5-2015

Theory and Experiments Exploring Behavioral, Financial, and Public Economics

Matthew John McMahon
University of Tennessee - Knoxville, mmcmaho2@vols.utk.edu

To the Graduate Council:

I am submitting herewith a dissertation written by Matthew John McMahon entitled "Theory and Experiments Exploring Behavioral, Financial, and Public Economics." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Economics.

William S. Neilson, Major Professor

We have read this dissertation and recommend its acceptance:

Donald J. Bruce, Jacob S. LaRiviere, Paul R. Armsworth

Accepted for the Council:
Dixie L. Thompson

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

Theory and Experiments Exploring Behavioral, Financial, and Public Economics

A Dissertation Presented for the

Doctor of Philosophy

Degree

The University of Tennessee, Knoxville

Matthew John McMahon

May 2015

“And on the pedestal these words appear:
‘I am Ozymandias, king of kings:
look on my works, ye Mighty, and despair!’
Nothing beside remains.”
–*Percy Bysshe Shelley*

I dedicate this dissertation to my family, friends, and mentors who have supported me throughout. I couldn't have done it without you.

I would especially like to thank my loving parents, John and Susan, and my sister, Jenny. Your help and guidance was beyond critical to my success. Thank you for making me who I am today. I owe it all to you.

I would also like to thank my grandparents– Kathleen Henderson, the late Neil Henderson, the late Charles "Speedy" McMahon, and the late Evelyn McMahon. Thank you for all her help over the years. You've taught me countless lessons and provided unwavering support.

I also I owe an eternal debt to my friends. I've been blessed with many close friends who have always stood by me and given me support. You have listened to far more than your fair share of complaints from me over the years. I can say with honesty that you have become family to me, and I could not have done this without you. Thank you.

To all my teachers and mentors over the year, thank you for helping me. Thank you to Ronnie Powell for encouraging my love of math from an early age, to Dr. Don Luther for nurturing my love for history and teaching me to think critically about everything, and to Benita Albert for reinforcing my love of math, teaching me the value of hard work, and being the best calculus teacher who ever lived. Thank you to Dr. Eric Marland for taking my passion for math in so many directions throughout college and to Dr. Jean-Pierre Courbois for nurturing my passion for economics. To Dr. Jake LaRiviere, thank you for developing my economic intuition; I view both economics and the world as a whole differently because of it.

Last, and certainly not least, I would like to extend a very special thank you to Dr. Bill Neilson. This dissertation is a testament to all the guidance you've provided throughout graduate school. You've been a wonderful teacher, both in and out of the classroom; you've helped teach me how to look at every issue from an economic perspective; and you've shown me that hard work does pay off. You always believed in me, and I wouldn't be here today if it weren't for all your help along the way.

Acknowledgments

Support for this dissertation came in many forms, and I again would like to thank everyone who has helped me. My committee members have all provided an enormous amount of feedback on this project and have guided me the whole way. Their insights have been invaluable. I would also like to thank the rest of the faculty in the UT Economics Department for their help and guidance. The department as a whole has been very supportive.

I would also like to thank my fellow graduate students for reviewing my work and attending my presentations despite having their own dissertations to work on. I would especially like to thank Clay McManus, Justin Roush, Becky Davis, Michael Craig, Yin Chu, Jens Schubert, Nick Simmons, Jing Li, Jill Welch, Bong-Kyun Kim, Langchuang Peng, and Jie Deng for spending much of their valuable time helping me with my work. Your input has been just as valuable.

Thank you to Sherri Pinkston, Kathy Zelem, Susan McGee, and Donna Kemper for their help navigating the giant bureaucracy at UT. You've saved me countless hours and headaches– I would never have gotten anything done without your assistance. Thank you for all your help.

Thanks to the University of Tennessee's Office of Research for providing a grant for my experiment, and a special thanks again to Dr. Bill Neilson for funding the other experiment.

Thanks to participants at the 2012 Colorado University Environmental and Reserouce Economics Workshop, the 2012 Meeting of the Southern Economic Association, the 2013 Northeast Workshop on Energy Policy and Environmental Economics, and the 2014 Midwest Economic Theory Conference, as well as to those who attended my presentation to the Economics Department at Centre College in December 2014. Additional thanks to any and all who have read and reviewed drafts of any of this work. I appreciate your time and insights.

Abstract

I study three questions which relate to one another only in that each explores facets of economics. First, I theoretically examine the conditions under which introducing an impure public good decreases total public provision. I introduce a central planner who can tax the private good to correct this and identify the market characteristics that typify this scenario. Second, I test the two standard competing dividend puzzle hypotheses using a laboratory experiment. Evidence from the lab, including variables unobservable in the field, reinforces empirical work supporting the outcome model over the substitute. Last, I obscure from dictators information regarding recipients' income sources in a standard laboratory dictator game. When other-regarding dictators are unaware how much of recipients' income is earned relative to luck-borne, they assume recipients are unlucky rather than lazy. This supports the social insurance literature over that on moral wiggle-room.

Preface

It has been shown that the introduction of a good that has portions of both private and public characteristics (an impure public good) can actually decrease the total level of the public characteristic provided. Using the standard impure public good model, I first isolate the conditions under which this occurs in a general equilibrium. I then introduce a central planner whose goal is to counteract this decrease. She chooses a tax rate for the purely private good, and spends the tax revenue generated to increase the provision of the public characteristic. In choosing the optimal tax rate, she minimizes the tax's deadweight loss subject to the total public characteristic given the tax and the impure good being at least as large as it was in the absence of both. I then identify the properties both of markets and of impure public goods that tend to necessitate such a situation, thus also identifying those which necessarily increase the resulting deadweight loss. It is suggested that such properties are harmful to society, either through the decrease in the public characteristic absent a planner or through the resulting increase in deadweight loss if a planner is present.

For nearly half a century, the literature on dividend payments has questioned why firms pay dividends at all. Two competing models have emerged to explain this. In low investor protection regimes, corporate insiders give dividends to signal trust. In high-protection regimes, outside investors have the power to demand dividends. Using a unique laboratory experiment, our results back the empirical literature, favoring the complement (outcome) model over the substitute model. An increase in outside shareholders' rights increases dividend payments, reducing overall efficiency. The efficiency loss moving from low- to high-protection treatments can occur via multiple channels, and we find significant evidence of losses through each. One benefit of laboratory analysis is that some of these channels are not observable outside the lab (such as insider expropriations). Additionally, we find that high investor protection causes a greater efficiency loss than dividend taxes do.

There is anecdotal evidence that people often treat income earned by effort differently than that gained by luck, yet economists often overlook this distinction. Furthermore, others' income sources may often be (at least partially) obfuscated. I adapt a common inequality aversion model to allow for income distinction by source, both for own income and others', and for uncertainty over others' sources. I empirically test resulting hypotheses in a dictator game experiment wherein subjects gain income via both effort and luck. Dictators know recipients' income by source in control, but only total income in treatment. I find that partially informed dictators treated wealth as fully informed dictators did luck, but not as they did earnings– nor as conditional expectations of recipient sources. This is evidence of social insurance rather than moral wiggle-room, and it also refutes an expected value approach.

Table of Contents

List of Tables

List of Figures

Chapter 1

Optimal Taxation with Impure Public Goods

1.1 Introduction

Consumers take both private and public characteristics of goods and services into account when making purchasing decisions. There are large and still expanding markets for many goods that are considered fair trade, environmentally conscious, and ethically produced. Everything from shade-grown and fair trade coffee beans to organic and locally grown food products to the prominent rise of ecolabels on appliances and automobiles exemplifies not just that consumers are keeping these public characteristics in mind when deciding which goods to purchase, but they also demonstrate that firms and advertisers are accutely aware of this phenomenon. Of course, all of the goods mentioned above also have private benefits; people enjoy drinking coffee, appliances save us time doing necessary chores, and automobiles have become the staple for transportation in the developed world. Such goods that exhibit both public and private characteristics have come to be known as green goods, impure public goods, or simply impure goods.

The original model of the private provision of (purely) public goods, designed by Bergstrom et al. (1986), provided the framework for the first models of impure public goods, built and analyzed by Cornes and Sandler (1984).[1] Later, Kotchen (2005), Kotchen (2006), and Kotchen (2007) put forth a simplified model, which has now become the standard. The heart of the model allows for consumers to have preferences over private and public characteristics, rather than over goods themselves, in an environment where there are both a purely private good and a purely public good, as well as an impure good that provides some of both characteristics.

Kotchen (2006) highlights what is perhaps the most interesting result from this model. The introduction of an impure public good into the market can actually decrease the overall provision of the public characteristic, depending on the market parameters and consumers' preferences. This result is even possible assuming *all* consumers positively value both characteristics enough to purchase some of each. In fact, the model even allows for the possibility that the introduction of the impure good can actually make *every* individual worse off than before.

I add a central planner to the model to correct for this reduction. The central planner taxes the purely private good and in turn uses that revenue to provide more of the public characteristic. The planner's goal is to choose the smallest possible tax rate that boosts the overall public characteristic back to where it was. This minimizes the deadweight loss required to achieve the original level of the public characteristic. I then identify the market and impure good properties that typify situations when such intervention is required.

It may not immediately be obvious how the introduction of a green good can actually decrease the total public characteristic present in the economy. Such a counterintuitive effect implies that the introduction of a product meant to improve air quality, for example, could instead actually make it worse. Consider a city in which residents can either drive their car or walk to get around town. Residents have a

[1] See also Steinberg (1987), McClelland (1989), Cornes and Sandler (1994), and Sandler and Cornes (1996).

preference for rapid transit, which is a private characteristic, and for air quality, which is a public characteristic. Their taste for rapid transit incentives driving, the purely private good, at the cost of pollution. On the other hand, their taste for air quality incentivizes walking, the purely public good, at the cost of transit time. Naturally, residents of this city vary in their consumption bundles of walking and driving, and thus also in their contributions to overall air quality and (their own level of) rapid transit.

Now suppose the city decides to introduce a mass transit bus system with the goal of increasing air quality.[2] The buses are faster than walking but slower than driving, although they provide less air quality than walking but more than driving. The introduction of the bus system affects the implicit prices of both rapid transportation and air quality.

The city's intended effect is for people to substitute away from driving and into riding the bus. For these individuals, there is a tradeoff that is easily illustrated using a classic Slutsky decomposition. Each mile travelled now effectively costs less, in terms of air quality sacrificed, which is the substitution effect. This increases air quality on the internal margin. However, because each mile effectively costs less, demand for travel increases, which is the income (or rebound) effect. This decreases air quality on the external margin. The net effect here is ambiguous.

To further complicate matters, because the bus is faster than walking, others will substitute away from walking and into riding the bus. For these individuals, both effects harm air quality. Add on top that each individual's own substitution patterns has spill-overs in air quality that affect everyone else, thereby further affecting everyone else's choices, and things quickly get complicated. Perhaps the one thing that is clear, however, is that there are multiple effects working in both directions, and so the net impact on air quality is ambiguous unless more market parameters are known.

[2] One could also imagine a number of other impure goods as examples here, such as hybrid cars, electric cars, light rail, or even carpool incentives.

Munro and Valente (2009) take this scenario the lab to examine it. They find that introducing an impure good that is a better substitute for the purely public good has little effect on total public contributions. Introducing one that favors the private good, on the other hand, has a significant negative impact on public contributions. Beyond this, little empirical work has been done on this issue.[3]

Additionally, while this potential problem has been recognized, little work has been done to examine possible solutions. It is well known that the general model of the private provision for public goods yields an inefficient underprovision of the public good. While Samuelson (1954) shows that taxing the private good in order to finance the public good yields inefficient results, such a simple tax-redistribution scheme has not been theoretically examined in an impure public good scenario. As Altemeyer-Bartscher et al. (2011) note, one of the key differences between pure and impure public goods is that the private aspect allows for a stronger alignment of incentives, and thus a simple tax-redistribution policy may yield better results in a market with an impure good.[4] Herein, I examine a similar question.

This paper takes a general-equilibrium approach to the standard impure public goods model and allows for a tax and redistribution policy that can correct for a negative impact that results from the introduction of an impure public good into the market.[5] The model operates like a two-stage game. First, the central planner chooses a tax rate for the purely private good, from which the revenue generated is then redistributed in the form of the purely public good. The planner optimizes this tax

[3]Similar complicated implicit price interactions do show up in other places, such as work by Bento et al. (2014) showing that the "free" policy of allowing hybrid cars in California's HOV lanes actually cost both carpoolers enough (in terms of increased travel time) to render the policy far less efficient that many other simpler alternatives.

[4]Even if such a policy yields a suboptimal result, however, it may still suffice as a second-best policy, due largely to the power of the price effects generated by a tax, the magnitude of which Bergstrom and Andreoni (1996) point out.

[5]It should be noted that in the original model introduced by Kotchen (2005) and in most subsequent extensions thus far, a general equilibrium analysis is left out. Most work tends to ignore the change in the amount of the public characteristic that the individual receives as a spill-in that results from the change in a static parameter. Such an analysis instead looks at a change in that individual's best-response function as a reaction to all others' choices. Bergstrom et al. (1986) do find the general equilibrium solution and statics for the standard two-good case.

rate by minimizing deadweight loss (DWL) in the economy, subject to the contraint that the total amount of the public characteristic in the economy given the presence of both the tax and the impure good is at least as great as in the absence of both.[6] In the second stage, consumers optimize their consumption bundles, accounting for the tax, the government-provided public good spill-in, and all other agents' public spill-in.

Given the market equilibria in the differenct scenarios, I look at how both the likelihood of and the magnitude of the gap created by the introduction of the impure good change with respect to the properties of the impure good and of the market as a whole. I also examine the properties of the impure good and the market that affect the size of the optimal tax rate, and thus also affect the amount of deadweight loss to society. The number of individuals who purchase both the private and the impure good negatively affects the optimal tax rate, while the number of individuals who purchase solely the impure good positively impacts it. The technology parameters, which capure the portion of the impure good entering either the public or the private characteristic, both negatively affect the optimal tax rate. These intuitive results are in line with those found regarding changes in the gap itself.

These results imply that policy makers should be wary of markets with impure public goods. Markets with large substitution effects away from the public good or small substitution effects away from the private good are both particularly at risk. These substitution effects manifest themselves based on the external margin, the internal margin, and the impure good's technology. That is, they are based on the number of people who substitute, the intensity of such individuals' substitution (as dictated by preference parameters), and the extent to which the impure good is a substitute for the purely public and the purely private goods. Both markets with immensely popular impure goods and those with unpopular private goods will require a larger tax in order to correct for the impure public good's effect. Additionally,

[6] Minimizing deadweight loss (excess burden) subject to raising a given minimum level of revenue has been the standard optimization function for central planners dating back to seminal work by Ramsey (1927), Diamond and Mirrlees (1971a), and Diamond and Mirrlees (1971b). For further discussion on the development of this within the optimal taxation literature, see Auerbach (1982).

markets where the impure good is a poor substitute for either the purely public or the purely private good will require a larger tax.

Of course, in actuality a tax on the purely private good is not the only possible policy instrument. Within the standard public good scenario, other solutions to provide the socially efficient level of the public good have been examined; for example, Kolmar and Wagener (2012) and Giebe and Schweinzer (2013) examine a lottery where paying taxes increases the probability of winning, and Kotchen (2013) suggests pursuing policies that turn purely public goods into impure goods by "creating" private characteristics. Additionally, more complicated tax schemes that better align incentives, such as those proposed by Bergstrom and Andreoni (1996) and Falkinger (1996), have been shown to induce the socially optimal level of the public good theoretically.[7] Experimentally, Falkinger et al. (2000) show that theoretically predicted results for one such tax scheme are a good prediction of observed behavior.[8] While this result lends hope that such policies may actually induce optimal public good provision on a larger scale, the administrative costs of instituting such a complex system of taxes are likely quite cumbersome.

In the specific context of climate stabilization as an impure public good, as is argued by Pearce (2000), there has been some work examining the role of taxes. Löschel and Rübbelke (2009) and Markandya and Rübbelke (2012) look at the potential for country-to-country transfer payments, although no direct tax on goods (or anything else) is applied. Altemeyer-Bartscher et al. (2011) and Altemeyer-Bartscher et al. (2014) build a two-nation model featuring only private goods that cause a negative externality and the public good of depleting that externality. Each government can independently set a tax rate on the private good and use the revenue generated to abate the negative externality. Unsurprisingly, these results are suboptimal compared

[7]For example, Falkinger et al. (1996), Nordhaus (2006), and Altemeyer-Bartscher et al. (2010) have looked at various alternative tax schemes in the context of climate stabilization.

[8]Although, results from an experiment by Koppel and Schulze (2013) indicate that the public characteristic of the impure good matters less than the pricing mechanism itself, at least in the context of a premium markup on the impure good.

to the reference scenario of a global central planner. In the context of this negative externality, the optimal tax rate is higher when abatement is realized to be an impure public good rather than a purely public good.

While these models are useful within the given context of climate stabilization, they have little external validity. It is not feasible to back out a more general context of impure goods that allow for positive or negative public characteristics and more open-form spillovers from these models. Thus, many important impure goods scenarios are left unanalyzed. This paper plays an important role by allowing for a broad-based analysis of a simple and feasible solution to the problem that is potentially created by the introduction of an impure public good in a very general context.

The remainder of the paper is organized as follows: Section 1.2 introduces the new variant of the standard model for impure public goods, featuring a central planner, and solves for the equilibrium level of the public characteristic. Section 1.3 uses comparative statics surrounding that equilibrium to prove related propositions. Those results are also discussed in this section. Section 1.4 concludes.

1.2 Basic Model

There are three possible scenarios in the economy. In the baseline scenario ("BL"), there are only two goods, a purely private good c and a purely public good d. In the Kotchen scenario ("K"), an impure public good g (for "green" good) has been exogenously introduced into the market. In the tax/central planner scenario ("τ"), there is now also a central planner who can tax the private good and use that revenue to boost the provision of the public good.

In all three scenarios, there are only two characteristics: the private characteristic X and the public characteristic Y. Agents, indexed by $i \in N$, get utility over the two characteristics, but face a budget constraint over the three goods. Using the language of Kotchen (2006), there are exogenous technology constraints that map between goods and characteristics. These constraints are scenario-specific.

In the BL scenario, there are only the two standard goods, so they map into the two characteristics straightforwardly. That is, $X_i = c_i$ and $Y_i = d_i$. Note that Y_i (d_i) only represents agent i's contribution to the overall public characteristic (good), which is denoted Y (d). Additionally, let Y_{-i} represent all other agents' combined contribution.

In the K scenario, there is now also an impure good g. Let there be some portion α of the impure good that goes to the private characteristic, and some portion β that goes to the public characteristic. That is, $X_i = c_i + \alpha g_i$ and $Y_i = d_i + \beta g_i$.

In the τ scenario, there is also central planner who can introduce a tax on the purely private good, denoted τ_c. The amount of tax revenue that agent i generates is then $\tau_c c_i$. Thus, $X_i = c_i + \alpha g_i$ and $Y_i = \tau_c c_i + d_i + \beta g_i$.

As mentioned, agent i's budget constraint is in terms of the three goods. For simplicity, assume that all *goods'* prices are unity before the tax, as Kotchen (2006) does. The exception here, of course, is that in the tax scenario, the private good c will tax a price $p_c = 1 + \tau_c$. Thus, the budget constraint can be written succinctly for all scenarios as

$$p_c c_i + d_i + g_i = w_i. \tag{1.1}$$

Also following Kotchen (2006), let $w_i = w_j \; \forall \; i, j$.

Let agent i's utility function be defined over the two characteristics. Note that agent i gets utility solely from her own private characteristic X_i, but from everyone's contribution to the public characteristic Y. That is, $U_i = U_i(X_i, Y)$. As will become clear later, an additional restriction is required here in order to isolate the exact conditions under which the introduction of the impure good causes the overall provision of the public characteristic to drop. Namely, an assumption must be made regarding the form of the utility function. Kotchen (2006) uses CES utility functions for the numerical examples that demonstrate and identify necessary but non-sufficient conditions for the decrease in the public characteristic to occur (Hicksian gross substitutability is necessary but not sufficient). In a similar vein, I have adopted a Cobb-Douglas utility function.[9] Agent i's utility is thus given by

$$U_i(X_i, Y) = X_i^{\theta_i} Y^{1-\theta_i}. \tag{1.2}$$

Recall that the consumers are effectively the second movers (in the presence of a central planner). In order to find the optimal tax rate in that scenario, then, backwards induction is required. Hence, finding the equilibrium demand functions for consumers is the first step to finding the general equilibrium for each scenario (and the only step for scenarios BL and K). After doing this, I will then turn back to the central planner's optimization problem.

Agent i's optimization problem is thus to maximize Equation (1.2) subject to the constraint of Equation (1.1). In order to solve this problem, it is necessary to either translate the utility function into goods space or translate the budget constraint into characteristics space. As Kotchen (2006) points out, it is far easier to visualize things in standard two-dimensional space, and so I will follow his lead and translate everything into characteristics space. This mapping effectively breaks apart Equation (1.1) into two separate constraints, creating a (potentially) kinked budget frontier. Additionally, rather than having agent i choose Y_i, it is simpler to visualize when letting her choose Y subject to another additional constraint: $Y \geq Y_{-i}$. Thus, agent i faces the following optimization problem, shown in its most general form:

$$\max_{X_i, Y} \; U_i(X_i, Y) \tag{1.3}$$

$$\text{subject to } \frac{1-\beta}{\alpha} X_i + Y \leq w_i + Y_{-i}, \tag{1.4a}$$

$$\frac{(p_c - 1) - \beta p_c}{(p_c - 1)\alpha - \beta} X_i + \frac{\alpha p_c - 1}{(p_c - 1)\alpha - \beta} Y = w_i + \frac{\alpha p_c - 1}{(p_c - 1)\alpha - \beta} Y_{-i}, \tag{1.4b}$$

$$\text{and } Y \geq Y_{-i}. \tag{1.4c}$$

[9]The Cobb-Douglas utility function is simply the limit of the CES utility function as the exponent on each characteristic approaches 0. That is, where $\lim_{\rho \to 0} [\theta_i X_i^{\rho} + (1-\theta_i) Y^{\rho}]^{\frac{1}{\rho}}$.

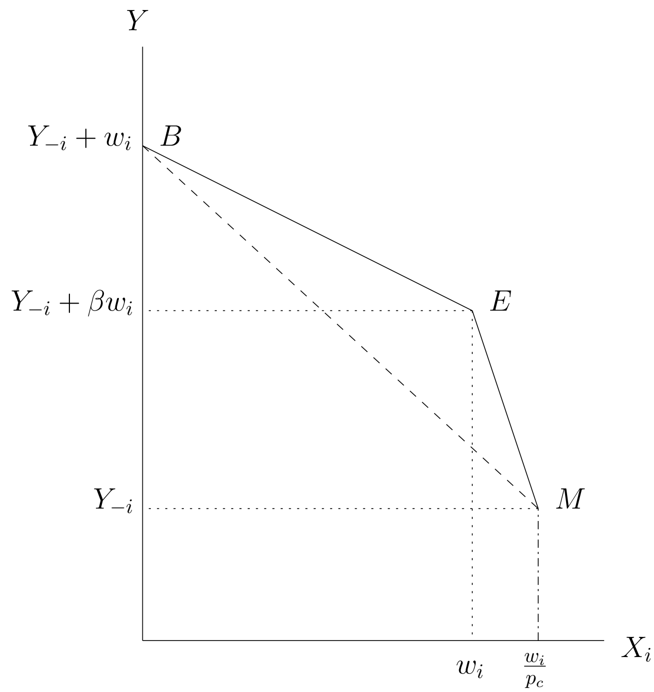

Figure 1.1: Budget frontier in characteristic space.

In Equations (1.4a) and (1.4b), the coefficients on X_i and Y can be seen as the implicit prices for characteristics X and Y, which differ across the two constraints. The reason for the notation will become more clear momentarily, but suffice for now to let the implicit characteristic prices in (1.4a) be called P_{X_d} and P_{Y_d}, and those in (1.4b) be called P_{X_c} and P_{Y_c}. Note that Equation (1.4b) reduces to $X_i + \frac{1-\alpha}{\beta} Y_i = w_i$ in scenarios BL and K, when $p_c = 1$. This is then identical to the optimization problem in Kotchen (2006).

As mentioned, by simplifying the problem into $X_i \times Y$ space, the budget frontier can be seen more clearly. Figure 1.1 illustrates this budget set. The diagram is easier to understand by thinking of what bundles of goods are purchased by individuals operating along different portions of the budget frontier.

Suppose agent i spends all her income, w_i, on the purely private good, d_i. That is, (c_i, d_i, g_i)=$(0, w_i, 0)$. Then she receives $X_i = 0$ and $Y = Y_{-i} + w_i$ (the public spill-in plus her own public contribution). This places her at point B.

Now suppose instead that i spends all her income on the impure good, g_i. That is, (c_i, d_i, g_i)=$(0, 0, w_i)$. She then receives $X_i = \alpha w_i$ and $Y = Y_{-i} + \beta w_i$. This places her at point E.

Now change the supposition once more so that she spends all her income on the private good, c_i. Predictably, she ends up at point M.

Last, suppose instead that she spends her income on a mixture of goods, say on both the purely public good and the impure good. She then falls on the budget frontier along line segment $\overline{BE}$. If she instead purchases a mixture of the private good and the impure good, then she falls along line segment $\overline{EM}$.

This discussion noticeably omitted any scenario in which agent i purchased some of both the private and the purely public goods. This omission is intentional. If the technology parameters α and β were such that any individual i would purchase both goods c_i and d_i simultaneously, then that necessarily implies that the impure good is completely dominated. That is, if $\alpha + \beta < 1$, then the point E is southwest of the dashed line segment $\overline{BM}$. If $\alpha + \beta = 1$, then point E lies along $\overline{BM}$. These represent scenarios in which the introduction of the impure good adds nothing, and thus the following assumption from Kotchen (2006) will be carried over here.

