informed audit treatments indicates that participants are clustered between sophistication levels 1 and 2.[24] This coincides with much of the experimental literature on level-k reasoning, which reports that most individuals hold levels of sophistication 3 and below (see, e.g., Ho et al. 1998; Arad & Rubinstein 2012). These patterns also suggest that the level-k model complements the Nash predictions and provides additional insights for understanding criminal behavior.

[24]The aggregated patterns from the repeated game sessions are similar to the one-shot games (Figure 4.5) and most players are also categorized between level-0 and level-3. Following the convention in the level-k literature, we focus on the one-shot treatments in order to provide the cleanest test of level-k reasoning and to exclude complications caused by reputation, learning, cooperation etc., which may confound the interpretation of the sophistication level.

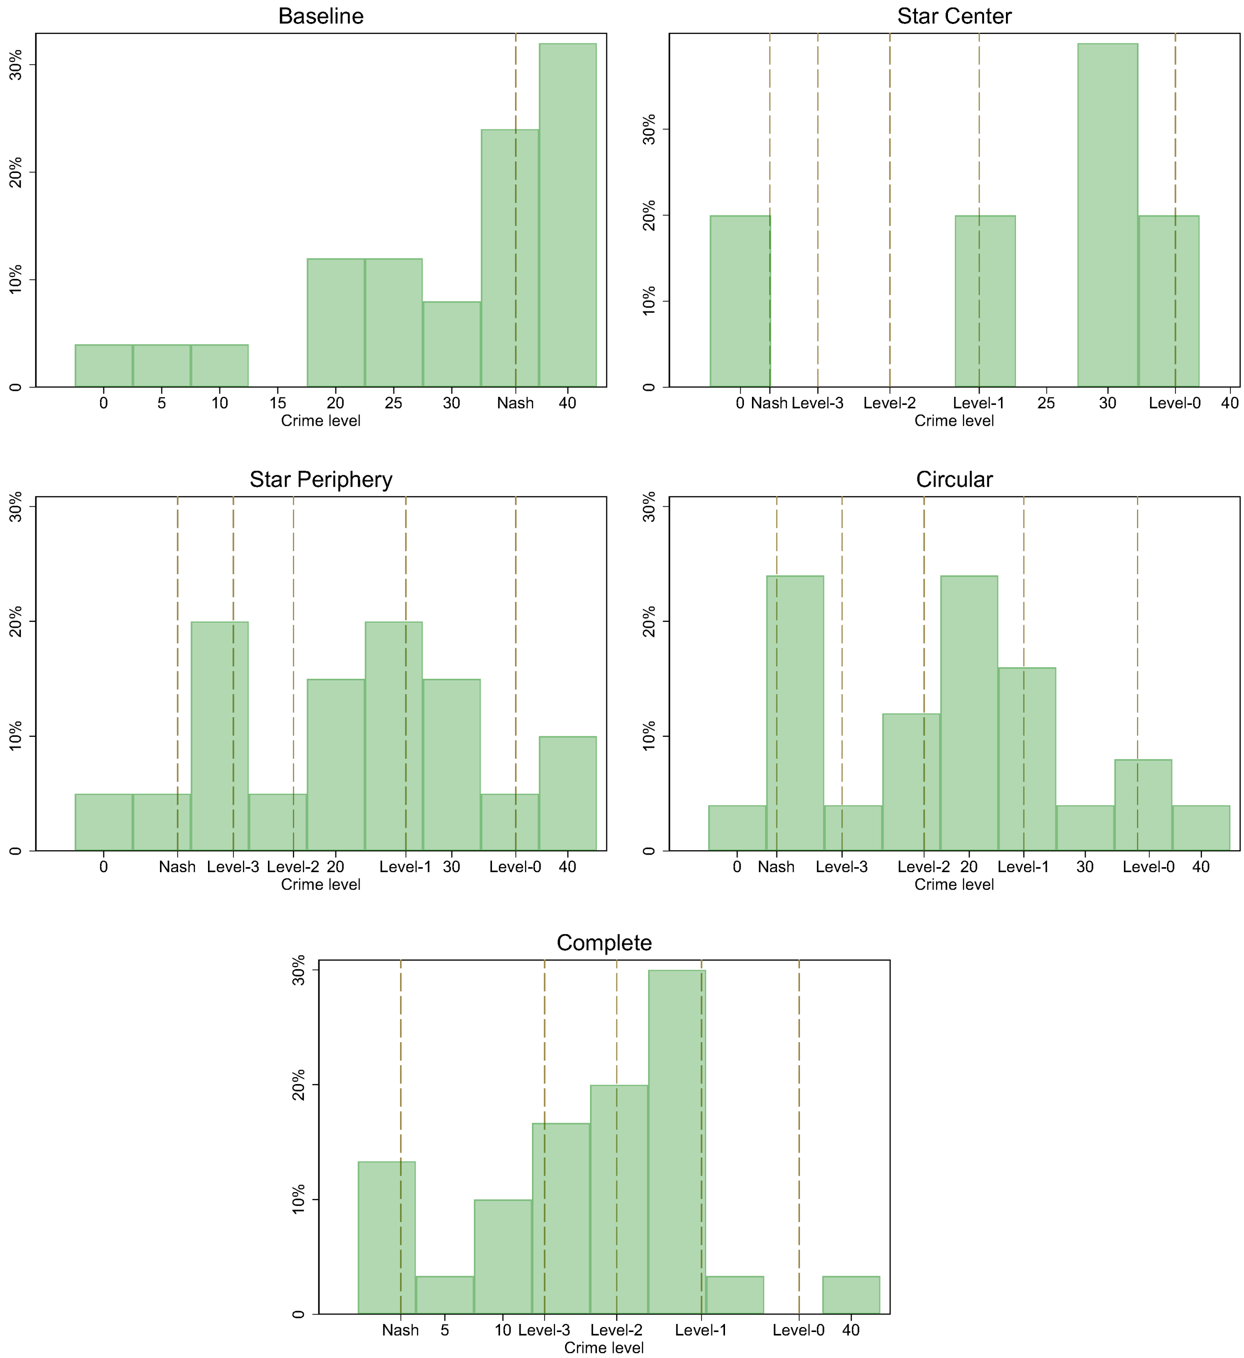

Notes: This figure plots the histogram of the crime levels in the one-shot treatments against the prediction from the level-k model. The title of each panel specifies the treatment and the position in the network.

Figure 4.5: Crime level and level-k reasoning

4.6. Conclusion

In this paper, we proposed a novel peer-informed audit stage added to the traditional baseline audit scheme to utilize the insider information in networks to further deter crime. This would be particularly important in cases in which it is important to apprehend the most severe offender or the key criminal in the group. Consistent with the Nash equilibrium, our experiment reveals that the new mechanism successfully lowers the crime levels for all the network structures we consider.

Interstingly, we find that crime levels are above what is predicted based on the Nash equilibrium and that there are no significant differences across different network treatments. We also find that participants with higher cognitive skills (GPA) perform closer to the Nash play and that certain participants report using strategies similar to those captured by level-k reasoning. With a level-k model adapted to our network setting, we investigate bounded rationality as a potential explanation for observed criminal activities. We find that the level-k model fits the patterns observed in the experiment, implying that naive or bounded-rational participants may impede the full effectiveness of the peer-informed audit mechanism to deter crimes through social networks. Nash predictions may, therefore, not be a good benchmark for mechanisms where one player's best response depends on others' choices as well. Policymakers could consider behavioral models in such cases, particularly when behavioral models yield quantitatively opposite predictions from Nash.

This paper is a first attempt to utilize insider information contained in social networks to deter crime. We highlight the peer-informed audit mechanism in a simple and stylized environment. Although ample evidence shows that laboratory experiments capture the essence of behavior in naturally occurring environments (see e.g., Benz & Meier 2008; Fehr & Leibbrandt 2011; Alm et al. 2015), more research needs to be conducted to understand deterrence using social networks.

We envision several avenues for future research. First, in the field, many social networks are endogenous and dynamic. The introduction of the mechanism may change the network structure (for example., some people may tend to hide their actions from peers). Second, while in theory, we assume no disutility in challenging peers, this may not always be the case. Some members in a crime group may be close friends or even come from the same family who may collude by not challenging at all. Moreover, potential retaliation from those who can infer the source of the challenge may also weaken the incentive to challenge others. It would be interesting to examine the effectiveness of the mechanism in such cases and how it could be modified. Hence, testing the mechanism in a field setup may provide further insights regarding the effectiveness of this mechanism and the development of new deterrence mechanisms. Third, it is interesting to extend the chain of the peer-informed audit process to take further advantage of insider information in networks. For example, the peer-informed audit process could continue as long as the peer informant can nominate another group member responsible for higher criminal activity. Finally, we foresee merit in additional research into level-k reasoning patterns within social network environments.

4.7. References

Abbink, K., Dasgupta, U., Gangadharan, L., and Jain, T. (2014). Letting the briber go free: an experiment on mitigating harassment bribes. Journal of Public Economics, 111: 17-28.

Agranov, M., Potamites, E., Schotter, A., and Tergiman, C. (2012). Beliefs and endogenous cognitive levels: An experimental study. Games and Economic Behavior, 75(2): 449-463.

Alm, J. (1991). A perspective on the experimental analysis of taxpayer reporting. The Accounting Review, 66(3):577-593.

Alm, J., Bloomquist, K. M., and McKee, M. (2015). On the external validity of laboratory tax compliance experiments. Economic Inquiry, 53(2): 1170-1186.

