How Behavioral Economics Can Shape Firm Strategy and Public Policy: Lessons from the Field and Laboratory

Thesis by

Matthew Chao

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY

Pasadena, California

2015

(Defended April 30, 2015)

ACKNOWLEDGEMENTS

I owe many thanks to my advisors at Caltech for supporting my research from the ground up. I am grateful to Colin Camerer for sharing his expertise and perspective through lab meetings and one-on-one discussions, and especially for imparting me with the desire to push the boundaries of social science methodology. I am also thankful to Colin for his advice and support throughout the job market process. Matt Shum has always been available to meet at a moment’s notice, and I can always count on him for great advice on both my broader research agenda and on the minute details in my papers. Mike Alvarez has always encouraged me to pursue my ideas, which has been a great help through those inevitable times of doubt. I also especially appreciate the time that Mike spent helping me build a relationship with a public radio station; without Mike’s efforts, the third chapter of my thesis would not have been possible. I am also indebted to Marina Agranov for willingly reading my many papers and providing advice on each and every one, despite being a late addition to my committee.

I owe special thanks to Ian Larkin, my co-author and unofficial advisor from my time as a research associate at Harvard Business School. I would not have made it this far without the many hours of advice from Ian about my research, my job market strategy, and all of the idiosyncrasies and nuances of the field that would have been difficult to learn anywhere else. I am also especially grateful to Ian for all of the time that he spent networking on my behalf when I was on the market, and I am honored that he has always been willing to stake his reputation on me.

I am thankful to the many other professors who encouraged me in my stops at both Dartmouth College and Harvard Business School. I appreciate the time that Bruce Sacerdote and Catherine Cramer spent advising me on graduate school options, even after I had already left Dartmouth. I am also thankful to Peter Coles and Al Roth for encouraging my interest in a research career during my time at Harvard Business School.

There are also many others who have provided invaluable advice and support. In particular, Michael Ewens, Anthony Fowler, Jonathan Chapman, and Neeru Bhardwaj have provided numerous brilliant suggestions that have greatly improved this thesis. Others that I am indebted to include (but are not limited to) Geoff Fisher, Ben Gillen, Alec Smith, Rahul Bhui, Euncheol Shin, Romann Weber, John Ledyard, Bob Sherman, Jean-Laurent Rosenthal, Alex Hirsch, Saurabh Bhargava, Randy Wurster, and many others. I am also grateful to Laurel Auchampaugh, Tiffany Kim, and Barbara Estrada for making sure that all administrative aspects of my research went smoothly.

I have also received support from many organizations and groups. I am thankful to the National Science Foundation, the Russell Sage Foundation, and the Moore Foundation for providing funds to support my projects. I am also thankful to Matthew Rabin, David Laibson, and the Russell Sage Foundation for the wonderful summer institute that they put together, which aided me immeasurably in my understanding of the field. I am especially grateful to the local public radio station that I collaborated with, and particularly their talented development team; I am honored that they were willing to place their trust in me to help them optimize their fundraising campaigns.

Last but certainly not least, I am grateful to the many friends and family who have supported me unconditionally from the very beginning. This starts with my parents, who always cared strongly about my education and funded my college education, only to watch me switch careers multiple times both during and after my undergraduate years. I am also thankful to my brother and sister, as well as to their spouses, for tolerating my many idiosyncrasies and providing support and encouragement when needed. I am also thankful to my many friends from ABRHS, Dartmouth, and Caltech who likewise have put up with me for many years and have helped me to reach this stage.

ABSTRACT

Incentives are not always economic or monetary in nature. Individuals are often influenced by *socially-based* incentives centered on how he or she wants to be perceived in a social setting, such as the desire to publicly adhere to a norm of fairness. Likewise, individuals can also be influenced by *cognitively-based* incentives centered on self-perception and self-attribution, such as the desire to convince oneself that he or she is altruistic. Behavioral economists have incorporated some of these concepts into standard economic models of decision-making, but there is much we still do not understand about the role of these *psychology-based* incentives in organizational strategy and public policy contexts.

This thesis examines the effects of several social and cognitive incentives across different settings. In Chapter I, I use a laboratory experiment to examine why individuals feel the need to reciprocate to gifts and favors, even when those gifts are from businesses looking to take advantage of our tendency to reciprocate. Specifically, I demonstrate that individuals reciprocate simply to the intent to give a gift or favor, regardless of the ulterior motives or actual utility resulting from the favor. In Chapter II, I use observational data to investigate how pharmaceutical firms exploit social incentives that invoke reciprocity in order to influence how physicians prescribe. I concurrently examine how regulators also use social incentives, but as a way to protect consumers from these manipulative marketing strategies. In Chapter III, I collaborate with a non-profit public radio station to test the interaction between psychological and economic incentives in a fundraising context. In particular, Chapter III uses a field experiment to show that economic incentives in a fundraising campaign can reduce donation rates by detracting attention from the psychological reasons for donating, and thereby inducing a different mindset in donors.

This thesis builds upon the field of behavioral economics in two ways. First, it uses experimental methods to extend our theoretical understanding of non-monetary, psychological incentives, including some of the mechanisms that drive reciprocity, social image, and motivation crowding out. Second, the thesis applies this knowledge toward understanding the effects of several commonly used marketing campaigns and regulatory policies.

TABLE OF CONTENTS

Introduction

Social scientists have documented many categories of non-monetary incentives that can influence decision-making through psychological mechanisms. These include socially-based incentives that influence individuals through how they want their actions to be perceived by others, as well as cognitive-based incentives that influence individuals through their desire to perceive themselves in a particular way. For instance, individuals may choose to donate to a charity both because they want others to perceive them as altruistic (a *social* incentive), and because they want to perceive themselves as altruistic (a *cognitive* incentive). Both social and cognitive incentives are often based off of commonly held societal norms, such as altruism, fairness, trust, reciprocity, and the like (Cialdini 2006). This thesis explores several examples of norm-driven psychological incentives, including how they can influence the effectiveness of certain types of marketing strategies and regulatory policies.

One such psychological incentive is an individual's desire to reciprocate to a gift or favor. Reciprocity refers to our desire to respond to an action with a similar action in kind. Some version of this norm has been documented in virtually every culture (Gouldner 1960; Mauss 1950) and may even have been evolutionarily selected for (Nesse 2009; Axelrod 1984). Chapter I of this thesis investigates the cognitive motives behind an individual's desire to reciprocate to a gift or favor. In particular, Chapter I tests whether individuals care primarily about the intent of the giver, as hypothesized by the economist Matthew Rabin (1993), or whether individuals instead care more about maintaining a self-image as a reciprocator. This represents the first economic study to actually test whether individuals reciprocate purely to the intent to give a gift, independent of any change in actual pay or outcome. In addition, it is the first study to simultaneously test whether self-image can motivate reciprocity either independently or in conjunction with the perceived intent of the giver.

Reciprocity can be an extremely useful tool for marketers and salespeople (Cialdini 2006). For instance, firms often give free gifts and product samples, which can induce consumers to reciprocate with a purchase even when the gifts are unrelated to what the consumer ends up purchasing (Beltramini 1992; 2000). This is particularly evident in the world of pharmaceutical marketing, where pharmaceutical salespeople regularly provide free meals, office supplies, and other small gifts to clinicians and their staffs; such gifts have been shown to increase the volume of marketed drugs prescribed by these physicians (Larkin, Ang, Chao, and Wu 2014).

Chapter II examines whether another psychological incentive, concern for one's public image, can mitigate the influence of pharmaceutical marketing on doctors. In 2009, Massachusetts began requiring that all free pharmaceutical gifts and meals to doctors be publicly disclosed in a consumer-searchable database. Chapter II uses doctor-level prescriptions data on doctors from multiple states to measure how mandated disclosure changed prescriptions for Massachusetts doctors; doctors in other states represent the counterfactual. Importantly, Chapter II examines changes in prescriptions that occur before patients gained access to the disclosures data, and before the pharmaceutical industry made any changes to how they marketed to doctors. Thus, any observed changes in prescription patterns would likely be driven by doctor-level response to disclosure. In particular, if prescriptions of marketed drugs decreased due to disclosure, then doctors likely stopped accepting free gifts and meals in order to avoid having anything to disclose. According to past research, this change in gifts and meals would likely lead to a decrease in marketed drug prescriptions (Larkin et al 2012). This public image mechanism would also be consistent with laboratory experiments on mandated disclosures of conflicts-of-interest (Sah and Loewenstein 2014).

Non-profits often also use gifts in fundraising campaigns, but gifts can invoke a very different set of incentives in this charitable context. Individuals donate to charities for many reasons: to signal to themselves that they are charitable, to support a cause they believe in, and so on. Introducing a gift can crowd out these intrinsic motives by directing attention away from these reasons for donating. Moreover, the gift can cause donors to direct their attention toward whether the gift item is worth the amount they are donating; this shift in attention can therefore cause donors to shift from a communal, pro-social mindset towards an economic, transactional mindset, further diluting the emphasis an individual places on their intrinsic motives for donating.

Chapter III examines whether offering conditional gifts in a fundraising campaign can decrease contributions. Chapter III tests for this by collaborating with a public radio station to run a direct-mail field experiment on past donors to the station. The experiment implements two treatment conditions and a control condition; the control condition solicits donations and offers no gift. In the treatment conditions, donors are offered either a gift that benefits themselves (a thermos) or a gift that benefits others (meals to a food bank) if they donate above a specified threshold. If the gift reduces donations only through decreasing the self-signal value of a donation, then only one of the two gift types (the thermos) should decrease donation rates. Note that because these gifts are conditional in nature, they do not invoke reciprocity as the gifts in the previous Chapters do. Instead, Chapter III investigates whether in this context these conditional gifts can invoke a set of cognitive incentives relating to attention, mindset, or self-signaling.

C H A P T E R   O N E

For it is in giving that we receive.
- St. Francis of Assisi

Introduction

The Advertising Specialty Institute estimates that more than $18.5 billion is spent annually in North America on promotional products (ASI 2012), which are often handed out as free gifts to prospective consumers. According to the social psychologist Robert Cialdini, this practice profits firms in part by inspiring consumers to reciprocate to these gifts. To support this interpretation, Cialdini (2006) cites the example of the Hare Krishna. This religious group became infamous in the 1970s and 1980s for soliciting in airports by first giving passers-by a flower or candy cane before requesting a donation. This strategy proved extremely effective at increasing donations, even though many donors gave grudgingly, did not appear to like the solicitor, and subsequently threw the flower into the nearest trash bin.[1]

Reciprocation to gifts in this context seems to run counter to standard economic thinking on reciprocity. These models emphasize that reciprocity-based social preferences hinge on the perceived intentions of the first-mover (Rabin 1993). However, it is debatable whether consumers would infer kind intentions behind a small gift from a salesperson or solicitor; the ulterior selfish motives of the gift should be obvious to most individuals and could prevent the gift from being viewed as kindly intended.

Alternatively, perhaps individuals simply feel the need to reciprocate in order to preserve their self-image (or public image) as someone who returns favors and adheres to the widely accepted norm of reciprocity. Anthropologists claim that some form of reciprocity norm inevitably develops in every culture (Gouldner 1960), perhaps because it is an evolutionarily favored outcome (Axelrod 1984). Since it is such a commonly held norm, individuals may reciprocate to gifts to avoid perceiving oneself as a non-reciprocator, which may lead to feelings of guilt. Economic models suggest that guilt can motivate

[1] Knowing this, the solicitors would periodically check the trash to replenish their supply of flowers (Cialdini 2006).

individuals to adhere to commonly held fairness norms (Battigalli and Dufwenberg 2004), and it is possible that guilt can cause adherence to reciprocity norms as well. If so, then contexts that can increase self-reflection or self-image concerns could increase the likelihood of reciprocation.[2]

Existing laboratory studies cannot distinguish between intentions versus self-image motivators of reciprocity. Previous studies have only measured reciprocation to intentions when those intentions are accompanied by a change in outcome. That is, they always implement an increase (or decrease) in pay and then vary whether this increase (or decrease) is due to chance or to an opposing player's actions (e.g., Charness and Rabin 2002; Falk, Fehr and Fischbacher 2008). This design therefore cannot determine if individuals are indeed reciprocating to the intention of the giver, or if they are instead reciprocating simply because they have received a favor and must maintain their self-image by returning the favor.

This study uses a simple modified dictator game to test between these motivators of reciprocity. I incorporate gift-giving actions into a standard dictator game (i.e., the Dictator Game with Gifts, or DGG). The non-dictator is endowed with a small gift that they can give to the dictator to try and influence the dictator's allocation. These gifts can be small or large in value and either monetary or non-monetary (i.e., pens or duffel bags) in nature. This setup mimics the ulterior motives behind promotional gifts from salespeople.

The first treatment tests whether intentions alone can motivate reciprocity in this setting. When the non-dictator chooses to give a gift, the intention-to-give is announced; however, the experiment varies whether the gift is actually allowed to exchange hands. In this way, I can measure whether intentions alone are enough to engender reciprocity, as posited by Rabin (1993). This represents the first study to test whether intentions can motivate reciprocity independent from a simultaneous change in outcome.

The second treatment measures whether self-image concerns can motivate reciprocation. This treatment increases the saliency of self-image by announcing ex-ante that dictators will see how their allocations compare to other dictators' choices. Subjects will be shown approximately how their

---

[2] Similarly, in more public settings such as the Hare Krishna example, public image concerns could motivate reciprocity; after all, social image has been shown to motivate people to adhere to other norms, such as fairness (Andreoni and Bernheim 2009) or impure altruism (Andreoni, Rao, and Trachtman 2012).

allocations compare (via a percentile rank score) to those of their peers, both when they receive a gift and when they do not. [3] Since these comparisons are only shown to each dictator individually, only self-image (and not public image) concerns are invoked. If individuals fear appearing as non-reciprocators to themselves, they should be motivated to reciprocate more upon receiving a gift.

In all sessions, I control for dictator preferences over wealth distribution. Since gifting an item to the dictator changes how wealth is distributed, wealth distributional preferences must be controlled for (Charness and Rabin 2002). To accomplish this, I implement control sessions where dictators are endowed with a gift item, but from the computer instead of from their partner. Reciprocity can then be measured by comparing how dictators allocate when they receive an item from their partner versus when they receive the same item from the computer.

I find that dictators reciprocated positively to gifts from other participants. Dictators allocated on average 30% more when they received the item from their partner instead of from the computer. This increased to 50% more on average when the gift was a large-value item. In addition, 60% of *selfish* dictators (those that always allocate $0 in rounds where no item is present) allocated a positive amount on average when receiving a gift; half of this shift was driven by changes in wealth distribution (i.e., when the gift came from the computer), and half was driven by reciprocity to the gift-giver (i.e., when the gift was from the other player).

Dictators also responded very strongly to declared intentions-to-give that resulted in no change in wealth distribution. In these instances, average dictator allocations approximately doubled relative to rounds with no declared intention-to-give.[4] In addition, 22% of *selfish* dictators allocated positive amounts when a gift was attempted, but prevented by the computer. Thus, simply the intent to give a gift was enough to engender reciprocity, and higher allocations, in many dictators.

[3]To avoid feedback and learning effects, subjects are shown these comparisons only at the end of the experiment. However, subjects are reminded at the beginning of each round that these comparisons will be made at the end.

[4] In rounds with no declared intention-to-give, dictators did not know if their partner was endowed with a gift or not; thus, there was no evidence of any confounding “punishment” effects for not declaring an intention-to-give.

Self-image did not seem to motivate reciprocity in this game. The incremental effect of a gift on allocation was the same regardless of whether self-image concerns were salient. In addition, there was no interaction effect between intentions and self-image; dictators responded to intentions similarly in both the self-image and standard treatments. The lack of any interaction between self-image and intentions is key; this suggests that dictators are *not* reciprocating to the intent to give a gift because their self-image will suffer if they do not.

Instead, self-image increased baseline allocations irrespective of gifts. That is, when self-image was more salient, dictators allocated more regardless of whether they received a gift or not. This experiment therefore finds that self-image may be a motivator of wealth distribution preferences, but that an intentions mechanism (and not self-image) is the more likely motivator of reciprocity in this setting.

This is the first study to directly test these reciprocity mechanisms. Although Rabin's intentions-based model was published back in 1993, no subsequent study measured whether people reciprocate to intentions independently from any change in outcome. As a result, these studies could not rule out the possibility that individuals reciprocate due to reasons other than just intentions, such as the desire to maintain one's image as an individual that returns favors and adheres to norms of reciprocity. In addition, this is the first study to test for reciprocity using actual items and gifts, which more closely resemble favors often encountered in daily life than the more abstract favors often used in reciprocity studies. Finally, no laboratory study specifically studied promotional gifts by examining whether gifts with ulterior motives could still induce positive reciprocity. The results from this study are consistent with the hypotheses that intentions matter even for promotional gifts with ulterior motives.

Literature Review

There is plenty of evidence suggesting that promotional gifts can be optimal for the firm. Beltramini (1992, 2000) collaborated with a construction products retailer and showed that small gifts ($20-$40 letter openers) improved consumer attitudes toward the firm and increased construction product purchases after the gift was sent. Small gifts also can increase support for a non-profit; Falk (2007) showed that inexpensive gifts from a charity can increase donations fourfold.[5] The Hare Krishna solicitation strategy is another example that was already discussed earlier in this paper (Cialdini 2006).

However, there is little evidence on whether and why these promotional products are so effective. Economic models propose that individuals derive utility from rewarding kind intentions and punishing unkind intentions (Rabin 1993; Dufwenberg and Kirchsteiger 2004; Falk and Fischbacher 2006). However, it is unclear whether intentions can apply to reciprocity to promotional gifts, since these may not be perceived as kindly intended. For instance, Cialdini (2006) notes that many donors to the Hare Krishna gave grudgingly and did not appreciate the gift (many tossed the item into the trash).

Instead, there may be alternative motivators of reciprocity at work in these contexts. In the Hare Krishna example, social image or self-image concerns could have motivated people to reciprocate out of reputational concerns. This would be consistent with anthropological studies which show that gift-giving practices entail social image concerns across many primitive (Mauss 1950) and modern-day cultures (Shen, Wan, and Wyer 2011). Social pressure and image concerns have also been shown to influence other norms, including altruism (Andreoni, Trachtman and Rao 2012; DellaVigna, List and Malmendier 2012) and equality (Andreoni and Bernheim 2009).

Image motivators of reciprocity could be founded upon guilt and guilt avoidance. Battigalli and Dufwenberg (2004) suggest that preferences for fairness could be driven by a desire to avoid guilt; in particular, they stress that others' expectations matter and a person feels "pressured" to be fair because it is common knowledge that others expect him to be fair. Similarly, Blanco and co-authors (2010) propose that, in the absence of complete information, people ask themselves what they would expect if their positions were reversed, and then assume others hold similar expectations. Thus, recipients of gifts may feel the need to reciprocate because they expect that others expect reciprocation, and their self-image will suffer (i.e., they will experience guilt) if they do not oblige.

---

[5] However, the gifts used in this field experiment were also designed to invoke emotional responses beyond just reciprocity; the charity benefited orphans, and the gift consisted of postcards designed by the orphans.

On the other hand, intentions may still matter, perhaps because we are evolutionarily hard-wired to view most gifts in a kind manner. Social scientists suggest that reciprocity is an evolutionarily favored outcome since it encourages cooperative and altruistic behavior (Nesse 2009; Axelrod 1984), and this may cause us to assume gifts in neutral or positive contexts to always be kindly intended. Promotional gifts are often handed out in a neutral or positive manner (and in the case of sales professionals, often by attractive personnel), and thus we may find it difficult to infer anything but positive intentions from such gifts.

This study represents the first to test between these motivators of reciprocity. Experiments in economics have shown that reciprocity occurs in many contexts, but they often shed little insight on the motivations behind reciprocity. For instance, studies on the Gift Exchange game show that offering higher-than-equilibrium wages can induce higher effort in workers (Akerlof 1982; Fehr, Kirchsteiger, and Riedl 1993, 1998), but do not discuss the reasons why individuals reciprocate. Malmendier and Schmidt (2012) examine whether cash gifts can influence agents in a principal-agent setup. Their study confirms that reciprocation to gifts can occur in anonymous one-shot settings, but it does not causally identify a mechanism behind why this reciprocation occurs. In addition, these studies used only cash gifts, which may not generalize to non-cash gifts; people perceive cash differently and do not always respond the same to cash versus non-cash gifts (Kube, Marechal, and Puppe 2012).

Experimental Design

Experimental Task

In all sessions, each subject played three variations of the dictator game, two of which involved gift-giving decisions. Subjects were seated in front of computers and randomly assigned by the computer to one of two roles, called role 0 and role 1. These corresponded, respectively, to the dictator and non-dictator in a standard dictator game. Subjects were told that for each task they performed, the computer would randomly and anonymously pair them with a subject of the other role. Subjects were informed that they would retain their role throughout Parts I-III of Section 1 of the experiment. Parts I-III each corresponded to one of the three versions of the dictator game in the experiment. The order in which subjects played each version was randomized across sessions, but every subject played all three versions. No feedback or pay information was given to subjects until the end of the entire experiment.