In order to focus on non-trivial cases, assume that (*i*) $0 < \alpha < 1$; (*ii*) $0 < \beta < 1$; and (*iii*) $\alpha + \beta \geq 1$. Let this jointly be know as "the technology assumption." Kotchen (2006) discusses the implications of this joint assumption in detail. Parts (*i*) and (*ii*) ensure that the impure good does not strictly dominate either pure good, and part (*iii*) allows for the impure good to weakly dominate any convex combination of the two pure goods. Kotchen (2006) and Kotchen (2007) show that if part (*iii*) holds strictly, then the impure good strictly dominates any convex combination of pure goods and there thus necessarily exists a unique Nash equilibrium. If $\alpha + \beta = 1$, then the budget set in characteristic space $(X_i \times Y)$ looks identical to that in the two-good scenario. While there are then an infinite number of unique three-good equilibria if the market has always been a three-good market, Kotchen (2006) and Kotchen (2007)

Table 1.1: Group Notation, by Consumption Bundle

In Figure 1.1	(c_i, d_i, g_i)	Notation	Size
Point B	$(0, +, 0)$	DD	n_{dd}
Along $\overline{BE}$	$(0, +, +)$	D	n_d
Point E	$(0, 0, +)$	G	n_g
Along $\overline{EM}$	$(+, 0, +)$	C	n_c
Point M	$(+, 0, 0)$	CC	n_{cc}

show that since a unique equilibrium first exists in a two-good market,[10] then the addition of the impure good will lead to a unique three-good Nash equilibrium.

In this paper, the budget frontier featuring $\overline{BM}$ represents baseline scenario, BL; the frontier featuring $\overline{BE}$ and $\overline{EM}$ when $p_c = 1$ represents the impure good/Kotchen scenario, K; and the frontier featuring $\overline{BE}$ and $\overline{EM}$ when $p_c = 1 + \tau_c$ represents the tax/central planner scenario, τ.

Given these natural divides, the set N of all consumers can be partitioned into 5 mutually exclusive, exhaustive subsets defined by their consumption bundles. Those consuming at point B, along $\overline{BE}$, at point E, along $\overline{EM}$, and at point M, are thus defined as belonging to the sets DD, D, G, C, and CC, respectively. These, along with corresponding consumption bundles and notation for the set sizes, are shown in Table 1.1. Note that the two sets that represent exterior corner solutions will be ignored, as discussed later. This simplifies the list to only sets D, G, and C. Table 1.1 lists all five groups solely for completeness's sake.

Conceptually, these sets need not be rigid to exogenous changes in the market. A shock that affects the shape of the budget frontier may, depending on consumers' preferences, cause some individuals to switch from one set to another. However, in order to make the model more tractable, a few restrictive assumptions are required.

First, assume that there is no individual i who, under any set of (α, β, w_i) parameters, contributes to only one characteristic. That is, $\forall i \in N$, $X_i > 0$ and $Y_i > 0$. This is a fairly standard assumption that simply eliminates (external) corner

[10] See Bergstrom et al. (1986) for proof.

solutions in characteristic space. Second, assume that any individual i in a given set k, for $k \in \{D, G, C\}$, always remains in that given set, regardless of changes in parameter values. In this way, the sizes of each set are fixed. Last, assume that there exists a single representative agent for each set. That is, agent i_c represents each agent $i \in C$, and so on. Let that notation also extend to goods and characteristics, such that g_{i_g} represents the amount of good g purchased by individual i, who is in set G, for example. (Of course, by the definition of being in set G, agent i_g is necessarily spending all of her income on good g, and thus $g_{i_g} = w_i$ for her.)

At this point, the agents' equilibrium demand functions can be found. Solving each representative agent's optimization problem yields their best-response demand functions for both X_i and Y, each as a response to the other agents'. This yields a system of equations that can be solved in the standard fashion, similarly to how one solves a problem simultaneously moving Cournot firms. In all, this yields 6 equilibrium demand functions (2 characteristics each for 3 representative agents) in each of the three scenarios.[11]

Finally, the total equilibrium level of the public characteristic can be found for each scenario– denoted $Y^{*^{BL}}$, $Y^{*^{K}}$, and $Y^{*^{\tau}}$–using the equation $Y^* = n_d Y^*_{i_d} + n_g Y^*_{i_g} + n_c Y^*_{i_c}$. This yields, for any given scenario,

$$Y^* = w_i \left[(1-\theta_d)(1-\theta_c) \left(\frac{n_d}{P_{Y_D}} + \frac{n_g}{P_{Y_G}} + \frac{n_c}{P_{Y_C}} \right) \right] \Gamma \tag{1.5}$$

where, in order to simplify notation, $\Gamma > 0$ is defined as

$$\Gamma^{-1} \equiv (1-\theta_c+\theta_c n_c)(1-\theta_d+\theta_d n_d) - \theta_c \theta_d n_c n_d. \tag{1.6}$$

[11]Note that there are limits on what each representative's θ value can be, so that each representative agent stays within the proper portion of the budget frontier. In order to assure that $X_{i_c} > \alpha w_i$, it must hold that $\theta_c > \frac{\alpha}{(\alpha)+(1-\alpha)(n_d+n_g+n_c)}$, and in order to assure that $X_{i_d} < \alpha w_i$, it must hold that $\theta_d < \frac{1-\beta}{(1-\beta)-(\beta)(n_d+n_g+n_c)}$.

First consider the baseline scenario with only the two pure goods. Without the impure good, all characteristic prices are simply unity for all individuals. Thus,

$$Y^{*^{BL}} = w_i \left[(1-\theta_d)(1-\theta_c)(n_d+n_g+n_c)\right]\Gamma. \tag{1.7}$$

Now consider the standard impure public good scenario, as originally presented by Kotchen (2006). Here, implicit prices vary across groups. For group D, $P_{X_d} = \frac{1-\beta}{\alpha}$ and $P_{Y_d} = 1$; for group G, $P_{X_g} = \frac{1}{\alpha}$ and $P_{Y_g} = \frac{1}{\beta}$; and for group C, $P_{X_c} = 1$ and $P_{Y_c} = \frac{1-\alpha}{\beta}$. Thus,

$$Y^{*^{K}} = w_i \left[(1-\theta_d)(1-\theta_c)\left(n_d+\beta n_g+\frac{1-\alpha}{\beta}n_c\right)\right]\Gamma. \tag{1.8}$$

Let $Y^{*^{DIFF}}$ be the difference between Equations (1.8) and (1.7). That is, $Y^{*^{DIFF}}$ is the increase in the total public good provision that stems from the introduction of the impure good. The introduction of the impure good lowers the overall provision of the public characteristic if and only if $Y^{*^{DIFF}} < 0$. So, for

$$Y^{*^{DIFF}} = \left[(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g+(\frac{\beta}{1-\alpha}-1)n_c\right)w_i\right]\Gamma, \tag{1.9}$$

it holds that $Y^{*^{DIFF}} < 0$ if and only if it holds that

$$\Phi \equiv \left((\beta-1)n_g+(\frac{\beta}{1-\alpha}-1)n_c\right) < 0. \tag{1.10}$$

Notice that the (negative) coefficient for n_g is the difference in the inverse of the new price of Y for $i \in G$ and the inverse of the baseline price of Y for $i \in G$. Similarly, the (positive) coefficient for n_c is the difference in the inverse of the new price of Y for $i \in C$ and the inverse of the baseline price of Y for $i \in C$. In fact, the same analogy applies for group D, however that price is unity in both scenarios, and thus the coefficient nets out completely.

Turning to the third scenario, where the central planner is present, the resulting equilibrium level of the public characteristic is

$$Y^{*^{\tau}} = w_i \left[(1-\theta_d)(1-\theta_c) \left(n_d + \beta n_g + \left(\frac{\beta - \alpha\tau_c}{1-(\alpha+\tau_c)} \right) n_c \right) \right] \Gamma. \tag{1.11}$$

Recall the goal of the central planner as the first mover: institute the smallest possible tax on the private good that pushes the public characteristic up to be at least as large as it was in the baseline scenario.[12] More formally, the central planner faces the following optimization:

$$\min_{\tau_c} \; \tau_c \tag{1.12}$$

$$\text{subject to } Y^{*^{\tau}} \geq Y^{*^{BL}}. \tag{1.13}$$

Solving the planner's new problem and isolating the optimal tax τ_c^* yields

$$\tau_c^* = \left(\frac{1-\alpha}{\alpha} \right) \left(\frac{(1-\beta)n_g + \left(1 - \frac{\beta}{1-\alpha}\right) n_c}{(1-\beta)n_g} \right). \tag{1.14}$$

Notice that the numerator of the second term is simply $-\Phi$, which again determines the sign. If the condition given in Equation (1.10) holds, then not only does that imply that Equation (1.9) is negative, it also necessarily implies that Equation (1.14) is positive. That is, if the introduction of the impure good decreases the provision of the public characteristic, then the optimal tax required to correct this decrease is necessarily positive, as expected.[13]

[12] Note that tax and deadweight loss are directly correlated, and so minimizing DWL, which is the actual goal, is identical to minimizing the tax, which is mathematically simpler. For completeness, note that DWL can be calculated using $\text{DWL} = n_c \left[\int_1^{1+\tau_c} \left(\frac{1}{1-\alpha p_c} X_{i_c}^{h*}(p_c, \bar{\cdot}) - \frac{\alpha}{1-\alpha p_c} w_i \right) \partial p_c - \tau_c \left(\frac{1}{1-\alpha(1+\tau_c)} X_{i_c}^{h*}(1+\tau_c, \bar{\cdot}) - \frac{\alpha}{1-\alpha(1+\tau_c)} w_i \right) \right]$, where the superscript $(^h)$ denotes the Hicksian demand function, and $(\bar{\cdot})$ denotes that all parameters supressed by $(\cdot)$ are held constant.

[13] The unrealistic but theoretically valid converse is also true.

1.3 Results

It is potentially relevant to policy-makers what types of properties typify situations wherein the introduction of the impure good decreases the total provision of the public characteristic. By being able to identify these properties, they can more swiftly move to counteract the negative effects. Some such parameters have directional effects, while others have magnitudinal effects. Propositions 1 and 2 summarize the former, while Propositions 3, 4, and 5 summarize the latter. Proofs of each are provided in Appendix A.

Proposition 1. *An increase in the number of individuals who consume both the private good and the impure good causes an increase in the difference between the public characteristic provision given the impure good and that given its absence. On the other hand, an increase in the number of individuals who consume solely the impure good causes a decrease in that difference.*

Consider first individuals who consume both the private good and the impure good. The introduction of the impure good into the market decreases the implicit price of the public characteristic for these consumers. The price effect pushes these consumers to substitute away from the private good toward the impure good, which in turn increases their contribution to the public characteristic. The resulting wealth effect bolsters this increase. Given that these consumers contribute more (per person) to the public characteristic when the impure good is present than in the baseline, increasing the number of these consumers increases the difference between the two scenarios.

Now consider individuals who consume solely the impure good. For such individuals, the introduction of the impure good into the market causes a decrease in the implicit prices of both the private and the public characteristics. If Equation (1.10) holds, then it must also hold that the price effect from the private characteristic is enough to outweigh that from the public characteristic to the extent that there is a net decrease in these individuals' contribution to the public characteristic. In such a scenario, increasing the number of these individuals thus increases the magnitude of the drop in the provision of the public characteristic. If instead the opposite were to hold, then the opposite (yet analogous) scenario would take place.

Proposition 2. *An increase in the portion of the impure good that maps into either the private or the public characteristic causes an increase in the difference between the public characteristic provision given the impure good and that given its absence.*

An increase in either of the technology parameters (α or β) increases the efficiency of the impure good. That is, the higher the portion of the impure good that enters either the private or the public characteristic, the closer the impure good is to a perfect substitute for the purely private good or the purely public good, respectively. This is simply a price effect; either technological gain decreases the implicit price of the public (private) characteristic for those who consume both the impure and the private (public) good. In turn, these price effects push consumers toward purchasing more the impure good. Because it comes through the efficiency channel, however, this push does not come at the expense of a decrease in the private (public) characteristic for those who purchase both the impure and the private (public) goods. Hence, the net effect of an increase in either technology parameter is an increase in the overall provision of the public characteristic given the presence of the impure good, and thus an increase in the difference between that level of provision and the unchanged baseline level.

By contrast, changes in some parameters cause purely magnitudinal effects.

Proposition 3. *An increase in the number of individuals who consume both the public good and the impure good causes a decrease in the magnitude of the difference between the total provision of the public characteristic given the impure good and that given its absence.*

Consider individuals who consume both the public good and the impure good. For such individuals, the introduction of the impure good decreases the implicit price of the private characteristic while leaving the implicit price of the public characteristic unchanged. Hence, the price effect of the introduction induces substitution away from the public characteristic in favor of the private characteristic. This puts positive pressure on the gap between the public characteristic provision given the impure good and that given its absence. However, this decrease in the implicit public characteristic price also has a positive wealth effect. The wealth effect works against the price effect and increases these individuals' demand for the public characteristic. This puts negative pressure on the public provision gap. The relative sizes of the market parameters and of the impure good parameters determine which of these competing effects dominates. If the introduction of the impure good causes a net increase in the total public characteristic provided, then the price effect dominates for these individuals. Thus, increasing the number of these individuals pushes the gap downward (approaching 0 in limit). If instead the introduction causes a net decrease in public provision, then the wealth effect dominates for these individuals. Here, increasing the number of such individuals then pushes the gap upward (again approaching 0 in limit).

Proposition 4. *An increase in the (Cobb-Douglas) taste parameter for the private characteristic either for individuals who purchase both the purely public good and the impure good or for those who purchase both the purely private good and the impure good causes a decrease in the magnitude of the difference between the total provision of the public characteristic given the impure good and that given its absence.*

Consider the scenario where the introduction of the impure good causes a decrease in the overall provision of the public characteristic (*i.e.*, where Equation (1.10) holds). An increase in an individual's taste parameter for the private characteristic causes substitution away from the public characteristic in favor of the private characteristic. The size of this substitution is directly proportional to the level of the public good provided. Because the magnitude of the public characteristic is smaller given the presence of the impure good that given its absence, there is a smaller reduction in public provision given the presence of the impure good, and thus there is a reduction in the magnitude of the difference between the public provision in the two scenarios. If instead the introduction of the impure good were to cause an increase in the provision of the public characteristic, then the signs would all flip accordingly, and the corresponding positive gap would also shrink in absolute value (and here in real value, too) as the individual's taste for the private characteristic increased.

Proposition 5. *An (equal) increase in each individual's wealth causes an increase in the magnitude of the difference between the total provision of the public characteristic given the impure good and that given its absence.*

Consider the scenario where the introduction of the impure good causes a decrease in the overall provision of the public characteristic (*i.e.*, where Equation (1.10) holds). An equal increase in each individual's wealth proportionally magnifies each individual's equilibrium provision of both characteristics. If the provision of the public characteristic is larger in the baseline scenario than in the impure good scenario, then a proportional increase in each of those levels of public provision only serves to proportionally magnify that gap. Intuitively, this static examines a pure wealth effect; with no price effect, no substitution across characteristics takes place. If instead the public provision were larger given the impure good than given its absence, implying a positive gap, then an increase in wealth would again only serve to magnify the size of the gap.

While Propositions 1 – 5 detail the goods and market characteristics that typify when intervention is required, they say nothing of the size of the tax required to fix the issue. Propositions 6 and 7 look at how the optimal tax rate, and therefore also the resulting DWL, are impacted by such characteristics. These propositions offer insight to and serve as basic guidelines for policy-makers facing this scenario in a given market. Proof of these propositions is given in Appendix A.

Proposition 6. *An increase in the number of individuals who consume both the private good and the impure good causes a decrease in the optimal tax rate. On the other hand, an increase in the number of individuals who consume solely the impure good causes an increase in the optimal tax rate.*

Consider first individuals who consume both the private good and the impure good. The effects such consumers have on the optimal tax rate is two-fold, and both of these effects work in the same direction.

First, the introduction of the impure good allows these individuals to substitute away from private good consumption and toward public characteristic consumption while still increasing their own private characteristic consumption. This movement occurs because the introduction of the impure good lowers these consumers' implicit price of the public characteristic without changing their implicit private characteristic price. That is, the price effect of the impure good's introduction causes these individuals to increase contribution to the public characteristic. Hence, an increase in the number of these consumers increases the total public characteristic provided given the impure good, which lowers the gap in the public provision generated by the impure good, thus implying a lower tax rate need to close that gap.

Second, the tax revenue from the private good depends solely on these individuals. The tax serves both as a nudge to push these consumers toward the public characteristic (via a price effect) and as a means to force some of their wealth spent on the private good back into the economy as a public good. Hence, an increase the number of these individuals increases both the tax base and the effect of the nudge, which in turn implies a lower optimal tax rate.

Now consider individuals who consume solely the impure good. For such individuals, the introduction of the impure good into the market causes a decrease in the implicit price of both characteristics. Given that the introduction of the impure good decreases the total provision of the public characteristic (*i.e.*, given that Equation (1.10) holds), it then must also hold that the price effect from the private characteristic is enough to outweigh that from the public characteristic to the extent that there is a net decrease in these individuals' contribution to the public characteristic. Increasing the number of these individuals thus increases the magnitude of the drop in the provision public characteristic, which in turn necessitates a greater optimal tax rate to close that gap.

Changes in the technological properties that map from the impure good into the two characteristics also affect the optimal tax rate.

Proposition 7. *An increase in the portion of the impure good that maps into either the private or the public characteristic causes a decrease in the optimal tax rate.*

An increase in either of the technology parameters (α or β) increases the efficiency of the impure good. That is, the higher the portion of the impure good that enters either the private or the public characteristic, the closer the impure good is to a perfect substitute for the purely private or the purely public good, respectively. This is simply a price effect; either technological gain decreases the implicit price of the public (private) characteristic for those who consume both the impure and the private (public) good. In turn, these price effects push consumers toward purchasing more of the impure good. Because it comes through the efficiency channel, however, this push does not come at the expense of a decrease in the private (public) characteristic for those who purchase both the impure and the private (public) goods. Hence, the net effect of an increase in either technology parameter is an increase in overall provision of the public characteristic (given the presence of the impure good), and thus there is a decrease in the optimal tax rate.

The tradeoffs can be succintly understood by examining Figure 1.2. For any given ratio $\eta \equiv \frac{n_c}{n_g}$, there is a bounded region for which $\tau_c^* > 0$ and the technology assumption both hold. The lower (southwest) boundary is always the red line $\alpha + \beta = 1$, as per the technology assumption. The upper (northeast) boundary is given implicitly; $\tau_c^* > 0$ holds if and only if $\frac{-(1-\alpha)(1-\beta)}{1-\alpha-\beta} > \frac{n_c}{n_g}$ is true. That is, the optimal tax rate τ_c^* is positive to the southwest of this boundary, and negative to the northeast of it (recall that $\frac{\partial \tau_c^*}{\partial \alpha} < 0$ and $\frac{\partial \tau_c^*}{\partial \beta} < 0$). In Figure 1.2, this upper boundary is shown for different levels of the ratio $\frac{n_c}{n_g}$. The pure blue line is the upper boundary when $n_c = n_g$. An increase in the numerator n_c shifts the upper boundary inward toward the origin, as indicated by the increasingly red lines. This movement illustrates the static $\frac{\partial \tau_c^*}{\partial n_c} < 0$. An increase in the denominator n_g pushes the upper boundary outward

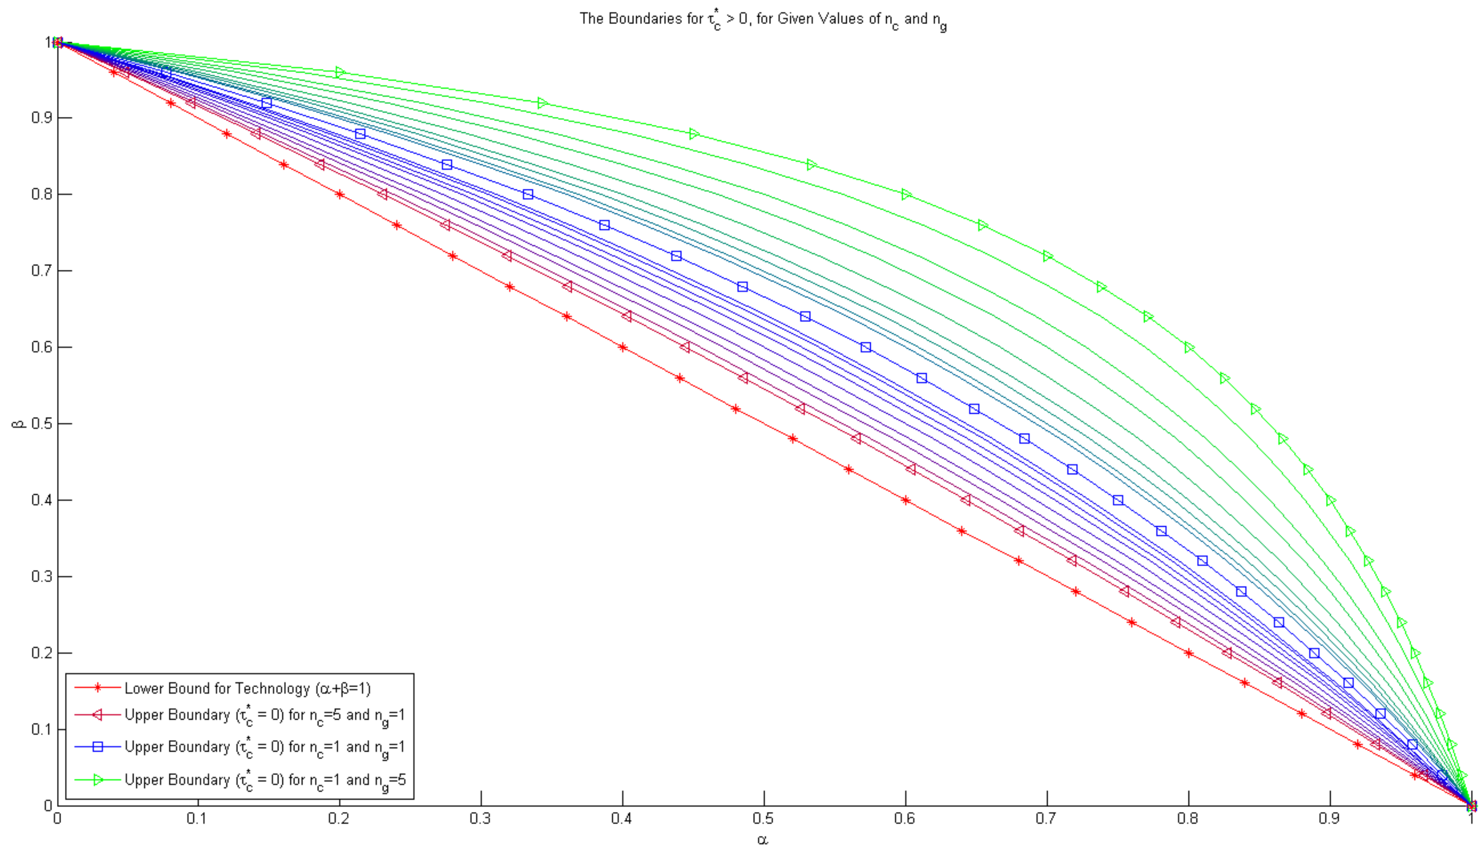

Figure 1.2: For any given ratio $\eta \equiv \frac{n_c}{n_g}$, there is a bounded region for which $\tau_c^* > 0$ and the technology assumption both hold. The lower (southeast) bound is always the red line $\alpha + \beta = 1$, given by the technology assumption. The upper (northeast) boundary represents the outer limit of where $\tau_c^* > 0$ for a given ratio η. The pure blue line is the upper boundary when $n_c = n_g$. As the numerator n_c increases, the boundary shifts inward toward the origin, as shown with increasingly red lines. As the denominator n_g increases, the boundary shifts outward away from the origin, as shown with increasingly green lines.

away from the origin, as indicated by the increasingly green lines. This movement illustrates the static $\frac{\partial \tau_c^*}{\partial n_g} > 0$. Note that as the upper boundary moves outward from the origin, the boundary itself is also limited by the top and right edges of the graph (recall that the technology requires both $\alpha < 1$ and $\beta < 1$).

1.4 Conclusion

Attention surrounding impure public goods is growing in both the theoretical and empirical literature. In the standard impure good model built by Kotchen (2006), an interesting theoretical result arises: the introduction of an impure good can actually lead to a decrease in the total provision of public characteristic in the economy. This decrease occurs because the effect from the decrease in the implicit price of the private characteristic can potentially outweigh that from the decrease in the implicit price of the public characteristc. That is, the substitution from the public good to the impure good potentially has a greater impact than that from the private good to the impure good.

In order to correct for this decrease, I introduce a central planner into the model. The central planner chooses a tax rate for the purely private good and then spends the tax revenue generated to increase the total public characteristic provision. Consumers take into account the tax, resulting government spill-in, and all other agents' spill-ins when making choosing their consumption bundles.

I first isolate the conditions under which the introduction of the impure good causes a net decrease in the (general) equilibrium provision of the public characteristic. Next, I look at how the likelihood and the magnitude of this gap change with respect to the properties of the impure good and of the market as a whole. I then solve the central planner's problem and find the optimal tax rate required to eliminate the gap. Finally, I examine both the market and impure good properties that affect the size of the optimal tax, and thus also affect the amount of deadweight loss to society.

The number of individuals who purchase both the private and the impure good negatively affects the optimal tax rate, while the number of individuals who purchase solely the impure good positively impacts it. The technology parameters, which capure the portion of the impure good entering either the public or the private characteristic, both negatively affect the optimal tax rate. These intuitive results are in line with those found regarding changes in the gap itself.

These results imply that policy makers should be wary of markets with impure public goods. Markets with large substitution effects away from the public good or small substitution effects away from the private good are both particularly at risk. These substitution effects manifest themselves based on the external margin, the internal margin, and the impure good's technology. That is, they are based on the number of people who substitute, the intensity of such individuals' substitution (as dictated by preference parameters), and the extent to which the impure good is a substitute for the purely public and the purely private goods. Both markets with immensely popular impure goods and those with unpopular private goods will require a larger tax in order to correct for the impure public good's effect. Additionally, markets where the impure good is a poor substitute for either the purely public or the purely private good will require a larger tax.

These results provide a first glimpse into how policy makers can correct for the social costs that can arise in markets with impure public goods. However, policy makers may have options other than a Pigouvian tax on purely private goods that they can use to help fix these situations. Indeed, depending on the market scenario, policy makers may consider placing a tax on the impure good or possibly subsidizing the purely public good, although the latter would require an external funding source. Beyond that, many other real-world scenarios offer additional options that are less easily built into this model, such as the example of hybrid cars being allowed in HOV lanes. As that example perfectly illustrates, though, each policy is worth careful analysis before implemenation. To this end, the model presented here is a first step.