Alm, J. and McKee, M. (2004). Tax compliance as a coordination game. Journal of Economic Behavior & Organization, 54(3): 297-312.

Andreoni, J. and Bernheim, D. (2009). Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects. Econometrica, 77(5):1607-1636.

Arad, A. and Rubinstein, A. (2012). The 11-20 money request game: A level-k reasoning study. American Economic Review, 102(7):3561-73.

Ballester, C., Calvo-Armengol, A., and Zenou, Y. (2006). Who's who in networks. wanted: The key player. Econometrica, 74(5):1403-1417.

Becker, G. S. (1968). Crime and punishment: An economic approach. In The economic dimensions of crime, pages 13-68. Springer.

Belleflamme, P. and Bloch, F. (2004). Market sharing agreements and collusive networks. International Economic Review, 45(2):387-411.

Bigoni, M., Fridolfsson, S.-O., Le Coq, C., and Spagnolo, G. (2012). Fines, leniency, and rewards in antitrust. The RAND Journal of Economics, 43(2):368-390.

Bigoni, M., Fridolfsson, S.-O., Le Coq, C., and Spagnolo, G. (2015). Trust, leniency, and deterrence. The Journal of Law, Economics, and Organization, 31(4):663-689.

Bloch, F. and Olckers, M. (2019). Friend-based ranking. Available at SSRN 3213311.

Calvo-Armengol, A., Patacchini, E., and Zenou, Y. (2009). Peer effects and social networks in education. The Review of Economic Studies, 76(4):1239-1267.

Calvo-Armengol, A. and Zenou, Y. (2004). Social networks and crime decisions: The role of social structure in facilitating delinquent behavior. International Economic Review, 45(3):939-958.

Carr, I. and Lewis, D. (2010). Combating corruption through employment law and whistleblower protection. Industrial Law Journal, 39(1):52-81.

ason, T. N., Friesen, L., and Gangadharan, L. (2016). Regulatory performance of audit tournaments and compliance observability. European Economic Review, 85:288-306.

Charness, G., Feri, F., MelÃ©ndez-JimÃ©nez, M. A., and Sutter, M. (2014). Experimental games on networks: Underpinnings of behavior and equilibrium selection. Econometrica, 82(5):1615-1670.

Chen, Z. and Rey, P. (2013). On the design of leniency programs. The Journal of Law and Economics, 56(4):917-957.

Coricelli, G., Joffily, M., Montmarquette, C., and Villeval, M. C. (2010). Cheating, emotions, and rationality: an experiment on tax evasion. Experimental Economics, 13(2):226-247.

Feltovich, N. and Hamaguchi, Y. (2018). The effect of whistle-blowing incentives on collusion: An experimental study of leniency programs. Southern Economic Journal, 84(4).

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

Friesen, L. (2012). Certainty of punishment versus severity of punishment: An experimental investigation. Southern Economic Journal, 79(2):399-421.

Gill, D. and Prowse, V. (2016). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. Journal of Political Economy, 124(6):1619-1676.

Gilpatric, S. M., Vossler, C. A., and McKee, M. (2011). Regulatory enforcement with competitive endogenous audit mechanisms. The RAND Journal of Economics, 42(2):292-312.

Hahn, Y., Islam, A., Patacchini, E., and Zenou, Y. (2017). Do friendship networks improve female education? IZA Discussion Paper.

Helsley, R. W. and Zenou, Y. (2014). Social networks and interactions in cities. Journal of Economic Theory, 150:426-466.

Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental "p-beauty contests". The American Economic Review, 88(4):947-969.

Jackson, M. O., Rogers, B. W., and Zenou, Y. (2017). The economic consequences of socialnetwork structure. Journal of Economic Literature, 55(1):49-95.

Jin, Y. (2018). Does level-k behavior imply level-k thinking? Available at SSRN 3138321.

Kaiser, B. A. (2007). The athenian trierarchy: Mechanism design for the private provision of public goods. The Journal of Economic History, 67(2):445-480.

KÃ¶nig, M. D., Liu, X., and Zenou, Y. (2019). R&d networks: Theory, empirics, and policy implications. Review of Economics and Statistics, 101(3):476-491.

Leibbrandt, A., Ramalingam, A., Saaksvuori, L., and Walker, J. M. (2015). Incomplete punishment networks in public goods games: experimental evidence. Experimental Economics, 18(1):15-37.

Levitt, S. D. and Venkatesh, S. A. (2000). An economic analysis of a drug-selling gang's finances. The Quarterly Journal of Economics, 115(3):755-789.

Lindquist, M. J. and Zenou, Y. (2019). Crime and networks: 10 policy lessons. Oxford Review of Economic Policy, forthcoming.

Marx, L. M., Mezzetti, C., and Marshall, R. C. (2015). Antitrust leniency with multiproduct colluders. American Economic Journal: Microeconomics, 7(3):205-40.

Miller, N. H. (2009). Strategic leniency and cartel enforcement. American Economic Review, 99(3):750-68.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. The American Economic Review, 85(5):1313-1326.

Nakajima, R., Tamura, R., and Hanaki, N. (2010). The effect of collaboration network on inventors' job match, productivity and tenure. Labour Economics, 17(4):723-734.

Oprea, R. (2020). What makes a rule complex? Working Paper.

Polinsky, A. M. and Shavell, S. (1979). The optimal tradeoff between the probability and magnitude of fines. The American Economic Review, 69(5):880-891.

Thomson, K. S. and Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. Judgment and Decision making, 11(1):99.

van Leeuwen, B., Ramalingam, A., Arjona, D. R., and Schram, A. (2019). Centrality and cooperation in networks. Experimental economics, 22(1):178-196.

van Leeuwen, B., Ramalingam, A., Rojo Arjona, D., and Schram, A. J. (2015). Authority and centrality: Power and cooperation in social dilemma networks. Available at SSRN 2573348.

4.8. Appendix

4.8.1. Derivation of Nash predictions

The baseline audit mechanism

The authority randomly audits a drug dealer j and fines the drug dealer by $f\left(\hat{e}_j\right)$. Drug dealer i's expected payoff is

$$\mathbb{E}\left[u_i\left(e_i, e_{-i}\right)\right] = \pi\left(e_i\right) - \theta \cdot \frac{1}{n} \cdot \mathbb{E}\left[f\left(e_i + \epsilon_i\right)\right].$$

In equilibrium, the marginal profit from committing an extra unit of crime is counterbalanced by the potential fine; thus, we have

$$\begin{aligned}\frac{\partial \mathbb{E}\left[u_i\left(e_i, e_{-i}\right)\right]}{\partial e_i} &= \pi'\left(e_i\right) - \frac{\theta}{n} \cdot \mathbb{E}'\left[f\left(e_i + \epsilon_i\right)\right] = 0 \\ \Leftrightarrow e^* &= \left(\frac{\theta f}{abn}\right)^{\frac{1}{b-1}} \approx 35.45.\end{aligned}$$

If the authority can identify two drug dealers without insider information

The authority randomly audits two drug dealers and fines the drug dealer with the higher signal. The expected payoff of player i playing e_i against everyone else playing e_{-i} is

$$\mathbb{E}\left[u_i(e_i, e_{-i})\right] = \pi\left(e_i\right)$$

$$
\underbrace{-\,\theta\frac{C_{n-1}^{1}}{C_{n}^{2}}\cdot\mathbb{E}\left[f\left(e_i+\epsilon_i\right)\cdot\mathbb{P}\left(\hat{e}_i>\hat{e}_j\middle|\,\epsilon_i\right)\right]}_{\text{Expected fine}}
$$

$$
=\pi\left(e_i\right)-\frac{2\theta}{n}\mathbb{E}\left[f\left(e_i+\epsilon_i\right)\cdot\mathbb{P}\left(\hat{e}_i>\hat{e}_j\middle|\,\epsilon_i\right)\right].
$$

In Nash equilibrium, drug dealer i needs to maximize her payoff against other drug dealers' crime e^*. The expected payoff function becomes

$$
\mathbb{E}\left[u_i(e_i,e_{-i})\right]=\pi\left(e_i\right)-\frac{2\theta}{n}\mathbb{E}\left[f\left(e_i+\epsilon_i\right)\cdot\mathbb{P}\left(e_i+\epsilon_i-e^*>\epsilon_j\middle|\,\epsilon_i\right)\right].
$$

We denote the gap between i's crime and other drug dealers' crime $\delta:=e_i-e^*$.