One of the three versions of the dictator game, the Gift Treatment, was designed to elicit a baseline level of reciprocation to gifts. To begin, the non-dictator was endowed with an item with three-fourths probability. If the non-dictator received an item, a picture of the item would appear on the recipient's screen. The dictator was not informed whether their partner received any item. The non-dictator could choose whether to give the item to their partner or keep the item for themselves. They made this choice by clicking on either a "yes" or a "no" radio button (see Figure 1a). If they select yes, the dictator was told that their partner had given them an item, and a picture of the item appeared on their screen (Figure 1b). If the item was not given, the dictator was told that they did not receive an item from their partner (Figure 1c); the dictator was also reminded that their partner may not have had an item to give.[6] The dictator then chose an allocation out of a $10 pie. While the dictator made his choice, their partner was asked how much they expected to be given. Subjects repeated this task eight times; each time was with a new, random and anonymous partner. Subjects were told that one of the eight rounds would be selected for payment.

The Intentions Treatment was designed to measure whether subjects reciprocated to nothing more than a person's intentions. To begin, the non-dictator was again endowed with an item with three-fourths probability. Similar to the Gift Treatment, the non-dictator chose whether to give this item to their partner. However, if the item was given, the computer would prevent the gift exchange with probability one-half; when this occurred, it would be announced to both players (see Figure 2). Dictators then chose how much of a $10 pie to allocate to their partner. Non-dictators were simultaneously asked how much they expected to be given. Thus, dictators at times made allocation decisions knowing that their partner intended to give them an item, but were prevented from doing so by the computer. Subjects repeated this

[6] This is to prevent dictators from punishing their partners for not giving them an item.

task eight times; each time was with a new, random and anonymous partner. Subjects were told that one of the eight rounds would be selected for payment.

The Control Treatment did not include any gift-giving and was designed to control for preferences over total wealth distribution. Receiving a gift would unbalance the distribution of wealth between both players; the Control Treatment isolated the effect that this change in wealth distribution had on dictator allocations. Since reciprocity and fairness are closely related concepts, it is important to design the experiment in a way that can separately estimate the effects of each (Charness and Rabin 2002). To begin, one of the two players was endowed with an item with three-fourths probability. If the dictator was the one endowed with an item, then their partner would be informed exactly what item the dictator received (see Figure 3). This is identical in wealth and information to the Gift Treatment when a gift item was given to the dictator by their partner.[7] As before, dictators were then asked to make their allocation decision, while non-dictators were asked what they expected to receive from the dictator. Subjects repeated this task eight times, each time with a new, random and anonymous partner. Subjects were told that one of the eight rounds would be selected for payment.

Session Types

I ran two different types of sessions in this study: Standard and Self-Image.[8] In both session types, subjects play the Gift Treatment, Intentions Treatment, and Control Treatment, as described in the previous section. In Standard sessions, subjects play each of these Treatments exactly as already described.

[7] In some rounds, the non-dictator is the one endowed with an item, and the dictator is informed that their partner has an item. This is identical in wealth and information to when intention-to-give is announced but the gift is prevented by the computer.

[8] In pilot tests, I attempted an anonymous Social Image session where dictator allocations were overridden by the computer to $1 with 33% probability. This meant dictators could allocate $1 and obscure whether the allocation came from the dictator or the computer. In this way, dictators could theoretically avoid public image concerns from not reciprocating. This design mimics Andreoni and Bernheim (2009), but unlike their study, my study maintained subject anonymity. As a result, dictator public image was not very strongly manipulated by this design, and the pilot test showed no effect of the treatment. A true public image treatment would need to violate anonymity, as in Andreoni and Bernheim (2009) or Tadelis (2011), and may still induce stronger reciprocity.

In the Self-Image sessions, dictators still play each of the Treatments described above, but they are also informed that at the end of the experiment, they will be told how their allocations compare to allocations made by dictators in previous sessions. They are instructed that this will be reported to them via a percentile rank (see Figure 4), and that they will receive separate rankings that detail how they allocated under different scenarios (i.e., when they received a gift, when they did not receive a gift, and so on). If self-image drives reciprocation to gifts, then dictators should allocate more in response to gifts than in the Standard sessions. Likewise, if self-image drives reciprocation to just the intent to give a gift, then allocations when a gift is attempted should also be higher in Self-Image sessions. Finally, note that if self-image increases concerns over fairness in wealth distributions, then base allocations in all rounds should be higher in the Self-Image sessions.

Items

In all sessions, there were three possible items that could be given in the experiment: a small-value item, a large-value item, and an actual cash item. The small-value item was a pen (market value in the $1–$2 range). The large-value item was a small duffel bag (market value in the $10–$12 range). The cash item consisted of two one-dollar bills. When a subject received an item, a picture of the item was shown; the same picture was used for all subjects and rounds.

Studies suggest that people infer different intentions behind cash versus non-cash gifts (Kube et al. 2013). For instance, a cash gift may be interpreted as a bribe, and it may not be seen as a kind intention. By using a cash gift and a non-cash gift of similar value (the pen), I can directly test whether people reciprocate more to the non-cash item.

In addition, the Hare Krishna example suggests that even an undesirable and value-less gift, such as a flower, can have significant effects on reciprocity. By using a small-value item that is relatively undesirable, I can test whether such a gift can induce reciprocation. In addition, by comparing reciprocation to a small- versus large-value item, I can measure the degree to which the value of the gift affects the degree of reciprocity displayed.

Questionnaire

At the end of the experiment, but before being informed of the outcomes from their dictator games, subjects filled out a questionnaire. This elicited basic demographics such as age, gender, and race. In addition, the questionnaire included a 20-question Machiavellianism scale (Christie and Geis 1970). If dictators rate high in Machiavellianism, they should show less concern over their partner's final pay. Subjects were also asked to estimate the value of the items in the experiment, and to rate how much they wanted each item on a 5-point Likert scale. Subjects were also asked to talk about what they thought the experiment was about, and what likely motivated each person's actions in the game. Finally, subjects were asked to estimate the number of economics experiments they have participated in, and whether they had ever played a similar game before.

Procedure

The sessions were run at CASSEL in the UCLA Public Affairs building and at SSEL in the Caltech Social Sciences building. Subjects were recruited via e-mail blasts to registered users in the CASSEL subject pool and the SSEL subject pool. Each CASSEL session recruited a maximum of 50 subjects, and sessions averaged 28 subjects. Each SSEL session recruited a maximum of 30 subjects and averaged 16 subjects.[9] Subjects were asked to take a seat in front of a computer. Each seat also contained a packet of handouts and a unique, printed ID number.

Subjects were welcomed and given a series of instructions that were also printed in their handouts (see Appendix A at the end of the thesis). First, they were told to open zLeaf by double-clicking on the icon, as the experiment was programmed in zTree (Fischbacher 2007). Subjects were then told that the experiment was on decision-making, and that items could also be earned in today's experiment; they were told the exact items they would have a chance to earn. They were also instructed that the study was anonymous and did not contain any deception. Next, they were told that each subject would be assigned

[9] Since the SSEL lab was smaller and consisted of a different subject pool, I ran control sessions at both SSEL and CASSEL. Thus, there are control groups of the proper type for each subject population and session size.

to either role 0 or role 1. Subjects then clicked on a button in zTree, and their role assignment was displayed to them. They were then told that this would be their role assignment for all of Section 1 of the experiment. They were then instructed to enter their unique, printed ID number into zTree, as they would later be paid according to those IDs.

The instructions for Part I of Section 1 was then read out loud to them. Part I could be the Gift Treatment, the Intentions Treatment, or the Control Treatment; the ordering was randomized across sessions. The instructions explained the exact decisions that each role would face, as well as the exact information that each role would have about the other. For instance, the non-dictators knew that if they kept the item for themselves, the dictator would not be able to tell that they had withheld an item from them (since both players knew that in some rounds, no item would be given to either player). Subjects were also informed the number of rounds they would play, and that one of those rounds would be selected for payment. A similar procedure took place when subjects reached Part II and then Part III of Section 1 of the experiment (Parts II and III consisted of the remaining two Treatments).

Upon finishing Part III of the dictator games, subjects were told that Section 2 would consist of a questionnaire that would require them to answer questions about both themselves and their experience in today’s experiment. When subjects completed the questionnaire, the computer selected three rounds for payment and informed the subjects of their earnings. In the Self-Image sessions, the dictators would also see information on where their allocations ranked relative to their peers from past sessions; this was shown just prior to the pay screen.

When all subjects finished the questionnaire, subjects were then paid according to their unique, printed ID number. Gift items that were earned during the experiment were also handed out by the lab manager (at CASSEL) or the experimenter (at SSEL). At this time, the experimenter also informed subjects that there was a minimum pay of $15 for the experiment, and that those with lower pay would have their earnings rounded up to the minimum. The minimum pay was implemented to maintain compliance with CASSEL and SSEL subject pay standards.

Results

The first three sessions were run at CASSEL at UCLA. These sessions were Standard sessions. CASSEL closed in June 2013; as a result, one Standard session and three Self-Image sessions were run at SSEL at Caltech. Statistical tests show that subjects from different subject pools behaved similarly; there were no significant differences in how non-dictators and dictators acted between subject pools. I thus combine sessions from different subject pools into one dataset. I analyze the Standard sessions first.[10]

Standard Sessions

There were 90 subjects across the three Standard experimental sessions at CASSEL, and 20 subjects in the Standard session at SSEL. In the CASSEL sessions, half of the subjects were assigned as dictators, yielding 45 dictators. In the SSEL session, approximately three-fourths of subjects were dictators, yielding 14 dictators (this was done to increase the number of dictators in the smaller SSEL subject pool; since instructions did not explicitly state that pairings were one-to-one, this did not necessitate any changes to the instructions). Dictators made 24 allocations (8 in each of the Gift, Intentions, and Control treatments) for a total of 1416 dictator allocations.

Summary Statistics

Figure 5a shows the distribution of dictator allocations by treatment and item. In the first graph, the gold histogram displays allocations when the dictator received no item from their partner, and no intention-to-give was announced; the black-outlined graph displays when their partner intended to give the dictator a gift, but the dictator did not end up with the item. This graph shows that the intention-to-give shifted allocations to the right. In the next three graphs, the gold histogram displays allocations from the Control treatment (when gifts come from the computer), while the black-outlined graph displays

[10] No sessions or subjects were dropped from the data. No stopping rule was used; I ran sessions until the subject pools at each institution were exhausted (i.e., until I was unable to recruit sufficient-sized sessions within a one-month timeframe).

allocations when the item is from the other player. These demonstrate that dictators allocated more when they received the item from their partner than when they received the item from the computer.[11]

The table in Figure 5a shows that gifts also shifted dictator types. Receiving an item from the computer caused approximately 40% of *selfish* dictators that typically allocate $0 (when no gift is given) to instead allocate a positive amount; this is a wealth distribution effect. Receiving a gift from their partner caused an additional 20% of *selfish* dictators to allocate a positive amount; this is a reciprocity effect. Additionally, simply the intent to give a gift caused approximately 20% of *selfish* dictators to allocate a non-zero amount. No dictator averaged allocations of $5 or more for any treatment-item combination. 20% of dictators allocated $0 in all rounds.

Figure 5b show average allocations in each treatment-item pairing, including rounds where no item was received. Dictators allocated $0.60 on average when they received no item from the computer, $0.64 when they received no item from their partner, and $1.30 when their partner tried to give them an item but the transaction was blocked. Thus, allocations nearly doubled in response to intentions-to-give relative to other rounds where the dictator also ended up with no item.

In the Control Treatment, when dictators received an item from the computer, they allocated more to their partner. When they received a pen, allocations averaged $1.17, an increase of $0.57 from when the dictator receives nothing. A $2 cash gift from the computer increases the average allocation to $1.54. Finally, a duffel bag from the computer increases the average allocation to $1.75. These are all wealth distribution effects.

When dictators received an item from their partner, they allocated more than when they received the same item from the computer. A pen from their partner raised average allocations to $1.50, an increase of $0.33 from when they receive the pen from the computer. When dictators received $2 from their partner, they allocated $2.31 on average, or $0.77 more than when they received $2 from the computer. Finally, when dictators received a duffel bag from their partner, they allocated $2.59, an

[11] Dictators allocated essentially identically in the Control and Gift treatments when they received no item and no intention-to-give was announced. These histograms are not shown, but they overlap more or less exactly.

increase of $0.84 over what they allocated when they duffel bag came from the computer. These are all reciprocity effects after controlling for wealth distribution preferences.

Figure 6 shows non-dictator giving behavior by treatment and item. In treatments where giving was allowed, non-dictators gave their items to their partners on average 44% of the time. The pen was gifted nearly 80% of the time it was received; the cash was gifted 15% of the time, and the duffel bag was gifted 38% of the time. Clearly, non-dictators preferred to hold onto the cash, suggesting that they generally did not expect the cash gift to yield an increase in allocation of greater than $2. Note that giving patterns were identical between both Gift and Intentions treatments, suggesting that the 50% chance of having a gift returned did not influence a non-dictator's likelihood of choosing to give a gift.

Regression Analysis (Standard Sessions)

1. Base regression specification

To measure how dictators differentially responded to each type of gift, I run a linear regression model with dollars allocated as the dependent variable (Model (1) in Table 1). I add indicators that identify when the dictator ended up with the small-value item, the large-value item, and the cash item. This captures the change in a dictator's allocation caused purely by having additional wealth in the form of the item. I interact the indicators for each item with another indicator that specifies whether the item was received from their partner (as opposed to from the computer/experimenter). I also add an indicator for rounds where a partner's intention to give was announced, but where the computer disallowed the gift (henceforth referred to as "returned gifts"). Finally, I add treatment indicators to ensure that, conditional on the above explanatory variables, allocation patterns were not different across treatments. I also assume subject random effects and I cluster standard errors by subject.[12]

[12]Random effects assumptions are met because all covariates are determined exogenously to the dictator (i.e., by their randomly selected partner or by the experimenter). Subject fixed effects yield almost identical results for all regressions in this paper, and the Hausman test comparing coefficients between the random effects and fixed effects models always suggests there is no systematic difference in their coefficients ($p>0.98$ for Model (1) discussed here).

2. Wealth effects

The regression results indicate that preferences over wealth distribution can have a significant influence on a dictator's allocation. In Model (1), the baseline allocation corresponding to when the dictator had no item was \$0.60. Simply possessing the small-value item increased allocations by another \$0.58, approximately doubling the amount allocated relative to possessing no item. Possessing the large-value item similarly increased allocations by \$1.14. Possessing the two-dollar cash item increased allocations by \$0.94, thus splitting the \$2 cash item between players fairly evenly. All three wealth effects are significant at the $p<0.01$ level, suggesting that dictators had clear wealth distribution preferences.

3. Reciprocity to gifts

Model (1) also indicates that, beyond simple wealth effects, receiving an item from their partner caused dictators to reciprocate. Receiving a large-value item from their partner increased dictator allocations by another \$0.62. This effect was in addition to the increase in allocation caused by simply possessing the large-value item, which was a \$1.14 increase. Receiving the cash item or small-value item from your partner instead of from the computer increased allocations by \$0.60 and \$0.25, respectively. These effects are significant at the $p<0.05$ level, implicating a clear reciprocity effect.

4. Reciprocity to intentions

Model (1) strongly suggests that intentions matter for reciprocation. When intention-to-give is announced but the gift is disallowed, dictators reciprocate by allocating more to their partner. In fact, the effect of a returned-gift is similar in magnitude to the reciprocity effects for receiving a large-value item, suggesting that intentions alone are enough to engender a large reciprocity effect. In Model (2), I replace the returned-gift indicator with an "intentions" indicator that takes a value of 1 for any instance where a gift is given or intention-to-give is announced. Not surprisingly, since the returned-gift indicator had such a large coefficient, this alternate regression suggests that intentions alone drive all of the observed reciprocity effects (after controlling for wealth distribution effects).

5. Punishment for not giving

Finally, Models (1) and (2) suggest that dictators did not punish their partners if they did not receive a gift. The treatment indicators are either insignificant or marginally positive, suggesting that dictators allocated similarly when they received no item in a round, regardless of whether they were in a Treatment where their partner was allowed to give them an item. Thus, dictators did not punish subjects for not giving a gift. Most likely, this is due to the fact that dictators were aware that in some rounds, their partner would not be endowed with any item to give (they were reminded of this fact in every round).

6. Tobit specifications

Model (3) runs Model (1) using a panel Tobit specification that is left-censored at 0.[13] Since dictators could not allocate less than $0, there is significant "piling" at 0 which could skew the results of a linear model. The Tobit specification corrects for this and yields very similar results. As in Model (1), wealth distribution, reciprocity, and intentions all significantly influence how a dictator allocates. Thus, these results are robust to correcting for data censoring.[14] However, it is important to note that there were 12 dictators (out of 59) that always allocated 0 in all 24 rounds. These 12 dictators have no within-group variation; in the random-effects Tobit specification, they are dropped from the within-group analysis, thus yielding biased results. For this reason, for future regressions I only show OLS specifications, although Tobit regressions yield similar results in all instances.

7. Non-dictator expectations

Model (4) examines whether controlling for the non-dictator's expectations over allocations can account for the observed wealth and/or reciprocity effects. Blanco et al. (2010) suggest that reciprocity is driven by acting according to how you believe others expect you to act.[15] However, Model (4) shows that

[13] Standard errors are bootstrapped (100 re-samples) and re-sampled by cluster.
[14] Tobit specifications of Model (2) are also consistent with Model (2).
[15] Directly soliciting dictator beliefs on partner-expectations would yield biased responses if solicited directly after the actual allocation decision. Likewise, soliciting dictator beliefs directly prior to the allocation decision would bias the allocation decision. Instead, I directly solicited non-dictator expectations after the non-dictator chose whether to give a gift.

partner-expectations of the allocation are not a predictor of the actual allocation; in addition, wealth and reciprocity effects are still significant after controlling for partner-expectations.

8. Effects by treatment type

Model (5) examines whether dictators responded differently to gifts in different treatments. Since gifts given in the Intentions treatment are returned with 50% probability, it is possible that dictators view those gifts as "less" kindly intended, since they were given with the knowledge that the item might not actually exchange hands. Thus, while Models (1)-(4) separately identified reciprocity effects for each item but pooled responses from all treatments, Model (5) instead pools all items into a single category and interacts receiving a gift with each of the non-Control Treatments.[16]

Model (5) yields wealth and reciprocity effects that are very similar to those in Model (1). In the Gifts Treatments, a gift from a partner increased allocations on average $0.44, while in the Intentions Treatments, a gift from a partner increased allocations by $0.43. The 50% possibility of a gift-return in the Intentions Treatment did not seem to cause dictators to perceive the gifts any differently.

Since Model (5) only differs from Model (1) in whether gift effects are collapsed across treatment or not, the other results from Model (1) also are evident in Model (5). The returned-gift effect is identical, as is the lack of punishment effects. Model (5)'s results are also robust to Tobit specifications.

Self-Image Sessions

I ran three Self-Image sessions in the SSEL lab at Caltech, yielding 44 subjects total. As before, since SSEL has a smaller subject pool than CASSEL, I paired each non-dictator with three dictators, so that three-fourths of subjects are dictators. This yielded a total of 31 dictators in Self-Image sessions and 744 dictator allocations. Since subjects behaved similarly in the SSEL Standard session as those in the three CASSEL Standard sessions, I use all four Standard sessions as a baseline for analyzing the Self-Image sessions.

[16] I also run regressions that interact each separate type of gift with each Treatment. This results in twelve different interaction terms (3 gift types x 4 treatment types). Results are generally consistent with Models (1)–(5), but due to small sample sizes for each interaction term, most do not reach critical statistical significance thresholds.

Basic Patterns of Behavior

Figure 7a displays the distribution of allocations by treatment and item. Results are similar to what was observed in the Standard Sessions. The first graph in Figure 7a shows that intentions-to-give caused dictators to increase their allocations relative to counterfactuals. The next three graphs demonstrate that dictators allocated more when they received the item from their partner than when they received the same item from the computer. The table shows that gifts can shift dictator types; 45% of dictators allocate nothing when they receive no item, while only 13% of dictators allocate nothing when they receive an item. As in the Standard Sessions, approximately half of this effect is due to wealth distribution preferences, and the other half is due to reciprocity.

Figure 7b shows the summary statistics across treatments and items for the Self-Image sessions. The average allocation across all treatments was $2.07, more than double the $0.99 average across the Standard sessions. The average allocation is still close to $2 after excluding outliers. Conversely, dictators showed similar reciprocity effects as those observed in the Standard sessions; self-image did not cause dictators to reciprocate more than dictators in the Standard sessions. Together, this suggests that image concerns in DGG affected wealth distribution preferences but not reciprocity.

Figure 8 examines non-dictator giving behavior by treatment and item. When giving a gift was allowed, non-dictators gave their items to their partners on average 58% of the time, an increase from the 44% average in the Standard sessions. This could be because the non-dictator is aware of the increased Self-Image saliency to the dictator; knowing this, they might believe that gifts will have a greater effect and be more likely to give the gift.

Regression Analysis (Both Session Types)

1. Regression specification

To measure whether Self-Image saliency can cause dictators to allocate differently, I re-run Model (1) but only on the Self-Image sessions – this is Model (6). Model (7) includes both Standard and

Self-Image sessions and adds an indicator for Self-Image. I interact this indicator with each covariate to separate wealth distribution and reciprocity effects in Standard versus Self-Image sessions.