Chapter 2

Shareholder Protection and Dividend Policy: An Experimental Analysis

2.1 Introduction

The literature on dividend payout policies follows the trends in how firms behave. For example, research agendas have grown around the question of why companies began moving away from dividends in the late 1970s, why share repurchases became more popular than dividends since the mid-1990s, and why firms currently hold so much cash without increasing dividends.[1] Before all these, and beginning with the seminal paper by Modigliani and Miller (1958), researchers studied why firms pay dividends in the first place. Likely explanations came from the agency literature which assumes that corporate insiders have means of expropriating cash holdings in ways that reduce shareholder value.[2] For example, they could invest in projects that yield personal benefits to the insiders but reduce overall earnings, they could dilute the

[1]See, for example, Fama and French (2001) for the decline in dividends, Grullon and Michaely (2002) and Skinner (2008) for the rise in repurchasing, Bates et al. (2009) for the abundance of cash on-hand, and Farre-Mensa et al. (2014) for a recent survey.

[2]The seminal paper in the agency literature is Jensen and Meckling (1976).

holdings of outside investors by issuing shares to insiders, or they could simply pay the insiders more with the excess cash. Dividends reduce cash holdings and so constrain insiders from expropriating them for their own benefit. Two competing mechanisms for dividend payouts emerged from this setting. One arises from shareholders having strong legal protection and demanding dividends in order to constrain insiders. The second emerges from insiders using dividends to build a reputation for not expropriating funds, in which case dividends substitute for shareholder legal protection.

La Porta et al. (2000) posit that if shareholder rights and dividends are complements dividend payout ratios should be higher in countries with stronger legal protection (the outcome model), but if they are substitutes dividend ratios should be higher with weak shareholder protection (the substitute model). Their empirical analysis of 4,100 corporations in 33 countries supports the outcome model. Many subsequent papers have used other data sets and other empirical strategies to determine whether dividends and corporate governance are substitutes or complements. We add to this literature by running an experiment in which treatments differ by the strength of shareholder protection.

Our experiment has two main treatments, both of which allow insiders the opportunity to expropriate funds for their own benefit. In the low-protection treatment insiders determine dividend payouts, and in the high-protection treatment outsiders do. This second treatment captures the idea advanced by La Porta et al. (2000) that in high corporate governance regimes investors can demand dividends. When we compare dividends across treatments we find that dividend/cash ratios are five times higher in the high-protection treatment, an effect with the same direction but a much larger magnitude as found by La Porta et al. Thus, the laboratory results provide more evidence favoring the outcome model over the substitute model.

There are particular reasons why laboratory experiments can add significant value to the literature on agency models of corporate dividend strategies. The explanations for dividends rest on the premise that unless otherwise controlled, insiders expropriate corporate funds for their own benefit, but empirical work based on financial data must, for the most part, proceed without observing insider expropriations. For example, La Porta et al.'s estimates came entirely from sales and dividend data along with national laws. Berkman et al. (2009) represents one exception, reporting direct evidence from Chinese firms of expropriation by large shareholders (tunneling) through loan guarantees, and Jiang et al. (2010) represents another based on Chinese intercorporate loans. The laboratory experiment allows us to observe insider behavior directly, and we find that the expropriation/cash-on-hand ratios are 20 percentage points higher in the high-protection treatment than in the low one.

The most striking contribution from the experimental approach is that we can observe efficiency of the investor-insider relationship and, moreover, can trace any inefficiencies to their sources. In the low-protection treatment, the outsider has 100 tokens and can invest any fraction of that in the firm. The amount retained by the investor does not grow in value, but the amount invested with the firm grows by 30% per period, and we refer to this new amount as the firm's cash on hand, or cash for short. In period 2 the insider can allocate any portion of the firm's cash back to the outsider (dividend) or to his own personal account (expropriation), and as with the outsider's personal account, funds in the insider's personal account do not earn interest. The outsider observes the dividend payment but not the expropriation, so from this point forward cannot infer the amount of cash on hand. The amount left as cash again grows by 30%, and in period 3 the insider can make another expropriation/dividend allocation. There is no interest earned between periods 3 and 4, but in period 4 the outsider can invest any funds in his personal account back with the firm. Periods 5 and 6 are the same as periods 2 and 3, with 30% growth in cash before the periods and expropriation/dividend allocations during the periods. Following the final period-6 expropriation/dividend decision, any remaining cash is allocated 60% to the outsider and 40% to the insider. The high-protection treatment is the same except that instead of the insider determining dividend payments, the outsider determines them.[3]

[3] In the high-protection treatment if the sum of the dividend demanded by the outsider and the expropriation demanded by the insider exceed the amount available in cash-on-hand, the insider's expropriation is allocated in full and the outsider receives the remainder.

This game generates three sources of inefficiency. Irrespective of the treament, surplus is maximized when the outsider invests all 100 tokens in the first period and nothing is removed until period 6. Funds taken out of the firm and placed in personal accounts do not grow, and so both dividends and expropriations reduce efficiency.[4] Furthermore, funds not invested by the outsider do not grow either, so uninvested funds are a third source of inefficiency.

Paradoxically, increased outsider protection reduces outsider investment, leading to markedly lower efficiency in that treatment. In the game combined earnings of the two players can reach as high as 286 tokens, and as low as 100 tokens if the outsider never invests, a 186-token difference. In the low-protection treatment the two parties capture 71 of the 186 possible difference, but in the high-protection treatment they only manage to accumulate 15 more tokens than the outsider started with. Combined earnings are a highly-statistically-significant 55 tokens *lower* in the high-protection treatment than in the low-protection one, which corresponds to a one-third decrease in efficiency. Most of this loss comes right at the beginning, with the outsider investing 17 fewer tokens, or 45% less, in period 1 of the high-protection treatment. Efficiency deteriorates still further from there, with losses from dividends, expropriation, and failure to reinvest all larger in the high-protection treatment. In our experiment, then, strong investor protection leaves the firm with far less money to work with than weak protection does.

Our diminished efficiency finding contrasts with the empirical literature on how legal and financial institution strengths correlate with firm value and economic growth. La Porta et al. (2002) show that increased shareholder protection against insider expropriation correlates with increased firm value,[5] and Knack and Keefer (1995) find that countries with better citizen protection against government expropriation have higher economic growth and income. On the other hand, Bae et al. (2012) find

[4]The one exception is a dividend paid in period 3, because cash-on-hand does not earn interest between periods 3 and 4 and the period-3 dividend payment can be reinvested in period 4, making the period-3 dividend potentially efficiency-neutral.

[5]See also Durnev and Kim (2005).

that firm value increases when corporate governance controls the expropriation ability of large shareholders, not insiders, and Gompers et al. (2010) find a similar result about controlling the expropriation ability of the privileged investors in dual-class firms. In these firms outsider investors receive one vote per share but insider investors receive more, and Gompers et al. (2010) find that firm value decreases with privileged investor voting rights.[6] In our high-protection treatment investors have the power to give themselves dividends, which can be considered a form of expropriation by large or privileged investors. We enhance the Bae et al. and Gompers et al. findings by showing that controlling shareholders not only increases firm value by reducing expropriation by the outsiders, but also by increasing their investment in the firm and by reducing *insider* expropriation as well. Because it allows for expropriation by both the investor and the insiders, our study suggests that expropriation by large shareholders may be a larger problem than expropriation by insiders.

Our experiment has relevance for the literature on dividend policy and taxes. We ran versions of both the low-protection and high-protection treatments in which dividends were taxed. While we found that taxes reduce combined earnings, the efficiency loss came directly from the taxes themselves, because we found no statistically significant effect on initial investment or dividends, but we did find that taxes led to an *increase* in expropriations. The non-effect of taxes on dividends is consistent with empirical studies of taxes and dividends, where the consensus result is that investor protection matters more than taxes for dividend policies.[7]

2.2 Experimental Design

There are at least two sets of critical features needed for an experiment to be able to provide evidence if dividends and investor protection are compliments or substitutes.

[6] Jordan et al. (2014) find that a substitute model holds for the privileged shareholders in dual-class firms, in that they use dividend payouts to commit to not expropriating cash-flow.

[7] See Chetty and Saez (2005) and Alzahrani and Lasfer (2012) for the interplay of tax rates and payout policy, as well as the survey by Farre-Mensa et al. (2014).

First, the experiment must embed the ability for subjects to build trust. In the substitute model, agents operating a firm on the inside use dividends as a costly signal thought to increase trust and secure future investment from principals investing from the outside. As a result, the experiment must give subjects enough time and repeated interactions to build trust during each run, along with a way for principals to respond to the signals.

Second, the experiment must have a mechanism to exogenously vary the level of shareholder protection. In order to maximize external validity, we use the hypothesized reason that dividends are paid within both the substitute and outcome models to motivate the way in which dividends are paid within the experiment. Varying this across treatments provides both a change in the level of outsider protection and a change in the mechanics of how dividends are paid, as desired. This mechanism is discussed in detail later.

We use a novel experimental design crafted with both of these important features in mind. The design mimics a principal investor purchasing shares in a firm and actual effort decisions of an employee at that company acting as the potential shareholder's agent. As in actual investment decisions, there are elements of trust, shirking, and patience which interact to affect decision making. To that end, the design bears some conceptual resemblance to both the trust game and the centipede game (Berg et al. (1995) and Rosenthal (1981)).

We recruited undergraduates for the experimental lab at the University of Tennessee, Knoxville to participate in the experiment. The experiment was run entirely using z-Tree.[8] We ran six sessions of 16 subjects each. Each subject plays 12 games, and each game has 6 periods. This totals to 576 observations at the game level. Subjects are paid based on exactly one of the twelve games they play, which is chosen at random on the session level. Tokens were earned at a rate of 10 tokens = \$1. This was paid on top of a \$10 participation fee.

[8]See Fischbacher (2007).

There were 6 treatments, set up in a 2×3 design. Subjects never varied across the binary dimension, but they rotated in the other dimension across the twelve games following the pattern *AAA BBB CCC AAA*. More details regarding the treatments are discussed later.

For each game, subjects are randomly and anonymously paired. Rematching occurs, with replacement, before each game. Within each pair, there are two roles, the outside investor (principal) and the corporate insider (agent). Each subject always plays the same role throughout the entire session.

Each player has a personal account. The amount of tokens in their account at the end of the game is the number of tokens they earn for that game. That is, for a given game, if that game is chosen as the payoff game for that session, then subjects are paid according to the number of tokens in their account at the conclusion of that game. In addition to the two accounts, there is also a "fund," which represents a firm, or more specifically, represents a firm's cash on hand. The firm periodically earns interest (as described below) at a rate of 30%.

The flow of the 6 periods is shown in Figure 2.1. Consider first our main treatment (low outsider protection with untaxed dividends, or "LU"). In the beginning, the outsider is endowed with 100 tokens in her account, while the insider has 0 in hers. The firm also has 0 tokens. In period 1, the outsider has the opportunity to invest any amount of her tokens in the firm.

Moving into period 2, the firm's fund earns 30% interest (rounded to the nearest whole number of tokens). The insider then has the opportunity to move tokens from the fund to his own account and to the outsider's account. Movements to his own account represent insider expropriations, and movements to the outsider's account represent dividend payments. It is crucial to note that the outsider cannot observe either the firm or the insider's account. From this point forward, the outsider can never infer either of these values.[9] Period 3 then operates exactly identically to period 2.

[9]This parallels investors' inability to check firms' books or insiders' shirking on demand.

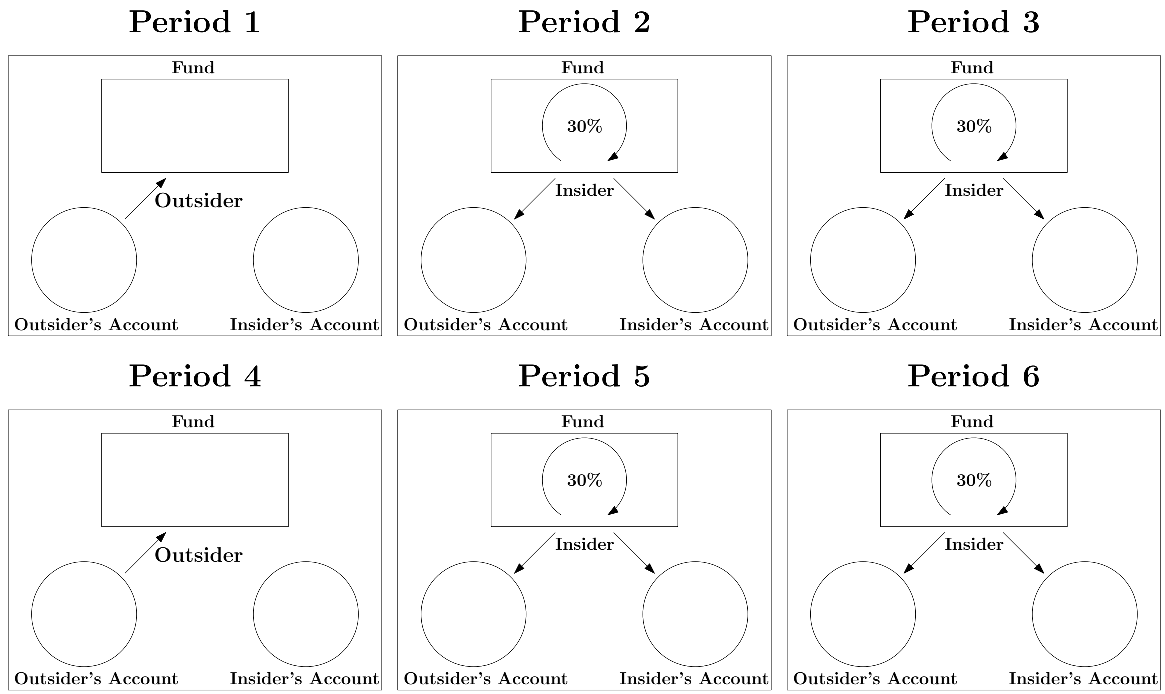

Figure 2.1: Visualization of the dividend treatment experimental design with low outsider protection

Period 4 is the reinvestment period. No interest is earned moving into period 4. Instead, the outsider again has the opportunity to invest tokens from her personnal account into the firm's fund. Recall that she cannot see the current value of the firm. She can, however, receive a signal about firm value in the form of period 3 dividends.[10] The outsider can invest any portion of her account, which at this point can range from 0 tokens (if all was inititally invested and no dividends were paid) to 169 tokens (if all was initially invested and all was repaid as period 3 dividends).

Periods 5 and 6 operate exactly as periods 2 and 3. Following period 6, any tokens remaining in the firm's fund are split according to a known 60:40 ratio in favor of the outsider. Again, this is all depicted in Figure 2.1.

Now consider the treatments, which again follow a 2×3 design. The treatment discriptions and corresponding abbreviations are given in Table 2.1. First varying in

[10]Period 2 dividends also send a signal, but sending that signal in period 2 rather than in period 3 does not make sense from the insider's perspective because they forgoe interest by doing so.

Table 2.1: Experimental Treatments

Dividend Treatment	Protection Treatment Low Protection	High Protection
No Dividends	LN	HN
Untaxed Dividends	LU	HU
Taxed Dividends	LT	HT

NOTE: All subjects remain either an insider or outsider for entire session. Protection treatment never varied within a session.

the binary dimension, we switch from low to high outsider protection. The mechanism through which this change occurs mimics the competing desctriptions by the substitute and outcome models regarding why dividends are paid. In the substitute model, corporate insiders pay dividends to signal trustworthiness. This was the situation we set up in the LU treatment previously described. In the outcome model, outside investors demand dividend payments. Hence, in the HU treatments, outsiders are able to directly take dividends from the firm.

It is critical to note that the outsider still cannot observe the value of the firm in the high-protection treatments. The investor can request any amount in dividends, and if that much is not available, then the total remaining value of the firm will instead be paid. In the case that the sum of the dividends demanded by the outsider and the expropriations taken by the insider together overdraws the firm's cash on hand, the insider receives first dibs. That is, the insider is paid the amount that he requested, and the remainder goes to the outsider. This mimics the ability of an agent to expropriate before equity investors are repaid.

Looking at the variation in treatment design along the trinary dimension, there is one mechanism that varies. There is again a difference in how dividends are paid. For both the low- and high-protection treatments, the tax treatments introduce a 25% tax on dividends paid in periods 2, 3, and 5. There is no tax in period 6 because dividends merely reflect a distributional conern at that point. These treatments are abbreviated as LT and HT, as shown in the bottom row of Table 2.1. This creates an efficiency loss in either using dividends to either send a signal or prematurely taking dividends.

Moving instead up to the top row of Table 2.1, treatments LN and HN completely eliminate the dividend channel in periods 2, 3, and 5. Again, they exist in period 6 only for distributional concerns. Because they serve as a baseline where no signaling is possible, the no-dividend treatments are sometimes referred to as the "baseline" treatments.

Each subject faced a total of 3 treatments across the 12 games they played. These varied only within the trinary dimension, cycling through one column of Table 2.1. Each of the 6 sessions started in a different treatment cell and worked downward, rotating back to the top following the tax treatment. Treatments were run in sets of 3 games. That is, calling treatments A, B, and C for simplicity, a subject starting on A would play the 12 games as AAA BBB CCC AAA.

At the end of each session, we administered a questionnaire to the subjects. The questionnaire did not affect payouts in any way. It consisted of a Holt-Laury risk ladder and standard demographic questions. This was done while payout envelopes were stuffed so that it added no time to the experiment. Average session times were less than 90 minutes and average payouts were $22. This includes the tokens earned within the game, which were exchanged at a rate of 10 tokens to $1, and a $10 show-up fee for all participants.

2.2.1 Hypotheses and Game Discussion

The protocols faced by the subjects are new to the laboratory, so the first task entails testing for evidence of rational responses. The design builds in two ways to test for within-treatment and across-treatment rationality. The first such test arises from the low-protection treatments where the insider makes all of the allocation treatments.

The insider can pay dividends to the outsider in periods 2, 3, 5, and 6. Recall that the cash on hand also grows at the beginning of these periods. Thus, paying out dividends too early creates lost potential income. For example, funds not paid out in period 2 grow by 30% in period 3. A dividend payout path (d_2, d_3, d_5, d_6) with $d_2 > 0$ generates less surplus than the alternative path $(d_2 - \epsilon, d_3 + \epsilon, d_5, d_6)$ that delays a payment of $\epsilon < d_2$ from period 2 to period 3, because the payout delay grows the surplus by $.3\epsilon$ in just one period. The same holds for delaying the period-5 dividend until period 6, as well as for an expropriation path (e_2, e_3, e_5, e_6) with $e_2 > 0$ or $e_5 > 0$. These considerations lead to the following hypothesis.

Hypothesis 1a. In the low-protection treatments, dividend payouts and expropriations are lower in periods 2 and 5 than in periods 3 and 6, respectively.

In order to understand why Hypothesis 1a specifically ignores high-protection treatments, it helps to examine the similarities between each protection regime and previous experiments. The low-protection treatment has much in common with the oft-studied trust game of Berg et al. (1995). In the trust game, player 1 begins with an endowment and can invest any fraction of that in the trust relationship. The amount invested is tripled, and then player 2 allocates the proceeds between himself and player 1. The low-protection treatment extends the trust game by (i) allowing the insider to allocate the proceeds before they have fully grown, (ii) providing the outsider an opportunity to invest more midway through the growth process, and (iii) identifying a default 60:40 outsider-insider split at the end of the game. Experiments on trust games have found initial investments averaging about half of the endowment and a slightly positive return on the initial investment (e.g. Berg et al. (1995)).[11] The

[11] Rietz et al. (2013) extend the trust game into a more dynamic setting by having player 1 invest with player 2, who can then invest any fraction of that with player 3. Investments are tripled at both stages, and at the end of the game player 3 first makes an allocation between herself and player 2, who then makes an allocation between herself and player 1. Greiner et al. (2012) take a different approach, with players participating in a series of trust games in which their endowments in period $t > 1$ are their accumulated earnings from the prior periods. Most closely related to our paper, Lunawat (2013) adapts a trust game by adding randomness to the interest rate that the insider can observe but the outsider cannot. The insider could disclose the actual interest rate in the disclosure treatment but not in the non-disclosure treatment. She finds that outsiders invest more in the non-disclosure treatment, providing another instance of increased corporate governance reducing efficiency.

design of the low-protection treatment was influenced by the trust game, and the 30% growth rate was chosen to yield a near tripling of the original investment.

The high-protection game is very different, and more closely resembles the centipede game of Rosenthal (1981) than the trust game. In the centipede game, players alternate choosing whether to continue the game or end it, with each continuation choice leading to an increase in the surplus. Whenever the game ends, the player whose turn it was to make a decision receives more than half of the available surplus, thereby alternating who would receive the larger share. The game has a fixed endpoint, and so backward induction leads players to try to end the game in their favor one period before their opponents do.[12] In the high-protection treatment of our investment game, the outsider has an incentive to guarantee his allocation through dividends before the insider expropriates it, and vice versa, and this should intuitively lead to backward unraveling and low initial investment. For this reason, Hypothesis 1a only examines low-protection treatments, where this incentive is not present.

The second gameplay-rationality hypothesis arises from the fact that in some treatments dividends are taxed but in others they are not. Taxes raise the price of dividends, so one would expect less use of dividends in the tax treatments than in the untaxed treatments. Because behavior may differ between the high-protection treatments and the low-protection ones, a test of a rational response to the price increase must compare dividends in the high-protection tax treatment to the high-protection untaxed treatment and in the low-protection tax treatment to the low-protection untaxed treatment. The baseline no-dividend treatments differ from the corresponding tax treatments in that dividend payouts can only occur in period 6 in the baseline but can occur in periods 2, 3, 5, and 6 in the tax treatments, so the baseline treatments cannot be used to test for this type of rationality.

Hypothesis 1b. Dividends are lower in the taxed dividend treatments than in the corresponding untaxed dividend treatments.

[12] For experimental evidence see McKelvey and Palfrey (1992) and Levitt et al. (2011).

Allowing for reinvestment in period 4 allows for an interesting test. The substitute model predicts that outsiders pay dividends to signal their trustworthiness. To accomodate such behavior, the game allows dividends twice before the reinvestment period. Without the reinvestment opportunity, the insider would have no reason to signal. Note that the insider can only send this signal in the low-protection treatment, when she controls dividend payments. We formalize this postulation into the following hypothesis, which omits period 6 because it pertains to distribution preferences.

Hypothesis 1c In the low-protection dividend treatments, dividends are higher in periods 2 and 3 than in period 5.

The next hypothesis is our main one, and it is driven by the literature started by La Porta et al. (2000) testing the outcome model against the substitute model. The outcome model implies that dividends are higher in the high-protection treatment, the substitute model predicts the opposite, and the preponderance of the empirical evidence favors the outcome model. Our hypothesis states that behavior in the lab mirrors that in the real world.

Hypothesis 2. Dividend/cash on hand ratios are higher in the high-protection treatments than in the corresponding low-protection treatments.

We look at dividends over cash on hand for two reasons. One is consistency with the empirical literature, which has looked at dividends as a fraction of cash on hand and as a fraction of cash flow. In our design, cash on hand corresponds to the amount invested in the firm, which is the maximal amount subjects could allocate to personal accounts in any of the payout periods 2, 3, 5, or 6. The design also generates cash flow in the form of interest on invested funds, but by design cash flow equals 30% of cash on hand, so analyzing that ratio separately cannot yield different results.

The second reason for using the dividend/cash on hand ratio is that it allows for differences in investment behavior across treatments. If the outsider invests less in one treatment than in another, the insider would have fewer funds to disperse in the first treatment. That might lead to lower dividends, which could lead to a false conclusion concerning the outcome model versus the substitute model based more on initial investment behavior than on subsequent payout behavior. Normalizing by cash on hand removes this concern.

The third set of hypotheses relates to efficiency. If the outsider invests all 100 tokens in the first period and subjects allow interest to accrue until period 6, their earnings would sum to 286 tokens.[13] Because neither personal account grows over time, failure to invest the full amount, dividends, expropriations, and failure to reinvest all represent leakages from the system that lead to efficiency losses.

We can measure the efficiency of the relationship by summing the tokens the two parties earn during the seven periods. The empirical literature is silent on efficiency because it is unobservable in the data, but there is no reason for efficiency to be the same across investor-protection regimes. The outcome model predicts higher dividends in the high-protection treatment, which would reduce efficiency there, but those leakages could be offset by differences in expropriations, reinvestment behavior, or initial investment. For this reason we offer the following null hypothesis.

Hypothesis 3a. The high-protection treatments generate the same total surplus as the corresponding low-protection treatments.

The experiment allows for tracing any differences in surplus back to the source of the leakage. The literature is again silent on how efficiency is affected via each leakage channel, so we also offer the related null hypothesis agnostically.

Hypothesis 3b. Leakages from initial investment, reinvestment, dividends, and expropriations are the same in the high-protection treatments as in the corresponding low-protection treatments.

[13] In the no-tax treatments the minimum-achievable surplus is 100 tokens which arises when the outsider invests nothing. In the tax treatments the minimum can be slightly lower and occurs in the following scenario. The outsider invests all 100 tokens in period 1, receives a dividend payout of the entire 130 in period 2, but the 25% tax on dividends reduces it to 98 tokens (after rounding). She reinvests all of this in period 4, it grows to 127 which she receives as a dividend in period 5, but after taxes it is only worth 95.

However, implications can be drawn about how these channels *might* behave given findings related to each. Bae et al. (2012) and Gompers et al. (2010) find that increases in the power of large and privileged shareholders correlate with reductions in firm value. In essence, large or privileged shareholders can expropriate funds in much the same way that insider managers can, leading to a contest to see who can get the funds out of the firm first. Backward unraveling then leads to the outsiders withholding investment in the first period, culminating in extremely low firm valuations. Within the context of our experiment, the high-protection treatments represent the increase in large shareholders' power, and the resulting game more closely resembles this backward unraveling. Initial investment might be expected to be lower in the high-protection treatment, thus increasing leakage.

Leakages through dividends directly mirror the effects on the ratio of dividends to cash on hand. By including cash on hand in the analysis, any change in initial investment is accounted for. Thus, an increase in the ratio implies a greater loss due to leakages, as the outcome model predicts high outsider protection will cause, while a decrease implies less leakages, as the substitute model predicts high protection will cause.

Insider expropriation is generally unobservable in financial data. For this reason, predicting the impact of outsider protection regimes on expropriations based on existing empirical literature is difficult, and thus our link to predicting how leakages relate to protection regimes is tenuous at best. Berkman et al. (2009) examine insider expropriation in the form of "tunneling" from Chinese firms through loan guarantees, and Jiang et al. (2010) examine it via Chinese intercorporate loans. Both find that larger expropriations are correlated with smaller firm size. In our experiment, firm size is directly related to initial investment, and thus their results suggest lower initial investment should spark higher insider expropriations. If we take at face value the previously discussed link suggesting higher outsider protection spurs lower initial investment, then our high-protection treatments should yield higher insider expropriations, and thus cause a larger loss in efficiency via this channel. Again, due to the multiple leaps of faith required to reach these conclusions, our null for Hypothesis 3b remains agnostic.