When $2q\geq\delta\geq0$, drug dealer i's crime level is no less than the other drug dealers, but the noise in the signal makes it still possible for i to avoid the fine. The expected fine from being audited or being challenged is

$$
\begin{aligned}
&\frac{f\theta}{2q^2n}\left(2e_iq^2-\frac{1}{2}e_i\delta^2\right)\\
&+\frac{f\theta}{2q^2n}\left(\frac{1}{3}\left(q-\delta\right)^3+\frac{1}{2}\left(\delta+q\right)\left(q-\delta\right)^2+\frac{1}{3}q^3-\frac{1}{2}\left(\delta q^2+q^3\right)\right)\\
&+\frac{2f\theta}{2qn}\left(\delta e_i+q\delta-\frac{1}{2}\delta^2\right).
\end{aligned}
$$

The marginal expected fine is

$$
\begin{aligned}
&\frac{f\theta}{2q^2n}\left(2q^2-\frac{1}{2}\left(\delta^2+e_i2\delta\right)\right)\\
&+\frac{f\theta}{2q^2n}\left(-\left(q-\delta\right)^2+\frac{1}{2}\left(\left(q-\delta\right)^2-2\left(\delta+q\right)\left(q-\delta\right)\right)-\frac{1}{2}\left(q^2\right)\right)\\
&+\frac{2f\theta}{2qn}\left(\delta+e_i+q-\delta\right).
\end{aligned}
$$

When $-2q\leq\delta\leq0$, drug dealer i's crime level is no more than other drug dealers, but the noise in the signal still makes it possible for i to pay the fine. The expected fine from being audited or being challenged is

$$
\frac{f\theta}{2q^2n}\left(\frac{1}{3}q^3+\frac{1}{2}\left(\delta+q+e_i\right)q^2+\left(e_iq+e_i\delta\right)q\right)
$$

$$+\frac{f\theta}{2q^2n}\left(\frac{1}{3}(q+\delta)^3-\frac{1}{2}(\delta+q+e_i)(q+\delta)^2+(e_iq+e_i\delta)(q+\delta)\right).$$

The marginal expected fine is

$$\frac{f\theta}{2q^2n}\left(q^2+q(q+e_i+\delta)+(e_iq+e_i\delta)\right).$$

In Nash equilibrium, when each player expects that othersemploy the same strategy (i.e., $e_i = e^*$), the marginal benefit from committing an extra unit crime is equal to the marginal costs; thus, we have

$$ab\left(e_i^*\right)^{b-1}-\frac{2f\theta}{2qn}(e^*+q)=0$$
$$\Leftrightarrow e^*\approx 6.58.$$

Peer-informed audit mechanism with a circular network

The authority randomly audits one drug dealer and the drug dealer can challenge one of the neighbors. In the second stage, the weakly dominant strategy for the audited drug dealer is to challenge the neighbor with the bigger signal. If we assume that everyone plays the weakly dominated strategy in the second stage, then the expected payoff of player i playing e_i against everyone else playing e_{-i} is

$$\begin{aligned}\mathbb{E}\left[u_i(e_i,e_{-i})\right]&=\pi(e_i)\\&\quad-\underbrace{\frac{\theta}{n}\cdot\mathbb{E}\left[f(e_i+\epsilon_i)\cdot\mathbb{P}\left(\hat{e}_i>\max\left\{\widehat{e_{i-1}},\widehat{e_{i+1}}\right\}\middle|\,\epsilon_i\right)\right]}_{\text{Expected fine from being audited}}\\&\quad-\underbrace{\frac{2\theta}{n}\cdot\mathbb{E}\left[f(e_i+\epsilon_i)\cdot\mathbb{P}\left(\hat{e}_i>\max\{\widehat{e_{i-1}},\widehat{e_{i-2}}\}\middle|\,\epsilon_i\right)\right]}_{\text{Expected fine from being challenged}}\\&=\pi(e_i)-\frac{3\theta}{n}\mathbb{E}\left[f(e_i+\epsilon_i)\cdot\mathbb{P}\left(\hat{e}_i>\max\{\widehat{e_{i-1}},\widehat{e_{i-2}}\}\middle|\,\epsilon_i\right)\right].\end{aligned}$$

In Nash equilibrium, drug dealer i must maximize her payoff against the crime of other drug dealers e^*. The expected payoff function becomes

$$\mathbb{E}\left[u_i(e_i, e^*_{-i})\right] = \pi\left(e_i\right) - \frac{3\theta}{n} \cdot \mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \max\{\hat{e^*_{i-1}}, \hat{e^*_{i-2}}\}\middle| \epsilon_i\right)\right].$$

We denote the gap between i's crime and the crime of other drug dealers as $\delta := e_i - e^*$.

When $2q \geq \delta \geq 0$, drug dealer i's crime level is no less than the other drug dealers, but the noise in the signal still makes it possible for i to avoid the fine. Then, the expected fine from being audited or being challenged is

$$\frac{3\theta}{n}\frac{1}{2q}\left(fq^2 + \frac{2qf}{3}\left(e_i - \delta - q\right)\right) - \frac{f}{\left(2q\right)^3}\left(\frac{1}{4}\delta^4 + \frac{1}{3}\left(e_i - \delta - q\right)\delta^3\right) + \frac{1}{2q}\left(f\delta e_i + qf\delta - \frac{1}{2}f\delta^2\right).$$

The marginal expected fine from being audited or being challenged is

$$\frac{3\theta}{n}\left[\frac{f}{\left(2q\right)^3}\left(\delta^3 + \left(e_i - \delta - q\right)\delta^2\right) + \frac{1}{2q}\left(f\delta + fe_i + qf - f\delta\right)\right].$$

When $-2q \leq \delta \leq 0$, drug dealer i's crime level is no more than other drug dealers, but the noise in the signal still makes it possible for i to pay the fine. Then, the expected fine from being audited or being challenged is

$$\frac{3\theta}{n}\frac{f}{\left(2q\right)^3}\left(\frac{1}{4}\left(2q + \delta\right)^4 + \frac{1}{3}\left(e_i - \delta - q\right)\left(2q + \delta\right)^3\right).$$

The marginal expected fine from being audited or being challenged is

$$\frac{3\theta}{n}\left[\frac{f}{\left(2q\right)^3}\left(\left(2q + \delta\right)^3 + \left(e_i - \delta - q\right)\left(2q + \delta\right)^2\right)\right].$$

When $\delta \geq 2q$, drug dealer i's crime level is so much higher than other drug dealers' crime levels such that it cannot escape the fine whenever she is challeged through an audit; when $\delta \leq -2q$, drug dealer i's crime level is so much less than other drug dealers' crime levels, such that it never pays the fine. These two cases clearly cannot be an equilibrium solution as i can reduce/increase crime level to increase her expected payoff.

In Nash equilibrium, when each player expects that others use the same strategy (i.e., $e_i = e^*$), the net marginal benefit from committing an extra unit crime is zero; thus, $\frac{\partial \mathbb{E}\left[u_i(e_i^*, e_{-i}^*)\right]}{\partial e_i^*} = \pi'\left(e_i^*\right) - \frac{3\theta}{n}\frac{1}{2q}\left(fe_i^* + qf\right) = 0$. The Nash equilibrium is the solution to

$$ab\left(e_i^*\right)^{b-1} - \frac{3\theta}{2qn}\left(fe_i^* + qf\right) = 0,$$

where the unique numerical solution is

$$e^* \approx 3.43.$$

Peer-informed audit mechanism within the star network

The authority randomly audits one drug dealer. If the center drug dealer is audited, then she can challenge all her peers; if a periphery drug dealer is audited, then he can only challenge the center drug dealer. In the second stage, the weakly dominant strategy for a periphery drug dealer, if audited, is to challenge the center drug dealer; while the strategy for the center drug dealer is to challenge the periphery drug dealer with the largest signal. Due to such asymmetry in the challenge possibility, we need to study the conditions for both center and periphery players.

The center player

For the center player, if we assume all periphery drug dealers play the weakly dominated strategy in the second stage, then the expected payoff of the center drug dealer i playing e_i against everyone else playing e_{-i} is

$$\begin{aligned}\mathbb{E}\left[u_i(e_i, e_{-i})\right] = \pi\left(e_i\right) & \\ & - \frac{\theta}{n} \cdot \underbrace{\mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \max\left\{\hat{e_{-i}}\right\} \middle| \epsilon_i\right)\right]}_{\text{Expected fine from being audited}} \\ & - \frac{\theta\left(n-1\right)}{n} \cdot \underbrace{\mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \hat{e}_j \middle| \epsilon_i\right)\right]}_{\text{Expected fine from being challenged}}.\end{aligned}$$

In the Nash equilibrium, drug dealer i must maximize her payoff against the crime level of periphery drug dealers e_p^*. We denote the gap between center drug dealer i's crime and any periphery drug dealer j's crime $\delta_c := e_i - e_p^*$. When $2q \geq \delta_c \geq 0$, the center drug dealer i commits no less than the periphery drug dealers, but the noise makes it possible for i to nevertheless avoid the fine. The expected fine from the directly audited by the authority is

$$\frac{\theta}{n} \cdot \mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(e_i + \epsilon_i > e_p^* + \max\left\{\epsilon_{-i}\right\}\middle| \epsilon_i\right) \mathbb{I}\left(-q < \epsilon_i < q - \delta_c\right) + f\left(e_i + \epsilon_i\right) \mathbb{I}\left(q - \delta_c < \epsilon_i < q\right)\right]$$
$$= \frac{f\theta}{n} \cdot \left[\frac{1}{n}\left(e_i - q - \delta_c\right) + \frac{2q}{n+1} - \frac{1}{(2q)^n}\left(\left(e_i - q - \delta_c\right)\frac{1}{n}\delta_c^n + \frac{1}{n+1}\delta_c^{n+1}\right) + \frac{\delta_c}{2q}\left(e_i + \frac{1}{2}\left(2q - \delta_c\right)\right)\right].$$

The marginal expected fine from being directly audited by the authority is

$$\frac{\theta}{n}\left[-\frac{f}{(2q)^n}\left(e_i - q\right)\delta_c^{n-1} + \frac{f}{2q}\left(e_i + q\right)\right].$$