2. Interactions between wealth distribution and self-image

Model (6) and (7) both indicate that, in accordance with the summary statistics, preferences over wealth distribution have a larger magnitude effect on a dictator's allocation when self-image is more salient. In Model (7), the baseline Self-Image indicator has an effect size of $0.93, suggesting that dictators allocate $0.93 more than in the Standard session when receiving no item.

3. Interactions between reciprocity and self-image

The regressions also indicate that gifts cause dictators to reciprocate. Model (6) shows that dictators in self-image sessions still respond to intentions; returned-gifts led to $0.58 greater allocations on average, which is very similar in magnitude to the effect sizes observed in the Standard sessions. In addition, receiving a large-value item from their partner increased dictator allocations by another $0.47.

Model (7) shows that all types of gifts induced reciprocity, but that the interactions between self-image and each gift item are not statistically significant. This suggests that reciprocity to gifts (after controlling for wealth distribution) are relatively similar between the Self-Image and Standard sessions.

4. Robustness checks

In unreported regressions, I run similar robustness checks as I do for the Standard sessions. Results do not change when including partner-expectations as a control, or when changing the intention-to-give indicator to the one used in Model (2). Finally, these results are all robust to Tobit specifications.

Regression Analysis (All Sessions, Subject Heterogeneity)

1. Regression specification

Social preferences are often highly heterogeneous (Erlei 2008). Model (8) runs a simple linear regression on a number of demographic variables known to correlate with selfishness, including gender and Machiavellianism. I also include indicators for ethnicity to test whether guilt-based cultures (i.e., Caucasians) are more likely to reciprocate. Lastly, I also include a threshold variable representing whether a subject has participated in five or more economics experiments.[21] Since reciprocity to gifts was not statistically different in Standard versus Self-Image sessions, I simply include Self-Image as an indicator but do not interact it with each of the reciprocity indicators.

2. Baseline effects on allocation

Dictator behavior is very heterogeneous along these different characteristics. Those who have more experience participating in economic experiments allocate marginally less; they may have learned to treat experiments as a selfish source of income. Not surprisingly, Machiavellianism is also a significant predictor of selfishness. When I break Machiavellianism into its sub-parts – see Model (9) – it is the Tactics score that is driving all of the explanatory power of Machiavellianism. Naturally, the Tactics sub-score should correlate with dictator behaviors more closely than the Morals or Views sub-scores. Even when controlling for these differences, wealth effects, gift reciprocity effects, and intentions effects are still present and highly significant.

3. Interaction effects with gifts

In unreported regressions, I interact each of these subject traits (gender, experience, race, and Machiavellianism[22]) with an indicator identifying whether the dictator received a gift from their partner. I run each trait independently of the others in order to preserve power. None are robustly significant. Thus, while these traits are predictive of lower allocations overall, they do not seem particularly predictive of propensity to reciprocate to gifts.

[21] Choosing a different threshold has no substantive effect on these results.

[22] For Machiavellianism and experience, I use binary indicators representing whether the variable is above a particular threshold.

Interpretation and Discussion

This study used a gift-giving task embedded within a dictator game to measure whether intentions and self-image can motivate reciprocity to gifts given with ulterior motives. This represents the first study to directly measure whether intentions that yielded no change in outcome can still induce reciprocal behavior. This is also the first study to test whether self-image concerns can increase reciprocal behavior. Finally, this study also sheds insight on how people interpret intentions behind a gift that is clearly given in order to induce reciprocation.

In all regression analysis, intentions very robustly predicted reciprocation. When intentions to give a gift were announced, allocations increased significantly even when the gift was not allowed to be given. The effect sizes were also generally large; dictators reciprocated to a returned-gift to the same degree that they reciprocated to the cash and large-value gifts. In addition, over 20% of dictators that normally allocated $0 in standard dictator games reciprocated positively to intentions to give a gift, even when those gifts yielded no change in outcomes.

This is consistent with the hypothesis people have a tendency to view gifts as kindly intended, even when they may be given with an ulterior motive. Social scientists have suggested that our mindset towards gifts and favors is evolutionarily selected for and it could be, to some degree, hard-wired in us (Axelrod 1984; Nesse 2009) to respond with a positive mindset to any sort of gift or favor.

The results did not provide evidence that self-image can motivate reciprocity. When self-image was more salient, dictators still reciprocated, but not to a greater magnitude than dictators in the Standard sessions. Importantly, self-image also did not cause dictators to change how they responded to the intent to give a gift; dictators reciprocated similarly to intentions in both Self-Image and Standard Sessions. This suggests that reciprocation to intentions is not likely being driven by a self-image motive; that is, the intent to give a gift is not being interpreted as a favor that needs to be reciprocated in order to preserve one's self-image. Instead, the intent is more likely simply interpreted as a positive or kind intention.

This study also derived very robust results on dictator preferences over wealth distribution. Dictators allocated more when they obtained an item and their partner did not. In addition, these wealth distribution preferences were affected by the saliency of self-image concerns. Dictators implemented more equitable allocations in all of the Self-Image sessions relative to the Standard sessions. These results are consistent with models that suggest guilt-aversion can drive fairness (Battigalli and Dufwenberg 2004), and it complements evidence that suggests social image concerns can also drive fairness in wealth distribution preferences (Andreoni and Bernheim 2009).

However, it is important to note that the observed self-image effects on wealth distribution preferences could also be interfering with self-image effects on reciprocity. The self-image treatment increased allocations in all rounds, and this effect could crowd out additional effects of self-image on reciprocity. Likewise, this increase in baseline allocations could push dictators toward a ceiling on allocation amounts that they refuse to go past. This could restrict how much a dictator could reciprocate upon receiving a gift.[23]

The gifts in the lab were meant to mimic, at some level, the gifts that marketers and non-profits use to induce reciprocity. The cash item generated robust reciprocation, while the small-value item generated weaker reciprocation despite the fact that subjects estimated its price to be approximately $2, the same monetary value as the cash item. This runs counter to Kube et al. (2012), which suggested that cash is not as effective a gift as a non-cash item of equal value. This discrepancy could simply be due to the difference in desirability of each item. It could also be because Kube et al. used a *wage increase* instead of an outright cash gift; future research can help determine whether people infer slightly different intentions from a wage increase versus a cash gift.[25]

The large-value item generated the largest and most robust reciprocation results. This indicates that the value of the gift does matter; however, on average subjects reciprocated less than a dollar more in

[23]Further tests can attempt to draw this distinction; however, fairness norms are often present in the same contexts as when reciprocity norms are present, and thus this interaction between norms may accurately represent some situations in the field where reciprocity plays a role.

[25] If so, this may carry implications for employee incentive compensation; for instance, perhaps year-end bonuses are more effective at conveying positive intentions and engendering worker reciprocity than a wage increase.

response to the large-value item relative to the small-value item. This was despite the fact the large-value item was estimated by subjects as worth $14 on average. Even after factoring in the additional amount that dictators allocated due to changes in wealth distribution, the large-value item was not cost-effective for the gift-giver in this particular context. This may be because the $10 pie was too small relative to the perceived value of the duffel bag. Nonetheless, this suggests that to optimize the use of promotional gifts, firms naturally must choose a gift of the right value relative to the desired reciprocal outcome.

These results as a whole affirm the intentions-based approach toward modeling reciprocity (Rabin 1993). Perhaps surprisingly, people reciprocate positively even to gifts that are given with an ulterior motive. In addition, intentions can induce reciprocation even when those intentions do not actually change or improve the intended recipient's welfare. Conversely, individuals do not seem to reciprocate to promotional gifts (or to intentions) only due to a need to preserve one's self-image as a reciprocator.

There are additional possible motivators of reciprocity that this study does not explore. First, public image (and not self-image) may still motivate reciprocity. In the Hare Krishna example, individuals may donate so that the solicitor and the others in the airport do not view the person as a non-reciprocator. Follow-up studies would need to violate anonymity in order to test this public image hypothesis. In addition, it is possible that dictators view the intention to give a gift as *unkind*, but still feel obligated to reciprocate because the intention to give a gift leaves the dictator indebted to their partner. This would also be consistent with the Hare Krishna anecdote. Future studies should attempt to test whether this "obligation aversion" can also drive reciprocity.

References

Advertising Specialty Institute. ASI Survey: Companies thanking employees with gifts averaging $42. 2012 Press Release. Retrieved on May 27, 2013 from www.reuters.com.

Akerlof GA. Labor contracts as partial gift exchange. *Quarterly Journal of Economics 1982; Vol 97(4): 543-569.*

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

Andreoni J, Rao J, Trachtman H. Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. Working Paper, 2012.

Axelrod, R. The Evolution of Cooperation. *Basic Books 2006; New York.*

Battigalli P, Dufwenberg M. Guilt in games. *AEA Papers and Proceedings 2007: 170-176.*

Beltramini R. Exploring the effectiveness of business gifts: A controlled field experiment. *Journal of the Academy of Marketing Science 1992; Vol 20(1): 87-91.*

Beltramini R. Exploring the effectiveness of business gifts: Replication and extension. *Journal of Advertising 2000; Vol 29(2): 75-78.*

Blanco M, Celen B, Schotter A. On blame-freeness and reciprocity: An experimental study. *Universidad del Rosario Faculty of Economics Working Paper No. 85, 2010.*

Charness G, Rabin M. Understanding social preferences with simple tests. *Quarterly Journal of Economics 2002; Vol 117(3): 817-869.*

Christie R, Geis FL. Studies in Machiavellianism. Academic Press: New York, 1970.

Cialdini R. Influence: The Psychology of Persuasion. HarperBooks: USA, 2006.

DellaVigna S, List J, Malmendier U. Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics 2012; Vol 127(1): 1-56.*

Dufwenberg M, Kirchsteiger G. A theory of sequential reciprocity. *Games and Economic Behavior 2004; Vol 47: 268-298.*

Erlei M. Heterogeneous social preferences. *Journal of Economic Behavior and Organization 2008; Vol 65(3-4): 436-457.*

Falk A. Gift exchange in the field. *Econometrica 2007; Vol 75(5): 1501-1511.*

Falk A, Fehr E, Fischbacher U. Testing theories of fairness – intentions matter. *Games and Economic Behavior 2008; Vol 62: 287-303.*

Falk A, Fischbacher U. A theory of reciprocity. *Games and Economic Behavior 2006; Vol 54(2): 293-315.*

Fehr E, Kirchsteiger G, Riedl A. Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics 1993; Vol 108: 437-460.*

Fehr E, Kirchsteiger G, Riedl A. Gift exchange and reciprocity in competitive experimental markets. *European Economic Review; Vol 42: 1-34.*

Field AJ. Altruistically Inclined? The Behavioral Sciences, Evolutionary Theory, and the Origins of Reciprocity. University of Michigan Press: Ann Arbor, 2006.

Fischbacher U. z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics; Vol 10(2): 171-178.*

Gouldner, Alvin. The Norm of Reciprocity. *American Sociological Review 1960; 25: 161-78.*

Kube S, Marechal M, Puppe C. The currency of reciprocity: Gift-exchange in the workplace. Working Paper No.377, Institute for Empirical Research in Economics 2010.

Larkin I, Ang D, Chao M, Wu T. The impact of pharmaceutical detailing on physician prescribing: Quasi-experimental evidence from academic medical center policy changes. Working Paper, 2012.

Malmendier U, Schmidt K. You owe me. Working Paper, 2012.

Mauss M. The Gift: Forms and Functions of Exchange in Archaic Societies. Presses Universitaires de France, 1950

Nowak M, Roch S. Upstream reciprocity and the evolution of gratitude. *Proc. R. Soc. B 2007; Vol 274: 605-609*

Nesse RM. Social selection and the origins of culture. In Evolution, Culture, and the Human Mind. Lawrence Erlbaum Associates: Philadelphia, PA, 2009.

Rabin M. Incorporating fairness into game theory and economics. *American Economic Review 1993; Vol (83): 1281-1302.*

Shen H, Wan F, Wyer R. Cross-cultural differences in the refusal to accept a small gift: The differential influence of reciprocity norms on Asians and North Americans. *Journal of Personality and Social Psychology; Vol 100(2): 271-281.*

Tadelis S. The power of shame and the rationality of trust. Working Paper, 2011.

Trivers R. The evolution of reciprocal altruism. *Quarterly Review of Biology 1971; Vol 46: 35-57.*

Figure 1. Gift Treatment – Screenshots

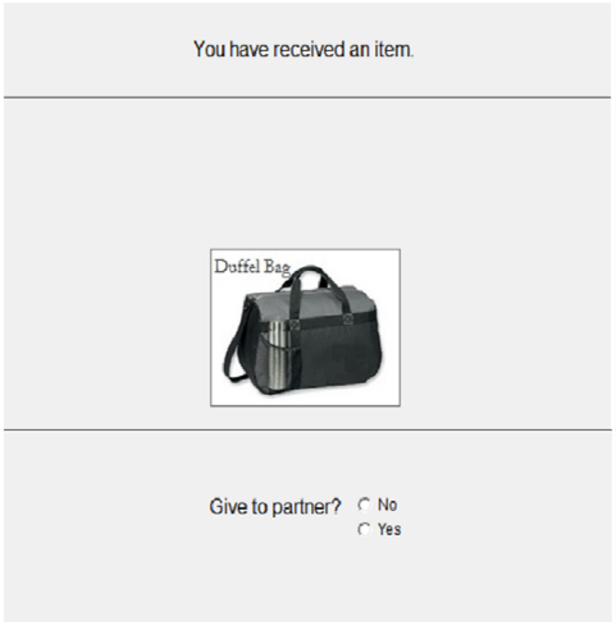

1a. Non-Dictator's Screen

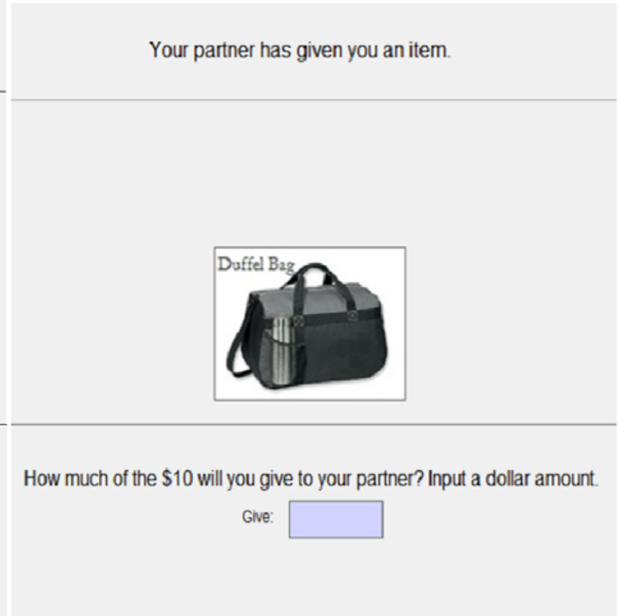

1b. Dictator's Screen

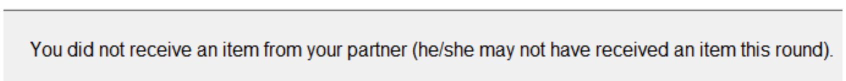

1c. Message to Dictator if No Gift

Figure 2. Intentions Treatment – Messages Upon Gift-Return

[Announced to both]: The computer has overridden an attempt to give a gift.

2a. Message to Dictator

[Announced to both]: The computer has overridden an attempt to give a gift.
[To you]: Thus, you will keep the item yourself.

2b. Message to Non-Dictator

Figure 3. Messages in Control Treatments

You did not receive an item.
Your partner received a pen; he/she knows you were informed of this.

3a. Message to non-Dictator about Dictator's item

You have received an item.
Your partner was told what item you received.

3b. Message to Dictator about what their Partner Knows

Figure 4: Social Comparison Information, Sample Screenshots (Self-Image Sessions)

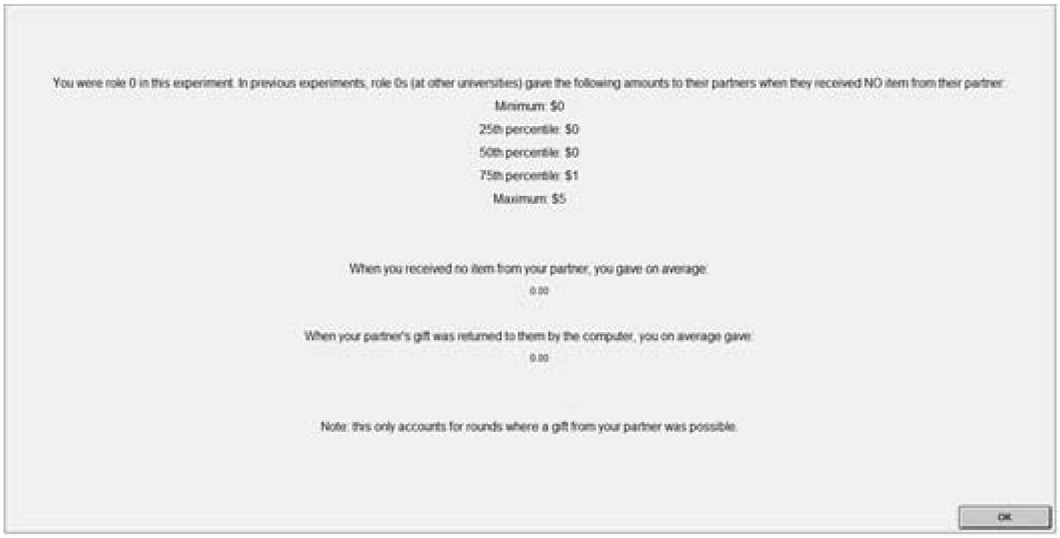

Figure 5a: Allocations by Item, Standard Sessions

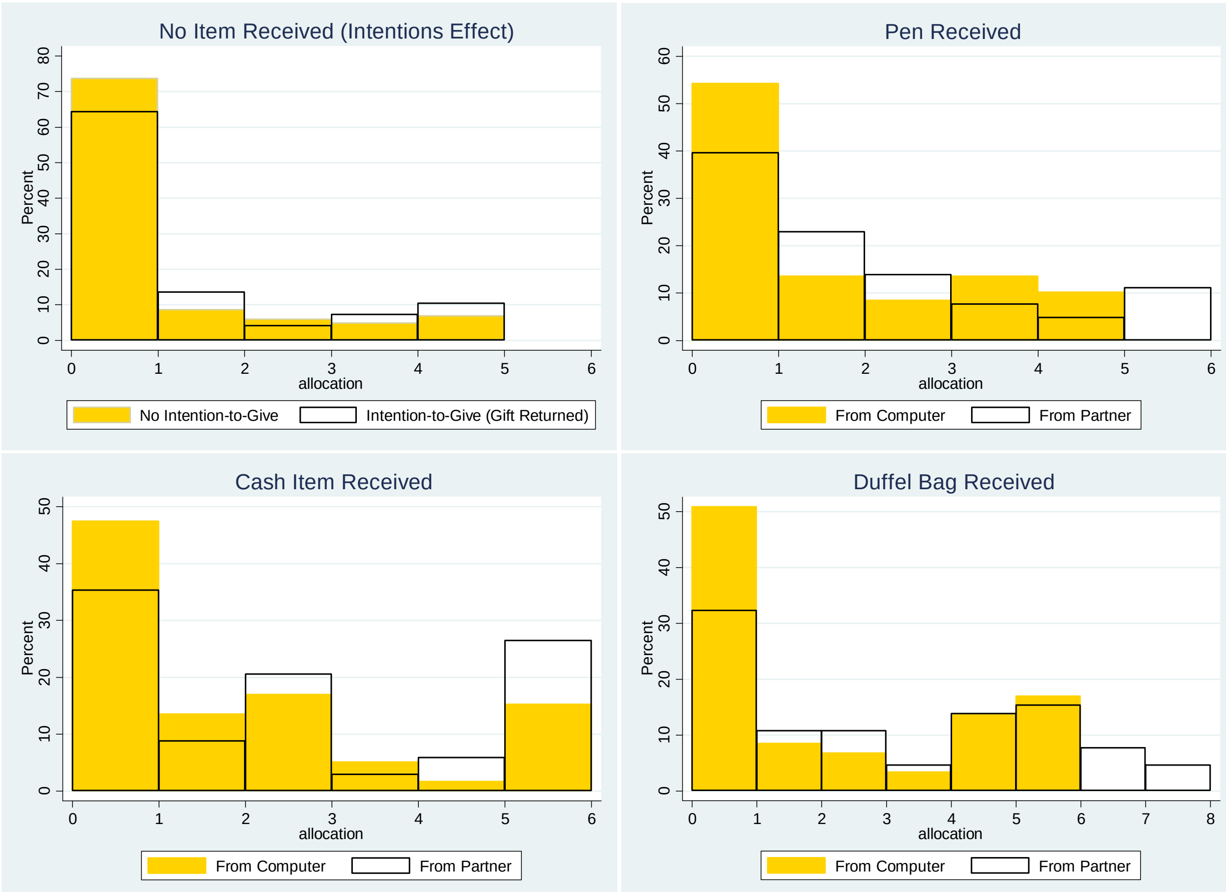

Number of Dictators Allocating $0 in all rounds
(when receiving no item, and when receiving an item)

"Selfish" Dictators	No Item	Received Item
Control	36 (61%)	23# (39%)
Gift*	32 (54%)	13 (22%)
Control + Gift	29 (49%)	12 (20%)
Returned Gift	28 (47%)	N/A^
Always Zero	12 (20%)	

*Includes rounds in the Return treatment where the gift was not returned.
#One subject here did not always give zero when receiving no item.
^When a gift is returned, the dictator naturally never ends up with an item.