Hypotheses 2, 3a, and 3b all relate to corporate payout policies. It would be possible to add a number of behavioral hypotheses. For example, fairness preferences would suggest that dividends and expropriations should be approximately equal in size, an aversion to being behind would suggest that the insider makes a sufficient period-6 expropriation to guarantee that he does not earn less than the outsider, and so on. Exploring these hypotheses might add to the behavioral literature, but it would detract from the focus of the paper on dividend payout policies, so we leave these avenues unpursued.[14]

The key to our design rests in switching control of dividends from the insiders to the outsiders. Beyond this there were a number of design choices that had to be made, and these included the length of the game, especially the choice to have a fixed endpoint for the relationship rather than allowing for one that would mimic an infinitely-repeated game. We chose to do this because an infinite number of rounds would allow for supergame strategies that would promote cooperation, which in turn would take the form of increased combined earnings. Dividends, expropriations, and reinvestment choices could all be driven by punishment strategies, and the existing literature in finance does not think about payout policy in this way. Consequently, we believe that our design isolates the theorized rationales for dividends more than an infinite-round version of the game would.

The prevailing theories predict a number of different behaviors, and we made design choices in order to allow these behaviors to occur and to be testable. As mentioned, we allow for dividends before the reinvestment period in order to test signaling, and we allow for two dividend and expropriation periods between investment periods to detect unraveling. Most importantly, though, our design had to capture the essences of weaker and stronger investor protection regimes. The choice here was driven primarily

[14] Breuer et al. (2014) explore behavioral issues in dividend payout policy using firm-level data on dividends and country-level data on behavioral characteristics like time preferences along with risk, loss, and ambiguity attitudes.

by the outcome model, which posits that dividends occur because strong shareholders demand them. Rather than having outsiders request dividends from the insiders, we chose to let the outsiders simply take them. Giving insider claims precedence over outsider ones in case the simultaneous claims exceeded cash on hand had the effect of giving insiders some control over how much outsiders could demand, but not too much control. In all we contend that the design captures the salient aspects of investor protection laws and allows for the behavior posited in empirical studies of the topic. We find evidence in line with the hypothesized rationales for dividend payments, and we also find some new results.

2.3 Results

Table 2.2 provides an overview of the variables of interest across the six experimental treatments. The labels "Low" and "High" refer to the low-protection and high-protection treatments, respectively, and the remaining labels refer to whether dividends can be paid and whether they are taxed in periods 2, 3, and 5. The second column of the table introduces the codes we use to refer to the different treatments, with LN referring to the low-protection, no-dividend treatment, and so on. The table aggregates dividends and expropriations from periods 2, 3, and 5, leaving out period 6 because those payments are more for allocating the final cash on hand differently than the prescribed 60:40 outsider-insider split. The final firm value represents the amount of cash on hand the two parties have available to allocate in periods 6 and 7, and this value comes from funds remaining invested in the firm at the end of five periods. The maximum that this value can attain is 286 when the outsider invests everything in the first period and neither party removes any funds prior to period 6, and the minimum value is 0, which occurs when the outsider invests nothing.

Hypothesis testing requires disaggregating by period, but still the aggregated data can reveal several striking patterns. Insiders do, in fact, expropriate funds, and they expropriate more in the high-protection treatments than in the corresponding

Table 2.2: Experiment summary statistics

Treatment	Code	Initial investment	Total dividend	Total expropriations	Final firm value
Low, no dividends	LN	38.4	-	6.3	108.5
Low, untaxed dividends	LU	38.2	2.0	7.7	107.8
Low, taxed dividends	LT	31.0	1.3	7.5	82.7
High, no dividends	HN	20.0	-	18.4	33.5
High, untaxed dividends	HU	20.8	13.2	23.1	6.6
High, taxed dividends	HT	19.1	10.1	20.9	8.3

Total dividends = dividends in periods 2, 3, and 5
Total expropriations = expropriations in periods 2, 3, and 5
Final value of the relationship = dividends and expropriations in period 6 + final amount split

low-protection ones. The primary question for the this paper relates to comparing dividends across the low-protection and high-protection treatments, and firms pay more aggregate dividends in the high-protection treatments than in the corresponding low ones. This is consistent with the outcome model and the findings in La Porta et al. (2000), but these numbers do not account for how much cash was on hand for dispersal each period. Taxing dividends leads to a small reduction in payouts. Finally, low-protection treatments have lower initial investment and produce less-valuable firms than high-protection ones do.

2.3.1 Gameplay Rationality

As argued in Section 2, the experiment allows for a straightforward prediction based solely on subject rationality. In particular, in the low-protection treatments the insider alone has control of how cash on hand is allocated, and can therefore time allocations to take advantage of the fact that cash on hand grows between periods. Consequently insider allocations through dividends and expropriations should be larger in period 3 than in period 2, and larger in period 6 than in period 5.

Testing such a hypothesis is a way to validate an experimental design that is new to the literature. To do so we estimate the following specification:

$$Payout_{ipgt} = \sum_{t \in \{N,U,T\}} \left[\beta_{2t} \mathbf{1}_{p \in \{2,3\}} + \beta_{3t} \mathbf{1}_{p=3} + \beta_{5t} \mathbf{1}_{p \in \{5,6\}} + \beta_{6t} \mathbf{1}_{p=6} \right] + \varepsilon_{ipgt}. \quad (2.1)$$

In equation (2.1) the dependent variable is the number of tokens paid either as dividends or expropriations. Because we restrict the sample to only the low-protection treatments, we use number of tokens paid, rather than tokens as a percent of cash on hand, for ease of interpretation. **1** is the indicator function, i indexes individuals, and t indexes treatments (N= no dividends, U= untaxed dividends, and T= taxed dividends). Subjects play each treatment several times, and g indexes the time the treatment is being played, which we refer to as a game. Finally, p indexes the period within the game. There are four periods in which dividends could be paid: $p = 2, 3, 5, 6$. The coefficient β_{2t} measure average dividends in tokens common to periods 2 & 3 of treatment t. β_{5t} is defined similarly for periods 5 & 6. The coefficients β_{3t} and β_{6t} capture the marginal effects of advancing one period on payouts. We estimate equation (2.1) via OLS including subject random effects.

Hypothesis 1a posits that the marginal effects for periods 3 and 6 are nonnegative, highlighting that subjects exploit the ability to earn interest before either paying dividends or expropriating tokens. Hypothesis 1a can be tested in estimating equation (2.1) by rejecting the joint null hypothesis: $H_0 : \beta_{3t} = \beta_{6t} = 0 \ \forall \ t$. Hypothesis 1b posits taxes on dividends decrease dividend payments. Hypothesis 1b can be tested in estimating equation (2.1) by rejecting the joint null hypothesis $H_0 : \beta_{3T} = \beta_{3U}, \ \beta_{6T} = \beta_{6U}$.

Table 2.3 shows results from estimating equation (2.1).[15] The first two columns have dividends as the dependent variable, and the third and fourth columns have

[15] The coefficients for pooled periods 5 and 6 had to be dropped for for both treatments LU and LT in column 2 due to a collinearity issue; when dropping games 1, 4, and 7, no dividends were paid out during period 5 in the low-protection treatment for either untaxed or taxed dividends.

Table 2.3: Treatment by Period Effects on Dividends and Expropriations

	Div		Exprop	
	(1)	(2)	(3)	(4)
Trimmed:	None	1, 4, 7	None	1, 4, 7
LN x (P2 or P3)			2.594*	1.958
			(1.277)	(1.378)
LN x P3			-0.010	0.125
			(1.894)	(2.255)
LN x (P5 or P6)			1.094	0.417
			(0.661)	(0.409)
LN x P6			60.542***	66.347***
			(10.892)	(13.739)
LU x (P2 or P3)	0.156	0.167	2.063**	0.847*
	(0.107)	(0.138)	(0.883)	(0.491)
LU x P3	1.260*	1.319	1.771	2.764*
	(0.692)	(0.908)	(1.511)	(1.585)
LU x (P5 or P6)	0.438	0.000	1.792**	1.056
	(0.431)	(.)	(0.675)	(0.746)
LU x P6	16.063**	14.986**	63.146***	68.639***
	(6.705)	(6.039)	(11.988)	(14.737)
LT x (P2 or P3)	0.365	0.139	3.521***	1.486**
	(0.221)	(0.092)	(1.211)	(0.650)
LT x P3	-0.083	0.069	-1.240	0.569
	(0.129)	(0.047)	(1.585)	(1.514)
LT x (P5 or P6)	0.677	0.000	1.719**	0.639
	(0.465)	(.)	(0.698)	(0.612)
LT x P6	4.042*	5.889**	43.333***	43.792***
	(1.971)	(2.712)	(9.332)	(10.858)
N	768	576	1152	864
r2	0.125	0.112	0.365	0.364

Standard errors in parentheses
Notes: All errors are clustered at the subject level. Only low-protection treatments are used.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

expropriations. Columns 2 and 4 remove the first instance that each game is played, that is, games 1, 4, and 7. For dividends we use only the untaxed and taxed dividend treatments since dividends cannot be paid in the no dividends treatment. Since expropriations are always available to the insider, we include all three low protection treatments. For the dividend portion of hypothesis 1a, we find that the average dividends are 1.3 tokens higher in period 3 relative to period 2 (significant at the 10% level). The effect is even more pronounced later in each game: period-6 dividends are 16 tokens higher than period-5 dividends (significant at the 5% level). Average dividends in periods 2 and 5 were not significantly different from zero.

For the expropriation portion of hypothesis 1a, the period-6 marginal effects are large and significant in each of the LN, LU and LT treatments (60.5, 63.1, and 43.3 respectively). Put another way, period-6 expropriations are roughly at least an order of magnitude larger than all other expropriation coefficients. However, unlike in the dividends regressions, in no treatment was the marginal effect of period 3 significantly different from zero. This is expected: with dividends, there is an opportunity to build trust with period-3 dividends. Conversely, with expropriations there is only an efficiency loss from expropriating in period 3. Insiders do, though, display some degree of impatience for expropriating tokens: in each of the LN, LU and LT treatments, average expropriations across periods 2 and 3 are significantly greater than zero, although magnitudes are small (2.6, 2.1, and 3.5 respectively).

Turning to the second gameplay rationality hypothesis 1b, in the taxed dividend treatments we observe that subjects only paid dividends significantly greater than zero in period 6. Hence, taxes appear to significantly influence dividend behavior in a way consistent with the hypothesis. We can reject the joint null hypothesis that taxes have no effect, that is, $\beta_{3T} = \beta_{3U}$ and $\beta_{6T} = \beta_{6U}$ using an F-test ($p = 0.0415$).

Hypothesis 1c predicts that, in the low-protection treatments, dividends in periods 2 and 3 are higher than in period 5. This hypothesis stems from the substitute model's postulation that insiders pay dividends as a means to signal trustworthiness to outside investors. This is not predicted in high-protection treatments, where investors choose their own dividends. Hence, in looking at hypothesis 1c, we both look for evidence to reject the null in low-protection treatments and check the corresponding high-protection treatments as a test for false-positives. We also examine the extent to which introducing a dividend tax, which effectively increases the cost of signaling, dampens that signaling.

Table 2.5 shows the period-by-period effects of each treatment on dividends over cash on hand. The first pair of columns represent one regression, with the first column representing the low-protection treatments and the second representing the marginal impact of switching to high protection. The second pair does the same but trims off games 1, 4, and 7.

For the LU treatment, we test the coefficient for period 3 against that for period 5. We can reject the null that the two are equal ($p = 0.074$), lending support to hypothesis 1c. When taxes are introduced in treatment LT, the effect goes away ($p = 0.487$). Additionally, unlike in treatment LU, there is no reason to see signaling in treatment HU. Testing the coefficient for period 3 against that for period 5 here, we cannot reject the null ($p = 0.494$). Thus, we find evidence for trust signaling in the untaxed low-protection treatment, as the substitution model predicts. However, either switching to the high-protection treatment or introducing a dividend tax effectively negates any traces of signaling.

In sum, we find gameplay results very consistent with rational gameplay and our first set of hypotheses. We also find evidence that insiders are impatient, even when there is no incentive for them to be. This impatience finding highlights a strength of our design: because we include a baseline treatment with no dividends across both the high and low protection scenarios, we can difference out this impatience effect to focus solely on how protection affects dividend payouts and efficiency losses from expropriations.

2.3.2 The La Porta et al. Hypothesis

Turning to the paper's primary question of how dividend payouts change with investor protection, the outcome model predicts dividend payout ratios will be higher in the high-protection treatment than in the low ones, and the substitute model predicts the opposite. The former has found more support in the literature, beginning with La Porta et al. (2000). These tests, using financial data, typically regress dividend/cash-on-hand ratios on control variables that include a measure of investor protection. We do the same here.

Figure 2.2 shows dividend payouts averaged across all periods as a percent of cash on hand for each of the six treatments. There are, by definition, no dividends in the LN and HN treatments. Dividends in treatments HU and HT are both significantly larger than in LU and LT, respectively. The differences are statistically significant. Consistent with hypothesis 1b, taxes appear to decrease dividend payouts. However, the effects of taxes are very much second order to the effects of shareholder protection treatments.

To test the precise magnitudes of moving from low to high shareholder protection, we estimate the following regression restricting the sample to only those treatments in which dividends were paid:

$$Dratio_{igt} = \beta_{LU}\mathbf{1}_{t\in\{LU,HU\}} + \beta_{HU}\mathbf{1}_{t=HU} + \beta_{LT}\mathbf{1}_{t\in\{LT,HT\}} + \beta_{HT}\mathbf{1}_{t=HT} + \epsilon_{igt}. \quad (2.2)$$

In equation (2.2), the variable *Dratio* is the ratio of dividends to cash on hand expressed as a percentage. β_{LU} and β_{LT} measure the average dividend payouts in the untaxed and taxed dividends low-protection treatments, respectively, aggregated across all periods. The coefficients β_{HU} and β_{HT} capture the marginal effects of switching to the untaxed and taxed dividend high-protection treatments. In line with Hypothesis 2, our data support the outcome model when $\beta_{HU} > 0$ and $\beta_{HT} > 0$. Such a finding would imply the marginal effect of switching to the high dividend treatment

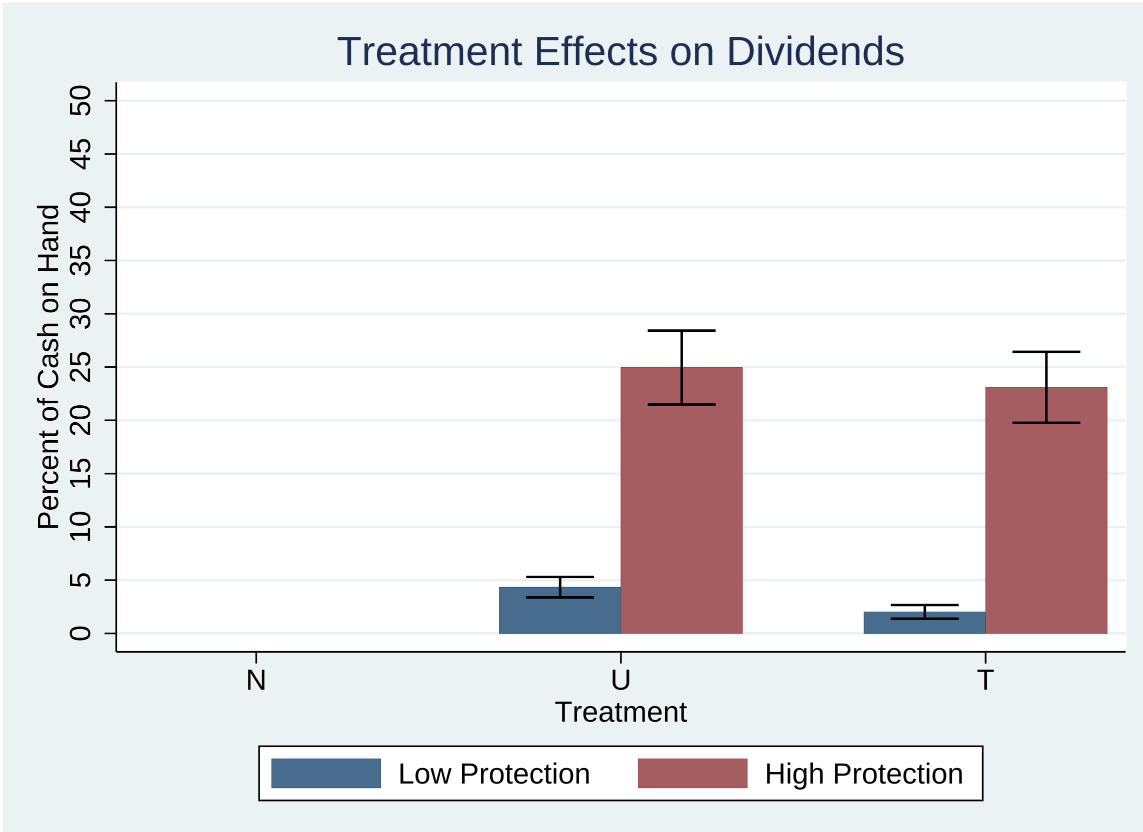

Figure 2.2: Treatment Effects on Dividends as a Percent of Cash on Hand

is to increase dividend payout. Conversely, a finding of $\beta_{HU} < 0$ and $\beta_{HT} < 0$ would support the substitute model.

We estimate equation (2.2) using OLS, with three modelling notes. First, we exclude individual fixed effects because there is no high versus low protection variation within subjects. This was an experimental choice to eliminate contamination across treatments. Second, we cluster standard errors at the subject level to allow for correlation in dividend behavior within subjects. Lastly, we multiply the dividend to cash on hand ratio by one hundred so that coefficients are in percentage points.

The first two columns of Table 2.4 show the estimated coefficients from equation (2.2). The first column uses the full sample, while the second column removes the first iteration of each treatment to eliminate any learning effects within a treatment. The third and fourth columns show results for the same estimating equation with

Table 2.4: Treatment Effects on Dividends and Expropriations (as a Percent of Cash on Hand)

	$\frac{\text{Div}}{\text{CoH}}$		$\frac{\text{Exprop}}{\text{CoH}}$	
	(1)	(2)	(3)	(4)
Trimmed:	None	1, 4, 7	None	1, 4, 7
LN or HN			19.820***	19.454***
			(1.789)	(2.202)
HN			15.877***	17.097***
			(3.650)	(4.647)
LU or HU	4.342***	3.855***	21.928***	21.231***
	(0.961)	(0.989)	(1.611)	(1.965)
HU	20.610***	21.861***	20.595***	20.988***
	(3.466)	(4.122)	(4.076)	(5.629)
LT or HT	2.026***	1.697***	23.334***	23.274***
	(0.642)	(0.619)	(2.995)	(3.055)
HT	21.075***	21.353***	21.054***	20.408***
	(3.328)	(4.643)	(4.586)	(5.584)
N	806	539	1256	852
r2	0.302	0.290	0.395	0.378

Standard errors in parentheses

Notes: All errors are clustered at the subject level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

expropriations as a percent of cash on hand as the dependent variable; we discuss those results at length below. The results show strong support for the outcome model, with dividend ratios a highly-significant 21 percentage points larger in the high-protection treatments than in the low ones (i.e., both $\hat{\beta}_{HU} > 0$ and $\hat{\beta}_{HT} > 0$). Looked at differently, high investor protection generates dividend ratios that are 5 times larger in the no-tax treatments and 10 times larger in the tax treatments when compared to low-protection regimes.

To investigate the results in a disaggregated way, Table 2.5 shows results from estimating equation (2.2) by period. That is, we estimate

$$Dratio_{ipgt} = \Sigma_{p \in \{2,3,5,6\}} \, \Sigma_{t \in \{U,T\}} \left[\beta_{Ltp} \mathbf{1}_{t \in \{Ltp, Htp\}} + \beta_{Htp} \mathbf{1}_{Htp} \right] + \epsilon_{ipgt}. \tag{2.3}$$

This allows us to compare the relative magnitudes of any dividends used as signals in the low protection treatments versus dividend behavior in the high protection treatments. Rows of the table are paired corresponding to periods of the game, with the first row in each pair pertaining to the relevant period of untaxed dividend treatments and the second pertaining to the taxed dividend treatment. The first column can be interpreted as the coefficient for the low-protection treatment and the second column shows the marginal effect of moving to the high-protection treatment. The third and fourth columns show robustness to excluding the first game each individual plays in each treatment.

Coefficients in the high-protection treatment columns are marginal effects, and so positive coefficients provide support for the outcome model. With one exception all of these coefficients are positive, highly significant, and large. The sole exception comes from period 6 dividends in the no-tax treatments, and this could be driven by the high period 6 dividends in the LU treatment rather than low ones in the HU treatment. Still, the disaggregated results support the outcome model, and the laboratory experiment aligns well with the results from studies using financial data.

The outcome model is driven by investors with high protection demanding higher dividends, while the substitute model is driven by firms paying dividends to build a reputation for good stewardship of the invested funds. Any evidence of using dividends to signal good stewardship would appear in Table 2.5, and signaling would lead to positive payouts in the low protection treatments in period 3, before the reinvestment opportunity. Table 2.5 shows that this occurs in the LU treatment, with insiders paying an average of 5.3% of the available cash as dividends in that period, but it does not occur in the LT treatment, with the insiders paying an insignificant 0.6%

Table 2.5: Treatment Effects on Dividends (as a Percent of Cash on Hand) by Period

Trimming:		None		1, 4, 7	
		L or H	H	L or H	H
P2	U	0.322*	25.150***	0.245	26.883***
		(0.191)	(4.199)	(0.175)	(4.828)
	T	0.889	23.034***	0.321	22.281**
		(0.589)	(7.064)	(0.229)	(8.509)
P3	U	5.345**	18.968***	5.338**	17.475**
		(2.494)	(6.802)	(3.112)	(7.908)
	T	0.611	29.334***	0.434	31.710***
		(0.375)	(5.035)	(0.315)	(6.663)
P5	U	0.534	29.801***	0.000	33.816***
		(0.512)	(5.851)	(0.000)	(7.408)
	T	1.491	15.815***	0.000	17.918***
		(1.144)	(4.655)	(0.000)	(5.529)
P6	U	10.949***	-2.207	9.668***	-9.668***
		(3.509)	(6.291)	(3.310)	(3.310)
	T	5.103***	15.734**	6.009	12.778
		(1.851)	(7.783)	(2.363)	(11.552)
N		806		539	
r2		0.326		0.327	

Standard errors in parentheses
Notes: All errors are clustered at the subject level. Each cell represents a single coefficient in the regression (or corresponding standard error). The row gives the cell's period and dividend treatment. The column gives its protection treatment. Treatments with no dividends in periods 2, 3, and 5 are excluded.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of cash as a period-3 dividend. The signaling effect present in the untaxed dividend treatment is fully dampened by taxing dividends, which increases the cost of signaling.

In sum, our results are very much consistent with the outcome model as the primary driver of dividend behavior. However, we also find evidence consistent with the proposed signaling rationale for the substitute model of dividends. Also consistent with the substitute model, increasing the cost of signals significantly decreases observed signaling behavior. While we find evidence of both models, dividends are almost two orders of magnitudes higher in the high protection treatments in periods where signaling is inconsistent with rational gameplay (e.g., periods 2 and 5). Dividends are over three times larger in the high protection treatment as the low protection treatment when they can act as signals.

2.3.3 Efficiency Hypotheses

Results from testing the first two sets of hypotheses show that our lab results are broadly consistent with empirical results from financial data in that the outcome model is the primary determinant of dividend payouts. The primary advantage of the experimental approach, though, is that we observe behavior and outcomes that remain hidden in the financial data. Expropriations are difficult to identify from financial data, of course, but so is total investment relative to possible investment: financial data show how much outsiders invested, but not how much they chose *not* to invest with the firm. Both expropriations and non-investment decisions are observable in the lab.

Observing expropriations and investment allow us to test how shareholder protection levels affect the total efficiency induced by low versus high shareholder protection.[16] Expropriations and non-investment decisions are two of the three

[16] Importantly, total efficiency in our laboratory experiment is conditioned on the particular type of remuneration scheme we used in the experiment. In the field, insider remuneration schemes could vary systematically with shareholder protection levels. We found little to no research linking executive compensation to shareholder protection level. This is a promising future line of research for both field and lab data.

sources of inefficiency in the outsider-insider relationship, the third being dividends. Tokens invested with the firm grow by 30% per period. However, tokens allocated to individual accounts do not. Therefore uninvested tokens are leakages from the system or, alternatively, are sources of forgone earnings. Uninveseted tokens, by the outsider, dividend payouts to the outsider in periods 2, 3, and 5, and expropriations by the insider in periods 2, 3, and 5 all lead to forgone earnings.[17]

Figure 2.3 shows the evolution and sources of cumulative forgone earnings for our two main treatments- the low and high protection untaxed dividend treatments- across each of the three channels averaged across all such games. The computations are as follows: Funds kept by the outsider in period 1 do not grow in period 2, and so the only forgone earnings in period 2 are 0.3 times the amount kept in the outsider's individual account in period 1. There is no growth between periods 3 and 4, so we combine these into one period for the purposes of the graph, and there are three potential sources of forgone earnings in period 3-4: forgone growth from lack of investment, forgone growth from funds paid as dividends in period 2, and forgone growth from funds paid as expropriations in period 2. These three sources (non-invested funds, dividend payouts, and expropriations) continue through periods 5 and 6. Thus, the bars in period 2 depict the forgone growth that would have occurred between periods 1 and 2, the bars in 3-4 add to that the forgone growth that would have occurred between periods 2 and 3, the period-5 bars add forgone growth between periods 4 and 5, and the period-6 bars add forgone earnings from growth between periods 5 and 6.

Three patterns emerge immediately from Figure 2.3. First, forgone earnings are substantial, and they are larger for the high-protection treatment than the low one. By inspection, Figure 2.3 rejects Hypothesis 3a: the high protection treatment creates lower surplus than the low protection treatment. We explore this hypothesis in greater detail below, but the difference is striking. Second, in the low-protection treatment nearly all of the efficiency loss stems from non-investment, while all three channels

[17] Note that tokens paid as dividends in period 3 could feasibly be reinvested by the outside in period 4 resulting in no efficiency loss. We almost never observed this behavior in the experiment, though.