The expected fine from being challenged by a periphery player is

$$\frac{\theta\left(n-1\right)}{n}\mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e_i} > \hat{e_p^*}\middle| \epsilon_i\right) \mathbb{I}\left(-q < \epsilon_i < q - \delta_c\right) + f\left(e_i + \epsilon_i\right) \mathbb{I}\left(q - \delta_c < \epsilon_i < q\right)\right]$$
$$= \frac{(n-1)}{n}\frac{f\theta}{4q^2}\left(\frac{1}{3}(2q)^3 + \left(e_i - q - \delta_c\right)2q^2 - \frac{1}{3}\delta_c^3 - \frac{1}{2}\left(e_i - q - \delta_c\right)\delta_c^2\right)$$
$$+ \frac{(n-1)}{n}\frac{f\theta}{2q}\left(e_i\delta_c + \frac{1}{2}\left(2q\delta_c - \delta_c^2\right)\right).$$

The marginal expected fine from being challenged is

$$\frac{(n-1)\,\theta}{n}\left[-\frac{f}{4q^2}\delta_c\left(e_i - q\right) + \frac{f}{2q}\left(e_i + q\right)\right].$$

When $-2q \leq \delta_c < 0$, drug dealer i commits no more crime than the periphery drug dealers, but the noise in the signal makes it still possible for i to pay the fine. The expected fine from being directly audited by the authority is

$$\frac{\theta}{n}\mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(e_i + \epsilon_i > e_p^* + \max\left\{\epsilon_{-i}\right\}\middle| \epsilon_i\right) \mathbb{I}\left(-q - \delta_c < \epsilon_i < q\right)\right]$$
$$= \frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{1}{n+1}\left(2q + \delta_c\right)^{n+1} + \left(e_i - q - \delta_c\right)\frac{1}{n}\left(2q + \delta_c\right)^n\right).$$

The marginal expected fine from being directly audited by the authority is

$$\frac{\theta}{n}\frac{f}{(2q)^n}\left(e_i+q\right)\left(2q+\delta_c\right)^{n-1}.$$

The expected fine from being challenged by another player is

$$\begin{aligned}&\frac{\theta\left(n-1\right)}{n}\mathbb{E}\left[f\left(e_i+\epsilon_i\right)\cdot\mathbb{P}\left(\hat{e_i}>\hat{e_p^*}\middle|\epsilon_i\right)\mathbb{I}\left(-q-\delta_c<\epsilon_i<q\right)\right]\\&=\frac{\theta\left(n-1\right)}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}\left(2q+\delta_c\right)^3+\left(e_i-q-\delta_c\right)\frac{1}{2}\left(2q+\delta_c\right)^2\right).\end{aligned}$$

The marginal expected fine from being challenged is

$$\frac{\theta\left(n-1\right)}{n}\frac{f}{(2q)^2}\left(e_i+q\right)\left(2q+\delta_c\right).$$

When $\delta_c \geq 2q$, the center drug dealer commits a lot more crime than periphery drug dealers that she is always liable for the fine; thus, the marginal benefit of committing an extra unit of crime is $\pi'\left(e_i\right)-f<0$. When $\delta_c<-2q$, the center drug dealer commits a lot less crime than periphery drug dealers that she is never liable for the fine; thus, the marginal expected benefit of committing an extra unit of crime is $\pi'\left(e_i\right)>0$.

In Nash equilibrium, the marginal expected benefit from committing an extra unit of crime is zero:

$$\frac{\partial\mathbb{E}\left[u_i(e_i,e_{-i}^*)\right]}{\partial e_i}=\begin{cases}\pi'\left(e_i\right)-f=0\text{ (impossible)} & \text{if }\delta_c>2q\\ \pi'\left(e_i\right)+\frac{\theta}{n}\left[\frac{f}{(2q)^n}\left(e_i-q\right)\delta_c^{n-1}-\frac{f}{2q}\left(e_i+q\right)\right] & \\ +\frac{(n-1)\theta}{n}\left(\frac{f}{4q^2}\delta_c\left(e_i-q\right)-\frac{f}{2q}\left(e_i+q\right)\right)=0 & \text{if }2q\geq\delta_c\geq 0\\ \pi'\left(e_i\right)-\frac{\theta}{n}\frac{f}{(2q)^n}\left(e_i+q\right)\left(2q+\delta_c\right)^{n-1} & \\ -\frac{(n-1)\theta}{n}\frac{f}{(2q)^2}\left(e_i+q\right)\left(2q+\delta_c\right)=0 & \text{if }-2q\leq\delta_c<0\\ \pi'\left(e_i\right)=0\text{ (impossible)} & \text{if }\delta_c<-2q.\end{cases}$$

A periphery player

Similarly, for a periphery drug dealer, if we assume that the center drug dealer plays the weakly dominated strategy in the second stage, then the expected payoff of the periphery drug dealer j playing e_j against the center drug dealer playing e_i and other periphery drug dealers playing $e_{-j\backslash i}$ is

$$\begin{aligned}\mathbb{E}\left[u_j(e_j,e_{-j})\right] = {} & \pi\left(e_j\right) \\ & -\underbrace{\frac{\theta}{n}\cdot\mathbb{E}\left[f\left(e_j+\epsilon_j\right)\cdot\mathbb{P}\left(\hat{e}_j>\hat{e}_i\middle|\,\epsilon_j\right)\right]}_{\text{Expected fine if being audited}} \\ & -\underbrace{\frac{\theta}{n}\cdot\mathbb{E}\left[f\left(e_j+\epsilon_j\right)\cdot\mathbb{P}\left(\hat{e}_j>\max\left\{\hat{e_{-j}}\right\}\middle|\,\epsilon_j\right)\right]}_{\text{Expected fine if being challenged}}.\end{aligned}$$

In Nash equilibrium, drug dealer j must maximize her payoff against the center player's crime e_c^* and the crime of each of the other periphery drug dealers e_p^*. We denote the gap between the crime of the periphery drug dealer j and the crime of the center drug dealer i by $\delta_p = e_j - e_c^*$ and the gap between drug dealer j and the crime of any of the other periphery drug dealer as $\delta := e_j - e_p^*$.

Expected fine from being directly audited

If a periphery drug dealer is audited, then challenging the center drug dealer is the only option and, thus, the crime levels of other periphery drug dealers are of no consequence.

When $2q \geq \delta_p \geq 0$ (regardless of δ), periphery drug dealer j commits more crime than the center, but the noise made it possible for j to still pass the fine to the center drug dealer. The expected fine for j from being directly audited by the authority is

$$\frac{\theta}{n}\left[\frac{f}{\left(2q\right)^2}\left(\frac{1}{3}\left(2q\right)^3+\frac{1}{2}\left(e_j-q-\delta_p\right)\left(2q\right)^2\right)-\frac{f}{\left(2q\right)^2}\left(\frac{1}{3}\delta_p^3+\frac{1}{2}\left(e_j-q-\delta_p\right)\delta_p^2\right)+\frac{f}{2q}\left(e_j\delta_p+q\delta_p-\frac{1}{2}\delta_p^2\right)\right]$$

The marginal expected fine from being directly audited by the authority is

$$-\frac{\theta}{n}\frac{f}{4q^2}\left(e_j-q\right)\delta_p+\frac{\theta}{n}\frac{f}{2q}\left(e_j+q\right).$$

When $-2q \leq \delta_p \leq 0$ (regardless of δ), periphery drug dealer j commit less crime than the center, but the noise made it possible for j to still pay the fine when audited. The expected fine from being directly audited by the authority is

$$\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}(2q+\delta_p)^3+(e_j-q-\delta_p)\frac{1}{2}(2q+\delta_p)^2\right).$$

The marginal fine from being directly audited by the authority is

$$\frac{\theta}{n}\frac{f}{(2q)^2}\left((2q+\delta_p)^2+(e_j-q-\delta_p)(2q+\delta_p)\right).$$

When $\delta_p \geq 2q$ (regardless of δ), drug dealer j is always liable for the fine if directly audited by the authority. In this case, the marginal expected fine is

$$\frac{f\theta}{n}.$$

When $\delta_p \leq -2q$ (regardless of δ), drug dealer j is never liable for the fine if directly audited by the authority. In this case, the marginal expected fine is

$$0.$$

Expected fine from being challenged

The only possibility that a periphery drug dealer j is challenged is that the center drug dealer i is audited and j has the largest signal among all her peers.

When $\delta_p \geq 2q$:

When $\delta_p \geq 2q$ and $\delta \geq 2q$, drug dealer j commits so much more crime than other periphery and the center drug dealers that j's signal is always the largest. The center drug dealer always tends to challenge j and j is always liable for the fine. The marginal fine from being challenged by the center player is

$$\frac{f\theta}{n}.$$

When $\delta_p \geq 2q$ and $0 \leq \delta \leq 2q$, drug dealer j commits so much more crime than the center drug dealers that j's signal is always larger than the center drug dealer; thus, once challenged, j is liable for the fine. However, the center drug dealer may not choose to challenge drug dealer j because it is possible for other periphery drug dealers to have a larger signal than j due to the draw of the noise. Then, the expected fine from being challenged by the center player is

$$\begin{aligned}
&\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q)^{n}+(e_j-q-\delta)\frac{1}{n-1}(2q)^{n-1}\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}\delta^{n}+(e_j-q-\delta)\frac{1}{n-1}\delta^{n-1}\right)\\
&+\frac{\theta}{n}\frac{f}{2q}\left(e_j\delta+\frac{1}{2}\left(2q\delta-\delta^2\right)\right).
\end{aligned}$$

The marginal expected fine from being challenged by the center player is

$$-\frac{\theta}{n}\frac{f}{(2q)^{n-1}}(e_j-q)\,\delta^{n-2}+\frac{\theta}{n}\frac{f}{2q}(e_j+q).$$