Figure 5b. Allocations by Item, Standard Sessions

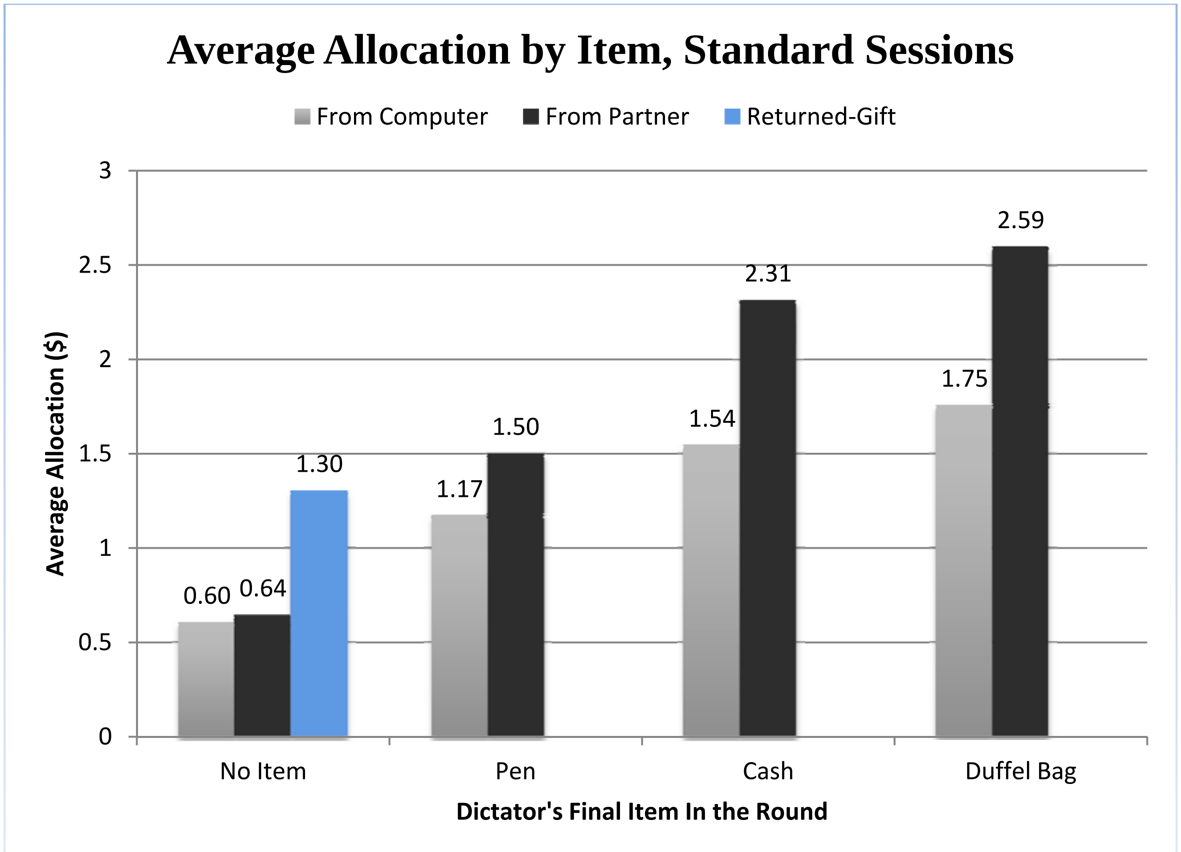

Gift From	Item	N*	Mean ($)
Computer	None	295	0.60
Computer	Pen	59	1.17
Computer	Cash	59	1.54
Computer	Duffel Bag	59	1.75
Partner	None	615	0.64
Partner	Pen	144	1.50
Partner	Cash	34	2.31
Partner	Duffel Bag	65	2.59
Partner	Gift Returned	86	1.30
		1416	1.00

*Note that there are multiple observations from each dictator within each group.

Figure 6: Non-Dictator Giving Rates, Standard Sessions

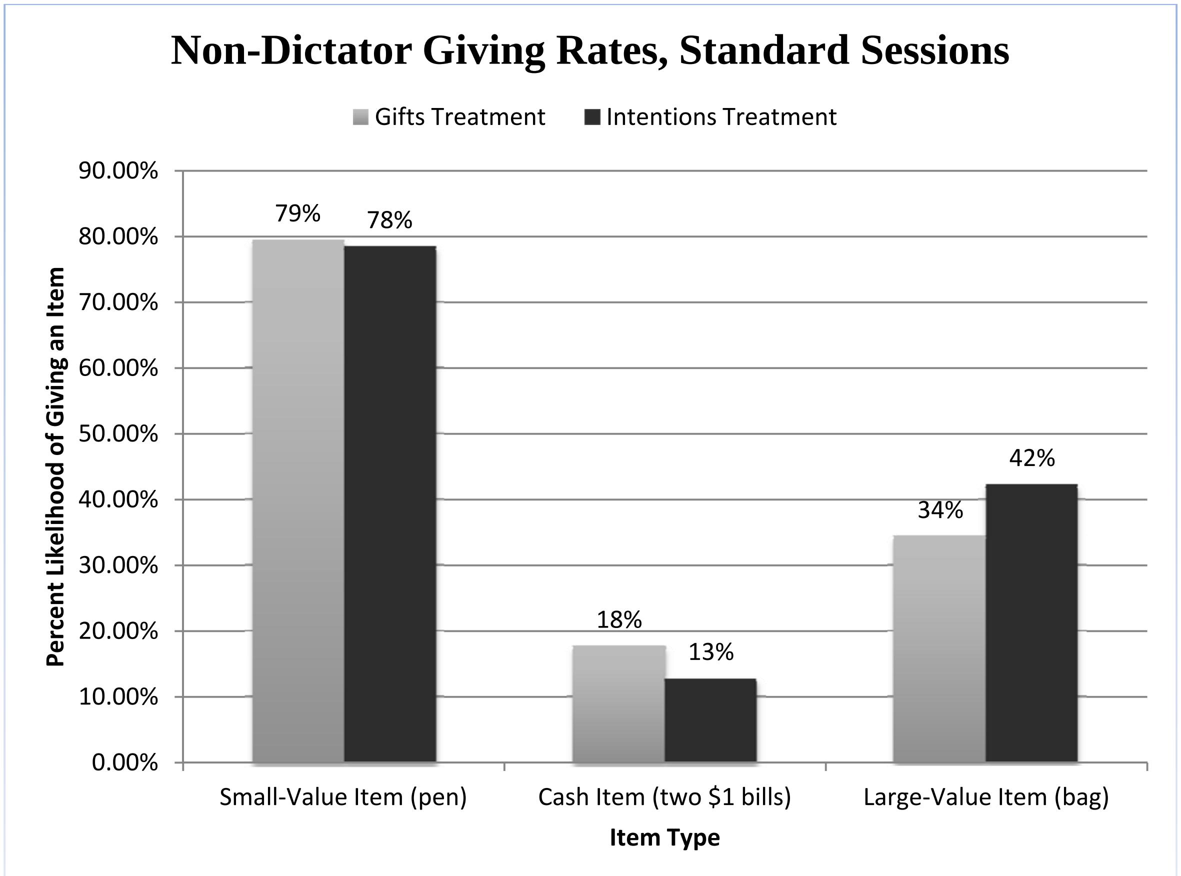

Non-Dictator Giving Behavior

Item	Treatment	Give	Keep	% Give
Small-Value Item (pen)	Gifts	81	21	79.4%
Small-Value Item (pen)	Intentions	80	22	78.4%
Cash Item (two $1 bills)	Gifts	18	84	17.6%
Cash Item (two $1 bills)	Intentions	13	89	12.7%
Large-Value Item (bag)	Gifts	35	67	34.3%
Large-Value Item (bag)	Intentions	43	59	42.2%
Total		134	172	43.8%

Table 1. Allocation by Type of Gift Received, Standard Sessions

DV = dollars allocated	**(1)**	**(2)**	**(3)**	**(4)**
Specification	OLS	OLS	Tobit^	OLS
Dictators (Rounds)	59 (1416)	59 (1416)	59 (1416)	59 (1416)
R^2	0.123	0.121	--	0.122
Treatment Indicators				
Gift	0.165*	0.145	0.455*	0.161
	(0.099)	(0.099)	(0.262)	(0.098)
Intentions	-0.027	-0.047	-0.153	-0.030
	(0.058)	(0.060)	(0.255)	(0.058)
Item Indicators				
Has small item	0.581***	0.488***	1.650***	0.588***
	(0.115)	(0.117)	(0.265)	(0.116)
Has cash item	0.939***	0.919***	2.227***	0.944***
	(0.177)	(0.176)	(0.364)	(0.177)
Has large item	1.135***	1.092***	2.441***	1.150***
	(0.218)	(0.215)	(0.389)	(0.218)
*Gift-from-Partner*Item Type Indicators*				
Given small item by partner	0.254**	-0.204	0.512*	0.261**
	(0.126)	(0.161)	(0.278)	(0.127)
Given cash item by partner	0.604***	0.077	1.051***	0.625***
	(0.192)	(0.297)	(0.395)	(0.194)
Given large item by partner	0.618**	0.107	0.792*	0.628**
	(0.276)	(0.297)	(0.4437)	(0.275)
Partner's gift was returned	0.620***		1.850***	0.634***
	(0.154)		(0.388)	(0.156)
Intention-to-give		0.551***		
		(0.142)		
Partner's E[Allocation]				-0.010
				(0.007)
Constant	0.604***	0.623***	-1.854***	0.628***
	(0.147)	(0.148)	(0.527)	(0.142)
Subject Random Effects	YES	YES	YES	YES

All OLS SEs are robust and clustered by Subject.
All Tobit SEs are bootstrapped (100 re-samples) and re-sampled by cluster.
^Tobit is left-censored at 0 (883 left-censored observations).
*p<0.10; ** p<0.05; *** p<0.01

Table 2. Allocation by Treatment, Standard Sessions

DV = dollars allocated	**(5)**
Specification	OLS
Dictators (Rounds)	59 (1416)
R^2	0.120
Treatment Indicators	
Gift	0.169*
	(0.096)
Intentions	-0.032
	(0.063)
Item Indicators	
Has small item	0.440***
	(0.125)
Has cash item	1.000***
	(0.185)
Has large item	1.214***
	(0.211)
*Gift-from-Partner * Treatment Indicators*	
Gift * Gift Treatment	0.436***
	(0.143)
Gift * Intentions Treatment	0.433***
	(0.158)
Partner's gift was returned	0.624***
	(0.159)
Constant	0.604***
	(0.148)
Subject Random Effects	YES

All SEs are robust and clustered by Subject.
*p<0.10; ** p<0.05; *** p<0.01

Figure 7a: Allocations by Item, Self-Image Sessions

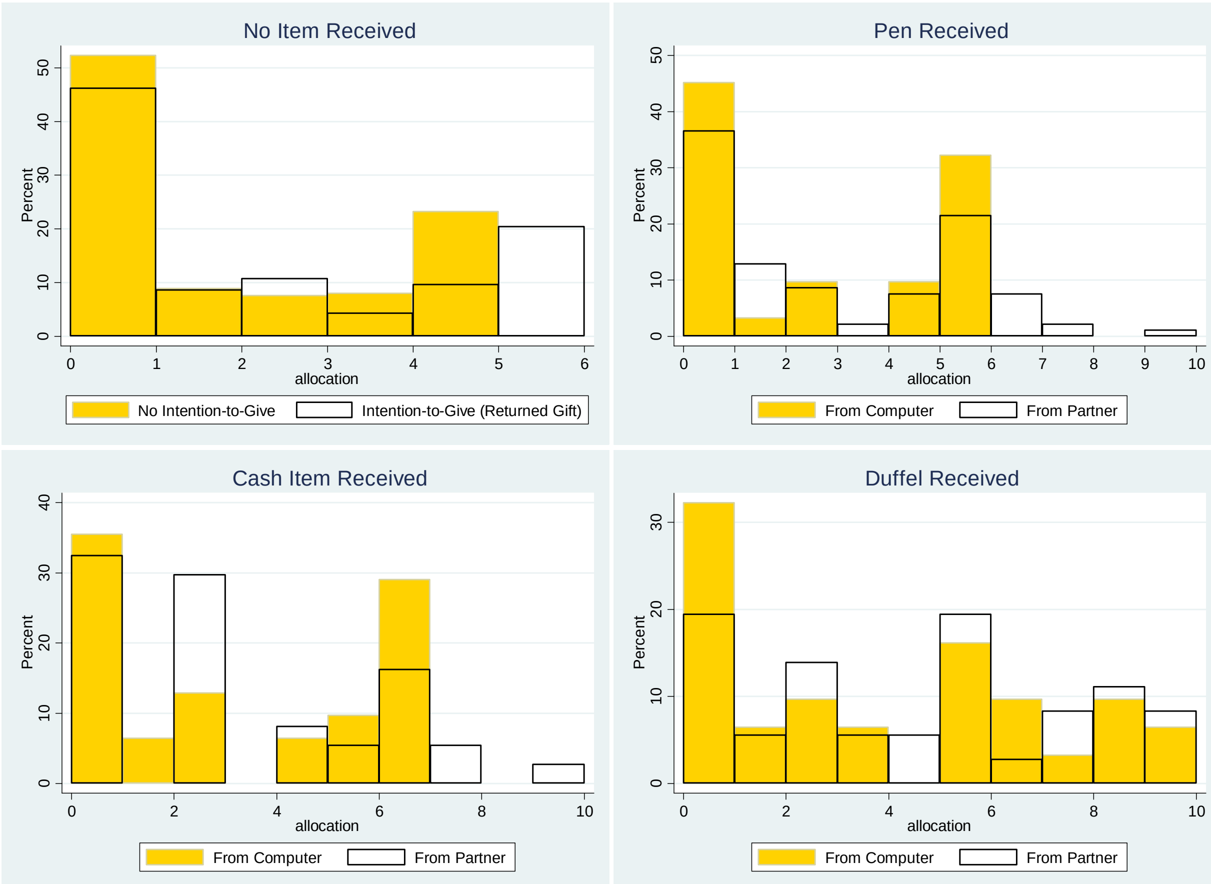

Number of Dictators Allocating $0 in all rounds
(when receiving no item, and when receiving an item)

"Selfish" Dictators	No Item	Received Item
Control	12 (39%)	9# (29%)
Gift*	14 (45%)	4 (13%)
Control + Gift	11 (35%)	4^ (13%)
Returned Gift	11 (35%)	N/A
Always Zero	3 (10%)	

*Includes rounds in the Return treatment where the gift was not returned.
#Three subjects did not always give zero when receiving no item.
^One subject did not always give zero when receiving no item.

Figure 7b Summary Statistics, Self-Image Sessions

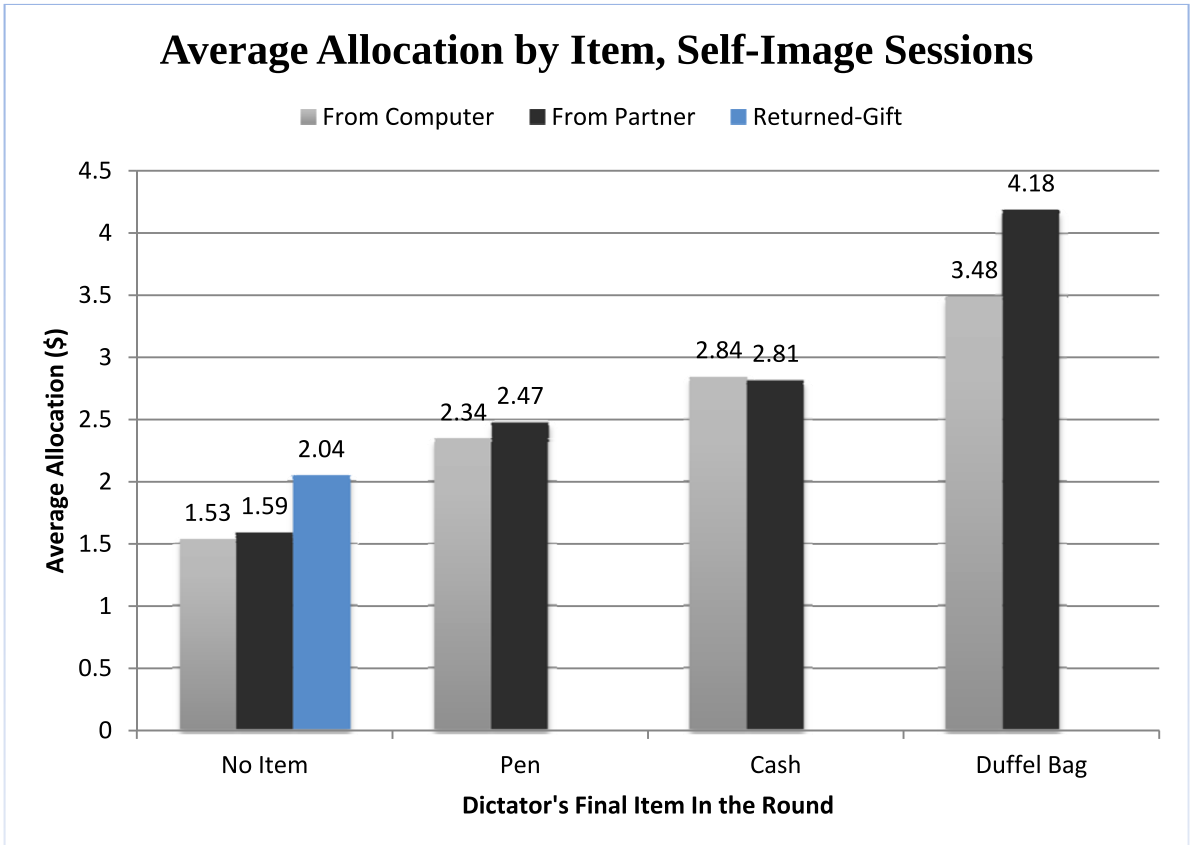

Gift From	Item	N	Mean ($)
Computer	None	155	1.53
Computer	Pen	31	2.34
Computer	Cash	31	2.84
Computer	Duffel Bag	31	3.48
Partner	None	268	1.59
Partner	Pen	93	2.47
Partner	Cash	37	2.81
Partner	Duffel Bag	36	4.18
Partner	Gift Returned	62	2.04
		744	2.07

Figure 8: Non-Dictator Giving Rates, Self-Image Sessions

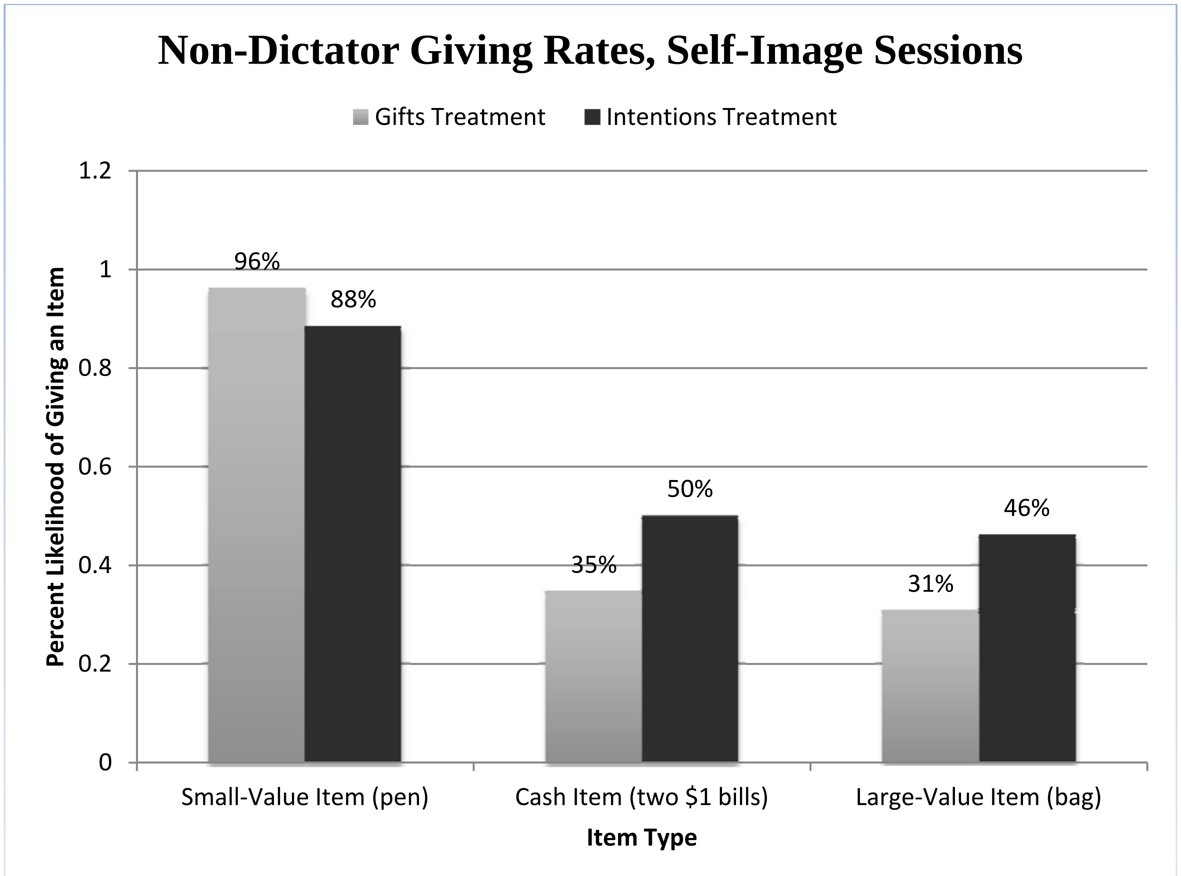

Item	Treatment	Give	Keep	% Give
Small-Value Item (pen)	Gifts	25	1	96.2%
Small-Value Item (pen)	Intentions	23	3	88.5%
Cash Item (two $1 bills)	Gifts	9	17	34.6%
Cash Item (two $1 bills)	Intentions	13	13	50.0%
Large-Value Item (bag)	Gifts	8	18	30.8%
Large-Value Item (bag)	Intentions	12	14	46.2%
Total		90	66	57.7%