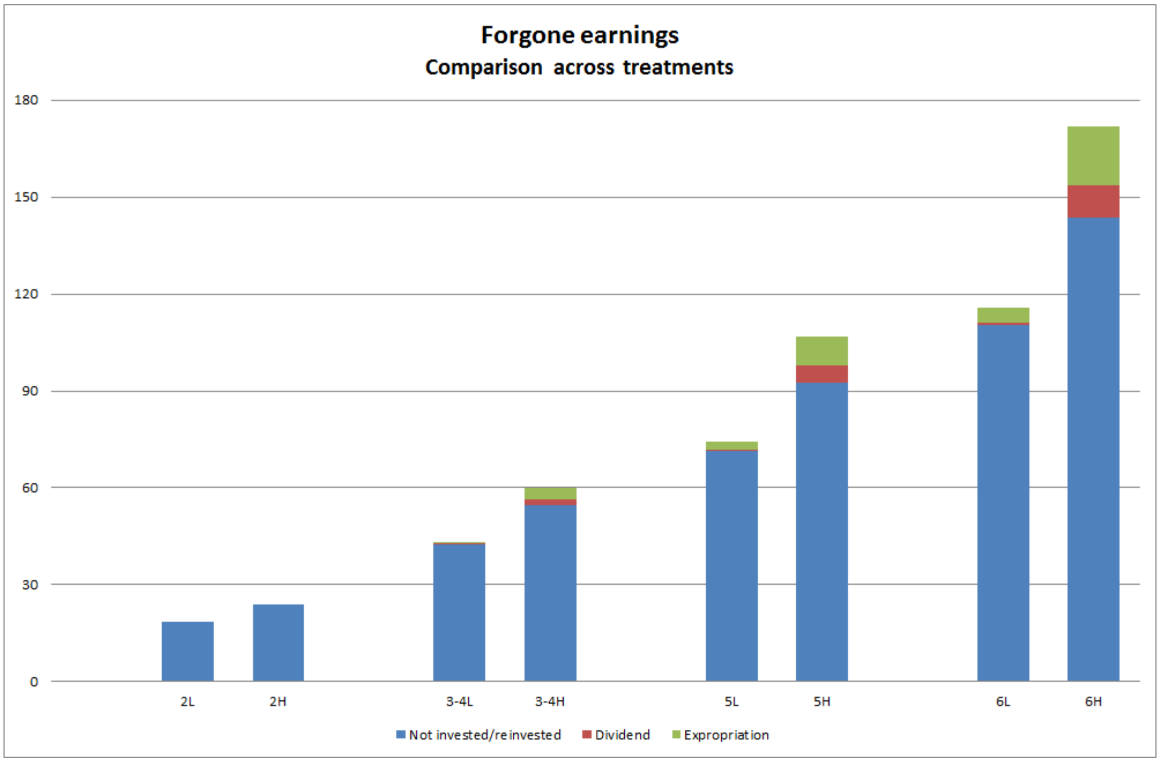

Figure 2.3: Efficiency loss by channel

matter in the high-protection treatment. Third, in both treatments most of the efficiency loss arises from lack of initial investment, and initial investment is sensitive to the investor protection regime.

Figure 2.3 provides evidence that total efficiency is higher in the low protection treatment. Our design allows us to investigate the composition of efficiency leakage across the high and low protection treatments. We first delve into hypothesis 3b, utilizing our experimental design that allows us to detect leakage through each channel, unlike with standard financial data. Hypothesis 3b agnostically postulates that leakages from investment, dividends, and expropriations are the same in high-protection treatments as the corresponding low-protection treatments. That is, for each leakage channel, within a given dividend treatment, the marginal impact of switching from low to high outsider protection is zero. We examine these each in turn, beginning with expropriation.

Figure 2.4 shows the effects of each treatment on expropriations as a percent of cash on hand. Analogous to equation (2.2), we estimate the following equation, now using expropriations over cash on hand as the dependent variable:

$$\begin{aligned} Expratio_{igt} =& \beta_{LN}\mathbf{1}_{t\in\{LN,HN\}} + \beta_{HN}\mathbf{1}_{t=HN} + \beta_{LU}\mathbf{1}_{t\in\{LU,HU\}} + \\ & \beta_{HU}\mathbf{1}_{t=HU} + \beta_{LT}\mathbf{1}_{t\in\{LT,HT\}} + \beta_{HT}\mathbf{1}_{t=HT} + \epsilon_{igt}. \end{aligned} \tag{2.4}$$

As before, the error term is clustered at the subject level. Note that here, because expropriations are possible in all treatments, no-dividend treatments are included. The blue bars in Figure 2.4 represent the estimates for the low-protection treatments, or $\hat{\beta}_{LN}$, $\hat{\beta}_{LU}$, and $\hat{\beta}_{LT}$, respectively. The coefficients on the high-protection terms in equation (2.4) represent the marginal impact of high-protection treaments over respective low-protection ones, so the red bars in Figure 2.4 are instead the sum of each dividend treatment's respective low- and high-protection coefficients. The black bars show each estimate plus or minus one standard error.

In general, an increase in expropriations implies a loss in efficiency. It is clear by looking at Figure 2.4 that, within each dividend treatment, switching from low to high investor protection causes a significant increase in expropriations as a percent of cash on hand. This increase is statistically significant, thus rejecting Hypothesis 3a as it pertains to leakage through the expropriation channel.

In order to see how the switch in protection treatment impacts expropriation leakages, we separate expropriations by period. To test for differences in for expropriations across the H and L treatments within each period, we estimate the period by period analog of equation (2.4):

$$Exp\ ratio_{ipgt} = \Sigma_{p\in\{2,3,5,6\}}\ \Sigma_{t\in\{N,U,T\}} \left[\beta_{Ltp}\mathbf{1}_{t\in\{Ltp,Htp\}} + \beta_{Htp}\mathbf{1}_{Htp}\right] + \epsilon_{ipgt}. \tag{2.5}$$

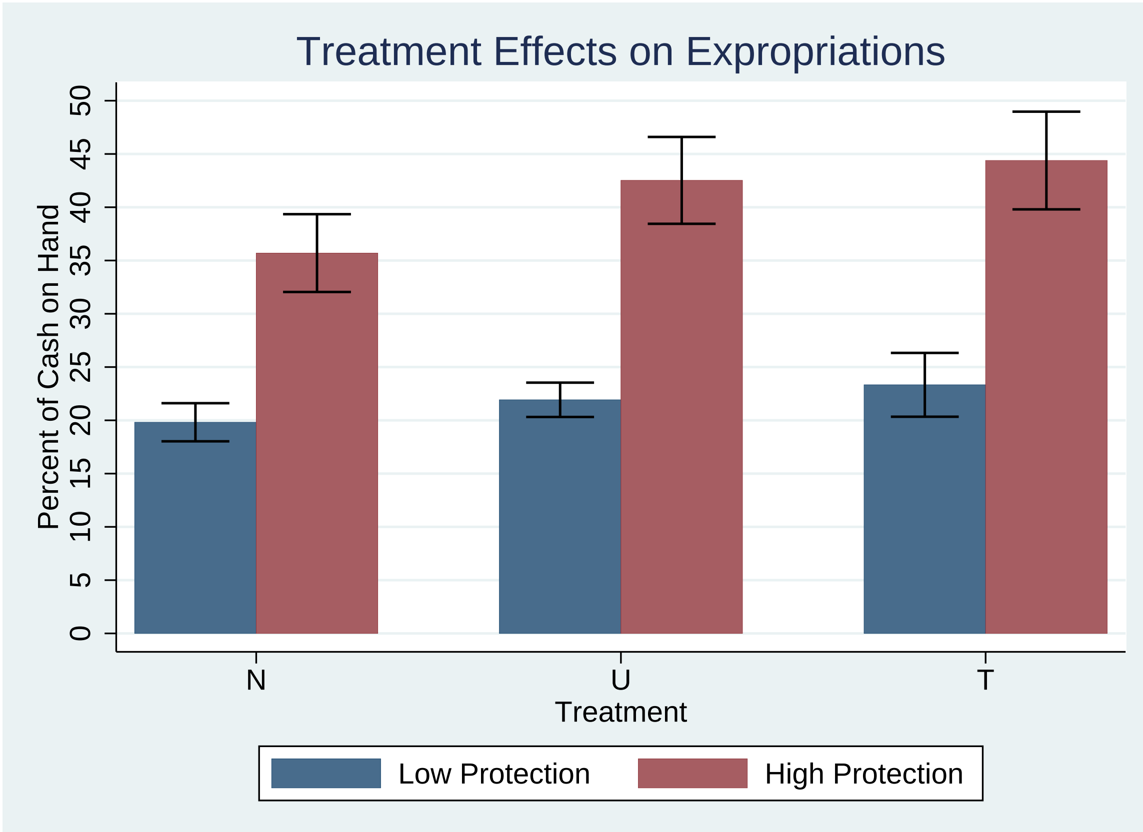

Figure 2.4: Treatment Effects on Expropriations as a Percent of Cash on Hand

The dependent variable in equation (2.5) is again expropriations as a percent of cash on hand. The coefficient of interest for each period and each dividend treatment (e.g., no dividends (N), untaxed dividends (U) and taxed dividends (T)) is β_{Htp}. β_{Htp} describes the marginal effect of moving from the low to high protection treatment on expropriations in dividend treatment t for period p. A significant coefficient rejects Hypothesis 3b: if $\hat{\beta}_{Htp} > 0$ and is significant, the high protection treatment has signficantly higher expropriations in treatment t for period p. As before, we cluster standard errors at the subject level.

Table 2.6 shows the results from estimating equation (2.5) via OLS. The first column of coefficients reports average expropriation rates in the different periods of the low-protection treatments, and the second column reports marginal effects for the high-protection treatments. The third and fourth columns demonstrate robustness to excluding the first game of each treatment.

Expropriation rates are positive and significant in every period of every treatment, and high protection leads to significantly more expropriation in all but three instances, period 3 of the no-dividend treatments, period 6 of the untaxed dividend treatments, and period 6 of the tax treatments. The picture that emerges from Table 2.6 is that insiders expropriate more when there is more investor protection.

Perhaps the most surprising result from Table 2.6 is that high protection leads to more expropriation even in the baseline treatments. In the baseline treatments, dividends can only be paid in period 6. Because the insider has primary claim to funds if both the insider and outsider claim them simultaneously in period 6, the outsider has no real control over allocations in either of these no-dividend treatments. Still, the insider expropriates in every period of the low protection treatment. This implies a similar impatience of the insider as observed of the outsider in the period-level dividend regressions in the previous subsection. This impatience is amplified in the high protection treatments: $\hat{\beta}_{HU2}$, $\hat{\beta}_{HU5}$ and $\hat{\beta}_{HU6}$ are all significant and large in magnitude. Providing the outsider with marginally more control leads to a substantial increase in insider expropriation behavior.

The difference in the marginal impact of introducing (untaxed) dividends in the high-protection treatment and that in the low-protection treatment can be more clearly seen using a standard difference-in-difference estimator. Again, it is split apart by periods:

$$\begin{aligned} Exp\ ratio_{ipgt} = \sum_{p\in\{2,3,5,6\}} \Big[\alpha_p + \beta_{Up}\mathbf{1}_{t\in\{LUp,HUp\}} + \\ \beta_{Hp}\mathbf{1}_{t\in\{HNp,HUp\}} + \beta_{DDp}\mathbf{1}_{t=HUp}\Big] + \varepsilon_{igpt} \end{aligned} \tag{2.6}$$

As before, the dependent variable is expropriations as a percent of cash on hand. There is now a dummy that provides a baseline for each period, α_p, as well as period-specific

Table 2.6: Treatment Effects on Expropriations (as a Percent of Cash on Hand) by Period

Trimming:		None		1, 4, 7	
		L or H	H	L or H	H
	N	7.025**	18.751**	5.808	17.475**
		(3.080)	(7.001)	(3.966)	(8.100)
P2	U	6.395**	31.799***	3.642	33.816***
		(2.974)	(5.459)	(2.969)	(7.941)
	T	12.546***	27.363***	9.142**	30.818***
		(4.602)	(7.140)	(4.428)	(8.248)
	N	7.422***	6.348	5.606*	7.214
		(2.603)	(4.327)	(3.099)	(4.993)
P3	U	6.715***	17.358**	6.228**	23.800***
		(2.011)	(6.926)	(2.593)	(8.327)
	T	7.283**	17.635***	6.269*	17.762**
		(3.207)	(4.875)	(3.182)	(6.794)
	N	0.927*	27.937***	0.476	30.407***
		(0.520)	(5.468)	(0.451)	(8.289)
P5	U	3.689***	45.818***	2.229*	44.132***
		(1.025)	(6.610)	(1.241)	(8.975)
	T	7.167**	51.252***	4.411	55.440***
		(2.768)	(7.594)	(2.974)	(8.279)
	N	63.319***	17.506**	65.622***	21.073***
		(5.623)	(6.766)	(6.145)	(7.689)
P6	U	68.686***	2.609	70.268***	3.194
		(4.631)	(9.582)	(4.633)	(11.827)
	T	66.533***	-2.578	72.852***	-13.078
		(6.198)	(9.771)	(7.187)	(12.518)
N		1256		862	
r2		0.628		0.642	

Standard errors in parentheses
Notes: All errors are clustered at the subject level. Each cell represents a single coefficient in the regression (or corresponding standard error). The row gives the cell's period and dividend treatment. The column gives its protection treatment.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7: Treatment Diff-in-Diff by Period on Expropriations over Cash on Hand

	Trimming:	None	1, 4, 7
		(1)	(2)
	P2	25.843***	29.986***
		(6.794)	(9.386)
	P3	23.407***	30.433***
Diff-in-Diff		(7.706)	(9.020)
	P5	36.775***	32.702***
		(8.053)	(11.796)
	P6	-58.396***	-64.048***
		(9.713)	(13.046)
N		869	612
r2		0.321	0.342

Standard errors in parentheses
Notes: All errors are clustered at the subject level. Tax treatments are excluded. The first column uses the full (remaining) sample. The second trims off games 1, 4, and 7.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

unconditional marginal impacts of both untaxed dividends, captured by β_{Up}, and high outsider protection, captured by β_{Hp}. The coefficients of interest are the period-specific diff-in-diff estimates, β_{DDp}.

Table 2.7 shows the estimates for the diff-in-diff coefficient for each period. Column 1 shows the results for the full sample (excluding tax treatments), and column 2 trims off games 1, 4, and 7 for each individual as a robustness check. Standard errors are again clustered at the individual level.

The diff-in-diff coefficient represents the difference in the marginal impact of introducing untaxed dividends in the high-protection treatments relative to that impact in the low-protection treatments. For periods 2 and 5, this coefficient measures the difference in the degree of impatience created by dividends between the high-protection and low-protection scenarios.[18] The first row shows that, for period 2,

[18] For period 3, low-protection dividends serve as a trust signal, while high-protection dividends do not. This difference is not present in periods 2 and 5, and so while impatience may play a role in

introducing dividends in the high-protection scenario increases expropriation (as a percent of cash on hand) by a highly significant 25 percentage points more than in the low-protection scenario. In period 5, the high-protection increase is a highly significant 36 percentage points higher than the low-protection increase.

The difference in increased impatience can potentially be explained as a backward unraveling that is present in the high-protection treatment but not in the low, although this still does not explain the early expropriations in the no-dividend treatments seen in Table 2.6. The mechanics and implications of the unraveling are discussed in Section 4.

Whereas estimating equations (2.5) and (2.6) above tests hypothesis 3.b in the context of expropriations, we estimate the following equation to test hypothesis 3.b for efficiency losses due to underinvestment:

$$Tokens_{igt} = \Sigma_{t \in \{N,U,T\}} \left[\beta_{Lt} \mathbf{1}_{t \in \{Lt,Ht\}} + \beta_{Ht} \mathbf{1}_{Ht} \right] + \epsilon_{igt}. \tag{2.7}$$

We estimate equation (2.7) using "tokens" as a measure of efficiency in three different ways– initial investment in period 1, the total number of tokens at the end of the game, and the cash on hand in period 6 after interest is earned but before any allocation decisions are made. As before, in each case we estimate average efficiency across the the low and high protection treatments in each dividend treatment t as β_{Lt}. The coefficient of interest is β_{Ht}, which is the marginal effect of moving to the high protection regime within a dividend treatment t. Each of these measures generates one observation per game.

Table 2.8 presents the results from estimating equation (2.7) by OLS with standard errors clustered at the subject level for each measure of efficiency: initial tokens invested, total combined tokens achieved, and final cash on hand. The first group of columns uses the full sample, while the second group trims out the first game in

explaining the period-3 coefficient, it may not be the sole factor. There is still a highly significant 23 percentage point difference in period 3, though.

Table 2.8: The Effects on Initial Investment, Total Funds, and Final Cash on Hand

Trimming:		None			1, 4, 7		
		(1)	(2)	(3)	(4)	(5)	(6)
		Init Inv	Total Tokens	Final COH	Init Inv	Total Tokens	Final COH
N	L or H	38.365***	170.469***	108.479***	38.167***	170.889***	109.847***
		(7.062)	(13.246)	(20.260)	(8.007)	(14.891)	(22.947)
	H	-18.385**	-42.802***	-74.958***	-24.111***	-51.403***	-87.292***
		(7.829)	(14.136)	(21.345)	(8.570)	(15.597)	(24.171)
U	L or H	38.167***	170.438***	107.802***	37.764***	171.278***	110.417***
		(7.418)	(13.270)	(20.310)	(8.316)	(15.035)	(23.086)
	H	-17.323**	-55.479***	-101.229***	-19.278**	-59.694***	-106.903***
		(8.289)	(13.528)	(20.490)	(9.062)	(15.164)	(23.126)
T	L or H	30.958***	153.198***	82.740***	26.528***	147.236***	75.292***
		(5.733)	(11.247)	(17.967)	(5.604)	(10.722)	(17.146)
	H	-11.885*	-40.635***	-74.396***	-10.611	-36.500***	-68.069***
		(6.464)	(11.527)	(18.258)	(6.379)	(11.046)	(17.556)
	N	576	576	576	432	432	432
	r2	0.432	0.878	0.434	0.403	0.874	0.427

Standard errors in parentheses
Notes: All errors are clustered at the subject level. The first group of columns uses the full sample. The second group trims off games 1, 4, and 7.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

each treatment. Outsiders can invest anywhere between 0 and 100 tokens in the first period, and combined payoffs can range from 100 if the outsider invests nothing to 286 if the outsider invests everything and there are no subsequent leakages.[19]

The first column of Table 2.8 shows that in our main untaxed dividend treatment, outsiders initially invest an average of 38 of their 100 tokens in the low-protection environment, but invest only 20 tokens in the high-protection one. This occurs despite the fact that they have more ability to recover their investment through their own actions in the high-protection treatments. Since $\hat{\beta}_{HU} = -17.32$, which represents the marginal impact of switching from low to high protection given untaxed dividends,

[19]In tax treatments, the lowest possible number of tokens achievable is 95, although we never observed this outcome.

and is significant at the 5% level, we reject the null hypothesis that the low and high protection treatments have the same efficiency level through the investment channel.

Comparing coefficient estimate rows 1 and 2 to rows 3 and 4, initial investment is identical in corresponding no-dividend and untaxed dividend treatments. In particular, outsiders invest about 38 tokens in both the LN and LU treatments, and invest about 20 tokens in both the HN and HU treatments. The switch from low protection to high protection only changes who makes the dividend payout decision. When these dividends are possible throughout the game, it makes sense that this change could impact investment decisions; it could affect the size and timing of later expropriations, potentially affecting the cash on hand available for dividends, thus decentivizing initial investment.

When dividends are only available as a means of final allocation, however, the reasoning for the decrease in initial investment is less clear. The likely explanation is that the switch from low to high outsider protection induces the attitude of a race to withdraw firm value sooner, even if it does not directly create one. Regardless of the reason, as shown in Table 2.8, the marginal impact of switching from low to high protection given only period-6 dividends, $\hat{\beta}_{HN} = -18.39$, is significant at the 5% level (when naïve and exposed games are pooled). Thus, we reject the null for hypothesis 3b when looking at the impact on initial investment for no-dividend treatments, as we did for dividend treatments.

Table 2.8 also shows that taxes reduce investment, although not significantly. For the low-protection treatment, taxes reduce initial investment by 7 tokens ($p = 0.101$), and in the high-protection treatment, taxes reduce investment by only 2 tokens ($p = 0.736$). There are several reasons why this could be the case. For example, the dividend treatments allow dividends in periods 2, 3, 5, and 6, while the no-dividend treatments prohibit dividends until period 6, essentially making tax rates infinite for the first three dividend periods. The tax treatments fall in between by making dividends expensive, but not impossible, in periods 2, 3, and 5. Thus there is some evidence for a behavioral result. Alternatively it could be that outsiders are risk averse and the tax treatments lead to uncertainty over which type of gameplay (i.e., a no-dividend equilibrium or a dividend equilibrium) will dominate. Further theoretical and empricial work are needed to investigate this result.

The second column of Table 2.8 reports the same regression for the subjects' combined final payoffs instead of initial investment. Reduced initial investment inherently leads to lower combined payoffs because only invested funds earn interest, but the regression on total tokens allows for a comparison that takes into account all leakages (i.e., including uninvested dividends and expropriations) through the first five periods.

Hypothesis 3a proposes the agnostic null that the total number of tokens generated is the same in both the low-protection and high-protection treatments within each dividend treatment. This implies that total leakages combined across all channels is the same across corresponding low-protection and high-protection treatments.[20] The literature has looked empirically at how dividends are affected, but that is just one channel through which leakages occur. Because investment counterfactuals and expropriations are generally not observable, overall efficiency has not been studied using field data. Our experiment, however, allows us to capture leakages through each channel, as well as overall efficiency. This is perhaps the biggest contribution our experiment makes to the literature.

Switching from low to high outsider protection causes a highly statistically significant decrease in the total number of tokens in the game within each dividend treatment. This corroborates the results surrounding hypothesis 3a, which showed that this switch decreased initial investment, increased premature dividends, and increased premature expropriation within all three dividend treatments. An increase in leakages through all three channels should, and does, yield a decrease in overall tokens.

[20]This does leave open the possibility that leakages from some channels increase while those from others decrease, as long as the total leakage is unchanged. The results for Hypothesis 3b examine the effects on each channel individually.

Comparing across dividend treatments, it is interesting to note a similarity between total tokens and initial investment. As before, the total number of tokens is nearly identical in the no-dividend and untaxed-dividend treatments within the low protection regime. However, while the decrease resulting from switching to the high-protection regime had the same marginal impact on initial investment for both of those dividend treatments, the decrease in total tokens is much more pronounced given untaxed dividends than when given no dividends. Each of the tax treatments again produced the smallest results within each respective protection regime.

The third column reports the same regression for cash on hand at the beginning of period 6, measured after interest was earned but before dividends or expropriations, which we interpret as the final value of the firm. Periods 6 and 7 are devoted solely to disbursing this cash, through dividends and expropriations in period 6 and the default split of remaining cash in period 7. Both untaxed low-protection treatments enabled firms to grow to a statistically-significant 108 tokens, but the untaxed high-protection treatments dropped these by a highly-significant 75 tokens in the baseline treatment and by a highly-significant 101 tokens in the dividends treatment. Thus, high investor protection reduces firm value, which contrasts with the empirical findings of La Porta et al. (2002). Taxes also reduce firm value in the low-protection treatment, but the reduction is not statistically significant.

2.4 Discussion and conclusions

Like much of the empirical literature, our experimental findings support the original results from La Porta et al. (2000) favoring the outcome model over the substitute model of agency motivations for dividend policy. In our study, dividends are higher in the high-protection treatments than in the corresponding low-protection ones. In the La Porta et al. (2000) setting, this was interpreted as outsiders demanding dividends to reduce insider access to cash on hand and thereby reduce expropriations, which is unobservable in their data.

A close look at the games corresponding to the low-protection and high-protection treatments reveals an alternative reason why this might occur. The low-protection treatments closely resemble the standard trust game, while the high-protection treatments more closely mirror the centipede game. Switching to the high-protection treatment gives the outsider more control, thereby reducing the amount of trust they must put in the insider. However, when this happens, both sides must now simultaneously trust each other to not remove tokens from the firm. This leads to a backward unraveling of payouts, meaning more of both expropriations and dividends are demanded prematurely, ultimately stifling initial investment.

The effects of this unraveling are thus twofold: the outsider reduces the total potential value of the firm by initially investing less, and both parties then race to remove the firm's funds before the other can. Initial investment (relative to outside options) and insider expropriation generally cannot be observed in financial data. In our experiment, however, these are observable. Both of the unraveling effects in turn *increase* insider expropriation, the opposite of what La Porta et al. (2000) posit as the outsider's intent. Holding initial investment constant, moving expropriations sooner means an identical expropriation comes from a smaller pot, increasing the expropriation percentage. Holding the expropriation timing constant, a decrease in initial investment implies a smaller pot in each period, thus lowering the expropriation percentage. While the outsider's intent may be to decrease insider expropriations by reducing insider access to cash on hand, in doing so this decreases firm value and thereby increases insider expropriation.

The effects to the pair of agents jointly are then threefold. Lower initial investment reduces the firm's potential earnings regardless of dividend and expropriation behavior. As discussed in Section 2.1, this falls in line with the literature showing a negative correlation between shareholder protection and firm size. As a percent of cash on hand, which is now smaller, both dividends and expropriations naturally increase. Additionally, the unraveling affect creates an increase in *premature* dividends and expropriations, both of which in turn hurt overall efficiency.

Indeed, the major discovery in this paper comes not from the new exploration of whether dividends are higher in high investor protection markets, but from the effect of investor protection on investment. While the existence of dividends has no impact on initial investment within either low- or high-protection treatments, switching who sets dividends from the insider to the outsider leads to a drop in initial investment of about 17 percentage points, which is close to half of the amount invested in the low-protection treatments.

Losses continue throughout the relationship. If the outsider invests the entire 100-token endowment and the two parties leave it untouched throughout the relationship, they can share 286 tokens at the end. In the baseline low-protection treatment they are able to achieve 60% of this amount, and allowing the insider to pay dividends throughout has no impact on this result. However, combined earnings fall by a quarter when we go from the low-protection baseline to the high-protection one, and they fall by a third when the outsider can demand dividends throughout the course of the relationship.