When $\delta_p \geq 2q$ and $-2q \leq \delta \leq 0$, drug dealer j commits so much more crime than the center drug dealers that j's signal is always larger than that of the center drug dealer; thus, once challenged, j is liable for the fine. However, the center drug dealer probably chooses to not challenge drug dealer j because it is likely that one of the other periphery drug dealers has a larger signal than that of j due to the draw of the noises. The expected fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q+\delta)^{n}+(e_j-q-\delta)\frac{1}{n-1}(2q+\delta)^{n-1}\right).$$

The marginal fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left((2q+\delta)^{n-1}+(e_j-q-\delta)(2q+\delta)^{n-2}\right).$$

When $\delta_p \geq 2q$ and $\delta \leq -2q$, drug dealer j is never challenged although j's signal is always higher than that of the center. This is because the other periphery drug dealers definitely have larger signals. In this case, the marginal fine from being challenged by the center player is

$$0.$$

When $0 \leq \delta_p \leq 2q$:

When $0 \leq \delta_p \leq 2q$ and $2q \leq \delta$, drug dealer j is always challenged if the center drug dealer is audited, but who is liable for the fine depends on the draw of noises. The expected fine from being challenged by the center player is

$$
\begin{aligned}
&\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}(q-\delta_p)^3+(e_j+q+\delta_p)\frac{1}{2}(q-\delta_p)^2+(e_jq+e_j\delta_p)(q-\delta_p)\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^2}\left(-\frac{1}{3}q^3+(e_j+q+\delta_p)\frac{1}{2}q^2-(e_jq+e_j\delta_p)q\right)\\
&+\frac{\theta}{n}\frac{f}{2q}\left(e_j+\frac{1}{2}q^2\right)\\
&-\frac{\theta}{n}\frac{f}{2q}\left(e_j(q-\delta_p)+\frac{1}{2}(q-\delta_p)^2\right).
\end{aligned}
$$

The marginal expected fine from being challenged by the center player is

$$
\frac{\theta}{n}\frac{f}{(2q)^2}(\delta_pq-e_j\delta_p)+\frac{\theta}{n}\frac{f}{2q}(1+e_j).
$$

When $0 \leq \delta_p \leq 2q$ and $\delta_p \leq \delta \leq 2q$, drug dealer j commits more crime than all other drug dealers, but the drug dealer that must be challenged by the center drug dealer and the one that is liable for the fine are uncertain. The expected fine from being challenged by the center player is

$$
\begin{aligned}
&\frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{1}{n+1}(2q)^{n+1}+(e_j+\delta_p-2\delta-q)\frac{1}{n}(2q)^n+(\delta_p-\delta)(e_j-\delta-q)\frac{1}{n-1}(2q)^{n-1}\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{1}{n+1}\delta^{n+1}+(e_j+\delta_p-2\delta-q)\frac{1}{n}\delta^n+(\delta_p-\delta)(e_j-\delta-q)\frac{1}{n-1}\delta^{n-1}\right)\\
&+\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}(q-\delta_p)^3+(e_j+q+\delta_p)\frac{1}{2}(q-\delta_p)^2+(e_jq+e_j\delta_p)(q-\delta_p)\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}(q-\delta)^3+(e_j+q+\delta_p)\frac{1}{2}(q-\delta)^2+(e_jq+e_j\delta_p)(q-\delta)\right)\\
&+\frac{\theta}{n}\frac{f}{2q}\left(e_jq+\frac{1}{2}q^2\right)\\
&-\frac{\theta}{n}\frac{f}{2q}\left(e_j(q-\delta_p)+\frac{1}{2}(q-\delta_p)^2\right).
\end{aligned}
$$

The marginal expected fine from being challenged by the center player is

$$-\frac{\theta}{n}\frac{f}{(2q)^n}\left(\delta^n+(e_j+\delta_p-2\delta-q)\,\delta^{n-1}+(\delta_p-\delta)(e_j-\delta-q)\,\delta^{n-2}\right)+\frac{\theta}{n}\frac{f}{2q}(q+e_j).$$

When $0\le\delta_p\le 2q$ and $\delta_p\ge\delta\ge 0$, drug dealer j commits more crime than all the other drug dealers, but the drug dealer that the center drug dealer must challenge and the drug dealer that is liable for the fine are uncertain. Then, the expected fine from being challenged by the center player is

$$\begin{aligned}
&\frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{(2q+\delta-\delta_q)^{n+1}}{n+1}+(e_j+\delta_p-2\delta-q)\frac{(2q+\delta-\delta_q)^n}{n}+(\delta_p-\delta)(e_j-\delta-q)\frac{(2q+\delta-\delta_q)^{n-1}}{n-1}\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{1}{n+1}\delta^{n+1}+(e_j+\delta_p-2\delta-q)\frac{1}{n}\delta^n+(\delta_p-\delta)(e_j-\delta-q)\frac{1}{n-1}\delta^{n-1}\right)\\
&+\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q)^n+(e_j-q-\delta)\frac{1}{n-1}(2q)^{n-1}\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q+\delta-\delta_p)^n+(e_j-q-\delta)\frac{1}{n-1}(2q+\delta-\delta_p)^{n-1}\right)\\
&+\frac{\theta}{n}\frac{f}{2q}\left(e_jq+\frac{1}{2}q^2\right)\\
&-\frac{\theta}{n}\frac{f}{2q}\left(e_j(q-\delta)+\frac{1}{2}(q-\delta)^2\right).
\end{aligned}$$

The marginal fine from being challenged by the center player is

$$-\frac{\theta}{n}\frac{f}{(2q)^n}\left(\delta^n+(e_j+\delta_p-2\delta-q)\,\delta^{n-1}+(\delta_p-\delta)(e_j-\delta-q)\,\delta^{n-2}\right)+\frac{\theta}{n}\frac{f}{2q}(q+e_j).$$

When $0\le\delta_p\le 2q$ and $-2q\le\delta\le 0$, drug dealer j commits more crime than the center drug dealer but less than other periphery drug dealers, but the drug dealer that the center drug dealer must challenge and the one that is liable for the fine are uncertain. The expected fine from being challenged by the center player is

$$\begin{aligned}
&\frac{\theta}{n}\frac{f}{(2q)^n}\left(\frac{(2q+\delta-\delta_q)^{n+1}}{n+1}+(e_j+\delta_p-2\delta-q)\frac{(2q+\delta-\delta_q)^n}{n}+(\delta_p-\delta)(e_j-\delta-q)\frac{(2q+\delta-\delta_q)^{n-1}}{n-1}\right)\\
&+\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q+\delta)^n+(e_j-q-\delta)\frac{1}{n-1}(2q+\delta)^{n-1}\right)\\
&-\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left(\frac{1}{n}(2q+\delta-\delta_p)^n+(e_j-q-\delta)\frac{1}{n-1}(2q+\delta-\delta_p)^{n-1}\right).
\end{aligned}$$

The marginal fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^{n-1}}\left((2q+\delta)^{n-1}+(e_j-q-\delta)(2q+\delta)^{n-2}\right).$$

When $0 \leq \delta_p \leq 2q$ and $\delta \leq -2q$, other periphery drug dealers always have larger signals and, thus, drug dealer j will never be challenged. The marginal fine from being challenged by the center player is

$$0.$$

When $-2q \leq \delta_p \leq 0$,

When $-2q \leq \delta_p \leq 0$ and $\delta \geq 2q$, drug dealer j commits so much more crime than other periphery drug dealers; thus, j is always challenged by the center drug dealer. However, j is not very likely to be liable for the fine as j commits less crime than the center drug dealer. The expected fine from being challenged by the center player is

$$\begin{aligned}&\frac{\theta}{n}\frac{f}{(2q)^2}\left[\frac{1}{3}q^3+(e_j+q+\delta_p)\frac{1}{2}q^2+(e_jq+e_j\delta_p)q\right]\\&-\frac{\theta}{n}\frac{f}{(2q)^2}\left[\frac{1}{3}(-q-\delta_p)^3+(e_j+q+\delta_p)\frac{1}{2}(-q-\delta_p)^2+(e_jq+e_j\delta_p)(-q-\delta_p)\right].\end{aligned}$$

The marginal expected fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^2}\left[q^2+(q+e_j+\delta_p)q\right]+\frac{\theta}{n}\frac{f}{(2q)^2}\left[(e_jq+e_j\delta_p)\right].$$

When $-2q \leq \delta_p \leq 0$ and $0 \leq \delta \leq 2q$, drug dealer j commits less crime than the center drug dealer but more than other periphery drug dealers; however, which drug dealer must be challenged by the center drug dealer and who is liable for the fine are uncertain. Then, the expected fine from being challenged by the center player is

$$\begin{aligned}&\frac{\theta}{n}\frac{f}{(2q)^n}\left[\frac{1}{n+1}(2q)^{n+1}+(e_j-2\delta+\delta_p-q)\frac{1}{n}(2q)^n+(\delta_p-\delta)(e_j-q-\delta)\frac{1}{n-1}(2q)^{n-1}\right]\\&-\frac{\theta}{n}\frac{f}{(2q)^n}\left[\frac{1}{n+1}(\delta-\delta_p)^{n+1}+(e_j-2\delta+\delta_p-q)\frac{1}{n}(\delta-\delta_p)^n+(\delta_p-\delta)(e_j-q-\delta)\frac{1}{n-1}(\delta-\delta_p)^{n-1}\right]\\&+\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}q^3+(e_j+q+\delta_p)\frac{1}{2}q^2+(e_jq+e_j\delta_p)q\right)\\&-\frac{\theta}{n}\frac{f}{(2q)^2}\left(\frac{1}{3}(q-\delta)^3+(e_j+q+\delta_p)\frac{1}{2}(q-\delta)^2+(e_jq+e_j\delta_p)(q-\delta)\right).\end{aligned}$$