Table 3. Allocation by Type of Gift Received, Self-Image Sessions (OLS)

DV = dollars allocated	**(6)**	**(7)**	**(8)**
Sessions	Self-Image	All	Self-Image
Dictators (Rounds)	31 (744)	90 (2160)	31 (744)
R^2	0.092	0.166	0.093
Treatment Indicators			
Gift	0.067	0.165*	0.072
	(0.107)	(0.099)	(0.110)
Intentions	-0.036	-0.027	-0.048
	(0.123)	(0.058)	(0.085)
Item Indicators			
Has small item	0.806***	0.582***	0.759***
	(0.249)	(0.115)	(0.230)
Has cash item	1.306***	0.939***	1.247***
	(0.312)	(0.177)	(0.289)
Has large item	1.952***	1.135***	1.915***
	(0.420)	(0.217)	(0.409)
Gift-from-Partner*Item Type Indicators			
Given small item by partner	0.102	0.253**	0.142
	(0.117)	(0.126)	(0.125)
Given cash item by partner	0.320	0.603***	0.329
	(0.250)	(0.192)	(0.254)
Given large item by partner	0.453	0.617**	0.404
	(0.282)	(0.275)	(0.296)
Partner's gift was returned	0.611***	0.619***	0.587***
	(0.234)	(0.154)	(0.224)
Partner's E[allocation]			0.025
			(0.016)
Session Indicator			
Self-Image		0.929***	
		(0.351)	
Interactions with Self-Image:			
Gift		-0.098	
		(0.145)	
Intentions		-0.012	
		(0.104)	
Has small item		0.225	
		(0.272)	
Has cash item		0.368*	
		(0.355)	
Has large item		0.816*	
		(0.468)	
Given small item by partner		-0.151	
		(0.171)	
Given cash item by partner		-0.286	
		(0.313)	
Given large item by partner		-0.163	
		(0.391)	
Partner's gift was returned		-0.009	
		(0.278)	
Constant	1.532***	0.604***	1.486***
	(0.381)	(0.147)	(0.328)
Subject Random Effects	YES	YES	YES

All SEs are robust and clustered by Subject. *p≤0.10; ** p≤0.05; *** p≤0.01

Table 4. Subject-level Characteristics

DV = dollars allocated	(8) All Sessions	(9) All Sessions
Groups / Subjects	90 (2160)	90 (2160)
R^2	0.256	0.265
Treatment Indicators		
Gift	0.133*	0.133*
	(0.075)	(0.075)
Intentions	-0.030	-0.030
	(0.048)	(0.049)
Session Indicator		
Self-Image	0.565	0.636
	(0.414)	(0.489)
Item Indicators		
Has small item	0.651***	0.651***
	(0.111)	(0.111)
Has cash item	1.062***	1.063***
	(0.157)	(0.158)
Has large item	1.406***	1.406***
	(0.204)	(0.204)
Gift-from-Partner * Item Type Indicators		
Given small item by partner	0.203**	0.203**
	(0.092)	(0.092)
Given cash item by partner	0.491***	0.491***
	(0.165)	(0.165)
Given large item by partner	0.582***	0.582***
	(0.202)	(0.202)
Partner's gift was returned	0.606***	0.606***
	(0.202)	(0.131)
Subject-level Traits		
Male	-0.024	0.014
	(0.292)	(0.291)
Machiavellianism	-0.048**	
	(0.022)	
Mach-Tactics		-0.093**
		(0.040)
Mach-Morals		0.070
		(0.127)
Mach-Views		-0.037
		(0.032)
5+ Past Studies Participated	-0.866***	-0.891***
	(0.299)	(0.305)
Caucasian (guilt-based culture)	0.016	0.064
	(0.413)	(0.397)
Constant	3.670 ***	3.751***
	(1.322)	(1.340)
Subject Random Effects	YES	YES

All SEs are robust and clustered by Subject.
*p<0.10; ** p<0.05; *** p<0.01

C H A P T E R   T W O

Introduction

The pharmaceutical industry spends up to $50 billion annually in the U.S. to build relationships with and market products to physicians (Gagnon and Lexchin 2008). Much of this money supports tens of thousands of sales representatives, who visit physicians' offices laden with free gifts and meals for physicians and their entire staffs (Rockoff 2012). Healthcare administrators have expressed concern over the impact of such activities on prescription patterns and normative patient care standards, and many have urged for stricter regulation of pharmaceutical marketing (Dana and Loewenstein 2003; Steinbrook 2009). For instance, non-profits such as The Prescription Project and physician groups such as the American Medical Student Association (AMSA) have campaigned for both mandatory and voluntary physician abstention from industry sponsored gifts, meals, and events.

State and federal governments have responded by increasing regulations on certain aspects of physician-industry interactions. In 2009, the state of Massachusetts implemented a physician payments *sunshine law* requiring that all pharmaceutical companies annually publish a list of Massachusetts-licensed doctors they provided gifts, meals, or consulting payments to. While other states have implemented marketing disclosure policies, Massachusetts is the first to implement a disclosure policy specific to industry-physician interactions, and where disclosure is to the public and not just to a government entity.[1] More recently, the U.S. Congress extended this type of disclosure to the national level via a provision within the Affordable Care Act. This provision required all pharmaceutical companies to publicly disclose all forms of compensation to U.S. physicians starting in 2013.

This study tests how the 2009 Massachusetts sunshine law affected doctor prescriptions in Massachusetts. I use a difference-in-differences design that examines pre- and post-sunshine law

[1] Minnesota, Washington D.C., and Vermont have public disclosure policies as well, but their disclosures information is not collated into an easily accessible electronic dataset. Minnesota and Vermont also ban most gifts, so disclosure in those states are limited to permissible, small-value gifts and items such as samples.

prescriptions in Massachusetts, while using prescriptions from doctors in other states as the control group. In total, I compare four years of prescriptions by 2,877 doctors in Massachusetts with four years of prescriptions by 2,878 doctors in four comparable states.

To date, there has been little empirical evidence on how these types of disclosure policies might affect prescribing. Laboratory studies suggest that disclosure can cause some individuals to opt out of conflicts-of-interest situations when opting out is possible (Sah and Loewenstein 2014). In a prescriptions context, this could imply that doctors will opt out of industry gifts and meals in response to disclosure. On the other hand, experiments also show that disclosure policies can cause those that disclose to feel morally licensed to be biased (Cain, Loewenstein and Moore 2005), since conflicts-of-interest are now common knowledge across all parties. Thus, disclosure to patients may allow some doctors to feel morally licensed to prescribe a marketed drug more often.

There could also be other negative consequences to stricter regulations on industry-physician interactions. Pharmaceutical companies publicly proclaim that these marketing activities are necessary for spreading product efficacy information to practitioners (e.g., PhRMA 2011). Without sales representatives and sponsored conferences, new drug information that is relevant to best practices could take longer to reach physicians; indeed, some academic studies have shown outcomes consistent with this hypothesis (e.g., Huang, Shum and Tan 2012).

I find that in aggregate, the disclosure law decreased prescriptions of branded, marketed drugs. More specifically, disclosure decreased branded drug prescriptions by 290,000 scripts in Massachusetts for just these doctors and these specific drugs, and for just the first 24 months after the law was passed. Importantly, disclosure also increased prescriptions of generic drugs, but only by approximately 100,000 scripts. This suggests that disclosure led to a *decrease* in total prescribing volume.[2]

These effects of disclosure are likely driven by doctor-level concern for their image. That is, doctors who did not want to harm their public image by disclosing conflicts-of-interest would stop

[2] For comparison, from 7/2007-7/2011, the dataset contains 10.7 million prescriptions, of which 4.2 million are of branded drugs.

accepting free gifts and meals. As demonstrated by previous literature (Larkin et al. 2012), these changes to their interactions with industry would then lead to changes in prescriptions. Consistent with this, doctors that had nothing to disclose post-law were more affected by disclosure than those who chose to accept gifts and meals post-law. This suggests that those with nothing to disclose likely changed their habits to a greater degree than those with payments to disclose. This mechanism is also consistent with laboratory evidence that disclosure can motivate some (but not all) agents to opt out of conflicts-of-interest when faced with mandatory disclosure (Sah and Loewenstein 2014). Importantly, doctors that accepted and disclosed gifts or meals still decreased prescriptions of branded drugs in response to disclosure, but just not to the same magnitude; thus, they displayed no evidence of moral licensing, at least in aggregate.

In addition, evidence in this paper suggests that pharmaceutical marketing influence is largely non-informational in nature. First, the results show that shocks to non-informational sources of influence such as free gifts and meals can alter prescriptions. Second, disclosure did not disproportionately affect newly released branded drugs (branded drugs that had been on the market for one year or less[3]), even though salespeople play a more informational role for these newer drugs. These results argue against industry's claims that salesperson influence occurs primarily through informational channels.

The rest of this paper is organized as follows. Section II reviews the literature on conflicts-of-interest, disclosure policies, and pharmaceutical marketing. Section III discusses the data sources and raw data trends. Section IV presents the main regression results on policy effects. Section V discusses the mechanisms that might be driving the observed effects. Section VI presents regression results that identify separate policy effects on Massachusetts doctors that disclosed meals or consulting payments after the sunshine law took effect. Section VII provides evidence on the informational versus non-informational influences of marketing. Section VIII concludes.

---

[3] Results are similar when using slightly different cutoffs for defining newly released branded drugs.

Relevant Literature

Biases from Conflicts-of-interest

There are multiple laboratory examples of how conflicts-of-interest can distort our perceptions and decision-making. Moore, Tanlu and Bazerman (2010) use a lab setting to show that when people are incentivized to advocate in favor of a position, they are later biased in their judgments when asked to be objective; this occurs even when the incentive is non-monetary and when the latter judgments are private and incentivized for objectivity. Babcock et al. (1995) demonstrate a similar effect in a judicial context.

These effects may be driven in part by motivated reasoning (Festinger 1954; Festinger and Carlsmith 1959), which refers to when people seek out and interpret information that confirms what they want to believe. In the context of conflicts-of-interest, people will be biased towards their favored position (Gneezy, Saccardo and van Vedhuizen 2013), especially when there is ambiguity in the environment (Haisley and Weber 2010). Malmendier and Schmidt (2012) further show that in a principal-agent setup, a simple monetary bribe from a third party will increase an agent's likelihood of aligning with that third party's preferences, even if that outcome is detrimental to the principal. Altogether, these studies suggest that conflicts-of-interest can significantly bias individual decision-making, both at one's own expense and at the expense of a client.

Conflict-of-interest Disclosure Policies

Sah and Loewenstein (2014) examine whether mandated or voluntary disclosure can cause agents to opt out of conflicts-of-interest when opting out is possible. This study implemented a simple adviser/advisee setup, where the adviser was given more information about a risky gamble and was responsible for providing advice on risky choices to the advisee. The advisers were allowed to choose between two incentive schemes, only one of which introduced a conflict-of-interest with their responsibilities to the advisee. With mandated disclosure, more than half of all advisers chose to opt out of the payscheme that introduced a conflict-of-interest between their own pay and their responsibilities to their advisee, even though this payscheme could yield higher payoffs to the adviser.

However, Cain, Loewenstein and Moore (2005) demonstrate how disclosure of conflicts-of-interest can fall short of correcting for these biases. In particular, the authors show that when conflicts-of-interest are disclosed, agents may exaggerate their bias in anticipation that others (e.g., the principal, aka the patient) will discount their advice. In a second, similar study by the same authors (2011), they demonstrate that agents believed it was more morally acceptable to exaggerate their advice when the conflict-of-interest is disclosed. This suggests that agents can feel morally licensed to be biased when their conflict-of-interest is known to the recipient of their advice.

Some medical studies have examined disclosure in medical publications. One such paper found that pharmacology studies with results supporting the use of calcium channel antagonists were much more likely to contain financial conflict-of-interest disclosures than studies with neutral or negative results (Stelfox, Chua, O'Rourke, and Detsky 1998). However, this suffers from endogeneity; do researchers tend to find results that support their sponsor, or do companies simply sponsor research that is more likely to find support for their products? A similar study by Perlis et al. (2005) suffers from the same endogeneity concerns. Finally, since disclosure is always mandated in these contexts, these studies do not directly test whether disclosure has any causal effect on each paper's findings and recommendations.

Pham-Kanter et al. (2012) used the 2004 Maine and West Virginia sunshine laws to measure the effect of disclosure on doctor prescriptions. They examined selective serotonin reuptake inhibitors (SSRIs), a class of antidepressants, as well as statins, which treat high cholesterol. They found that these disclosure policies did not cause any changes to the rates of branded drugs prescribed relative to neighboring states. However, these states' disclosure policies did not require public disclosure of payments; instead, disclosure was only to a state government department. Thus, doctors knew that their patients had no information on physician-industry interactions, and the effects measured by this study have no relation to *public* disclosure. These results may not generalize to the effects of the Massachusetts or federal disclosure laws, both of which require disclosure via a publicly searchable database.

Biases from Pharmaceutical Marketing

Researchers have speculated that doctors with conflicts-of-interest are likely biased in their decision-making (Dana and Loewenstein 2003). Within the medical literature alone, enough has been published to require three successive meta-analyses (Lexchin 1993; Wazana 2000; Spurling et al. 2010). More than half of the studies examined in these meta-analyses have indicated that marketing activities can alter prescription decisions or related aspects of patient care; however, over one-third of those studies found no such effect. In addition, many of these medical studies are small-sample cases that focused on a small group of physicians and drugs (e.g., Orlowski and Wateska 1992; Loertscher et al. 2010; Chren and Landefeld 1994), and almost all only looked at correlations. Thus, their results are difficult to generalize.

The marketing literature also struggles to handle endogeneity in doctor-industry relationships. These studies obtain data on physician-level prescriptions over time, as well as physician-level measures of marketing exposure over time. They then run regressions to estimate the effect of marketing exposure on prescriptions (e.g., Gonul et al. 2001; Mackowiak and Gagnon 1985; Manchanda and Chintagunta 2004; Mizik and Jacobson 2004; Rizzo 1999). However, this methodology is confounded since pharmaceutical companies selectively target only a subset of physicians. The companies track physician prescriptions on a bi-weekly basis and change their marketing tactics in response to changes in a physician's prescription patterns (Fugh-Berman and Ahari 2007). Researchers attempting to use actual data on marketing and prescriptions cannot determine whether changes in marketing exposure led to changes in prescribing, or if changes in prescribing led to changes in marketing exposure.

Most recently, a paper by Larkin et al. (2012) used a quasi-experimental methodology to circumvent these endogeneity and small-sample constraints. The authors treated changes in academic medical center policies that governed marketing activities as exogenous to physician-industry interactions. Therefore, changes in prescription patterns (relative to controls) could be directly attributed to the effects of bans on various marketing activities.

I use a difference-in-differences approach akin to Larkin et al. (2012), but I instead evaluate the effects of the Massachusetts disclosure law on prescriptions. I also identify doctor-level effects by incorporating data made available through the Massachusetts disclosures database. This lets me identify doctors that accepted payments from industry and measure whether this sub-group of doctors responded differently to disclosure. Finally, by examining drug-level differences, such as whether newly released drugs are differentially affected, I can evaluate whether the observed effects are driven by informational or non-informational sources of influence. This evaluates the accuracy of industry claims that salesperson influence is solely through informational channels, such as the distribution of clinical trials results.

Data

Data Overview

The base dataset can be broken down into five main components: (1) physician–hospital affiliations over time for all physicians affiliated with any academic medical center (AMC) in one of six metropolitan regions (MSAs), one of which is in Massachusetts;[4] (2) monthly prescriptions by each physician from January 2006 to December 2012 for all drugs within nine drug classes; (3) drug characteristics such as whether a drug is branded and when a generic version became available for a branded drug; (4) the list of doctor names published by the state of Massachusetts for having accepted industry payments any time between July 2009 and December 2012; and (5) details on all academic medical center (AMC) marketing policies. Each of these components is listed in more detail in Figure 1.

(1) Physicians and Affiliation Data

The data consists only of physicians who are affiliated full-time with an academic medical center (AMC). I make this choice for several reasons. First, since AMC-level policies on marketing activities are publicly available, using AMC-affiliated doctors allows me to control for the marketing activities that are allowed or not allowed at every physician's workplace. I can therefore evaluate whether AMC policies are complements or substitutes to the Massachusetts disclosure policy. Second, full-time attending

[4] These MSAs are chosen simply because they contain the most AMCs within each region.

physicians at AMCs are generally salaried clinical faculty members whose patients are derived through referrals from primary care physicians. Thus, both income and patient volume for these doctors are relatively independent from their actual reputation or public image. If the results suggest that these doctors are averse to disclosing industry conflicts-of-interest, I can better isolate the reasons why these particular doctors might care (or not care) about their reputation. Finally, I choose these doctors because their 2006-2009 prescriptions data were already available to me through a previous paper (Larkin et al. 2012).

Next, I obtained a list of all physicians that were full-time affiliated with any of the 35 AMCs in these MSAs.[5] This data was obtained from IMS Health, a leading pharmaceutical market research firm. IMS identified all hospitals that were owned or otherwise governed by these AMCs.[6] IMS also surveys hospitals on a quarterly basis in order to obtain physician–hospital affiliations data. IMS identified all attending-level (FTE) physicians affiliated with any of these hospitals for at least one quarter between January 2006 and June 2009. In total, 9,998 physicians met these criteria and are initially included in this dataset. Physicians did not switch affiliations very frequently in the dataset (only 12% switch at any point in the data); for affiliations data corresponding to July 2009 – December 2012, I simply carry forward the affiliations data from June 2009.[7]

(2) Prescriptions Data

IMS Health also provided physician prescriptions data. IMS compiles monthly doctor-level prescriptions data by purchasing data directly from retail pharmacies. In total, IMS purchases data on approximately 75% of the retail market and projects the rest using geographic and demographic variables. This is the most comprehensive dataset available on the market, and IMS is widely considered the best source of prescription data by pharmaceutical companies and academic researchers alike. However, since

[5] Physicians that did not write prescriptions for drugs in this study are excluded.
[6] Note also that, per IMS, no hospital-AMC affiliations changed during this timespan.
[7] Additional affiliations data for 2010-2012 would require additional funding to purchase; in addition, IMS does not hold on to historical affiliations data, as affiliations data is only useful to their pharmaceutical clients if it is current information. This implies that doctors newly affiliated with these AMCs after July 2009 are not included in the data.

the data represents only the retail market, prescriptions filled at hospital pharmacies or by mail are not included. Additionally, prescriptions that are written but not filled are not observed in the data. IMS provided total prescriptions filled for every physician-drug-month in the dataset.

The dataset included monthly prescriptions data for 377 drugs in nine drug classes of varying sizes: statins, antihyperglycemics, proton pump inhibitors, antihypertensives, sleep aids, antidepressants, anxiolytics, antihyperactives and antipsychotics. These drug classes were specifically chosen based on physician feedback, which identified these classes as a composition of drugs that are heavily marketed and not marketed at all. For each of these drug classes, all drugs within the class (according to IMS classifications) were included except for drugs that were rarely, if ever, prescribed during this time span.

This prescription data was purchased in two batches. The first, purchased in 2010, covered 218 drugs from January 2006 – June 2009 and was used for the study Larkin et al. (2012). The second batch, purchased in summer 2013, covered all original 218 drugs that were still commonly prescribed (some were no longer produced or commonly prescribed by June 2009); in addition, this second batch added new drugs from these drug classes that entered the market around or after June 2009 and were commonly prescribed, thus bringing the total number of drugs in the study to 377.

(3) Additional Drug Data

Additional details about each drug are obtained through the FDA's database on FDA-approved drugs. This data identifies whether a drug is branded, and if so, whether and when a generic version of it became available on the market. In this paper, to proxy for whether a drug is marketed in a particular month, I simply identify whether the drug is branded and whether a generic version was on the market in that month. This measure is highly correlated with whether a drug was actually marketed by salespeople in that month.[8] Additional data provided by the FDA database include the FDA approval date and manufacturer(s) for each drug.