This finding raises the question of whether strong investor protection laws really are beneficial for economic growth. In our experiment, the subjects in the high-protection treatments left the lab significantly less wealthy than those in the low-protection ones. Most studies relating wealth to legal protection look at overall protection, not just investor protection. Countries with strong investor protection likely also have strong legal systems, because a strong general set of laws probably precedes the adoption of strong specific ones to protect shareholders. Clever identification strategies would be needed to untangle the answer using country-level data, but the laboratory data suggest that it would be worth the trouble.

Beyond the standard context, our experimental design and results can be exported to fit any situation where one party must devote funds to a long-term relationship with opportunities for one or both parties to withdraw them early. Such situations arise in many situations. Industrial R&D projects can take years to complete, always with the chance that the company will pull the plug before the project reaches fruition.

Governments also face long investment projects, such as building nuclear power plants, designing new military weapons, fighting climate change, or sending humans to Mars. Taxpayers must pay for these projects as they go along, but different parties have both the opportunity and the incentive to divert funds to other uses. Different settings have different implications for the timing of surplus generation, but our current experiment suggests that long-term relationships are more successful when the payer of the funds has fewer opportunities to reclaim them.

Chapter 3

Better Lucky than Good: The Role of Information in Other-Regarding Preferences

3.1 Introduction

Theoretical models accounting for other-regarding preferences (ORP) have been shown to more accurately predict observed equilibrium behavior, especially in a controlled laboratory setting. Seminal work by Andreoni (1989) and Andreoni (1990), for example, illustrates how the existence of even one player with warm glow utility in a public goods game can alter the theoretically predicted equilibrium such that it more accurately predicts observed behavior. More recently, within the context of personnel economics, Neilson and Stowe (2010) examine how status-seeking preferences and inequality aversion affect workers' equilibrium effort levels, in turn impacting employers' equilibrium piece-rate contracts.

Early work seeking to illustrate this unselfishness experimentally was designed specifically to induce such behavior. For this reason, many early experiments, such as those by Kahneman et al. (1986) and Hoffman et al. (1996), gifted dictators the money that they were choosing how to split. However, individuals often treat money differently depending on what "account" they consider it to be in, as illustrated by Thaler (1985). For instance, money endowed by luck may be treated differently than that earned by exerting effort. Perhaps most cleanly demonstrated by Cherry et al. (2002), dictators tend to give less liberally when they've worked to earn their money. The experimental literature is silent, however, on how dictators act when the source of income is unclear.

More intuitively, consider a dictator who is unsure what portion of income the recipient actually *earned* relative to how much was effectively gifted, even if the total income is known. Analogous situations arise every day: A beggar on the street who claims to be a wounded Vietnam War veteran surely gets more donations than one who claims to be lazy, whether or not the claim is true. Students claim to be sick, even if they're actually skipping class to surf, knowing their story is relatively unverifiable. A car salesman, who is paid a piece-rate, may get a small paycheck one month either because he was unlucky or because he was lazy. His friend or coworker is much more likely to buy him a beer after work if he was unlucky rather than lazy, although that might not be perfectly observable. The role of information, or the lack thereof, is potentially large here.

I first build a theoretical model to support relevant hypotheses. I then test how individuals act when others' income sources are obscured relative to a full-information baseline. I build a laboratory experiment where income is partially earned and partially determined by a stochastic luck shock. In doing so, I introduce a novel real-effort task, which I argue more accurately captures effort than most commonly used tasks. Income differences naturally arise, and these in turn naturally vary with regards to source composition.

Subjects are then randomly and anonymously paired for a one-shot dictator game. Control dictators know the breakdown of the recipient's income by source, while treatment dictators know only the recipient's total. Comparing control donations to those in the treatment shows whether uninformed dictators treat recipients' total income similarly to how perfectly informed dictators treat earned income or to how they treat endowed income. That is, I examine whether, in the face of obscured information, dictators assume recipients are lazy or give them the benefit of the doubt. For those who do exhibit some degree of other-regarding behavior, I find support of the latter.

In order to test the difference between how partially and fully informed dictators treat recipients, the experiment requires a design that allows for variation in income sources. Income is gained and potentially lost through two rounds preceeding the dictator game. Subjects first have thirty minutes which they can split between browsing the internet and earning tokens by performing a real-effort task. The outside option of surfing the internet allows effort to vary on the external margin. The real-effort task used, solving CAPTCHAs, is a novel approach, and it allows effort to vary greatly on the internal margin.[1] Additionally, because earnings are based solely on one's own performance, there is no possibility of inducing a competitive spirit, as may arise in the tournament payoff structures commonly used (for example, by Erkal et al. (2011)).

I argue that solving CAPTCHAs is a more accurate measure of effort than many other commonly used real-effort tasks. There is likely little variation in typing skills for college students, the pool from which the experimental subjects were drawn. All remaining variation in earnings must thus stem from differing levels of effort. The extent to which the task captures effort is also likely obvious to all subjects, which is necessary. In this way, I also contribute to the mechanical design aspect of the experimental literature as a whole.

To motivate each hypothesis, I turn to both theory and empirics. Theoretical ORP models, most notably those of inequality aversion by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), fail to distinguish between earned income and that bequeathed by luck. I expand this literature by adapting the Fehr and Schmidt (1999)

[1] CAPTCHAs are text that have been distorted in a way that renders them unreadable to computerized text scanners. See Figure 3.1 for an example.

model ("FS") to account for this distinction. I then set up the dictator's objective function and to motivate and formalize the hypotheses.

Due to the nature of the problem and resulting model, the theory is agnostic as to predicting which competing null hypothesis will be observed in the lab. The hypotheses are mutually exclusive, and each is supported by a different economic literature.

The moral "wiggle-room" concept traces its roots to work done by Dana et al. (2007). They create uncertainty in the relationship between subjects' actions and the subsequent outcomes. This allows for subjects to self-justify less kind actions via that disconnect.

Extending that intuition to the context of my experiment, partially informed dictators may potentially use that lack of information as a means of creating a similar disconnect. They can tell themselves "I'm sure he was lazy, not unlucky" to self-justify lower donations. This is the moral wiggle-room hypothesis.

In contrast, the social insurance literature suggests that, in the absence of perfect information, individuals may assume others are unlucky rather than lazy in hopes that others would judge them similarly if the situation were reversed. Social insurance is a much more widely studied concept, likely because it can be seen impacting everything from charitable donations to tax policy (see section 3.2 for more background and discussion). In the context of this experiment, then, the social insurance hypothesis posits that partially informed dictators will treat recipients' total income just as fully informed dictators treat recipients' endowed income.

There also exists a third possible hypothesis. Dictators and recipients all know the distribution of random shocks to income that serve as the luck component. The fact that this is common to all subjects is also fully known. Because these shocks can be negative but total income cannot, partially informed dictators can back out conditional expectations of recipients' earnings and luck shocks in some cases. Perhaps more in line with standard rational expectations theory than any other economic literature, it may be that partially informed dictators break apart recipients' total income into the expected values of earnings and endowment, conditional on recipient total income, and then give donations as if these represent actual recipient earnings and endowments. This is the conditional expectations hypothesis.

Section 3.2 examines the relevant literatures, including those supporting each hypothesis. Section 3.3 describes the experimental design and introduces the relevant theory to motivate and formalize the hypotheses. Summary statistics and empirical results are shown and discussed in section 3.4. Section 3.5 concludes.

3.2 Related Literature

Each proposed hypothesis here is motivated by a different economic literature, which is uncommon. As mentioned before, Dana et al. (2007) first propose the concept of moral wiggle-room. In the laboratory, they contrast a baseline in which dictators' actions directly affect recipients' outcomes with treatment sessions where the two are less than perfectly linked. The disconnect created allows self-justification by the dictators by creating some moral wiggle-room. The application here extends their notion into the context of uncertainty over how others' income was acquired, either through effort or luck, creating the foundation for the moral wiggle-room hypothesis.

Social insurance is the notion that a luckier individual is willing to give to a less lucky one in the hopes that if the situation were to reverse in some future time period that the reverse transaction would occur. There is a much larger literature on social insurance, largely because it applies to and is observable in a wider variety of settings. For example, Varian (1980) was the first to illustrate that income redistribution taxes serve as a form of social insurance. Given the size of the literature on redistributive taxes, it likely provides the deepest look at social insurance.

From a theoretical perspective, Alesina and Angeletos (2005) build a model showing that two nations with differing beliefs about how much luck affects income, and which base their redistributive tax schemes around those beliefs, each create a self-fulfilling system. The result is that each nation's own *ex ante* beliefs in turn cause those beliefs to become an *ex post* reality.

Building on this concept, there have been a few related experimental studies in the tax and redistribution literature. Balafoutas et al. (2013) modify a traditional public goods game to allow for unequal endowments, which are either earned by performance on a general knowledge quiz or determined exogenously, depending on the treatment. In half the sessions, subjects then vote on the percentage of income that they must put into a public fund to be redistributed equally. Despite quiz performance (arguably) being largely correlated to exogenous ability, and hence not a great real-effort task, they do find evidence that there is less support for redistribution and less cooperation when tokens are earned. I attempt to mitigate this potential design criticism in my experiment by implementing a task which I argue is more purely a measurement of effort.

Other studies, such as Krawczyk (2010), have used *ex ante* redistribution voting based on variation in randomly assigned "probabilities of winning" (*i.e*, of getting the higher payout) to capture fairness attitudes. Similarly, Eisenkopf et al. (2013) find evidence that unequal access to study material before a quiz (*i.e.*, random and unequal *ex ante* probabilities of success) which determines income elicits redistribution attitudes comparable to those found when luck alone determines income. Krawczyk (2010) argues that the intuition of such approaches hinges on luck determining how upwardly mobile an individual is.[2] However, luck also affects many other aspects of income that occur later in life. I argue that my experiment more accurately represents realistic luck shocks by allowing luck to potentially act as a negative shock rather than either solely as a proxy for income mobility, as Krawczyk (2010) does, or as a necessarily positive rank/endowment determinant, as Balafoutas et al. (2013) do.

There is some evidence in the literature, though, that suggests income determinants do not affect distribution/redistribution attitudes. Cabrales et al. (2012) design an

[2] This idea that fairness of distribution should be based on one's equality of opportunity dates back at least as far as Alexis de Tocqueville.

experiment of repeating rounds where, within each round, subjects decide whether to exert high or low effort. While there is a non-zero probability that choosing high effort will result in a low payout, the converse is not true, which allows for free-riders in the subsequent public goods game. The game is characterized by a voting mechanism over redistribution schemes. They find that the threat of a Hobbesian low payoff environment[3] is not enough to sustain a Rousseaunian social insurance policy,[4] even when the effort is replaced with randomness. The design still relies on a voting mechanism, however, which may not fully capture other-regarding concerns. I contend that my experiment is much more broadly applicable by allowing for agents to make more consequential decisions for which they know the outcomes with certainty.

In this vein, Erkal et al. (2011) abstract away from the voting environment and instead allow for voluntary donations in a public goods game. After a within-group tournament determines relative earnings based on performance in solving codes using an encryption table, there is a potential for each individual to lose 30% of their earnings due to a shock before the public goods game takes place. This shock is only severe enough to switch the ordering between the first and second of the four participants in a given group. They find a significant difference in first- and second-ranked players when performance determines rank/income, but not when randomness does. One potential confound, however, is that there may be a bias towards decreased donations in the earnings treatment because the earned payouts are determined by a tournament, which may artificially induce a spirit of competitiveness early in the experiment. My experiment avoids this potential bias by not allowing any player's earned income to affect any other players' during the real-effort stage.

While the discussion surrounding income redistribution is certainly applicable to the general disparity in attitudes surrounding effort and luck, it is a particularly relevant example in another sense, too. The disparity of beliefs can be exacerbated in

[3]In a Hobbesian world, any "sucker" who exerted effort would be exploited by free-riders.

[4]In a Rousseaunian world, players create a social contract where redistribution and high effort are compatible in order to avoid the threat of a reversion to the Hobbesian world. This is analogous to sustained cooperation in a repeated prisonners' dilemma game.

reality because empirically disentangling the relative roles of effort and luck is difficult. This situation highlights how, given the difference in attitudes between effort and luck, such lack of information can yield two drastically different outcomes. If one buys that wealth determinants in the two regions are identical (or at least nearer than the respective tax code differences would suggest), then *at least* one region must have incorrect beliefs. Thus, there is a cost to social welfare for having imperfect information and a corresponding benefit to truthfully revealing that information.

The concept of social insurance applies beyond the tax structure Varian (1980) examines, however. More broadly defined, any act of income redistribution in which luckier individuals are systematically giving to less lucky ones in the presence of future uncertainty is a means of social insurance. Many less coordinated acts are also forms of social insurance, such as charitable donations. Cairns and Slonim (2011) examine charitable donations taken during Sunday Masses at a Catholic Church to examine substitution patterns across charities. While they stop short of examining the differences between more luck-driven causes (*e.g.*, the 2005 Indonesian Tsunami) and arguably less luck-driven ones, they certainly hint at such disparities and likely could have at least done some rudimentary empirical examination. Indeed, from a mechanical standpoint, charitable donations are perhaps the most directly related to the experiment used here.

3.3 Experimental Design

The experiment has 3 stages. Each stage's instructions are given to subjects just prior to the beginning of that stage. There is a control group and a treatment group, which allows for identification of the average treatment effect. There is also natural within-treatment variation for both the control and the treatment, as discussed below, which allows for examination of within-treatment effects.

In Stage 1, each player is given 30 minutes on a computer with access to the internet. The web browser is open for them when the time starts, and they may

Figure 3.1: Example of a CAPTCHA

allocate their time however they wish between two activities: freely surfing the internet and completing a real-effort task for payment.

The real-effort task is solving CAPTCHAs. CAPTCHAs are text that have been distorted in a way that renders them unreadable to computerized text scanners (see Figure 3.1 for an example). CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) are employed by many websites as a method to ensure that humans, not "bots," are responsible for certain actions on their sites. Solving a CAPTCHA is required to do many common online activites, such as buying concert tickets or creating email accounts.

This experiment uses Google's reCAPTCHA service.[5] For every 5 CAPTCHAs a player solves, she earns 1 token (without fractions or rounding). At the end of the experiment, tokens were exchanged at a rate of 2 tokens to $1, rounded up. Each subject has a real-time count of the number of CAPTCHAs she has solved and of the number of tokens she has earned shown on her screen. All earnings information is completely private at this stage.

An important feature of CAPTCHAs is that solving them fits the definition of a real-effort task extremely well. There is very little variation in the skill required to solve CAPTCHAs,[6] as long as subjects are moderately sufficient typists. Given

[5] The reCAPTCHA service, the most widely used CAPTCHA service on the Internet, also offers many societal benefits. It gives users two words to solve per CAPTCHA, where one is a known control word and the other is a treatment word unknown to the program. There are many projects around the world that attempt to digitize old print magazines, newspapers, and books using digital scanning software. When this software fails to recognize a word, it becomes a treatment word for reCAPTCHA. After a threshold of users who correctly identify the control word have matching answers for the treatment word, their answer becomes accepted as that word. Any software able to beat the reCAPTCHA system, then, is also an advance in digital scanning software, which can then enhance old print digitalization.

[6] See http://www.google.com/recaptcha/static/reCAPTCHA_Science.pdf.

the young and fairly homogeneous age group of the subjects, there is likely very little heterogeneity in typing ability. Thus, payment is rewarded almost exclusively on effort. This is in contrast to more traditional alternatives, such as word hunts, photo hunts, or mazes. Many of these are less routine skills and depend more on individual ability, which is more a product of luck rather than effort.

Additionally, solving CAPTCHAs is a moderately realistic analogy to a workplace environment. Firms that wish to gain access to sites blocked by CAPTCHAs have increasingly resorted to paying humans to solve them in order to bypass the filter.[7]

Stage 2 adds an element of luck to each player's earned income from the first stage. Each player faces the same lottery of 5 equally likely outcomes of the random variable ℓ_i. The support of ℓ_i is $L_i = \{-20, -10, 0, 10, 20\}\ \forall\ i$. Each player's realized value of ℓ_i tokens is added to (or removed from) her Stage 1 earnings ($w_i = e_i + \ell_i$). This total "wealth" of tokens is carried into Stage 3. The only common knowledge at this point is that each player has faced the same initial task and the same subsequent lottery.

Stage 3 begins by randomly pairing all players into i,j pairs to play a one-time dictator game. Each player in the pair is assigned as either dictator (i) or recipient (j) such that the dictator is the wealthier of the two ($w_i > w_j$).

In control pairs, each dictator knows her own wealth and its breakdown by earnings and luck. She also knows her matched recipient's wealth and its breakdown. In treatment pairs, however, while each dictator does know her own wealth and its breakdown, she knows only her matched recipient's wealth, but not its breakdown.

In both the control and the treatment, the recipient has full knowledge over both players' incomes and income sources, although the dictator doesn't know this. This information allows the recipient to play a hypothetical (non-incentive-compatible) dictator game by specifying how much he would give if he instead were the other

[7]For example, see the 2010 New York Times article on the subject: http://www.nytimes.com/2010/04/26/technology/26captcha.html?_r=0.

player. This maintains anonymity of roles in the lab, since otherwise only dictators would be using the keyboard and mouse during this stage.[8]

This design gives rise to only one observation per pair of subjects, which can be expensive. Unfortunately, this is by necessity; if players repeated the game, then it would allow for dictators to diffuse responsibility. That is, a dictator may self-justify giving less by convincing herself that even if the recipient is unlucky this time, *surely* he will be luckier next time. The impact of luck is thus mitigated when multiple draws take place. This potential self-justification is the intertemporal analogy to the contemporaneous imperfect information effect of the moral wiggle-room concept that the across-treatment variation is designed to examine. However, this cannot be perfectly known when looking forward across time, as it can be across contemporaneous treatments.

3.3.1 Theory and Hypotheses

The model constructed here is an adaptation of the commonly used model of inequality aversion originally developed by Fehr and Schmidt (1999) (the "FS" model).[9] Let there be n players indexed by $i \in \mathbb{Z}$ s.t. $1 \leq i \leq n$. Let $\mathbf{w} = (w_1, ..., w_n)$ represent a vector consisting of each player i's income, and let $U_i(\mathbf{w})$ be player i's utility, which takes the form

$$U_i(\mathbf{w}) = w_i - \alpha_i \frac{1}{n-1} \sum_{j \neq i} \max\{w_j - w_i, 0\} - \beta_i \frac{1}{n-1} \sum_{j \neq i} \max\{w_i - w_j, 0\}. \quad (3.1)$$

[8]The CAPTCHAs in Stage 1 were programmed using Parse. The client side of the program is located at http://uteconexperiment.herokuapp.com/. Stages 2 and 3 used z-Tree (Fischbacher (2007)).

[9]Note that while inequality aversion may not actually be the channel through which ORP are operating, the point of this paper is not to identify the "true" channel(s). Rather, this model is commonly used to describe ORP in experiments in an intuitive way. Further, empirical analysis, such as that done by Lazear et al. (2012), has shown the variables identified by FS as one potential set of statistically significant determinants of giving in dictator games.

Equation (3.1) is the standard FS utility function. The second term of the right-hand side of Equation (3.1) represents the utility loss from disavantageous inequality, while the third represents the utility loss from advantageous inequality.

Now suppose that income is determined by a (weakly) convex combination of two distinct channels: effort and luck. All wealth is then partitioned into two mutually exclusive, exhaustive subsets representing the two respective determinants. Mathematically, agent i has wealth $w_i = e_i + \ell_i$, where $\mathbf{e} = (e_1, ..., e_n)$ is the complete vector of earnings and $\boldsymbol{\ell} = (\ell_1, ..., \ell_n)$ is the complete vector of luck-borne wealth. If individual i has different preferences over e_i and ℓ_i, then if $e_i = \ell_i$, it does not necessarily follow that $U_i(e_i) = U_i(\ell_i)$.

Consider individual i, who is Fehr-Schmidt inequality averse and who also has differing preferences over e and ℓ. Her utility is not a function solely of her own earnings and luck, but also both of the difference between her and each other individual's earnings and of the difference between her and each other individual's luck-borne wealth. This is analogous to how the standard FS model treats general wealth.

In the standard FS model, there are three possible scenarios for each (i, j) pair: $w_i > w_j$, $w_i = w_j$, and $w_i < w_j$. By the definition of inequality aversion, and hence also by the setup of the model, the utility function need only account for scenarios one and three, since scenario two simply drops out by subtraction. In the interest both of brevity and of tailoring this model specifically to fit the dictator game used here, the adapted model will consider comparison to only one other individual, j. Furthermore, since the dictator is always wealthier overall, only the $w_i > w_j$ scenario is considered by design.

The adapted utility function, defined in general form, is thus

$$U_i(e_i, e_j, \ell_i, \ell_j) = u_i(e_i, \ell_i) - \theta_i v_i(e_i - e_j) - \phi_i v_i(\ell_i - \ell_j) \tag{3.2}$$

for individual i. Note that because $w_i > w_j$ is implicitly assumed by experimental design, it necessarily holds that either the dictator has earned more ($e_i > e_j$) or that the dictator has gotten luckier ($\ell_i > \ell_j$), or both.

The first term of the right-hand side of Equation (3.2) captures the direct wealth effect for both earnings and luck.[10] The second term represents the utility lost due to the pair's gap in earnings, and the third term represents that lost due to their gap in luck. These are referred to as the "earnings gap effect" and the "luck gap effect," respectively. While the function $v_i(\cdot)$ is identical for the earnings and luck gaps, the preceding coefficients differ, which allows for different marginal effects.

A few basic assumptions regarding Equation (3.2) are required. First, let $u_i(e_i, \ell_i)$ be increasing and weakly concave over both e_i and ℓ_i. Second, let it hold that $\theta_i > 0$ and $\phi > 0$; Additionally, let $v_i(\cdot)$ be increasing and weakly concave when the argument is positive and decreasing and weakly concave when the argument is negative. This ensures that individual i is actually averse to inequality over both earnings and luck.[11] Last, let the FS assumption of behindness aversion also hold here, for both earnings and luck. That is, for any $\epsilon > 0$, $v_i(\epsilon) < v_i(-\epsilon)$.

Now consider the optimization problem that dictator i faces. Rather than simply choosing $z_i \in [0, e_i + \ell_i]$ to donate, the mental accounting literature suggests that dictator i's donation may be driven by two separate considerations. She chooses x_i to give based on both her and the recipient's earnings, and she chooses y_i to give based on both her and the recipient's luck. Thus, her total donation is the sum of two separate but related donations ($z_i = x_i + y_i$). Of course, the experiment still restricts the total donation such that $0 \leq z_i = x_i + y_i \leq e_i + \ell_i$.

[10] The theory remains agnostic as to whether this term treats the two inputs as integrated or segregated, as defined by Thaler (1985), although minimal assumptions regarding it are made in the next paragraph.

[11] This is a direct extension of the assumption that $0 < \beta_i \ \forall \ i$ in Equation (3.1), as set forth by Fehr and Schmidt (1999).

In stage 3 of the experiment, then, dictator i is choosing (x_i, y_i) based on her utility function $U_i(x_i, y_i; e_i, e_j, \ell_i, \ell_j)$. That is, she faces the optimization problem

$$\max_{x_i, y_i} \; u_i(e_i - x_i, \ell_i - y_i) - \theta_i v_i\left((e_i - x_i) - (e_j + x_i)\right) - \phi_i v_i\left((\ell_i - y_i) - (\ell_j + y_i)\right) \quad (3.3)$$

$$\text{s.t. } e_i + \ell_i \geq x_i + y_i \geq 0. \quad (3.4)$$

The solution to this yields the optimal donations from each account, which sum to the total optimal donation ($x_i^* + y_i^* = z_i^*$). In the experiment, only z_i^* is directly observable, and so all analyses must revolve around it. For simplicity, assume an interior solution.[12]

So far the model has relied on the assumption that agent i has complete information about agent j's sources of wealth. However, this is often not the case in realistic scenarios. Consider instead a situation where agent i knows agent j's total wealth w_j, but doesn't know the breakdown between the two sources of it. As previously discussed, there are three competing hypotheses that could explain this, each supported by a different literature.

The moral wiggle-room story told by Dana et al. (2007) describes how individuals are willing to pay a cost to keep others from knowing that they avoided a situation in which they could give donations. In a similar vein, it follows that some individuals may donate less in the presence of imperfect information regarding income sources. Such individuals use the moral wiggle-room argument that "it wasn't their fault" that they were uninformed as self-justification for giving less. This would imply that partially informed dictators treat the total wealth gap between them and the recipient in the same way fully informed dictators treat the earned income gap. This is formalized into Hypotheses 1, the moral wiggle-room hypothesis.

[12] No wealth-constrained donations were observed in the experiment. Donations of 0 tokens imply either a lack of ORP or a status-seeking individual, neither of which are of interest here.

Hypothesis 1. *The marginal effect on dictator i's donation from an increase in the earnings gap when j's wealth decomposition is known is equal to that from an increase in the total wealth gap when it is not* $\left(\frac{dz_i^*}{d(e_i-e_j)}|_{e_j,\ell_j,f_j,w_j} = \frac{dz_i^*}{d(w_i-w_j)}|_{f_j,w_j}\right)$.

Theories in the social insurance literature, on the other hand, argue that some individuals are willing to give others the benefit of the doubt. This is a means to ensure that if they were ever to get unlucky in the future that society would be more likely to provide them with a similar safety net. The implication here is that partially informed dicatators treat the total wealth gap in the same way fully informed dictators treat the luck gap. This is formalized into Hypthesis 2, the social insurance hypothesis.

Hypothesis 2. *The marginal effect on dictator i's donation from an increase in the luck gap when j's wealth decomposition is known is equal to that from an increase in the wealth gap when it is not* $\left(\frac{dz_i^*}{d(\ell_i-\ell_j)}|_{e_j,\ell_j,f_j,w_j} = \frac{dz_i^*}{d(w_i-w_j)}|_{f_j,w_j}\right)$.

Last, there is a third alternative approach. While the dictator does not know the recipient's actual earned or luck-borne income, she does have knowledge of the underlying luck distribution, $f_j(\ell_j)$. She can use this information to construct expectations about the recipient's luck and earnings conditional on his luck distribution and total wealth.[13] Hence, the traditional Bayesian approach suggests that when choosing how much to donate, the dictator treats the conditional expectations of the earnings and luck differences given imperfect information identical to how she would treat the realized differences given perfect information. This is formalized into Hypothesis 3, the conditional expetations hypothesis.