The marginal fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^2}\left(q^2+(q+e_j+\delta_p)\,q\right)+\frac{\theta}{n}\frac{f}{(2q)^2}\left(e_jq+e_j\delta_p\right).$$

When $-2q\le\delta_p\le 0$ and $0\ge\delta\ge\delta_p$, drug dealer j commits less crime than all other drug dealers, but the drug dealer that the center drug dealer must challenge and the drug dealer that is liable for the fine are uncertain. The expected fine from being challenged by the center player is

$$\begin{aligned}=&\frac{\theta}{n}\frac{f}{(2q)^n}\left[\frac{(2q+\delta)^{n+1}}{n+1}+(e_j-q+\delta_p-2\delta)\frac{(2q+\delta)^n}{n}+(\delta_p-\delta)(e_j-q-\delta)\frac{(2q+\delta)^{n-1}}{n-1}\right]\\ &-\frac{\theta}{n}\frac{f}{(2q)^n}\left[\frac{1}{n+1}(\delta-\delta_p)^{n+1}+(e_j-q+\delta_p-2\delta)\frac{1}{n}(\delta-\delta_p)^n-(e_j-q-\delta)\frac{1}{n-1}(\delta-\delta_p)^n\right].\end{aligned}$$

The marginal fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^n}\left[(2q+\delta)^n+(e_j+\delta_p-2\delta-q)(2q+\delta)^{n-1}+(\delta_p-\delta)(e_j-q-\delta)(2q+\delta)^{n-2}\right].$$

When $-2q\le\delta_p\le 0$ and $-2q\le\delta\le\delta_p$, drug dealer j commits less crime than all other drug dealers, but the drug dealer that must be challenged by the center drug dealer and the one that is liable for the fine are uncertain. Then, the expected fine from being challenged by the center player is

$$\frac{\theta}{n}\frac{f}{(2q)^n}\left[\frac{(2q+\delta)^{n+1}}{n+1}+(e_j-2\delta+\delta_p-q)\frac{(2q+\delta)^n}{n}+(\delta_p-\delta)(e_j-q-\delta)\frac{(2q+\delta)^{n-1}}{n-1}\right].$$

The marginal fine from being challenged by the center player is

$$\frac{f}{(2q)^n}\left[(2q+\delta)^n+(e_j-2\delta+\delta_p-q)(2q+\delta)^{n-1}+(\delta_p-\delta)(e_j-q-\delta)(2q+\delta)^{n-2}\right].$$

When $-2q\le\delta_p\le 0$ and $\delta\le -2q$, drug dealer j commits so much less crime than other periphery drug dealers that j is never challenged. Then, the marginal expected fine from being challenged by the center player is

$$0.$$

When $\delta_p \leq -2q$, drug dealer j commits so much less crime than the center drug dealer that j is never liable for the fine regardless of other periphery players' strategy. Then, the marginal fine from being challenged by the center player is

$$0.$$

In Nash equilibrium, the center player maximizes her payoff with the expectation that all the peripheral players play e_p^* and all peripheral players maximized her payoff expecting that the center player plays e_c^* and her peer peripheral players play the same strategy as she plays (i.e., $e_j = e_p^*$), the marginal benefit from committing an extra unit of crime is zero, so $\frac{\partial \mathbb{E}\left[u_i(e_k^*, e_{-k}^*)\right]}{\partial e_k^*} = 0$ for all $k \in \{1, 2, ..., n\}$. The Nash equilibrium is the solution to the system of equations for both center and periphery players where the unique numerical solution is

$$\begin{cases} e_c^* \approx 2.40 \\ e_p^* \approx 6.37. \end{cases}$$

Peer-informed audit mechanism with complete network

The authority randomly audits one drug dealer and the drug dealer can challenge the remaining drug dealers. In the second stage, the weakly dominant strategy for the audited drug dealer is to challenge the drug dealer with the largest signal. If we assume that everyone plays the weakly dominated strategy in the second stage, then expected payoff of player i playing e_i against everyone else playing e_{-i} is

$$\begin{aligned} \mathbb{E}\left[u_i(e_i, e_{-i}^*)\right] &= \pi\left(e_i\right) \\ &\quad - \theta \cdot \underbrace{\frac{1}{n} \cdot \mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \max\left\{\hat{e_{-i}}\right\} \middle| \epsilon_i\right)\right]}_{\text{Expected fine from being audited}} \\ &\quad - \theta \cdot \underbrace{\frac{n-1}{n} \cdot \mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \max\left\{\hat{e_{-i}}\right\} \middle| \epsilon_i\right)\right]}_{\text{Expected fine from being challenged}} \\ &= \pi\left(e_i\right) - \theta \mathbb{E}\left[f\left(e_i + \epsilon_i\right) \cdot \mathbb{P}\left(\hat{e}_i > \max\left\{\hat{e_{-i}}\right\} \middle| \epsilon_i\right)\right]. \end{aligned}$$

In Nash equilibrium, drug dealer i must maximize her payoff against the crime of other drug dealers e^*. We denote the gap between i's crime and other drug dealers' crime $\delta := e_i - e^*$.

When $2q \geq \delta \geq 0$, drug dealer i commits no less crime than other drug dealers, but the noise in the signal still makes it possible for i to avoid the fine. The expected fine from being directly audited by the authority is

$$\begin{aligned}
&f\theta\left(\left(e_i - q - \delta\right)\frac{1}{n} + \frac{2q}{n+1}\right)\\
&-\frac{f\theta}{(2q)^n}\left(\left(e_i - q - \delta\right)\frac{1}{n}\delta^n + \frac{1}{n+1}\delta^{n+1}\right)\\
&+\frac{f\theta}{2q}\left(e_i\delta + \frac{1}{2}\left(2q\delta - \delta^2\right)\right).
\end{aligned}$$

The marginal expected fine from being audited or being challenged is

$$-\frac{f\theta}{(2q)^n}\left(\left(e_i - q - \delta\right)\delta^{n-1} + \delta^n\right) + \frac{f\theta}{2q}\left(e_i + q - \delta\right).$$

When $-2q \leq \delta \leq 0$, drug dealer i commits no more crime than other drug dealers, but the noise in the signal still makes it possible for i to pay the fine. Then, the expected fine from being challenged by another player is

$$\frac{f\theta}{(2q)^n}\left[\frac{1}{n+1}\left(2q+\delta\right)^{n+1} + \left(e_i^k - q - \delta\right)\frac{1}{n}\left(2q+\delta\right)^n\right].$$

The marginal expected fine from being audited or being challenged is

$$\frac{f\theta}{(2q)^n}\left(\left(2q+\delta\right)^n + \left(e_i - q - \delta\right)\left(2q+\delta\right)^{n-1}\right).$$

In Nash equilibrium, when each player expects that others employ the same strategy (i.e., $e_i = e^*$), the marginal benefit from committing an extra unit of crime is zero, so

$\frac{\partial \mathbb{E}\left[u_i\left(e_i^*, e_{-i}^*\right)\right]}{\partial e_i^*} = \pi'\left(e_i^*\right) - \frac{f\theta}{2q}\left(e_i^* + q\right) = 0$. The Nash equilibrium is the solution to

$$ab\left(e_i^*\right)^{b-1} - \frac{f\theta}{2q}\left(e_i^* + q\right) = 0,$$

where the unique numerical solution is

$$e^* \approx 1.16.$$

4.8.2. Instructions and quiz questions

Here, we provide the instructions for the treatment with complete network and one-shot interactions. The instructions for the other treatments differ in terms of the rule to determine the payer of the fine and the nature of repetition (one-shot v.s. repeated) and is available upon request.

Instructions

Introduction

Thanks for your participation! Please put away your cell phones and other distracting devices during the experiment. It is important that you make all your decisions on your own without any distraction. If you have a question at any time during the experiment, please raise your hand and an experimenter will come by to answer it.

In the experiment, you will be part of a group consisting of five people-- you and four others. The other people in your group are sitting in this room. Your identity and the identities of the other group members will remain anonymous throughout the experiment. You and your group members will make some decisions and you will be paid based on these decisions. After making the decisions, you will be asked to fill a short questionnaire, which concludes the experiment. Your earnings are denominated in Tokens, which will be exchanged at a rate of 4 Tokens = 1 dollar. The money you earn will be paid privately to you, in cash, at the end of the experiment. The amount of money you earn depends partially on the decisions that you make and, thus, you should read the instructions carefully. In order to improve the understanding of the instructions, you will take a computerized quiz and earn four Tokens for each correct answer you provide. The task will begin after the quiz.

What are the decisions you will make?

You are given an initial endowment of 60 Tokens and you will be asked to make decisions in two stages.

Stage one (select the production level)

Your decision (and those of others in the group) is to select an amount of output to produce and to enter it into the computer. The output you select must be between 0 and 40. Your decision will be entered on a screen like the one below.