(4) Massachusetts Disclosures

Data on Massachusetts doctors that accepted industry payments in July 2009-December 2012 are publicly available and were downloaded from the Massachusetts Department of Health and Human Services. This dataset includes the name of each doctor and the type of compensation they received (e.g., consulting fees, gifts, or food), as well as the valuation of what they received. This data was merged to the IMS dataset using doctor name; since doctor name is not necessarily a unique identifier, this introduces noise into the data, but it should not bias the results.

Table 1 summarizes the disclosures data. In total, there are 4,278 doctors in my data that were affiliated with a Massachusetts AMC after the sunshine law was implemented, although only 2,877 remain after data exclusions (see the subsequent sub-section titled "Data Exclusions"). Of these, 430 doctors accepted a meal of $50 or greater at least once and 370 doctors accepted consulting payments at least once. Conditional on having accepted a meal, the average doctor-year value of meals accepted is $235. Consulting payments are much higher; conditional on having accepted a consulting payment, the average doctor-year value of consulting payments was $12,107.

Note the following two caveats on Table 1. First, in 2009, only 6 months of data are included, since the policy took effect in July 2009. Second, in 2012, there are fewer entries; while this may partly be because doctors are accepting fewer meals and consulting payments over time, it is mostly due to a change in reporting requirements. For 2012 and onwards, the Massachusetts state legislature relaxed the requirements on what types of payments have to be reported; in particular, payments already reported to the federal government due to the federal sunshine law (which was supposed to begin tracking payments in 2012, although they did not start until 2013) did not have to be reported to Massachusetts.

---

[8] I have monthly data on the number of salespeople assigned to each drug from January 2006-June 2009. This measure correlates almost perfectly with whether a drug is branded with no generic version on the market. I use the latter as a proxy for whete a drug is marketed in a given month, since I do not have sales force data for 2010-2012.

(5) AMC Policies

In total, the 9,998 physicians in the dataset held full-time affiliations with 35 different AMCs that were located in or near the six chosen MSAs. I obtained dated hard copies of policies for all 35 AMCs through the conflict-of-interest policy database maintained by the Institute of Medicine as a Profession (IMAP). This organization, whose mission is to shape medical professionalism, regularly submits requests to all U.S. medical universities asking them to provide copies of their marketing policies. In private phone interviews with over 25 hospital administrators, I confirmed the information contained in many of these policy documents. I classify each policy according to the restrictions they place on a variety of marketing activities, including restrictions on accepting gifts, consulting contracts and free meals. These are listed in Table 2. Many of these policies followed a general set of recommended rules provided by the Association of American Medical Centers (AAMC) in a 2008 publication.

The minimum policies in place at every AMC are determined by PhRMA self-regulation guidelines that were published in 2002.[9] Among other restrictions, the PhRMA regulations restricted gifts to $100 per salesperson visit and also banned purely recreational gifts such as golf outings and concert tickets. These trade regulations theoretically represent the rules that governed marketing behavior at all AMCs regardless of AMC policy.[10]

The Massachusetts doctors in the dataset are mostly affiliated with one of four AMCs. These are the medical centers owned or operated by Boston University, University of Massachusetts, Tufts University, and Harvard University. Each of these implemented restrictions on industry-physician interactions at different points in time. BU and UMass implemented restrictions before the sunshine law took effect. The remaining two, Tufts and Harvard, implemented restrictions after the sunshine law was already in effect. Some of the Massachusetts physicians in the data are affiliated with a hospital owned by Partners Healthcare, a Massachusetts non-profit organization that is affiliated (but not owned by) Harvard and Tufts, and which also implemented their own marketing restrictions in 2009. Each of these

---

[9] PhRMA is a trade group including all major U.S. pharmaceutical manufacturers.
[10] Physicians and administrations indicated that most salespeople were obedient to these guidelines.

institution's policies varied in the types of restrictions placed on marketing activities; for instance, Harvard only bans personal gifts to doctors, while UMass bans all gifts, even those that are educational (e.g., medical textbooks) or professional (e.g., surgical shears) in nature.

Data Exclusions

When analyzing the data, I exclude several types of observations from the data. Primarily, this is done to keep the analysis and interpretation as simple and clean as possible. *Importantly,* statistical significance results in this paper are virtually all robust to analyzing the full dataset; in fact, disclosure effects are often larger in magnitude and greater in statistical significance when using the entire dataset.

First, I drop *doctor:drug-class* pairings that are almost always zero. For instance, the psychiatrists in the dataset will prescribe many drugs from the antidepressants class, but they virtually never prescribe any drug from the statins class. Likewise, cardiologists do not prescribe antidepressants frequently, but do prescribe many statins. I drop all doctor:drug-class pairings where the doctor prescribed fewer than 100 scripts of an entire drug class over all 84 months of the data. This is an extremely low threshold; a doctor would essentially have to prescribe no more than a single script per month of an entire drug class to fall under this threshold. This eliminates pairings such as psychiatrist-statins or cardiologist-antidepressants, and this reduces the dataset from approximately 24 million doctor-drug-month observations and 10,000 doctors to slightly over 16 million observations and 7,400 doctors. In unreported tests I determine that policies do not affect how much of a particular drug class a doctor prescribes.

Second, I drop doctors who, according to IMS Health designations, have full-time attending affiliations to multiple AMCs for at least one quarter in the sample. These doctors may split their time between many hospitals, making it more difficult to assess the exact marketing regulations they are subject to at all times. This reduces the data to approximately 14 million observations and 6,300 doctors.

Third, I drop all doctors who switch hospitals or states at some point in the dataset. This rules out the possibility that doctors switched affiliations in response to any policies. This reduces the data to approximately 13.5 million observations and 5,755 doctors. 2,877 of these doctors are in Massachusetts, and 2,878 are in other states.[11]

Finally, I focus all analysis on just the 24 months before and after the sunshine law was implemented. I exclude the months January 2006 through June 2007 in part because these coincide with the passing and implementation of Romney Care in Massachusetts. This bill required all Massachusetts residents to be covered under a health insurance policy, which could theoretically have some effect on how drugs were prescribed by doctors or filled by pharmacies. In addition, by focusing only on the two years before and after the treatment event, I reduce the likelihood of other events impacting the fixed differences assumption in the model. Importantly, the effects of the sunshine law are almost always of similar or greater magnitude and higher statistical significance when utilizing the full seven years of data.

Data Summary and Parallel Trends

Figure 2 displays the raw data for branded drug prescriptions for the two years before and after the sunshine law was implemented. Note that branded drugs refer only to drugs that do not have a biochemically equivalent generic version on the market for a given month.[12] The two lines represent the treatment and control groups. The MA group consists of doctors that are always affiliated with a Massachusetts AMC and are thus subject to the disclosure law as of July 2009. The Other group consists of doctors affiliated with AMCs in Pennsylvania, Illinois, New York, and California; based on their hospital affiliations data, none of these doctors were ever subject to any state disclosure requirements.[13]

---

[11] I deliberately increased the number of Massachusetts doctors by including Partners Healthcare doctors, a Massachusetts private hospital group affiliated with Harvard and Tufts, but not owned by them like the other hospitals in the dataset. I do this because Partners issued their own hospital-level marketing policies, which I have data on. In addition, these doctors are similar to the other doctors in that they are clinical faculty and full-time affiliated with an AMC. All main results are robust to excluding Partners Healthcare doctors.

[12] Branded drugs that have a generic version already on the market are classified with the generics, since they are virtually never marketed and are also rarely prescribed. Even if a doctor prescribed such a drug, the pharmacist would likely replace it with the generic version (most insurance companies often require this), and the filled-prescriptions data from IMS would only reflect a prescription for the generic.

[13] Note that if I break this graph out by individual state, New York is an outlier; branded drug prescriptions are relatively flat over time, and do not decrease the way the other states do. Since I cannot identify why New York differs, I have no exogenous reason for excluding it from the data. However, all results are robust to dropping New York (and indeed, dropping any single counterfactual state) from the analysis.

Since I have limited the data to doctors that never changed affiliations, the number of doctors in each group is constant; I therefore simply plot the total number of branded prescriptions that each group prescribed in each month.

Figure 2 provides a visualization of the parallel trends assumption necessary for analysis (see Figure 2B for a blow-up of the relevant trends). Figure 2 and 2B demonstrate that, over the one-year period prior to the treatment event, fixed differences in branded prescriptions is plausible between Massachusetts doctors and those in other states. Figure 2 also overlays a bar plot of the exact differences in each month; from August 2007 through August 2008, differences appear relatively fixed.

Figure 3 displays the full trends over the full seven-year period. Note that in 2006 and 2007, the difference in total branded prescriptions between Massachusetts doctors and other doctors is shrinking. This is simply because this set of Massachusetts doctors contains a significantly higher number of doctors prescribing psychiatric drugs than the set of control doctors. This is likely due to a difference in the distribution of doctor specialization between groups.[14] This difference in distribution matters because in 2006 and 2007, many highly prescribed psychiatric drugs went off patent, including several highly prescribed drugs: Xanax XR, Zoloft, Effexor, Wellbutrin XL, Ambien, Ambien Pak, Focalin, Paxil CR, and several others. Due to the difference in the distribution of doctor specializations in the treatment and control groups, branded prescriptions in Massachusetts are much more affected by these patent expirations. Figure 3B plots the non-psychotherapeutic branded prescriptions over time between both groups, and this illustrates that the trends are much more parallel when these psychotherapeutics are removed.

By similar logic, the trends in 2006-2007 are also parallel when the graphs are limited to branded drugs that remain on patent throughout the entire dataset. This analysis abstracts away from the effect of drugs being removed or added, but it also prevents the analysis from considering the full set of drugs within each class. Figure 4 displays prescriptions for just these drugs, and it further confirms that the

[14] 2134 out of 2877 MA doctors (74%) prescribe at least one psychiatric drug class in this dataset, while only 1912 out of 2878 doctors in other states (64%) do so. As before, these figures are after dropping doctor:drug-class pairings where doctors prescribed fewer than 100 scripts of the entire class during this seven-year period.

discrepancy across groups in Figure 3 from 2006-2007 was due to differences in the number of doctors prescribing branded drugs that came off patent in those years. In addition, Figure 4 demonstrates that the decreasing trend in branded prescriptions over time (seen in Figure 3 in both treatment and control groups) was also due to drugs coming off patent. Importantly, all results in this paper are robust to including only this subset of drugs. Figure 4B magnifies the graph for the two years before and after the policy.[15]

If there is a policy effect on branded prescriptions, these graphs should show a change in *slope* around the time of policy implementation, but not a discontinuous jump in branded prescriptions. This is because the prescriptions data includes prescription refills, which constitute the vast majority of total prescriptions filled at retail pharmacies. Patients already on medication that works for them will continue to stay on that medication regardless of changes to disclosure requirements (especially for the conditions treated by the drugs in this data, e.g., hypertension, depression, high cholesterol, etc.). Thus, refills are not likely to be affected by disclosure. Instead, disclosure will primarily affect first-time prescriptions, thus leading to a change in slope over time to branded prescriptions but not a discontinuous jump.

Figures 2 through 4 are suggestive of a change in slope in response to disclosure. The figures show relatively parallel trends in the year or two before the policy is implemented. In the time between when the law is signed (August 2008) and implemented (July 2009), branded prescriptions in Massachusetts begins to fall relative to the counterfactual. This likely reflects that in response to the law, doctors begin to cut off industry ties sometime in this eleven month period. This makes sense since doctors likely heard of the impending disclosure requirements sometime during this time period. For instance, the state released a summary of discussion points in December 2008 and announced the finer points of the policy in May 2009. Note also that the law was discussed during legislature sessions that began in late January 2008 and continued through the end of the sessions in July 2008, so the medical community may have been aware of the proposals even prior to the law being signed.

[15] Figure 4B still excludes branded drugs that come off patent any time during the full seven-year dataset.

Figure 5 provides the same graph as Figure 2, except it examines prescriptions of generic drugs instead of branded drugs. Since generics show a fairly fixed difference both pre- and post- sunshine law, the raw data suggests that mandated disclosure is more likely to affect how often branded drugs are prescribed rather than how often generic drugs are prescribed.

Other Identification Concerns

To avoid confounds, it is important to be aware of other changes to healthcare in Massachusetts that occurred around this time. There are two particular Massachusetts-specific changes to consider.

1. *Romney Care*

The most well-known Massachusetts-specific healthcare initiative in this time period is what is known as Romney Care, named after the Massachusetts governor who signed the law. This bill was signed in April 2006 and required that nearly all Massachusetts residents obtain a minimum level of health insurance coverage.[16] To help implement this, the bill created an independent public authority, the Commonwealth Health Insurance Connector, to act as an insurance broker and offer subsidized private insurance plans to residents. This bill was driven in part due to rising costs of insurance, as well as concern over free-riders who did not have insurance but would use emergency room services for non-emergency medical care. In 2010 the state also began restricting residents to an open enrollment period for purchasing insurance through the Connector (so that individuals could not as easily game the system).

Importantly, these changes do not seem to have greatly affected prescription patterns. Figures 2-4 show that in 2006-2008, during the first years of Romney Care and before the disclosure law was passed, prescription patterns between Massachusetts and the control states are relatively fixed after accounting for the distribution of physician specialties across states. This suggests that the effects of Romney Care on health insurance sign-ups did not significantly influence the rate of branded drug prescriptions in Massachusetts.

[16] The bill also called for the state to provide free health insurance to those earning less than 150% of the federal poverty level, and it also required employers with more than 10 full-time employees to provide insurance to employees.

In addition, Romney Care is not likely to influence prescriptions in a manner that could be confused with any effects of the disclosure policy. Even if Romney Care affected prescriptions, it would likely be in the first years of the bill, in 2006-2007, and not around 2008-2009 when the disclosure policy was signed and implemented.[17] Moreover, an increase in health insurance coverage should theoretically *increase* branded drug prescriptions, which is not necessarily the same directional effect as might be expected for the disclosure policy. This is because increased insurance coverage makes the expensive, branded drugs more affordable to patients; in addition, the increase in insurance coverage could also increase patient visits to doctors in general, thus increasing drug prescription volume. These all suggest against Romney Care being the driver of any observed decrease in prescriptions in response to the disclosure policy.

2. Alternative Quality Contracts

There is one other significant change to Massachusetts healthcare during this time period that is worth noting. In 2009, Blue Cross Blue Shield of Massachusetts launched a new payment arrangement, known as the Alternative Quality Contract (AQC); these contracts stipulated fixed payments for the care of a patient over a specified time period, and they connect payments to quality goals and a five-year budget (Chernew et al. 2011). In particular, providers could receive quality bonuses for staying under budget. In 2009, seven provider organizations in Massachusetts entered into these contracts, and another four joined in 2010. However, even after 2010, this covered only 1600 primary care physicians and 3200 specialists (Chernew et al. 2011), which represent a small fraction of the physicians in the state.

These changes are unlikely to lead to changes in prescriptions. Research suggests that AQCs in Massachusetts did not have any impact on the use of either branded or generic drugs (Afendulis et al. 2014). This study used a difference-in-difference approach comparing drug prescription usage by Massachusetts doctors belonging to providers that enrolled or did not enroll in an AQC in 2009. Since

[17] As will be demonstrated, results are all robust to excluding the years when Romney Care would most likely have caused Massachusetts prescriptions to change relative to controls (2006 and 2007).

AQCs showed no effect on prescriptions between these two groups of Massachusetts doctors, it is likely that all analysis in this paper would be robust to including only doctors in Massachusetts that were not subject to an AQC.

Aggregate Policy Effects

Statistical Method: Doctor-Month Data

I first use a panel OLS specification to evaluate whether the sunshine law influenced prescriptions. For ease of interpretation, the first specification ignores drug-level data and collapses data down to the doctor-month level. In this specification, the dependent variable (DV) is simply the total number of branded scripts written by that doctor in a given month. The independent variable of interest is an indicator for whether a sunshine law was in effect for that doctor-month. This is represented as:

$$R_{it} = \beta_0 + \beta_1 * sunshine_{st} + \lambda_1 * X_t \tag{1}$$

where i represents the doctor, t represents the month, s represents the state, and X_t represents control variables including doctor and month fixed effects. These fixed effects help control for the fact that doctor specialty is not distributed equally across treatment and control groups. R_{it} represents branded prescriptions in that doctor-month, and $sunshine_{st}$ is the indicator for the sunshine law. That is, it takes a value of 1 if a sunshine law was in effect for that doctor-month. For doctors in Massachusetts, $sunshine_{st}$ takes a value of 0 in all months prior to July 2009, and 1 for all months from July 2009 onwards. For doctors not in Massachusetts, it always takes on a value of 0.

Standard errors are clustered at the AMC level. Although the disclosure policy and thus treatment is at the state level, I only have five states in the data and clustering over so few states could lead to bias (Cameron and Miller 2014). Instead, I choose to cluster at the AMC level, which can also partially account for institution-level correlations in prescriptions (e.g., philosophical practices, administrative influences, patient demographics, etc.). However, all results are robust to clustering at the doctor or state level instead. All regressions include the same data exclusions specified earlier, although results are also robust to including any or all of those data.

Results: Doctor-Month Analysis

Table 3 displays regression results for the doctor-month specification. Model (1) suggests that the sunshine law caused doctors to prescribe 2.25 fewer branded scripts per month. This represents approximately 10% of the average branded scripts per doctor-month. Model (2) shows the same regression but on generic scripts per month. Disclosure may cause generic scripts to increase by approximately 2.5 scripts per month, but the result is only marginally significant. In addition, since generics are prescribed at much higher volumes than branded drugs,[18] this increase in generic drugs would only represent an approximately 4% change to average generic scripts per doctor-month. Thus, disclosure seems to affect branded, marketed drugs more than it affects generic drugs.

Statistical Method: Doctor-Drug-Month Data

I now incorporate drug-level data into the analysis. The primary DV of interest is the number of scripts of a specific drug prescribed by a particular doctor in a given month. The main independent variable of interest is whether a sunshine law is in place. The full regression model is represented as:

$$R_{ijt} = \beta_0 + \beta_1 * sunshine_{st} + \beta_2 * marketed_{jt} + \beta_3 * (marketed_{jt} * sunshine_{st}) + \lambda_1 * X_t \qquad (2)$$

where i represents the doctor, j represents the drug, t represents the month, s represents the state, and X_t represents a vector of control variables. As before, the *sunshine* variable is an indicator for the 2009 Massachusetts sunshine law.

The *marketed* variable represents whether a drug is marketed by its manufacturer in a given month. For simplicity, I use a binary indicator that represents whether the drug is branded with no generic

[18] On average, generic prescriptions represent approximately 75% of the total prescriptions in this dataset, with the ratio increasing in favor of generics over time (due to branded drugs coming off patent over time).

version available in that month. This represents whether a firm is incentivized to market that particular drug, and it correlates very strongly with whether a drug is actually marketed by sales reps or not.

The *marketed*sunshine* interaction separates the effect of the sunshine law on marketed drugs from the effect on non-marketed (i.e., generic) drugs. The coefficient for *sunshine* (β_1) measures the effect of policies on non-marketed drugs, and the coefficient for *marketed*sunshine* (β_3) measures the additional effect that policies have on marketed drugs. To properly interpret how the sunshine law affects a marketed drug, I will take the linear combination of the coefficients for *sunshine* and *marketed*sunshine*.

The remaining variables are a set of controls. This includes month fixed effects as well as doctor*drug fixed effects. State and AMC fixed effects are not included, since the sample is limited to doctors that never switch affiliations (so doctor is collinear to state and AMC). Since including both fixed effects and a lagged dependent variable can complicate identification (Angrist and Pischke 2008), the base specifications use fixed effects but no lag. Standard errors are clustered at the AMC level.

In a follow-up regression, I run the same specification but replacing the DV with a drug's marketshare. That is, for each observation, I divide the total prescriptions of that drug by the total number of drugs prescribed by that doctor in that month. This analysis adjusts the policy effect to account for differences in total prescribing volume across doctors.

Results: Doctor-Drug-Month Data

Table 4 (Model 3) displays the fixed effects panel OLS results from the above specification. The sunshine law significantly affects the prescriptions of both branded and generic drugs. Mandated disclosure causes generic drugs to increase on average by 0.07 scripts per doctor-month for *every* generic drug in the dataset prescribed by that doctor. Likewise, branded drug prescriptions decreased on average by 0.56 scripts per doctor-month for *every* branded drug the doctor prescribes. This suggests that the policy led to approximately 290,000 fewer prescriptions of branded drugs for just the Massachusetts doctors and drugs in this dataset; generics increased by approximately 100,000 prescriptions. This implies that the disclosure policy caused doctors in Massachusetts to prescribe *fewer* drugs overall. This in turn could imply that meals and payments from pharmaceutical companies caused doctors to prescribe branded drugs to patients they otherwise would not have prescribed anything to.

I also run a number of unreported robustness checks. Results are still statistically significant at the same thresholds when clustering standard errors by state or doctor. Results are robust to excluding each of the four Massachusetts AMCs in turn, excluding each of the counterfactual states in turn, or limiting the analysis to drugs that are always marketed or always generic throughout the entire dataset. Results are also robust to including all seven years of data. In placebo tests, I assign the sunshine law to another state/month pair, and these placebo tests show no negative effects of disclosure on prescriptions.