Hypothesis 3. *The marginal effect on dictator i's donation from an increase in the earnings gap when j's wealth decomposition is known is equal to that from an increase in the conditional expectation of the earnings gap when it is not*

[13]It is known that $f(\ell) = \frac{1}{5} \; \forall \; \ell \in L = \{-20, -10, 0, 10, 20\}$ and $w \geq 0$ for all subjects. If dictator i observes $0 \leq w_j < 10$, then $f(\ell_j|w_j) = 0 \; \forall \; \ell_j \in \{10, 20\}$ and $f(\ell_j|w_j) = \frac{1}{3} \forall \; \ell_j \in \{-20, -10, 0\}$. If dictator i observes $10 \leq w_j < 20$, then $f(\ell_j|w_j) = 0 \; \forall \; \ell_j \in \{20\}$ and $f(\ell_j|w_j) = \frac{1}{4} \forall \; \ell_j \in \{-20, -10, 0, 10\}$. Note that this requires no knowledge of or assumptions regarding either the recipient's effort or his effort distribution.

$\left(\frac{dz_i^*}{d(e_i-e_j)}|_{e_j,\ell_j,f_j,w_j} = \frac{dz_i^*}{d(e_i-E[e_j|f_j,w_j])}|_{f_j,w_j}\right)$.

Correspondingly, the marginal effect on dictator i's donation from an increase in the luck gap when j's wealth decomposition is known is equal to that from an increase in the conditional expectation of the luck gap when it is not $\left(\frac{dz_i^*}{d(\ell_i-\ell_j)}|_{e_j,\ell_j,f_j,w_j} = \frac{dz_i^*}{d(\ell_i-E[\ell_j|f_j,w_j])}|_{f_j,w_j}\right)$.

I test each of these three hypotheses empirically. The results in section 3.4 add empirical evidence to a previously purely theoretical contention.

3.4 Results

3.4.1 Data Summary

The experiment ran over the course of two days in late May, 2013. In total, eight sessions of the experiment were run. Seven sessions had 8 participants, and one session had 6. In total, this yielded 62 students, or 31 observations (pairs). Of these, 16 observations were in the treatment, while 15 were in the control.

Table 3.1 shows some basic summary statistics about subjects' earnings, luck, and total wealth levels, all before the dictator game occured. It also shows summary statistics for dictators' donations. Notice that, on average, dictators had higher earnings, luck, and wealth than recipients. This is to be expected, however, given the dictator selection process.[14]

Figure 3.2 illustrates the frequency of each donation, split into control and treatment. Notice that 8 of the 15 control pairs donated 0 tokens, while the same holds for 9 of the 16 treatment pairs. The highest observed donation was 10 tokens for each.

[14] It is ironic that the highest earning subject got unlucky in both his actual luck outcome and his pairing, since he ended up losing tokens and then becoming a recipient. This serves to punctuate the fact that luck can actually play a pivotal role in this experiment, just as it can in many realistic situations.

Table 3.1: Summary Statistics (Before Donation)

Tokens (Before Donation)	Mean	Std. Dev.	Min	Max
	All Subjects		n=62	
Earnings	54.04	13.32	20	80
Luck	-0.33	14.25	-20	20
Wealth	53.73	21.51	10	95
	Recipients		n=31	
Earnings	48.58	13.03	20	80
Luck	-6.13	13.34	-20	20
Wealth	42.45	18.54	10	78
	Dictators		n=31	
Earnings	59.52	11.36	31	75
Luck	5.48	12.87	-20	20
Wealth	65.00	18.28	21	95
Donations	1.84	2.99	0	10

Figure 3.3 shows earnings and total wealth (before the dictator game), split by dictators and their respective recipients, as well as the respective donations. The data points are arranged by the dictator's wealth, in increasing order. The key point illustrated by this graph is that donations are not merely increasing as dictator wealth increases; there is clearly something more affecting donations.

3.4.2 General Model Discussion

Because of the small sample size, many models are considered, and the strengths and weaknesses of each are discussed. Some are used, some are outright dismissed, and some serve as robustness checks. This section contains that modeling discussion, while section 3.4.3 highlights the results for the main specifications chosen.

To see what else is affecting donations, this section will consider a series of models that examine the previous hypotheses, beginning with Table 3.2. Note that columns 1, 2, and 3 follow Lazear et al. (2012) in clustering using standard errors at the session level to account for any possible session-specific distinctions that may have occured (*e.g.*, weather, time of day, number of subjects, *etc.*).

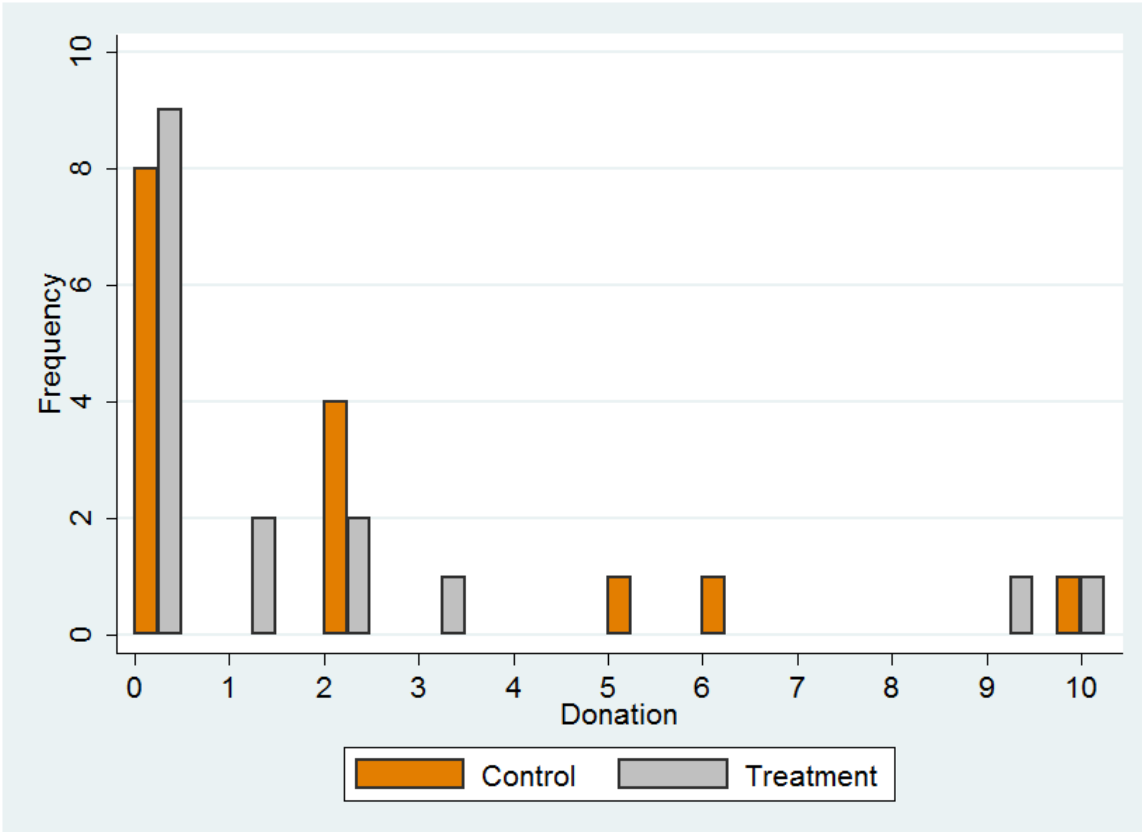

Figure 3.2: Frequency of Amount of the Tokens Donated, by Control and Treatment

The first model examines donations as a function solely of what each given dictator knows, including both own levels and differences for each source. The empirical specification is

$$\begin{aligned} Don_{ijs} =& \beta_0 + \beta_1 e_i + \beta_2 \ell_i + \beta_3 \mathbb{1}_{ij}\{C\}(e_i - e_j) + \\ & \beta_4 \mathbb{1}_{ij}\{C\}(\ell_i - \ell_j) + \beta_5 \mathbb{1}_{ij}\{T\}(w_i - w_j) + \beta_6 \mathbb{1}_i\{M\} + \epsilon_{ijs} \end{aligned} \tag{3.5}$$

where e represents earned income, ℓ represents luck-borne income, w represents total income, i indexes the dictator, j indexes the recipient, s indexes the session, C indicates control pairs, T indicates treatment pairs, and M indicates male dictators. Recall that both the control and the treatment dictators know their own earnings and luck levels, but only those in the control know the gaps for earnings and luck. Thus, the sources'

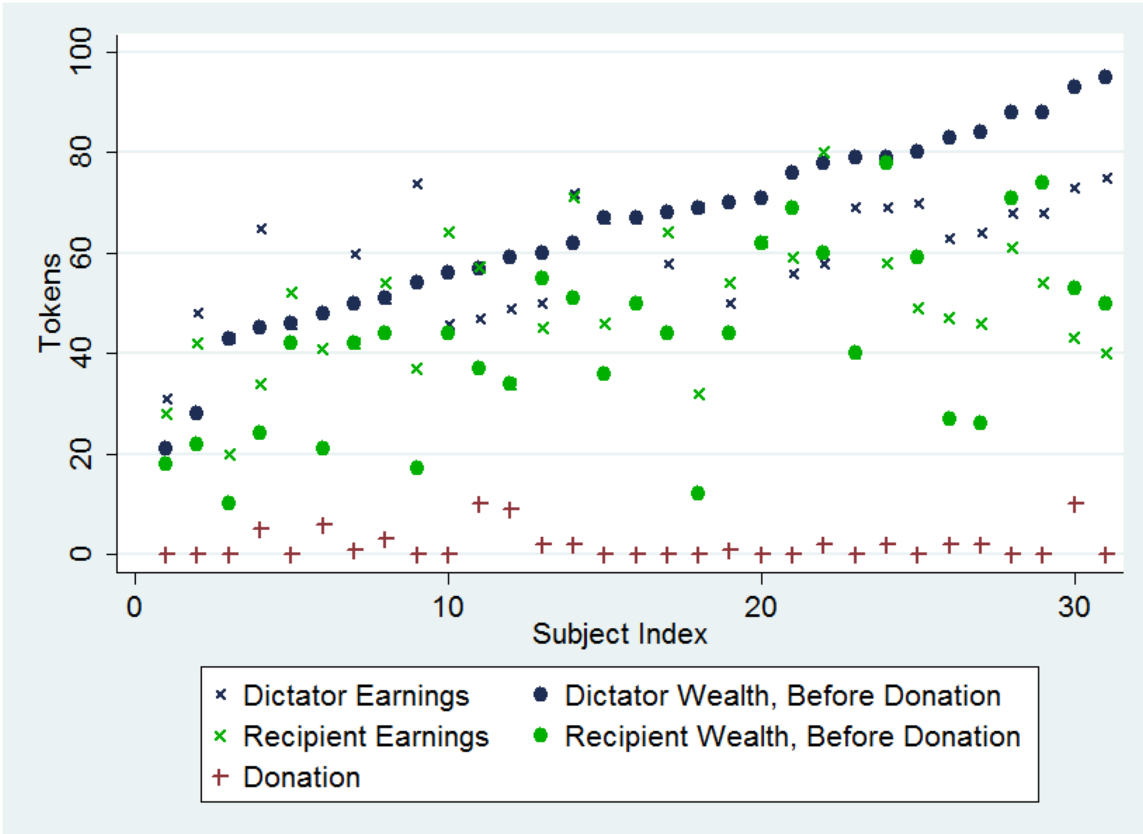

Figure 3.3: Wealth (Before Donation), Earnings, and Donation, for Dictator-Recipient Pairs

gaps are used for the control group only, while the treatment relies on the gap in total wealth levels. Running this specification with all the demographic variables (most not reported here), and all corresponding subsets, shows only the gender dummy as statistically significant. These results, with male as the only included demographic, are reported in column 1 of Table 3.2.

Comparing the empirical specification shown in Equation 3.5 to the theoretical specifications from section 3.3.1, it is first apparent that the marginal impact of increasing either the dictator's earnings or luck is not captured in a single coefficient. Rather, the marginal impact of increasing the dictator's earnings is $\frac{\partial Don_{ijs}}{\partial e_i} = \beta_1 + \beta_3$ if i is in the control, and $\frac{\partial Don_{ijs}}{\partial e_i} = \beta_1 + \beta_5$ if i is in the treatment. For increases in luck, $\frac{\partial Don_{ijs}}{\partial \ell_i} = \beta_2 + \beta_4$ if i is in the control, and $\frac{\partial Don_{ijs}}{\partial \ell_i} = \beta_2 + \beta_5$ if i is in the treatment.

Table 3.2: Regression Results

		(1)	(2)	(3)	(4)
Sample:		Full Sample	Donation>0	Donation>0	Donation>0
	(H_0)	Donation	Donation	Donation	Donation
Dict Earn		-0.0354	-0.00679		-0.00679
		(0.0569)	(0.0596)		(0.0839)
Dict Luck		-0.00586	-0.207***	-0.193***	-0.207**
		(0.0256)	(0.0445)	(0.0482)	(0.0738)
Earn Gap C	(+)	0.0360	0.0952		0.0952
		(0.0698)	(0.123)		(0.0855)
Luck Gap C	(+)	0.0646	0.181***	0.164***	0.181**
		(0.0453)	(0.0406)	(0.0433)	(0.0581)
Wealth Gap T	(+)	0.0346	0.250***	0.208***	0.250**
		(0.0321)	(0.0651)	(0.0442)	(0.0787)
Male		-3.181***	-5.809***	-4.932***	-5.809**
		(0.710)	(1.224)	(1.173)	(1.663)
Constant		4.992	4.078	4.268**	4.078
		(3.524)	(4.003)	(1.190)	(4.586)
Observations		31	14	14	14
Adjusted R^2		0.144	0.512	0.546	0.512

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Notes: Standard errors ar clustered by Session in columns 1, 2, and 3. Covariates denoted "C" indicate interaction with a control dummy, and those denoted "T" indicate interaction with a treatment dummy.

This is entirely due to the empirical specification; by using income gaps to better match the theoretical specification, the marginal impact of own wealth is necessarily split.

On the other hand, the marginal impact of increasing the recipient's income, through either source, is captured by a single coefficient. That is, in the control, an increase in j's earned income is $\frac{\partial Don_{ijs}}{\partial e_j} = -\beta_3$, and an increase in j's luck-borne income is $\frac{\partial Don_{ijs}}{\partial \ell_j} = -\beta_4$. In the treatment, and an increase in j's total income (from either source) is $\frac{\partial Don_{ijs}}{\partial w_j} = -\beta_5$.

Column 1 of Table 3.2 uses exactly this specification. Notice that no variables in column 1 are statistically significant, with the exception of the gender dummy (as previously noted). This specification indicates that males generally give 3.18 tokens less (p=0.003) than their corresponding female counterfactuals.

While the near complete lack of results may seem disheartening at first, breaking down the behavior more categorically, based on preferences, provides a different angle that offers better insight.

Before delving into the mechanics of the approach, let's take a step back and discuss the intuition at play. The goal is to capture how donations by other-regarding individuals are affected by changes in other variables. It may be, however, that not all individuals are other-regarding. Furthermore, it may be that some are other-regarding in a competitive way ("status-seeking" individuals), which would imply a negative donation, were that possible here. For solely self-interested or status-seeking individuals, changes in either their own or others' effort or luck will not affect the observed level of donation, since they will not donate tokens regardless. Hence, such individuals serve only to inject noise when pooled with other-regarding (non-status-seeking) individuals.

There are multiple approaches to tackling this issue. Perhaps the most complete approach is to first construct a two-stage Craggit (or Hurdle) model that first predicts which subjects will donate (*i.e.*, are other-regarding and non-status-seeking), and then use that to help explain donations in a second stage regression (*e.g.*, using predicted probability or propensity score matching). I explored many binary-choice models for the first stage (some are shown in Table A.1 in Appendix B) in an attempt to predict whether or not subjects would donate, but nothing significant surfaced. Whatever the characteristics may be that predict the presence of such preferences, they do not seem to have been picked up by any experimental decisions or any questionnaire responses.

The second approach to addressing the issue is to use a Tobit model, where donations are censored at a lower bound of 0 (shown in Table A.2 in Appendix B).[15] After various attempts at modeling donations using a Tobit approach, again nothing significant surfaced. One potential explanation for this is that there is no way to mechanically distinguish between solely self-interested individuals and status-seeking individuals.

I also use a Poisson distribution model for count data as a third approach to the issue (shown in Table A.3 in Appendix B). The properties of the Poisson distribution are such that it naturally fits well with data containing a large amount of zeros as observations. While slightly more covariates are significant, these specifications still predict little of the variation.

The lack of results using any of the previous approaches necessitates another, less complete approach. Since I'm not able to accurately predict whether or not an individual donates, I use specifications with the data restricted only to those who did indeed donate. While this approach does not pick up any factors that affect the external margin of the donation decision, it does allow for an examination of the driving factors on the internal margin.

Column 2 of Table 3.2 shows the results from the most straightforward specification. All three gap coefficients have the expected signs (more discussion on this follows in section 3.4.3). Luck, wealth, and the gender dummy coefficients are all statistically significant ($p<0.01$ for each), while the earnings coefficients and the constant are statistically insignificant.

There is, however, a potential issue with this regression; there are 7 covariates and only 7 clusters (session five had no positive donations). I used two further specifications as robustness checks to identify any potential issues. The first robustness check is the specification in column 3, which is identical to that in column 2 except that all statistically insignificant covariates, other than the constant, have been removed. The

[15]Note that the Tobit model is a special case nested in the Craggit model with added coefficient restrictions.

coefficients' signs are all unchanged, and all four statistically significant coefficients from specification 3 are still significant at the 1% level. Interestingly, the point estimate of the constant does not change much, but its p-value dropped (from p=0.348 to p=0.012), making it now statistically significant at the 5% level. Furthermore, the adjusted R^2 value increased when the difference between the dictator's earnings and the control's earnings were dropped, which indicates that specification 3 is a better fit.[16]

The second robustness check is simply to eliminate the clustering, as shown in column 4. This increases the number of degrees of freedom. All statistically significant variables in specification 2 are still significant ($p<0.05$). This also serves to address the issue that clustered errors have been shown to typically be slightly smaller and unreliable in situations with this few clusters.[17]

Because the gender dummy is significant in nearly every specification, I also run additional models to check that the shift-in-mean effect the dummy variable represents in Table 3.2 is an accurate specification. When interacted with other covariates, significance is lost uniformly, as can be seen in Table A.5 in Appendix B. This result suggests that males and females treat the various income-related factors identically on the margin, but that males have a lower baseline for giving on the whole.

3.4.3 Hypotheses Results

While Table 3.2 succinctly illustrates the main factors driving donations in an intuitive way, not all of the hypotheses laid out in section 3.3.1 can be tested using its specifications. Instead, Tables 3.3 and 3.4 will guide the discussion through each of the

[16] Regressing with just own earnings as a robustness check yielded insignificant estimates for earnings and the constant. Doing so with just the control's difference in earnings yielded an insignificant earnings estimate but a significant (p=0.042) constant. The only sign change was that own earnings was slightly positive. Neither result is shown here.

[17] MacKinnon and White (1985) use the residual-variance estimator HC3, which approximates a jackknife estimator to correct for this in small samples. I follow them for specifications 2 and 3, and I find that all statistically significant variables remain so ($p<0.1$), except for dictator's own luck in specification 2 (p=0.158). These results are shown in Table A.4 in Appendix B.

Table 3.3: Regressions for Hypotheses Tests

	Donation (1)	Donation (2)
Dict Earn	-0.007	-0.013
	(0.060)	(0.059)
Dict Luck	-0.207***	-0.242***
	(0.045)	(0.054)
Earn Gap C	0.095	-0.038
	(0.123)	(0.076)
Luck Gap C	0.181***	0.141**
	(0.041)	(0.043)
Wealth Gap T	0.250***	
	(0.065)	
Expected Earn Gap T		0.393***
		(0.075)
Expected Luck Gap T		0.690***
		(0.165)
Male	-5.809***	-3.969**
	(1.224)	(1.340)
Constant	4.078	6.528*
	(4.003)	(3.348)

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: "C" denotes control pairs, and "T" denotes treatment pairs.

Table 3.4: Hypotheses Tests

Hypothesis	H_0	F-stat	p-value
(1) Moral Wiggle-Room	$\frac{\partial Don_{ij}}{\partial(e_i^C - e_j^C)} = \frac{\partial Don_{ij}}{\partial(w_i^T - w_j^T)}$	3.93^{*}	(0.095)
(2) Social Insurance	$\frac{\partial Don_{ij}}{\partial(\ell_i^C - \ell_j^C)} = \frac{\partial Don_{ij}}{\partial(w_i^T - w_j^T)}$	0.95	(0.367)
(3) Conditional Expectations	$\frac{\partial Don_{ij}}{\partial(e_i^C - e_j^C)} = \frac{\partial Don_{ij}}{\partial(e_i^T - E[e_j^T])}$	14.42^{***}	(0.009)
	$\frac{\partial Don_{ij}}{\partial(\ell_i^C - \ell_j^C)} = \frac{\partial Don_{ij}}{\partial(\ell_i^T - E[\ell_j^T])}$	9.86^{**}	(0.020)
	Jointly:	7.24^{**}	(0.025)

$p-$values in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

hypotheses. In Table 3.3, column 1 shows the empirical specification for Hypotheses 1 and 2, while column 2 shows the empirical specification for Hypothesis 3.[18]

The null for Hypothesis 1 represents the moral wiggle-room hypothesis, detailed in column 1 of Table 3.3 and row 1 of Table 3.4. If the null cannot be rejected, then there is some support for the moral wiggle-room hypothesis. However, the F-test seen in row 1 of Table 3.4 shows that the null can be rejected at the 10% level ($p = 0.0948$). This rejection is evidence that imperfectly informed dictators do not treat wealth gaps in the same way informed dictators do earnings gaps. Rather, they do account for the possibility that recipients are unlucky.

Hypothesis 2, or the social insurance hypothesis, stands in direct contrast to the moral wiggle-room hypothesis. If the null cannot be rejected, then there is evidence that dictators treat ambiguously sourced wealth as they would luck, supporting the social insurance hypothesis. Unlike with Hypothesis 1, here the null cannot be rejected with significance ($p = 0.3666$), as seen in row 2 of Table 3.4. This failure to reject is

[18]The specification in column 2 does not cluster standard errors, since it has more covariates than there are session clusters.

evidence in support of the notion that there naturally exists some social insurance when the source of wealth is ambiguous.

The second column of Table 3.3 shows the specification used to examine Hypothesis 3, the conditional expectations hypothesis. The F-test results are shown in the rows 3, 4, and 5 of Table 3.4. The intuition here draws from of a standard Bayesian conditional expected value approach. A failure to reject the null would suggest that dictators make donations based on their expectations of recipients' income decomposition, conditional on their luck distribution and total wealth. That is, dictators treat these conditional expectations of luck and earnings gaps exactly as they would if these were the actual realized gaps. However, each of the two conditions is separately rejected, as shown in rows 3 and 4 ($p = 0.0090$ and $p = 0.0201$, respectively), as is the joint test ($p = 0.0252$) shown in row 5. This is evidence that dictators do not split the wealth gap into conditional expectations of source gaps in the decision-making process.

The implications of Hypotheses 1, 2, and 3 together provide an insight into how other-regarding dictators treat recipients' ambiguously sourced wealth. There is evidence against the theory that dictators use the lack of information as a means of self-justification for giving less by assuming recipients are just lazy. There is also evidence against the Bayesian expectation approach; dictators do not seem to decompose the wealth gap into conditional expectations of the earnings and luck gaps. There is, however, support for the remaining potential explanation. Partially informed dictators do *not* seem to treat their total wealth gap any differently than fully informed dictators treat their luck gap. That is, other-regarding dictators tend to innately provide some level of social insurance in the face of imperfect information regarding recipients' income sources. This evidence supports the story that people tend to help others out in hopes that the reverse would happen if the shoe were on the other foot.

3.5 Conclusion

There is ample anecdotal evidence of common situations in which individuals seem to treat wealth differently depending on whether they worked to earn it or simply got lucky. This distinction extends to individuals' perceptions of how others came to be in their own respective situations, too. Despite this evidence, such a distinction has gone largely ignored in the economics literature.

Furthermore, in most applicable realistic scenarios, the salience of others' income sources is less than perfect. It is generally unclear as to whether people tend to treat such obfuscation moreso as if the income were luck-based, as is suggested by the social insurance literature; as if it were effort-based, which is more analogous to the moral wiggle room story; or as a decomposition into conditional expectations for each source, as Bayesian theory suggests. Given that a large portion of these scenarios seem to be characterized by imperfect information, the implications of ignoring the lack of clarity are potentially large. Within the few studies that have distinguished income by source, however, the trend has been to do so with complete salience.

This paper examines each of these issues. I develop a theoretical preference structure that allows for the separation both of one's own and of others' income by source. The new structure is based on the classic model of inequality aversion developed by Fehr and Schmidt (1999). I nest the model within a dictator game framework wherein each player has both a luck component and an earnings component of her wealth. From this, I motivate relevant empirically testable hypotheses comparing partially and fully informed dictators' behavior. I then run a laboratory experiment to test these hypotheses.

With even a quick glance at the data, it is immediately obvious that there are approximately an equal number of people who make a positive donation and those who give nothing. This separation suggests that while there are many individuals who do exhibit other-regarding behavior, there are also many who appear solely self-interested (or status-seeking). While pooling the two types together produces no useful evidence, limiting the analysis to only the subsample who do donate provides some meaningful results.

There is statistically significant evidence that, in the presence of imperfect information surrounding the recipient's income sources, other-regarding dictators tend to treat the pair's wealth gap differently than if it had been earned, but not differently than if it had been acquired via luck. This implies that in the presence of source uncertainty, other-regarding dictators tend to give recipients the benefit of the doubt by assuming that the recipient is unlucky rather than lazy. Furthermore, there is also evidence that imperfectly informed dictators treat the conditional expectations of the earnings and luck gaps differently than perfectly informed dictators treat the known earnings and luck gaps, respectively. This result suggests that dictators give more under uncertainty than they would if their conditional expectations were realized and known. Overall, these results lend credence to the story of social insurance for other-regarding individuals.

In the future, theoretical models and empirical analyses should carefully consider the implications of assuming both that income preferences are identical across sources and that such information is perfectly known to all agents. Being able to empirically parse apart other-regarding and status-seeking individuals from their respective counterparts, whether based on individual characteristics or through observing some type of behavior (other than simply donation behavior), would be a strong next step. Additionally, just as empirical evidence has supported reference dependence, it may be that there are non-linearities for earnings and luck. With a larger sample, it would be possible to examine the data for any such non-linearities, instead of only picking up the average effects as I am constrained to doing here.