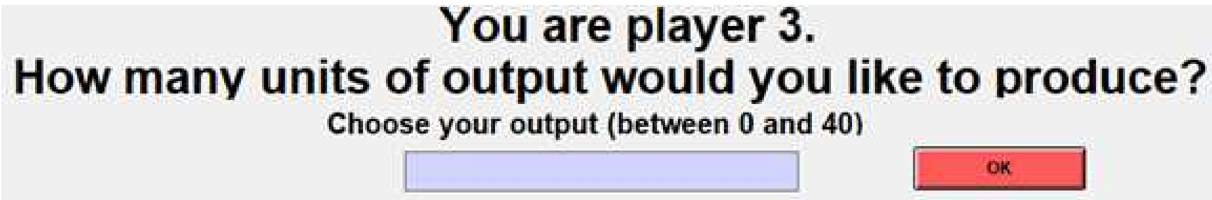

The output generates a profit for you. The relationship between the output and the profit is listed in the graph below. More output always generates more profit for you, but at a decreasing rate.

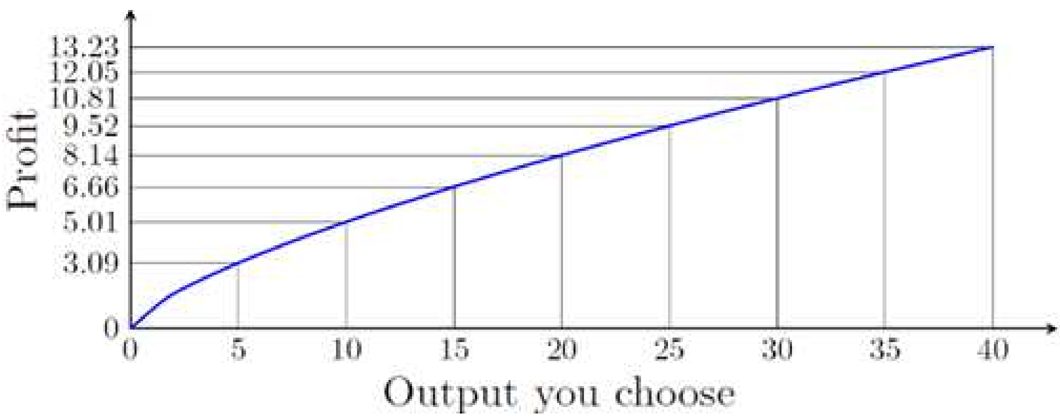

We also provide a calculator like the one below to help you calculate the profit. You need to type the output level in the blue box and then click "Calculate". The calculator will calculate the corresponding profit for you. For example, if you produce 10 units, then the profit is 5.01 Tokens.

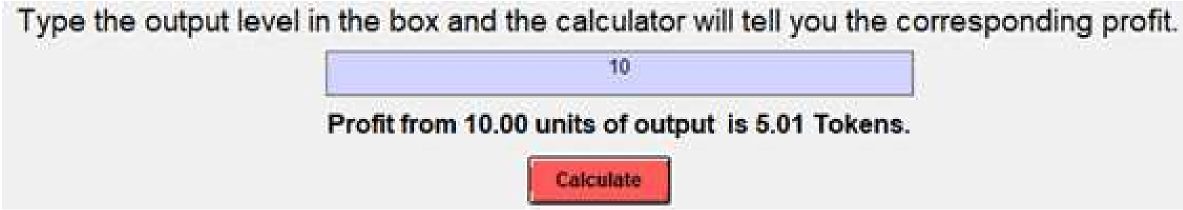

However, the more you produce the more likely that you need to pay a fine in stage two. You can consider the fine to be a penalty on your production.

Stage two (determine who must pay the fine)

In stage two, you (and the other participants in your group) see four signals informing you of the output levels of the other members of your group. Each signal is the sum of the true output level of that player plus a random number between [-10,10]. Each number between [-10,10] has the same chance to be chosen. For example, if player 4 produces 10 units in stage one and the random number chosen for player 4 is -5.5, then the signal for player 4 is 10+(-5.5) =4.5. Everyone except player 4 can see this number. Player 4 sees her output instead.

The computer then randomly audits 1 member in your group. Everyone is equally likely to be audited. If you are audited, then you will see the left panel of the picture on the next page. You can either pay the fine by selecting the top button or choose to challenge another player in your group to avoid paying the fine yourself by selecting one of the 4 buttons below the top button. If you click "I want to pay the fine", then you are liable for the fine. If you challenge another player, then whoever has the higher signal is liable for the fine. The fine amount is:

$$\text{Fine} = 1.2 \times (\text{the } \underline{\text{signal}} \text{ of the player liable for the fine}).$$

If you are not audited, then you see the right panel of the picture below and make no decision.

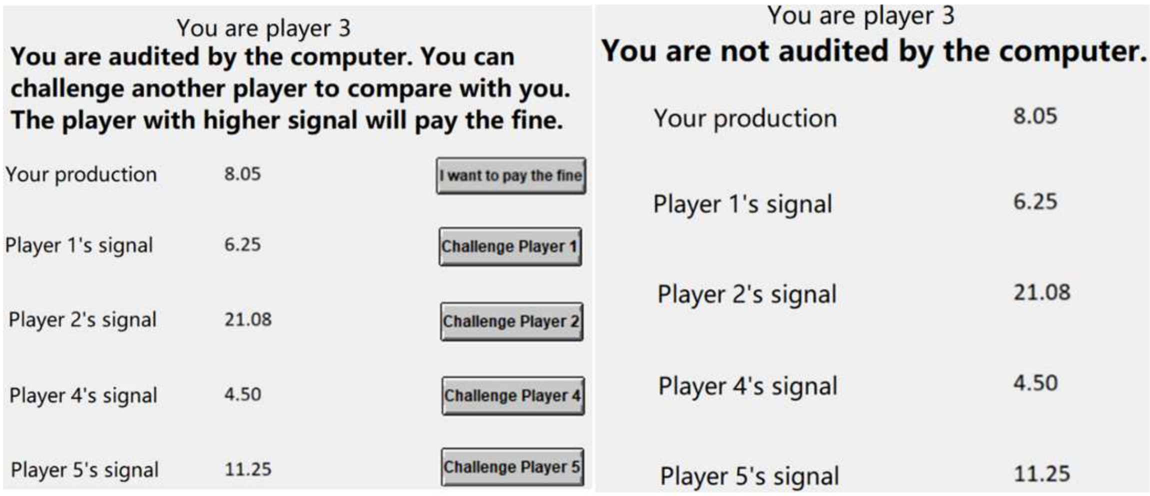

Payment

You will have an earning of:

Tokens earned from answering the quiz questions

+ Your endowment

+ Profit from production

- The fine you have to pay (if any)

Tokens will be converted to Australian dollars at an exchange rate of 4 Tokens=$1.

Quiz questions

Below, we list the quiz questions for the treatment with complete network (one-shot game). The quiz questions for the other treatments only differ in terms of the description of the rule of challenge and the number of repetitions and are available upon request. Participants are informed whether the answer is correct immediately after submitting the answer. If the participant provides a wrong answer, we display the correct answer and explain the reasoning behind it.

Q1: If you produce 22 units, how much profit (in Tokens) do you earn? You need to use the calculator below to solve this question.

Answer: 8.70.

Q2: Can you see your own signal in stage 2? Yes/No/Depends

Answer: No.

Q3: If you produced 5 units and the computer picks 5 as the random number for you, then what is your signal visible to the rest of your group?

Answer: 10.

Q4: If your production level is 40, what is the lowest signal that you can possibly get? A: 30/ B: 40/ C: 25/ D: 45

Answer: A.

Q5: Suppose you are player 5. You produced 8 units of output in stage 1 and were audited by the computer in stage 2. Suppose the signals of player 1 through 4 are -1, 5, 12, and 16, respectively. Which player should you challenge if you want to avoid paying the fine?

Answer: Player 4.

Q6: Suppose player 3 produces 0.5 units of output, player 2 produces 7.5 units; and player 3's signal is 5 and player 2's signal is -0.5. If player 3 is audited by the computer and challenges player 2, who should pay the fine?

Answer: Player 3.

Q7: Suppose player 1 is liable for the fine and her signal is 20. What is the amount she will pay as the fine (in Tokens)?

Answer: 24.

4.8.3. Repeated interaction treatments

The repeated game treatments inform our understanding of how the peer-informed audit mechanism influences criminal activity under repeated interaction among players, a feature that is relevant in numerous practical settings that motivate our study. For example, a group of criminals (in a mafia) may commit a series of crimes together over years. In order to capture this realistic feature and test the robustness of the peer-informed audit mechanism, we invite another 200 participants to play the repeated version of the production game. In the repeated game, the treatments and the rule of the game remains the same as the one-shot game, except that the same group of participants play the game 20 times and one randomly chosen round is selected for payment. In each of repeated interaction treatments, we conducted 10 groups with 5 subjects in each group. As each group is regarded as an independent observation in these treatments, we have a total of 10 independent observations for each treatment.

Since the stage game Nash equilibrium is unique and the finite number of rounds (20) is announced to the subjects, the repeated play of the game is likely not to change the predictions of the model. Our main goal with the repeated game treatments is to examine the robustness of the treatment effects rather than test the distinctive features of the repeated game theory.