I also run a number of robustness checks to account for any concerns over the fixed differences assumption between treatment and control states. First, results are robust to including a lag covariate with no fixed effects, or using an Arellano Bond GMM estimator that includes both a one-month doctor-drug lag as well as doctor, drug, and month fixed effects. Second, results also hold when including linear doctor-drug level trends as a covariate (in addition to the already included doctor-drug fixed effects). Third, the results are robust to using only subsets of the data where the parallel trends assumption is more closely satisfied. As mentioned, this includes using only drugs that are always generic or always on patent during this timeframe. This also includes using only the single counterfactual state, Pennsylvania, that most closely parallels Massachusetts in terms of pre-treatment trends. Results are also robust to using a control group consisting only of counterfactual doctors that are matched to a treatment doctor according to a nearest-neighbor match algorithm, where the match is assigned according to pre-treatment prescription patterns that account for both prescription volume and the distribution of prescriptions across drug classes.

Model (3B) estimates the same specification but using drug marketshare as the DV. Under these assumptions, the disclosure policy still has a positive effect on generics and a negative effect on branded drugs. In particular, for each doctor subjected to disclosure, branded drugs decrease in marketshare by 0.5% per branded drug on average, while generic drug marketshare increases by 0.2% per generic drug on average. The same robustness checks generally also hold when using drug-marketshare as the DV.

Statistical Method – Sunshine Law and AMC Policies

Since doctors in this dataset are all affiliated with an AMC, I can estimate disclosure effects after controlling for AMC policies on marketing activities. I use the following regression specification:

$$
\begin{aligned}
&\text{[Main Effects]} && P_{ijt} = \beta_0 + \beta_1 {*} AMC_{ut} + \beta_2 {*} sunshine_{st} + \beta_3 {*} marketed_{jt} + \\
&\text{[Interactions]} && \beta_4 {*} (marketed_{jt} {*} AMC_{ut}) + \beta_5 {*} (marketed_{jt} {*} sunshine_{st}) + \beta_6 {*} (AMC_{ut} {*} sunshine_{st}) + \qquad (2) \\
&\text{[Double-Interaction]} && \beta_7 {*} (marketed_{jt} {*} sunshine_{st} {*} AMC_{ut}) + \lambda_1 {*} X_t
\end{aligned}
$$

where i represents the doctor, j represents the drug, t represents the month, u represents the university/AMC, s represents the state, and X_t represents a vector of control variables. The *sunshine* and *marketed* indicators are the same as in the previous specification.

The *AMC* variable is an indicator that codes for the AMC policies. This represents a vector of indicators representing whether meals bans or consulting restrictions are in place for a given AMC-month. The *marketed*AMC* interaction separates the effect of each AMC policy on marketed drugs from their effect on non-marketed drugs.

The last two interaction terms measure the additional effect on prescriptions of having both AMC and sunshine policies in place. The coefficient for *AMC*sunshine* (β_6) represents the additional effect (beyond just β_1 + β_3) for months where both AMC and disclosure policies in place. Similarly, the coefficient for *marketed*AMC*sunshine* (β_7) represents the additional effect on marketed drugs of having both policies in place. As before, the specification includes both month fixed effects and doctor*drug fixed effects. Standard errors are again clustered at the AMC-level.

Results – Sunshine Law and AMC Policies

Table 5 displays the fixed effects panel OLS results using the specification described. The linear combinations of relevant coefficients are listed below the table. Model (4) suggests that the sunshine law effects are still negative and significant after controlling for both meals bans and consulting restrictions. The estimated effects of disclosure are smaller in comparison to those from Model (3), but the effects increase to even larger magnitudes than those in Model (3) when both AMC and disclosure policies are in place for a given data point. This suggests that AMC policies and disclosure laws may be complementary and additive in nature. The meals ban also has a slight negative effect on branded prescriptions and a marginal positive effect on generic prescriptions. This is consistent with results from Larkin et al. (2012).

Consulting restrictions had no impact. These restrictions only banned honoraria (payments for unspecified purposes) while also requiring that consulting contracts be approved by the AMC; thus they have less impact than outright bans on activities such as free meals or gifts. In addition, fewer doctors participate in consulting, which will reduce the observed aggregate effects of consulting restrictions (per Campbell et al. 2010, only 14% of doctors self-report accepting speaking or consulting payments, compared with self-report rates of over 70% for free gifts and meals).

Possible Mechanisms

The observed changes in prescriptions could be caused by changes in behavior in one or more different groups. The response could be driven by changes in doctor behavior (e.g., whether they accept meals or payments), in industry marketing behavior (e.g., whether sales reps offer meals and payments to doctors), in patient behavior (e.g., whether patients demand generics or switch doctors in response to disclosed payments), or a combination of the above. However, the most likely mechanism is through changes in doctor behavior.

These observed effects are not likely driven by patient-driven demand for changes. Importantly, the first set of data (for meals and payments from July 2009 – December 2009) was not published until November 2011. The analysis in the previous section only extends to data from July 2011, and thus all of the observed effects occur *prior* to any public release of data. These changes in prescriptions are not a response to any actual usage of the data, whether by patients, colleagues, hospitals, or the media.

The effects could theoretically be driven by changes in the industry's marketing approach. That is, pharmaceutical companies could have reduced meals and payments to doctors in response to disclosure requirements. However, this would only occur if industry anticipated that disclosure would reduce the effectiveness of these tactics.[19] In addition, since salespeople are generally assigned to a small geographical radius, such a change would require a reduction in the sales force in Massachusetts and/or a relocation of sales reps from Massachusetts to other states. There is no anecdotal evidence that pharmaceutical companies made any Massachusetts-specific alterations to detailing strategies, sales force size, or salesperson geographic allocation in response to this policy. In fact, scraped data from the online professional network LinkedIn show that pharmaceutical sales force size in Massachusetts trended similarly as pharmaceutical sales force sizes in other states during this timeframe.

Most likely, the observed policy effects are driven by doctors' fears of being perceived as biased or unethical in the eyes of patients, colleagues, or the public. As a result, some doctors may respond to disclosure by cutting off ties to industry in order to avoid having to disclose them; such a response would be consistent with laboratory evidence (Sah and Loewenstein 2014). This change could then lead to the observed changes in prescriptions in response to disclosure. Such a hypothesis would suggest that those doctors willing to accept and disclose payments or meals may not be as affected by the policy; I investigate this hypothesis in the next section.

It is important to note that doctor-level image concerns are separate from doctor-level concerns over how their reputation might affect their long-term income. Since the doctors in this dataset are full-time attendings at an academic medical center, they are generally salaried employees whose income is relatively independent of the number of procedures or tests they perform, or the number of patients that they see. In addition, patients for these doctors are driven through referrals through their AMC, so their patient volume is not as affected by their reputation as private practitioners' might be. Finally, doctors in general seem to care about how their patients perceive them ethically, independently from any impact on earnings, as evidenced by doctor message board responses to other disclosure efforts (Wen 2014).

[19] In addition, if industry changed their meals and payments activities prior to November 2011 (when the first set of data was released) then they anticipated a change in doctor response independent from release of the data.

Policy Effects on Massachusetts Doctors with Disclosed Industry Ties

Procedure

The Massachusetts payments data (as summarized in Table 1) identifies which doctors accepted meals or consulting fees after the sunshine law was passed. Unfortunately, the data only identifies payments to a doctor at the year-level. In addition, due to relaxed reporting requirements in 2012 (described earlier), there are fewer payments reported for 2012 than for previous years. Finally, the 2009 data only contains payments from six months of data, since the law was implemented on July 1, 2009. It is therefore difficult to use the timing of these payments to evaluate the effect of payments on prescriptions trends. Instead, I use this data to classify doctors as a “consulting doctor” if they ever appear in the data for accepting consulting fees at any time post-law. I similarly classify doctors as a “meals doctor” if they appear in the data for having accepted meals.

Raw Data Trends

Figure 6 displays the total branded prescriptions for Massachusetts doctors that accepted meals and Massachusetts doctors that did not accept meals, according to the disclosures data. This is compared to doctors in other states that were not subject to mandated disclosure. Figure 7 plots the same for doctors that accepted or did not accept consulting payments. Since the total number of doctors in each group is different, I plot average branded drugs prescribed per doctor instead of just total branded drugs prescribed by each group. This measure first sums each doctor’s total prescriptions that are branded in a given month. For each doctor group, it then finds the average branded drugs prescribed per month by equally weighting each doctor’s average.

In Figures 6 and 7, there is very little discernible difference in sunshine law effects on those that accepted meals or payments and those that did not. Of course, these graphs do not control for various fixed effects or other control variables, and they also ignore all of the drug-level information in the data.

Regression Results

I run the same base regression as in Model (3), but I include interaction terms that account for whether a doctor is a “meals doctor” or a “consulting doctor.” Results are displayed in Table 6. Standard errors are again clustered at the AMC level.

Model (5) shows the regression for meals doctors. As before, the regression demonstrates that doctors are significantly affected by disclosure. Specifically, doctors that did not have any meals to disclose decreased their branded prescriptions by 0.53 scripts per branded drug per month relative to counterfactual doctors. On the other hand, doctors that showed up in the database as having accepted meals responded to disclosure by reducing branded prescriptions by 0.73 scripts per branded drug per month.

However, this difference in magnitude may simply be because meals doctors prescribe more branded drugs on average, which is certainly apparent from Figure 6. I therefore run an alternate specification in Model (5B) that uses drug-level marketshare as the dependent variable, similar to Model (3B). These results suggest that meals-doctors are actually affected slightly less in terms of marketshare; branded drugs drop in marketshare by 0.4% per branded drug for meals doctors, and by 0.5% per branded drug for doctors that had no meals to disclose.

Moreover, the estimated effect of disclosure on the *non*-meals doctors is biased downward by the fact that this group includes doctors that never interacted with industry even pre-disclosure law. These doctors likely experienced no change in prescriptions due to disclosure, and these doctors would reduce the estimated average effect of the disclosure law on the entire group of non-meals doctors. When factoring this in, it appears that doctors that used to accept meals but stopped due to disclosure are likely much more affected by disclosure than those that used to accept meals and then continued to accept some meals post-disclosure.

Model (6) shows a similar trend for doctors that accepted consulting payments. The sunshine law led to decreases in branded drug prescriptions of 0.52 scripts per month-doctor-drug for those that had no consulting to disclose, and 0.84 scripts for consulting doctors. Model (6B) suggests that consulting doctors also showed a larger drop in branded drug marketshare in response to disclosure than those that had no consulting to disclose. However, as before, the majority of doctors in the no-consulting group likely had no consulting arrangements even pre-disclosure law; recall that only 14% of doctors self-reported accepting such payments (Campbell et al. 2010). These doctors would significantly downward bias the estimated effect of disclosure on the entire group of non-consulting doctors. It is likely that the effect of disclosure is in actuality much larger for those that used to accept consulting but stopped in response to disclosure than for those that continued to accept consulting post-disclosure.

Procedure – High and Low Types

I also define “high-meals” and “high-consulting” doctors” doctor types. These represent the 50% of doctors who accepted the highest total values of meals or consulting payments within this dataset. I then run the same regression specifications but using indicators for “high-meals” or “high-consulting” status instead of simply indicators for “meals doctors” or “consulting doctors.”

Results – High and Low Types

Models (7) and (8) in Table 7 show the regression results for high-meals doctors and high-consulting doctors, respectively. Models (7B) and (8B) show the equivalent regressions using drug-marketshare as the DV. These results show that high-meals and high-consulting doctors respond to disclosure fairly similarly as the full group of meals doctors and consulting doctors.

Interpretation

Taken altogether, the analysis using the actual state disclosure data suggests that doctors that were willing to disclose meals or consulting payments still decreased prescriptions of branded drugs relative to doctors in other states. Thus, there was no evidence of moral licensing among these doctors, at least when measured in aggregate. However, these doctors were likely also less affected by the policy than their Massachusetts peers who had nothing to disclose. This latter conclusion hinges on the assumption that those with nothing to disclose consist of a heterogeneous mix of doctors: those that never accepted industry payments even pre-policy and thus would not alter their prescriptions in response to disclosure, and those that cut off industry payments in response to disclosure, likely leading to large changes in their prescriptions.

Informational versus Non-Informational Influence

There is debate in the literature on whether meals and other forms of industry payments are informational or non-informational sources of influence on doctors. Anti-marketing groups naturally claim that gifts or meals have no informational content and are thus a source of non-informational persuasion. Industry argues that these meals are just a ticket for getting a salesperson into the door, and any influence from meals is a result of the information that the salespeople ultimately disperse to doctors, such as new clinical trials results (Carlat 2007).

This paper can take advantage of drug-level differences to test whether these marketing influences are informational in nature. In particular, new-to-the-market drugs require more information dissemination from manufacturers to doctors, since doctors are less familiar with the idiosyncrasies and details of these drugs. There is also likely a higher volume of clinical trials results being released for these newer drugs, leading to more information for salespeople to disseminate. If disclosure influences new-to-the-market branded drugs differently than older branded drugs, this can shed insight on the degree to which this marketing influence is informational or non-informational in nature.

In all specifications reported in this section, standard errors are again clustered at the AMC-level.

Old vs. New Branded Drugs

Within the full seven years of this dataset, a number of branded drugs were introduced to the market. 58 of the branded drugs in this dataset were brought to market sometime during this time frame, with 38 of those hitting the market in the middle of the four-year timespan around when the disclosure law was implemented (July 2007 to July 2011). I test whether disclosure differentially affects these new drugs relative to older branded drugs. I define a new drug as one that has been on the market for one year or less, although alternate cutoffs (e.g., 18 months, 24 months) also yield similar results.

Table 8 runs linear models that are similar to Model (3), except that an additional indicator is included for whether a drug is new to the market for that month (i.e., it entered the market during the time period of this dataset and has been on the market for 12 months or less). This indicator is also interacted with the sunshine law indicator.

Model (9) suggests that newer marketed drugs are *less* affected by the sunshine law than older branded drugs. Model (9B) shows a similar finding when using drug-level marketshare as the DV; older branded drugs drop in marketshare by 0.6% in response to disclosure, while new-to-the-market drugs only drop 0.3% in marketshare. Importantly, average script volume for new-to-the-market drugs is similar (and in fact slightly larger) than average script volume for older branded drugs, so these effects are not driven by simple differences in average total volume prescribed. Altogether, these results are consistent with gifts and consulting payments serving as *non-informational* sources of influence.

Discussion

This paper uses a state-level policy change to evaluate whether mandated public disclosure of industry-related conflicts-of-interest can alter how physicians prescribe. The results suggest that public disclosure reduced the prescriptions of branded drugs and increased the prescriptions of generic drugs, but yielding a net *decrease* in prescription volume. These changes likely are a consequence of changes in doctor-industry interactions that were brought about by the mandated disclosure requirements.

Specifically, these disclosure requirements may have invoked social image concerns in doctors. Doctors are trained in the spirit of the Hippocratic Oath (i.e., “First, do no harm”), and this norm helps enforce a desire in doctors to avoid appearing biased or unethical in the eyes of patients, colleagues, and the public. Mandated disclosure may interact with this norm and encourage doctors to abstain from meals and payments to avoid having to disclose such conflicts-of-interest. Since these free meals and payments have been shown to influence how doctors prescribe (Larkin et al. 2012), abstaining from these activities would then lead to the changes in prescriptions observed in this paper. In addition, these doctors' concerns over their public image are not necessarily tied to concerns over the effects of reputation on long-term income, since the doctors in this sample are salaried employees whose income does not depend on the number of patients they see or the number of procedures or tests they perform.

Alternative non-image or non-reputation explanations for the effect of disclosure are possible, but unlikely. Patient-driven responses could not have driven the observed results, since the state released the first batch of data with a two year lag, and the observed effects were evident even prior to release of the data. The pharmaceutical industry could have reduced payments to doctors in response to disclosure requirements, but this would only occur if the manufacturers anticipated that doctors would respond to disclosure; in addition, there is no evidence suggesting that industry responded in this way, since observational evidence suggests that sales force size in Massachusetts did not change in response to this policy.

These social image effects are consistent with behavioral economics literature on public image and norm adherence. Lab experiments have shown that social image can cause individuals to conform to norms of fairness (Andreoni and Bernheim 2009) and trust (Tadelis 2011). It stands to reason that social image concerns may also cause doctors to adhere to ethical norms that were emphasized to them during their training. It remains to be seen whether disclosure can reduce biases from conflicts-of-interest where ethical norms may not be as firmly entrenched through training (such as perhaps financial advising or political contexts). Indeed, some studies suggest that medical versus non-medical contexts can be relevant when measuring the effect of disclosure on biases (Sah, Loewenstein and Cain 2011).

This paper also finds that industry influences on doctors appear to be largely non-informational in nature. New-to-the-market branded drugs are not disproportionately affected by the sunshine law, despite the fact that these drugs require more information dissemination from salespeople to doctors. This evidence does not accord with arguments by industry, which claim that meals are only a way to get a foot in the door and not a direct source of influence unto themselves. Altogether, these results suggest that doctors are influenced by industry meals and consulting payments, and in at least partially a non-informational way, but that disclosure can be effective at reducing this influence by invoking image concerns and causing some doctors to opt out of industry meals and payments.

References

Afendulis CC, Fendrick AM, Song Z, Landon BE, Safran DG, Mechanic RE, Chernew ME. The impact of global budgets on pharmaceutical spending and utilization: Early experience from the alternative quality contract. *Inquiry 2014; 51:1-7.*

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

Angrist J, Pischke J. Mostly Harmless Econometrics. *Princeton University Press, 2009.*

Babcock L, Loewenstein G, Issacharoff S, Camerer C. Biased judgments of fairness in bargaining. *American Economic Review 1995; 85(5): 1337-1343.*

Cain D, Loewenstein G, Moore D. The dirt on coming clean: Perverse effects of disclosing conflicts-of-interest. *Journal of Legal Studies 2005; 34: 1-25.*

Cameron CA, Miller DL. A practitioner's guide to cluster-robust inference. *Working paper, 2013.*

Campbell E, Rao S, DesRoches C, Iezzoni L, Vogeli C, Bolcic-Jankovic D, Miralles P. Physician professionalism and changes in physician-industry relationships from 2004 to 2009. *Arch Intern Med 2010; 170(20): 1820-1826.*

Carlat D. Dr. Drug Rep. The New York Times, November 25, 2007. Accessed on May 30, 2013 at http://www.nytimes.com.

Chernew ME, Mechanic RE, Landon BE, Safran DG. Private-payer innovation in Massachusetts: The 'alternative quality contract.' *Health Affairs 2011; 30(1): 51-61.*

Chren MM, Landefeld CS. Physicians' behavior and their interactions with drug companies. A controlled study of physicians who requested additions to a hospital drug formulary. *JAMA 1994; 271(9): 684-689.*

Christakis NA, Fowler JH. Contagion in prescribing behavior among networks of doctors. Marketing *Science 2010, Articles in Advance, 1-4.*

Dana J, Loewenstein G. A social science perspective on gifts to physicians from industry. *JAMA 2003; 290(2): 252-255.*

Festinger L. A theory of social comparison processes. *Human Relations 1954; 7: 117-140.*

Festinger L, Carslmith J. Cognitive consequences of forced compliance. *J Abnorm Social Psych 1959; 58: 203-210.*

Fugh-Berman A, Ahari S. Following the script: How drug reps make friends and influence doctors. *PLoS Med 2007; 4(4): 150.*

Gagnon M-A, Lexchin J. The cost of pushing pills: A new estimate of pharmaceutical promotion expenditures in the United States. *PLoS Med 2008; 5(1): 1.*

Gneezy U, Saccardo S, van Veldhuizen R. Bribery: Greed versus reciprocity. *Working Paper, 2013.*

Gonul F, Carter F, Petrova E, Srinivasan K. Promotion of prescription drugs and its impact on physicians' choice behavior. *J of Marketing 2001; 65: 79-90*.