Bibliography

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American Economic Review*, 95(4):960–980. 72

Altemeyer-Bartscher, M., Markandya, A., and Rübbelke, D. T. (2011). The private provision of international impure public goods: The case of climate policy. *Basque Centre for Climate Change Working Paper*, (2011-09). 4, 6

Altemeyer-Bartscher, M., Markandya, A., and Rübbelke, D. T. (2014). International side-payments to improve global public good provision when transfers are refinanced through a tax on local and global externalities. *International Economic Journal*, 28(1):71–93. 6

Altemeyer-Bartscher, M., Rübbelke, D. T., and Sheshinski, E. (2010). Environmental protection and the private provision of international public goods. *Economica*, 77(308):775–784. 6

Alzahrani, M. and Lasfer, M. (2012). Investor protection, taxation, and dividends. *Journal of Corporate Finance*, 18(4):745–762. 29

Andreoni, J. (1989). Giving with impure altruism: applications to charity and ricardian equivalence. *The Journal of Political Economy*, pages 1447–1458. 68

Andreoni, J. (1990). Impure altruism and donations to public goods: a theory of warm-glow giving. *The economic journal*, pages 464–477. 68

Auerbach, A. J. (1982). The theory of excess burden and optimal taxation. *NBER Working Paper*, (w1025). 5

Bae, K.-H., Baek, J.-S., Kang, J.-K., and Liu, W.-L. (2012). Do controlling shareholders' expropriation incentives imply a link between corporate governance and firm value? theory and evidence. *Journal of financial Economics*, 105(2):412–435. 28, 39

Balafoutas, L., Kocher, M. G., Putterman, L., and Sutter, M. (2013). Equality, equity and incentives: An experiment. *European Economic Review*, 60:32–51. 73

Bates, T. W., Kahle, K. M., and Stulz, R. M. (2009). Why do us firms hold so much more cash than they used to? *The Journal of Finance*, 64(5):1985–2021. 25

Bento, A., Kaffine, D., Roth, K., and Zaragoza-Watkins, M. (2014). The effects of regulation in the presence of multiple unpriced externalities: evidence from the transportation sector. *American Economic Journal: Economic Policy*, 6(3):1–29. 4

Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. *Games and economic behavior*, 10(1):122–142. 30, 35

Bergstrom, T. and Andreoni, J. (1996). Do government subsidies increase the private supply of public goods? *Public Choice*, 88:295–308. 4, 6

Bergstrom, T., Blume, L., and Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49. 2, 4, 12

Berkman, H., Cole, R. A., and Fu, L. J. (2009). Expropriation through loan guarantees to related parties: Evidence from china. *Journal of Banking & Finance*, 33(1):141–156. 27, 39

Bolton, G. E. and Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *American economic review*, pages 166–193. 70

Breuer, W., Rieger, M. O., and Soypak, K. C. (2014). The behavioral foundations of corporate dividend policy a cross-country analysis. *Journal of Banking & Finance*, 42:247–265. 40

Cabrales, A., Nagel, R., and Mora, J. V. R. (2012). It is hobbes, not rousseau: an experiment on voting and redistribution. *Experimental Economics*, 15(2):278–308. 73

Cairns, J. and Slonim, R. (2011). Substitution effects across charitable donations. *Economics Letters*, 111(2):173–175. 75

Cherry, T. L., Frykblom, P., and Shogren, J. F. (2002). Hardnose the dictator. *American Economic Review*, pages 1218–1221. 69

Chetty, R. and Saez, E. (2005). Dividend taxes and corporate behavior: evidence from the 2003 dividend tax cut. *The Quarterly Journal of Economics*, 120:791–833. 29

Cornes, R. and Sandler, T. (1984). Easy riders, joint production, and public goods. *The Economic Journal*, 94(375):580–598. 2

Cornes, R. and Sandler, T. (1994). The comparative static properties of the impure public good model. *Journal of Public Economics*, 54(3):403–421. 2

Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80. 71, 72, 81

Diamond, P. A. and Mirrlees, J. A. (1971a). Optimal taxation and public production i: Production efficiency. *The American Economic Review*, pages 8–27. 5

Diamond, P. A. and Mirrlees, J. A. (1971b). Optimal taxation and public production ii: Tax rules. *The American Economic Review*, pages 261–278. 5

Durnev, A. and Kim, E. (2005). To steal or not to steal: Firm attributes, legal environment, and valuation. *The Journal of Finance*, 60(3):1461–1493. 28

Eisenkopf, G., Fischbacher, U., and Föllmi-Heusi, F. (2013). Unequal opportunities and distributive justice. *Journal of Economic Behavior & Organization*, 93:51–61. 73

Erkal, N., Gangadharan, L., and Nikiforakis, N. (2011). Relative earnings and giving in a real-effort experiment. *The American Economic Review*, pages 3330–3348. 70, 74

Falkinger, J. (1996). Efficient private provision of public goods by rewarding deviations from average. *Journal of Public Economics*, 62(3):413–422. 6

Falkinger, J., Fehr, E., Gächter, S., and Winter-Ebmer, R. (2000). A simple mechanism for the efficient provision of public goods: Experimental evidence. *The American Economic Review*, pages 247–264. 6

Falkinger, J., Hackl, F., and Pruckner, G. J. (1996). A fair mechanism for efficient reduction of global co2-emissions,". *FinanzArchiv / Public Finance Analysis*, 53(3):4. 6

Fama, E. F. and French, K. R. (2001). Disappearing dividends: changing firm characteristics or lower propensity to pay? *Journal of Financial economics*, 60(1):3–43. 25

Farre-Mensa, J., Michaely, R., and Schmalz, M. C. (2014). Payout policy. *Available at SSRN 2400618*. 25, 29

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pages 817–868. 70, 78, 80, 94

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178. 30, 78

Giebe, T. and Schweinzer, P. (2013). Consuming your way to efficiency: public goods provision through non-distortionary tax lotteries. Technical report, CESifo Working Paper. 6

Gompers, P. A., Ishii, J., and Metrick, A. (2010). Extreme governance: An analysis of dual-class firms in the united states. *Review of Financial Studies*, 23(3):1051–1088. 29, 39

Greiner, B., Ockenfels, A., and Werner, P. (2012). The dynamic interplay of inequality and trustan experimental study. *Journal of Economic Behavior & Organization*, 81(2):355–365. 35

Grullon, G. and Michaely, R. (2002). Dividends, share repurchases, and the substitution hypothesis. *The Journal of Finance*, 57(4):1649–1684. 25

Hoffman, E., McCabe, K., and Smith, V. L. (1996). Social distance and other-regarding behavior in dictator games. *The American Economic Review*, pages 653–660. 68

Jensen, M. C. and Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs, and ownership structure. *Journal of Financial Economics*, 3(4):305–360. 25

Jiang, G., Lee, C. M., and Yue, H. (2010). Tunneling through intercorporate loans: The china experience. *Journal of Financial Economics*, 98(1):1–20. 27, 39

Jordan, B. D., Liu, M. H., and Wu, Q. (2014). Corporate payout policy in dual-class firms. *Journal of Corporate Finance*, 26:1–19. 29

Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness and the assumptions of economics. *Journal of business*, pages S285–S300. 68

Knack, S. and Keefer, P. (1995). Institutions and economic performance: Cross-country tests using alternative institutional measures. *Economics & Politics*, 7(3):207–227. 28

Kolmar, M. and Wagener, A. (2012). Contests and the private production of public goods. *Southern Economic Journal*, 79(1):161–179. 6

Koppel, H. and Schulze, G. G. (2013). The importance of the indirect transfer mechanism for consumer willingness to pay for fair trade products evidence from a natural field experiment. *Journal of Consumer Policy*, 36(4):369–387. 6

Kotchen, M. J. (2005). Impure public goods and the comparative statics of environmentally friendly consumption. *Journal of Environmental Economics and Management*, 49(2):281–300. 2, 4

Kotchen, M. J. (2006). Green markets and private provision of public goods. *Journal of Political Economy*, 114(4):816–834. 2, 7, 8, 9, 10, 11, 14, 23

Kotchen, M. J. (2007). Equilibrium existence and uniqueness in impure public good models. *Economics Letters*, 97(2):91–96. 2, 11

Kotchen, M. J. (2013). Voluntary-and information-based approaches to environmental management: A public economics perspective. *Review of Environmental Economics and Policy*, 7(2):276–295. 6

Krawczyk, M. (2010). A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution. *Journal of Public Economics*, 94(1):131–141. 73

La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (2000). Agency problems and dividend policies around the world. *The Journal of Finance*, 55(1):1–33. 26, 37, 42, 47, 64, 65

La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (2002). Investor protection and corporate valuation. *The Journal of Finance*, 57(3):1147–1170. 28, 64

Lazear, E. P., Malmendier, U., and Weber, R. A. (2012). Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics*, 4(1):136–163. 78, 84

Levitt, S. D., List, J. A., and Sadoff, S. E. (2011). Checkmate: Exploring backward induction among chess players. *American Economic Review*, 101(2):975–90. 36

Löschel, A. and Rübbelke, D. T. (2009). Impure public goods and technological interdependencies. *Journal of Economic Studies*, 36(6):596–615. 6

Lunawat, R. (2013). An experimental investigation of reputation effects of disclosure in an investment/trust game. *Journal of Economic Behavior & Organization*, 94:130–144. 35

MacKinnon, J. G. and White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29(3):305–325. 90, 110

Markandya, A. and Rübbelke, D. T. (2012). Impure public technologies and environmental policy. *Journal of Economic Studies*, 39(2):128–143. 6

McClelland, R. (1989). Voluntary donations and public expenditures in a federalist system: Comment and extension. *The American Economic Review*, 79(5):1291–1296. 2

McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, pages 803–836. 36

Modigliani, F. and Miller, M. H. (1958). The cost of capital, corporation finance and the theory of investment. *The American economic review*, pages 261–297. 25

Munro, A. and Valente, M. (2009). Green goods: are they good or bad news for the environment? evidence from a laboratory experiment on impure public goods. MPRA Paper 13024, University Library of Munich, Germany. 3

Neilson, W. S. and Stowe, J. (2010). Piece-rate contracts for other-regarding workers. *Economic Inquiry*, 48(3):575–586. 68

Nordhaus, W. D. (2006). After kyoto: Alternative mechanisms to control global warming. *The American Economic Review*, pages 31–34. 6

Pearce, D. (2000). *Policy frameworks for the ancillary benefits of climate change policies*. Centre for Social and Economic Research on the Global Environment. 6

Ramsey, F. P. (1927). A contribution to the theory of taxation. *The Economic Journal*, pages 47–61. 5

Rietz, T. A., Sheremeta, R. M., Shields, T. W., and Smith, V. L. (2013). Transparency, efficiency and the distribution of economic welfare in pass-through investment trust games. *Journal of Economic Behavior & Organization*, 94:257–267. 35

Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1):92–100. 30, 36

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, pages 387–389. 4

Sandler, T. and Cornes, R. (1996). The theory of externalities, public goods and club goods. 2

Skinner, D. J. (2008). The evolving relation between earnings, dividends, and stock repurchases. *Journal of Financial Economics*, 87(3):582–609. 25

Steinberg, R. (1987). Voluntary donations and public expenditures in a federalist system. *The American Economic Review*, pages 24–36. 2

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science*, 4(3):199–214. 69, 80

Varian, H. R. (1980). Redistributive taxation as social insurance. *Journal of Public Economics*, 14(1):49–68. 72, 75

Appendix

Appendix A

Recall that

$$\Gamma \equiv (1-\theta_c+\theta_c n_c)(1-\theta_d+\theta_d n_d)-\theta_c\theta_d n_c n_d.$$

Proof of Proposition 1:

Proof. Consider Equation 1.9:

$$Y^{*^{DIFF}} = \frac{(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g+(\frac{\beta}{1-\alpha}-1)n_c\right)w_i}{\Gamma}.$$

Differentiation with respect to n_c yields

$$\begin{aligned}\frac{\partial Y^{*^{DIFF}}}{\partial n_c} =& \frac{(1-\theta_d)(1-\theta_c)w_i}{\Gamma^2}\Bigg[(1-\theta_d)(1-\theta_c)\left(\frac{\beta}{1-\alpha}-1\right)\\ &+(1-\theta_c)\theta_d n_d\left(\frac{\beta}{1-\alpha}-1\right)-(1-\theta_d)\theta_c n_g(\beta-1)\Bigg].\end{aligned} \tag{6}$$

Equation 6 is necessarily positive. Differentiation of Equation 1.9 with respect to n_g yields

$$\frac{\partial Y^{*^{DIFF}}}{\partial n_g} = \frac{-(1-\theta_d)(1-\theta_c)(1-\beta)w_i}{\Gamma}. \tag{7}$$

Equation 7 is necessarily negative. □

Proof of Proposition 2:

Proof. Consider Equation 1.9:

$$Y^{*^{DIFF}} = \frac{(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g+(\frac{\beta}{1-\alpha}-1)n_c\right)w_i}{\Gamma}.$$

Differentiation with respect to α yields

$$\frac{\partial Y^{*^{DIFF}}}{\partial \alpha} = \frac{(1-\theta_d)(1-\theta_c)\left(\frac{\beta n_c}{(1-\alpha)^2}\right) w_i}{\Gamma}. \tag{8}$$

Equation 8 is necessarily positive. Differentiation of Equation 1.9 with respect to β yields

$$\frac{\partial Y^{*^{DIFF}}}{\partial \beta} = \frac{(1-\theta_d)(1-\theta_c)\left(n_g + \frac{n_c}{1-\alpha}\right) w_i}{\Gamma}. \tag{9}$$

Equation 9 is necessarily positive. □

Proof of Proposition 3:

Proof. Consider Equation 1.9:

$$Y^{*^{DIFF}} = \frac{(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g + (\frac{\beta}{1-\alpha}-1)n_c\right) w_i}{\Gamma}.$$

Differentiation with respect to n_d yields

$$\frac{\partial Y^{*^{DIFF}}}{\partial n_d} = \frac{-(1-\theta_d)(1-\theta_c)^2\theta_d\left((\beta-1)n_g + (\frac{\beta}{1-\alpha}-1)n_c\right) w_i}{\Gamma^2}. \tag{10}$$

Let Φ be defined as $\left((\beta-1)n_g + (\frac{\beta}{1-\alpha}-1)n_c\right)$, exactly as earlier. If $\Phi > 0$, then Equation 10 is negative. If $\Phi < 0$, then Equation 10 is positive. □

Proof of Proposition 4:

Proof. Consider Equation 1.9:

$$Y^{*^{DIFF}} = \frac{(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g + (\frac{\beta}{1-\alpha}-1)n_c\right) w_i}{\Gamma}.$$

Differentiation with respect to θ_d yields

$$\frac{\partial {Y^*}^{DIFF}}{\partial \theta_d} = \frac{-(1-\theta_c)\left((1-\theta_c)n_d + (1-\theta_d)\theta_c n_c\right) w_i \Phi}{\Gamma^2}. \tag{11}$$

If $\Phi > 0$, then Equation 11 is negative. If $\Phi < 0$, then Equation 11 is positive. Differentiation of Equation 1.9 with respect to θ_c yields

$$\frac{\partial {Y^*}^{DIFF}}{\partial \theta_c} = \frac{-(1-\theta_d)\left((1-\theta_d)n_c + (1-\theta_c)\theta_d n_d\right) w_i \Phi}{\Gamma^2}. \tag{12}$$

If $\Phi > 0$, then Equation 12 is negative. If $\Phi < 0$, then Equation 12 is positive. □

Proof of Proposition 5:

Proof. Consider Equation 1.9:

$${Y^*}^{DIFF} = \frac{(1-\theta_d)(1-\theta_c)\left((\beta-1)n_g + (\frac{\beta}{1-\alpha}-1)n_c\right) w_i}{\Gamma}.$$

Differentiation with respect to θ_d yields

$$\frac{\partial {Y^*}^{DIFF}}{\partial \theta_d} = \frac{(1-\theta_d)(1-\theta_c)\Phi}{\Gamma}. \tag{13}$$

If $\Phi > 0$, then Equation 13 is positive. If $\Phi < 0$, then Equation 13 is negative. □

Proof of Proposition 6:

Proof. Consider Equation 1.14:

$$\tau_c^* = \left(\frac{1-\alpha}{\alpha}\right)\left(\frac{(1-\beta)n_g + \left(1-\frac{\beta}{1-\alpha}\right)n_c}{(1-\beta)n_g}\right).$$

Differentation with respect to n_c yields

$$\frac{\partial \tau_c^*}{\partial n_c} = \left(\frac{1-\alpha}{\alpha}\right)\left(\frac{1-\frac{\beta}{1-\alpha}}{(1-\beta)n_g}\right). \tag{14}$$

Equation 14 is necessarily negative. Differentiation of Equation 1.14 with respect to n_g yields

$$\frac{\partial \tau_c^*}{\partial n_c} = -\left(\frac{1-\alpha}{\alpha}\right)\left(\frac{\left(1-\frac{\beta}{1-\alpha}\right)n_c}{(1-\beta)n_g^2}\right). \tag{15}$$

Equation 15 is necessarily positive. □

Proof of Proposition 7:

Proof. Consider Equation 1.14:

$$\tau_c^* = \left(\frac{1-\alpha}{\alpha}\right)\left(\frac{(1-\beta)n_g + \left(1-\frac{\beta}{1-\alpha}\right)n_c}{(1-\beta)n_g}\right).$$

Differentation with respect to α yields

$$\frac{\partial \tau_c^*}{\partial \alpha} = \frac{-(n_g+n_c)}{\alpha^2 n_g}. \tag{16}$$

Equation 16 is necessarily negative. Differentiation of Equation 1.14 with respect to β yields

$$\frac{\partial \tau_c^*}{\partial \beta} = \frac{-n_c}{(1-\beta)^2 n_g}. \tag{17}$$

Equeation 17 is necessarily negative. □

Appendix B

Appendix B contains tables showing regression results for extraneous models mentioned in the paper. Table A.1 shows the results for models predicting whether or not a dictator donated a positive number of tokens. The dependent variable "DonDummy" equals 1 if the dictator donated a positive number of tokens, and equals 0 otherwise. Columns 1, 3, and 5 are probit models, while columns 2, 4, and 6 are linear probability models. Columns 1 and 2 use a more general specification, fully separating control and treatment dictator luck and earnings. Columns 3 and 4 follow the main specification from Table 3.2, also used to test Hypotheses 1 and 2. Columns 5 and 6 follow the specification used to test Hypothesis 3.

Table A.2 shows the results from various Tobit models with a censored lower bound at a donation of 0 tokens. Column 1 again follows the more general specification outlined above. Column 3 follows the main specification from Table 3.2, also used to test Hypotheses 1 and 2. Column 5 follows the specification used to test Hypothesis 3.

Table A.3 shows the results from various Poisson count data models. The columns represent the Poisson analogs to those in Table A.2, respectively. The bottom portion shows the hypothesis tests for Hypotheses 1-3.

Table A.4 shows the main specification, column 2 in Table 3.2, estimated using the residual-variance estimator HC3, as in MacKinnon and White (1985).

Table A.5 shows the main specification, column 2 in Table 3.2, with the addition of interactions between a gender dummy and each other covariate.

Table A.1: Binary Variable Prediction Models

	(1)	(2)	(3)	(4)	(5)	(6)
Model:	Probit	Lin. Prob.	Probit	Lin. Prob.	Probit	Lin. Prob.
Dict Earn			-0.00538	-0.00225	-0.0123	-0.00399
			(0.0269)	(0.0106)	(0.0287)	(0.0108)
Dict Earn C	-0.0250	-0.00836				
	(0.0390)	(0.0142)				
Dict Earn T	-0.00513	-0.00175				
	(0.0313)	(0.0118)				
Dict Luck			0.00459	0.00226	0.0122	0.00463
			(0.0223)	(0.00869)	(0.0233)	(0.00904)
Dict Luck C	0.0155	0.00608				
	(0.0411)	(0.0155)				
Dict Luck T	0.0159	0.00556				
	(0.0322)	(0.0125)				
Earn Gap C	0.0357	0.0122	0.00386	0.00193	0.0285	0.00931
	(0.0502)	(0.0189)	(0.0301)	(0.0121)	(0.0376)	(0.0143)
E Earn Gap T					-0.0310	-0.00962
					(0.0307)	(0.0110)
Luck Gap C	0.0253	0.00903	0.0224	0.00783	0.0272	0.00967
	(0.0265)	(0.00985)	(0.0241)	(0.00880)	(0.0243)	(0.00901)
E Luck Gap T					-0.0850	-0.0269
					(0.0735)	(0.0262)
Wealth Gap T	-0.0332	-0.0104	-0.00915	-0.00287		
	(0.0336)	(0.0125)	(0.0224)	(0.00847)		
Male	-1.235**	-0.434*	-1.032*	-0.381*	-1.273**	-0.441*
	(0.616)	(0.228)	(0.550)	(0.209)	(0.614)	(0.218)
Constant	1.420	0.971	0.775	0.787	0.773	0.753
	(1.938)	(0.692)	(1.547)	(0.593)	(1.624)	(0.595)
Observations	31	31	31	31	31	31
Adjusted R^2		-0.067		-0.019		-0.021

Standard errors in parentheses

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

Table A.2: Tobit Models

	(1)	(2)	(3)
Dep Var:	Donation	Donation	Donation
Dict Earn		-0.0356	-0.0424
		(0.103)	(0.106)
Dict Earn C	-0.00571		
	(0.135)		
Dict Earn T	0.00659		
	(0.114)		
Dict Luck		-0.0268	-0.0185
		(0.0874)	(0.0905)
Dict Luck C	-0.109		
	(0.147)		
Dict Luck T	0.0561		
	(0.126)		
Earn Gap C	0.0113	0.0466	0.0750
	(0.173)	(0.114)	(0.138)
E Earn Gap T			0.0143
			(0.114)
Luck Gap C	0.156*	0.121	0.129
	(0.0899)	(0.0807)	(0.0842)
E Luck Gap T			-0.0529
			(0.269)
Wealth Gap T	-0.0270	0.0419	
	(0.127)	(0.0860)	
Male	-5.518**	-5.718**	-6.003**
	(2.196)	(2.076)	(2.244)
Constant	2.427	3.957	3.834
	(6.566)	(5.720)	(5.810)
sigma			
Constant	4.152***	4.243***	4.263***
	(0.866)	(0.879)	(0.887)
Observations	31	31	31

Standard errors in parentheses

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

Table A.3: Poission Count Data Models

	(1) Donation	(2) Donation	(3) Donation
Dict Earn		-0.0265	-0.0265
		(0.0188)	(0.0189)
Dict Luck		0.00483	0.00473
		(0.0171)	(0.0172)
Dict Earn C	-0.00913		
	(0.0221)		
Dict Earn T	-0.0172		
	(0.0211)		
Dict Luck C	-0.0343		
	(0.0259)		
Dict Luck T	0.0369		
	(0.0241)		
Earn Gap C	0.0190	0.0392*	0.0382*
	(0.0246)	(0.0214)	(0.0222)
E Earn Gap T			0.0285
			(0.0184)
Luck Gap C	0.0455***	0.0374***	0.0363***
	(0.0157)	(0.0126)	(0.0139)
E Luck Gap T			0.0336
			(0.0430)
Wealth Gap T	0.0185	0.0263*	
	(0.0209)	(0.0136)	
Male	-1.899***	-2.048***	-2.033***
	(0.416)	(0.411)	(0.419)
Constant	1.419	2.174**	2.208**
	(1.079)	(0.965)	(0.986)
Observations	31	31	31
Pseudo R^2	0.3179	0.2931	0.2933

Standard errors in parentheses

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

Hypothesis		Test Stat	$p =$
H1 - Model (2):	$\frac{\partial z_i}{\partial(e_i^C - e_j^C)} = \frac{\partial z_i}{\partial(w_i^T - w_j^T)}$	0.42	0.5163
H2 - Model (2):	$\frac{\partial z_i}{\partial(\ell_i^C - \ell_j^C)} = \frac{\partial z_i}{\partial(w_i^T - w_j^T)}$	0.70	0.4034
H3 - Model (3):	$\frac{\partial z_i}{\partial(e_i^C - e_j^C)} = \frac{\partial z_i}{\partial(e_i^T - E[e_j^T])}$	0.13	0.7139
	$\frac{\partial z_i}{\partial(\ell_i^C - \ell_j^C)} = \frac{\partial z_i}{\partial(\ell_i^T - E[\ell_j^T])}$	0.00	0.9552
	Jointly:	0.16	0.9116

Table A.4: HC3 Estimator

	(1) Donation
Dict Earn C	-0.00679
	(0.110)
Dict Luck C	-0.207
	(0.131)
Earn Diff C	0.0952
	(0.226)
Luck Diff C	0.181*
	(0.0887)
Wealth Diff T	0.250*
	(0.117)
Male	-5.809**
	(2.158)
Constant	4.078
	(6.032)
Observations	14
Adjusted R^2	0.512

Standard errors in parentheses

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

Table A.5: Main Specification from Table 3.2, with Gender-Covariate Interactions

	(1)
Dep Var:	Donation
Sample:	Donation>0
Dict Earn	0.0907
	(0.151)
Dict Earn Male	0.110
	(2.956)
Dict Luck	-0.139
	(0.205)
Dict Luck Male	0.581
	(8.173)
Earn Gap C	-0.0769
	(0.263)
Earn Gap Male C	0.00424
	(3.317)
Luck Gap C	0.217
	(0.117)
Luck Gap Male C	-0.644
	(7.375)
Wealth Gap T	0.187
	(0.0961)
Wealth Gap Male T	-0.826
	(11.96)
Constant	-1.216
	(7.259)
Observations	14
Adjusted R^2	0.394

Standard errors in parentheses

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

Vita

Matthew John McMahon was born in Knoxville, Tennessee on January 21, 1988. He grew up in nearby Oak Ridge, Tennessee, and graduated from Oak Ridge High School in May 2006. He then attended Appalachian State University in Boone, North Carolina, where he received a Bachelor of Science degree in Mathematics and a Bachelor of Arts degree in Economics in May 2010. He moved to Knoxville, Tennessee the following July to study at the University of Tennessee. He received his Master of Arts degree in Economics in December 2011 and his Doctor of Philosophy degree in Economics in May 2015, both from the University of Tennessee. His work covers a broad range of fields in economics, although his main focus lies in behavioral, environmental, and experimental economics. He has accepted a position as an Assistant Professor at the University of Arkansas at Little Rock, where he will begin work in August 2015.