Results for the repeated game

Figure 4.6 reports the crime level over 20 rounds for the baseline and the network treatments combined. Our mechanism successfully reduces crime even when the game is repeated for 20 rounds. The level reduced from 26.66 units in the baseline to 18.73 units in the network treatments – a 30% drop ($p < 0.001$). Figure 4.7 provides a comprehensive view on the evolution of crime levels in all treatments. The new mechanism significantly decreases the crime level under all networks we consider. Compared to the baseline, where the average crime is 26.66 units, the mechanism reduces the levels to 17.71 ($p = 0.007$), 20.34 ($p = 0.013$), and 18.14 ($p = 0.003$) units under star, circular, and complete networks, respectively. The crime level of the center player is less than the periphery players in

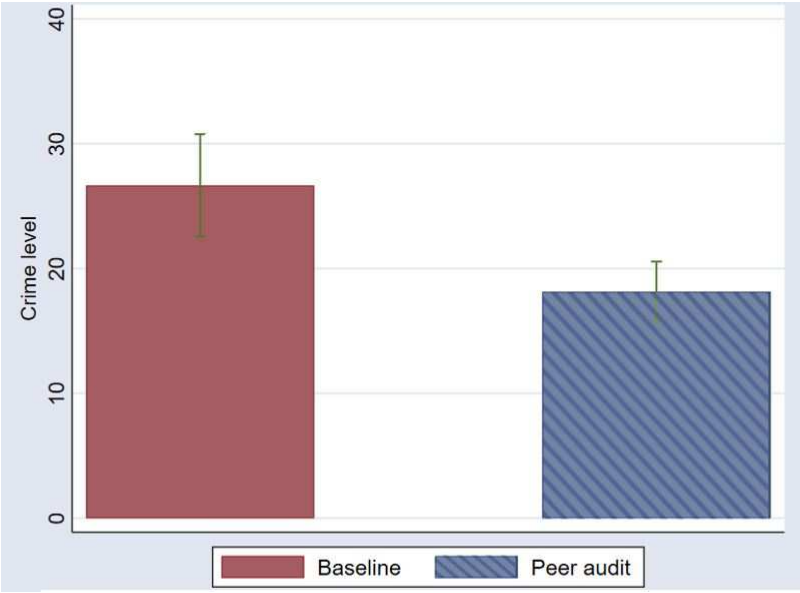

Notes: This figure plots the average crime levels (over 20 rounds) for the baseline and network treatments combined. The error bars indicate 95% confidence intervals.

Figure 4.6: Overview for the repeated games

the star network, which is qualitatively in line with the theory for the one-shot game, despite being marginally insignificant at the 10% level (p = 0.15). The average crime of the center player is 14.8 units, which is not significantly different from 18.0 units of the periphery players (p > 0.1). However, similar to the results described in the previous section, there are no treatment differences across the different network structures (p > 0.1 in all of the three pairwise comparisons). The crime levels in different peer-informed audit treatments are also significantly larger than theoretically predicted in Section 2 (p<0.001 in all three cases).

In order to examine the robustness of the results, we perform a regression analysis similar to that of the one-shot treatments. In addition to the independent variables discussed in the previous section, we also include a variable, referred to as round, to control for time effects in the model:

$$\begin{aligned}\text{Crime}_{i,t} = c &+ \alpha_1 \mathbb{I}_i^{\{\text{treatment=star center}\}} + \alpha_2 \mathbb{I}_i^{\{\text{treatment=star periphery}\}} + \alpha_3 \mathbb{I}_i^{\{\text{treatment=circular}\}} \\ &+ \alpha_4 \mathbb{I}_i^{\{\text{treatment=complete}\}} + \beta_1 \mathbb{I}_i^{\{\text{gender=male}\}} + \beta_2 \text{Understand}_i \\ &+ \beta_3 \text{GPA-RA}_i + \beta_4 \text{GPA-PIA}_i + \sum_{t=2}^{20} \gamma_t \mathbb{I}_{i,t}^{\{\text{round=t}\}} + \epsilon_{i,t}.\end{aligned}$$

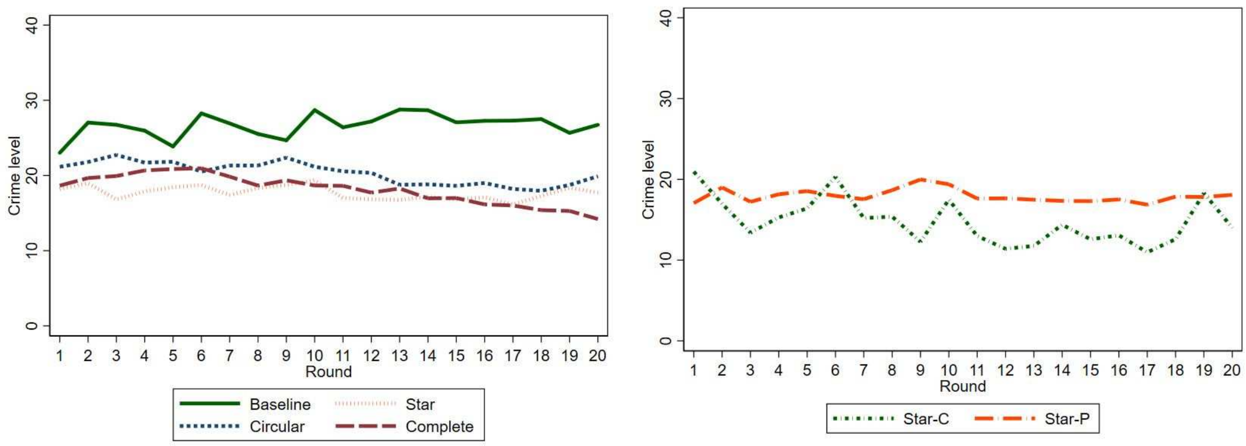

Notes: The left panel of the figure plots the average crime levels in each of the 20 rounds for the all four treatments. The right panel depicts the average crime level in each of the 20 rounds for the center and periphery players in the star treatment.

Figure 4.7: Crime levels in repeated game treatments

Table 4.3 reports the estimation results. The results for repeated games are robust to all the control variables we consider. The peer-informed audit mechanism successfully deters crime under all network structures we consider and the magnitude is both economically and statistically meaningful. There is also a significant difference ($p<0.1$ in all model specifications) between the center and periphery players in the star network in all four specification we consider. However, there are no significant differences between treatments across the star, circular, or complete networks (more specifically, there is no significant difference in any of the three pairwise comparisons of $\alpha_s := \left(\frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2\right)$, α_3 and α_4).[25] Overall, these set of results are similar to those based on one-shot interaction treatments reported in Section 4 of the paper.

4.8.4. Demographic variables

We report the demographic variables across treatment in Table 4.4. In general, the demographic variables are balanced across treatments from F-tests.

[25] We also performed a similar regression where we include a time trend for each treatment. We find that the crime levels move towards the corresponding Nash predictions, but the effect is weak. Even in the last round, crime levels remain significantly different from the corresponding prediction ($p < 0.01$). Importantly, the main findings are robust to controlling for the treatment specific time trends.

	Gender (male=1)	Asian (Asian=1)	GPA>70% (GPA>70%=1)
Baseline (one-shot)	0.4	0.76	0.68
Star (one-shot)	0.48	0.92	0.76
Circular (one-shot)	0.56	0.8	0.56
Complete (one-shot)	0.4	0.8	0.84
Baseline (repeated)	0.66	0.78	0.72
Star (repeated)	0.54	0.78	0.66
Circular (repeated)	0.48	0.86	0.72
Complete (repeated)	0.54	0.84	0.72
F_stat	1.08	0.58	0.83
p_value	0.37	0.77	0.57

Notes: This table presents the average demographic variables for each treatment. We also report the F-statistic and p-value testing the null hypothesis that there is no difference in the variable across all treatments.

Table 4.4: Balance check

	Model 1	Model 2	Model 3	Model 4
Star center (α_1)	-11.87***	-11.23***	-12.02***	-12.18***
	(2.47)	(2.49)	(2.58)	(2.60)
Star periphery(α_2)	-8.23***	-8.02***	-8.03***	-7.90***
	(2.54)	(2.56)	(2.58)	(2.59)
Circular (α_3)	-6.32***	-5.87***	-6.01***	-5.96**
	(2.17)	(2.27)	(2.33)	(2.35)
Complete (α_4)	-8.52***	-8.22***	-8.44***	-8.43***
	(2.02)	(2.11)	(2.14)	(2.17)
Gender (β_1)		2.46**	2.47**	2.58**
		(1.10)	(1.10)	(1.15)
Understand (β_2)			0.84	0.85
			(0.53)	(0.52)
GPA-RA (β_3)				-0.06
				(1.18)
GPA-PIA (β_4)				-0.97
				(0.73)
Constant (c)	25.59***	21.50***	19.84***	19.64***
	(1.84)	(2.70)	(2.80)	(2.82)
Control for time effects ($\sum_{t=2}^{20} \gamma_t \mathbb{I}_{i,t}^{\{\text{round=t}\}}$)	YES	YES	YES	YES
$H_0:\ \frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2 = \alpha_3$	p=0.19	p=0.17	p=0.16	p=0.16
$H_0:\ \frac{1}{4}\alpha_1 + \frac{3}{4}\alpha_2 = \alpha_4$	p=0.76	p=0.75	p=0.76	p=0.78
$H_0:\ \alpha_3 = \alpha_4$	p=0.19	p=0.17	p=0.15	p=0.15
$H_0:\ \alpha_1 = \alpha_2$	p=0.02	p=0.07	p=0.04	p=0.02

Notes: This table presents the estimation results for whether our mechanism reduces the crime level under different network structures. There are 4000 observations (200 participants×20 rounds) for each regression. We report the standard errors under each point estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All standard errors are clustered at the group level. As reported in the table, we also performed t-tests to study the treatment effect caused by different network structures and by different positions in the star network.

Table 4.3: Regressions for repeated game treatments