Haisley E, Weber R. Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior 2010; 68(2): 634-645.*

Huang G, Shum M, Tan W. Is advertising informative? Evidence from contraindicated drug prescriptions. *Working Paper, 2012.*

Iyengar R, Van den Bulte C, Valente TW. Opinion leadership and social contagion in new product diffusion. *Marketing Science 2011; 30(2), 233-248.*

Larkin I, Ang D, Chao M, Wu T. The impact of pharmaceutical detailing on physician prescribing: Quasi-experimental evidence from academic medical center policy changes. *Working Paper, 2012.*

Lexchin, J. Interactions between physicians and the pharmaceutical industry: What does the literature say? *Can Med Assoc J. 1993; 149:1401-1407.*

Loertscher LL, Halvorsen AJ, Beasley BW, Holmboe ES, Kolars JC, McDonald FS. Pharmaceutical industry support and residency education. *Arch Intern Med 2010; 170(4): 356-362.*

Loewenstein G, Cain D, Sah S. The limits of transparency: Pitfalls and potential of disclosing conflicts-of-interest. *American Economic Review: Papers and Proceedings 2011; 101(3): 423-428.*

Mackowiak J, Gagnon JP. Effects of promotion on pharmaceutical demand. *Soc Sci Med 1985; 20(11): 1191-1197.*

Malmendier U, Schmidt K. You owe me. *NBER Working Paper No.18543, 2012.*

Manchanda P, Chintagunta P. Responsiveness of physician prescription behavior to salesforce effort: An individual level analysis. *Marketing Letters 2004; 15(2-3): 129-145.*

Mizik N, Jacobson R. Are physicians 'easy marks'? Quantifying the effects of detailing and sampling on new prescriptions. *Management Sci 2004; 50(12): 1704-1715.*

Moore D, Tanlu L, Bazerman M. Conflict-of-interest and the intrusion of bias. *Judg and Dec Making 2010; 5(1): 37-53.*

Mullainathan S, Washington E. Sticking with your vote: Cognitive dissonance and political attitudes. *American Economic Journal: Applied Economics 2009; 1(1): 86-111.*

Nair H, Manchanda P, Bhatia T. Asymmetric social interactions in physician prescribing behavior: The role of opinion leaders. *Working Paper, 2006.*

Orlowski JP, Wateska L. The effects of pharmaceutical firm enticements on physician prescribing patterns. There's no such thing as a free lunch. *Chest 1992; 102(1): 270-273.*

Perlis RH, Perlis CS, Wu Y, Hwang C, Joseph M, Nierenberg A. Industry sponsorship and financial conflict-of-interest in the reporting of clinical trials in psychiatry. *Am J Psychiatry 2005; 162: 1957-1960.*

Pham-Kanter G, Alexander GC, Nair K. Effect of physician payment disclosure laws on prescribing. *Arch Intern Med; 172(10): 819-821.*

PhRMA. New survey emphasizes value of biopharmaceutical company engagement with healthcare providers [press release], March 29, 2011. Accessed on May 30, 2013 at http://www.phrma.org/media/releases

Rizzo JA. Advertising and competition in the ethical pharmaceutical industry: The case of antihypertensive drugs. *J Law and Econ 1999; 42(1): 89-116.*

Rockoff JD. Drug reps soften their sales pitches. *The Wall Street Journal*, Jan 10, 2012. Accessed on May 30, 2013 at http://online.wsj.com/.

Sah S, Loewenstein G. Nothing to declare: Mandatory and voluntary disclosure leads advisors to avoid conflicts-of-interest. *Psych Science 2014; forthcoming.*

Sah S, Loewenstein G, Cain D. Insinuation anxiety: Fear of signaling distrust after conflict-of-interest disclosures. SSRN 2011. *Available at SSRN: http://ssrn.com/abstract=1970691*

Samuelson W, Zeckhauser R. Status quo bias in decision making. *Journal of Risk and Uncertainty 1988; 1: 7-59.*

Sanchez JM, Self D. Gennder bias and moral decision making: The moral orientations of justice and care. *Journal of Medical Humanities 1995; 16(1): 39-53.*

Spurling GK, Mansfield PR, Montgomery BD, et al. Information from pharmaceutical companies and the quality, quantity, and cost of physicians' prescribing: A systematic review. *PLoS Med. 2010; 7(10).*

Steinbrook R. For sale: Physicians' prescribing data. *NEJM 2006; 354(26): 2745-2747.*

Steinbrook R. Controlling conflict-of-interest – proposals from the Institute of Medicine. *NEJM 2009; 360(21): 2160-2163.*

Stelfox HT, Chua G, O'Rourke K, Detsky AS. Conflict-of-interest in the debate over calcium-channel antagonists. *N Engl J Med 1998; 338:101-106.*

Tadelis S. The power of shame and the rationality of trust. *Working Paper, 2011.*

Wazana, A. Physicians and the pharmaceutical industry: Is a gift ever just a gift? *JAMA. 2000; 283(3):373-380.*

Wen, L. What your doctor won't disclose. *Ted Talks, Nov 2014.* Accessed on 2/3/15 from *www.ted.com.*

Figure 1. Data Summary (Prior to Data Exclusions)

Data Type	Data Included (Full Eight Years)	Data Selection Based on:	Data Excluded
Geographic Regions	6 Metropolitan Regions: Chicago, Boston, Northern California, Southern California, New York, Pennsylvania	AMSA's map of AMCs in the U.S.	All other U.S. Regions
Academic Medical Centers	35 AMCs in these Metropolitan Regions	IMAP Policy Database; Interviews with AMC Staff	AMCs outside of these Metropolitan Regions
Hospitals	282 hospitals and other medical centers owned or directly governed by AMCs	Quarterly IMS survey data from Jan 06 – Jun 09	All other hospitals and medical centers in the regions
Drugs	377 drugs in nine drug classes	IMS drug classifications; 2006-2012 IMS prescriptions data	All other drug classes; All drugs with rarely prescribed during 2006-2012
Doctors	All full-time attending doctors at AMC-owned hospitals for at least one quarter in Jan06 - Jun09	Quarterly IMS survey data from Jan 06 – Jun 09	All doctors with multiple affiliations, doctors that switch affiliations in 2006-2009, and doctors that are not full-time attendings
Measures of Marketing	Whether a drug is branded or generic; the months in which a branded drug had a generic version on the market	FDA drug databases	Other measures of marketing (e.g., journal advertisement spending, sample packs, etc.)
Payments to Doctors	All payments/gifts etc. of $50+ to Massachusetts-licensed physicians from Jul 2009 – Dec 2012	Disclosures made as part of the Massachusetts Physician Payments Sunshine Law	Payments to doctors not licensed to Massachusetts; payments before July 2009; payments under $50.

Table 1: Payments Data Summary (Massachusetts Doctors)

	Accepted Meals		**Accepted Consulting**	
	Total Doctors* (out of 2877)	**Average Dollars per Doctor (SD)**	**Total Doctors* (out of 2877)**	**Average Dollars per Doctor (SD)**
2009	215	$209 (204)	198	$9,724 (14,385)
2010	213	$241 (328)	239	$12,676 (23,779)
2011	160	$253 (336)	220	$13,064 (22,100)
2012	28	$297 (269)	27	$16,752 (18,158)
2009-2012	430 (616 obs)	$235 (291)	370 (684 obs)	$12,107 (20,707)

*****There were 4278 doctors (out of 9998) in this data affiliated with a Massachusetts AMC from July 2009 onwards. 2877 of these remain after the data exclusions specified in the Data section.

Table 2. AMC Policy Changes

University	On-site Gifts	Meals	Consulting	Date of change
NYU	1	1	0	3/29/2005
U Penn	1	1	0	9/27/2006
Stanford	0	1	0	10/1/2006
U Chicago	1	1	0	7/1/2007
UC Davis	2	1	0	7/1/2007
UC SF	1	1	0	7/15/2007
Boston Univ.	1	1	0	8/10/2007
Loyola	0	0	1	10/15/2007
Illinois at Chicago	1	1	0	12/1/2007
Mount Sinai	1	1	1	1/1/2008
Pittsburgh	2	1	1	2/25/2008
UC Irvine	2	1	0	3/12/2008
UC SD	2	1	0	3/12/2008
Rochester	1	1	0	6/25/2008
UC - system	2	1	0	7/1/2008
U Mass	2	1	0	7/1/2008
Columbia	1	1	0	1/1/2009
Rush	2	1	1	1/29/2009
Stony Brook	1	1	0	3/23/2009
Temple	1	1	0	6/15/2009
Rosalind Franklin	2	1	0	6/17/2009
U Nevada	1	0	0	7/6/2009
Southern Calif.	1	1	1	9/1/2009
Partners Health	2	2	1	10/1/2009
UC LA	2	1	0	10/7/2009
NY Med. Coll.	0	2	0	1/10/2010
Northwestern	2	1	0	1/25/2010
SUNY Upstate	2	1	0	2/1/2010
Tufts	1	0	0	4/5/2010
SUNY Downstate	1	1	0	4/21/2010
Stanford	2	1	1	7/22/2010
Boston Univ.	1	1	1	1/7/2011
Yeshiva	0	1	0	6/15/2011
Harvard	1	1	0	7/1/2011
Rochester	2	1	0	10/10/2011
Cornell	-	-	-	-
UMDNJ/RWJ	-	-	-	-
NYCOM	-	-	-	-
Western Univ.	-	-	-	-
Thomas Jefferson	-	-	-	-

Legend:

On-Site Gifts:
0: PhRMA regulations; <$100 per visit
1: Personal gifts banned
2: Personal and educational gifts banned (e.g., medical textbooks, etc.)

Meals:
0: PhRMA regulations; <$100 per visit
1: Banned on campus except for ACCME*
2: Banned everywhere except for ACCME*

*ACCME = Accreditation Council for Continuing Medical Education

Consulting:
0: No restrictions
1: No honoraria^; contracts require department approval; pay must be fair market value for specified services.

^honoraria = payments for unspecified purposes

Figure 2: Total Branded Rx (Four Years)

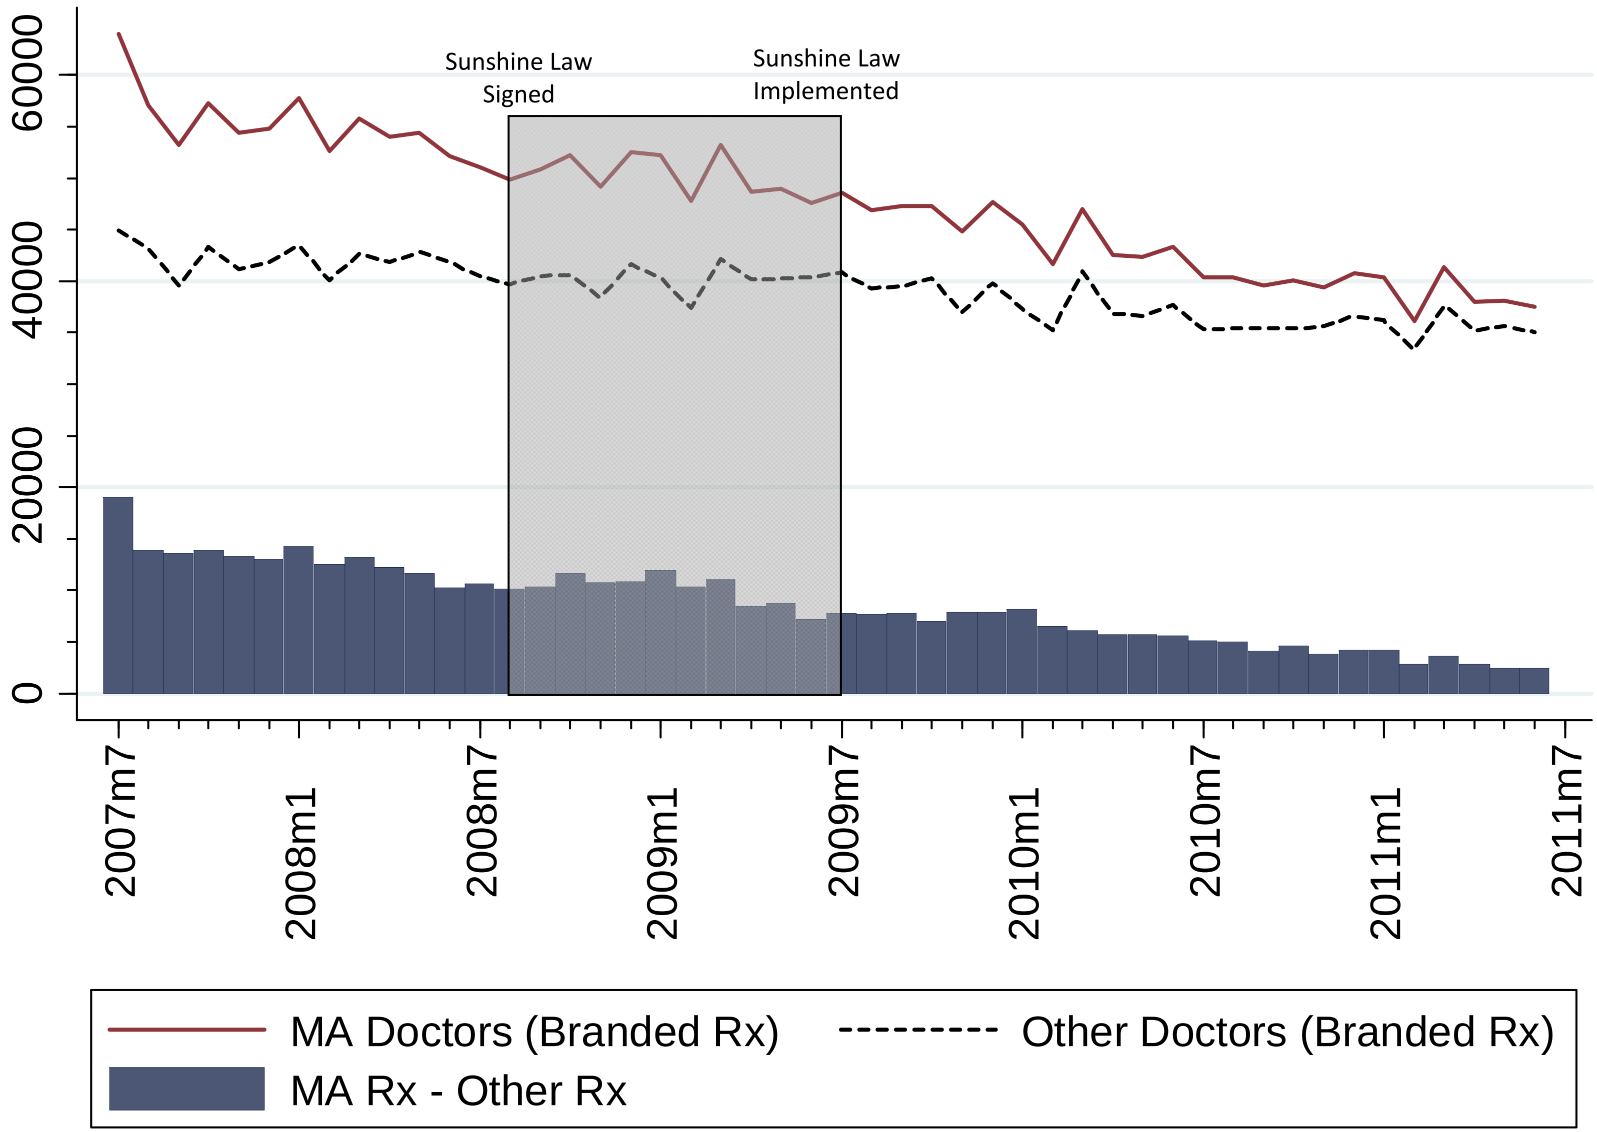

The “MA” group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The “Other” group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 2B: Total Branded Rx (Close-Up)

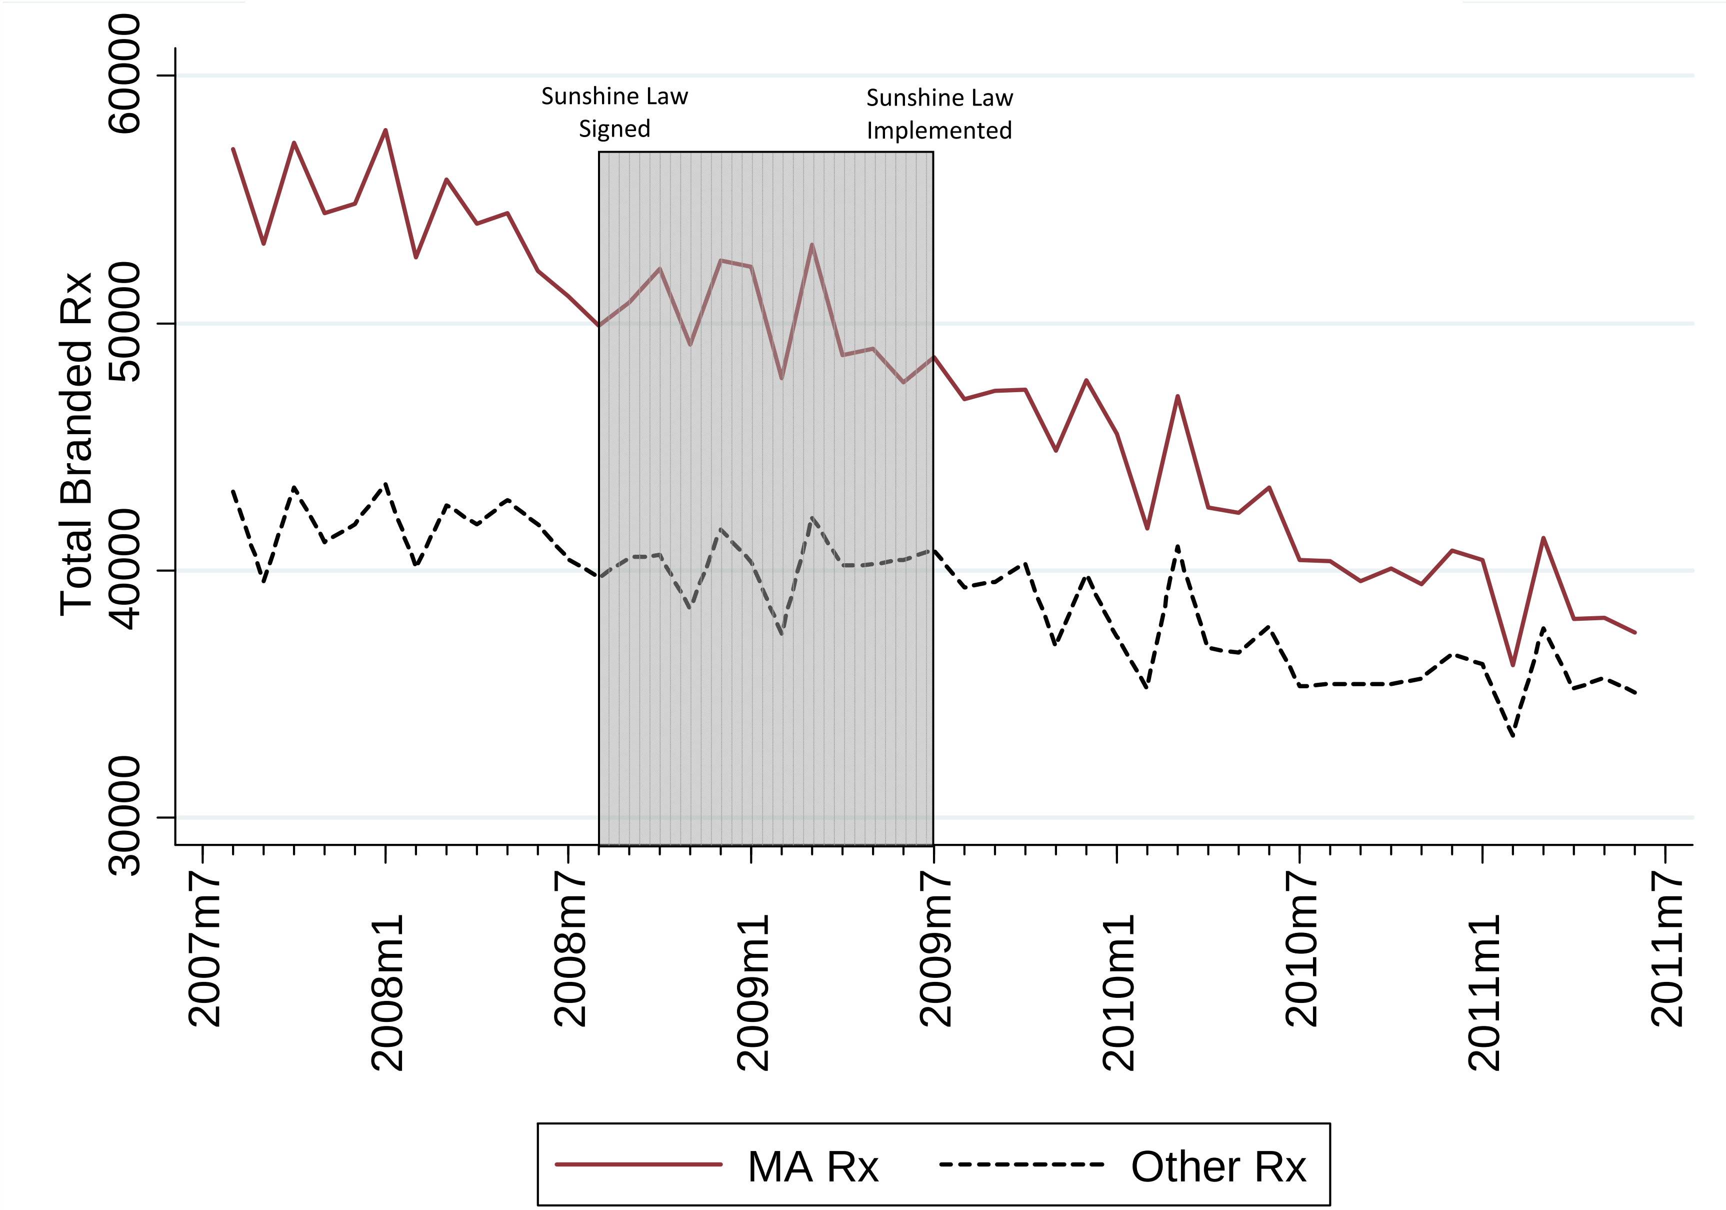

The "MA" group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The "Other" group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 3: Total Branded Rx (Seven Years)

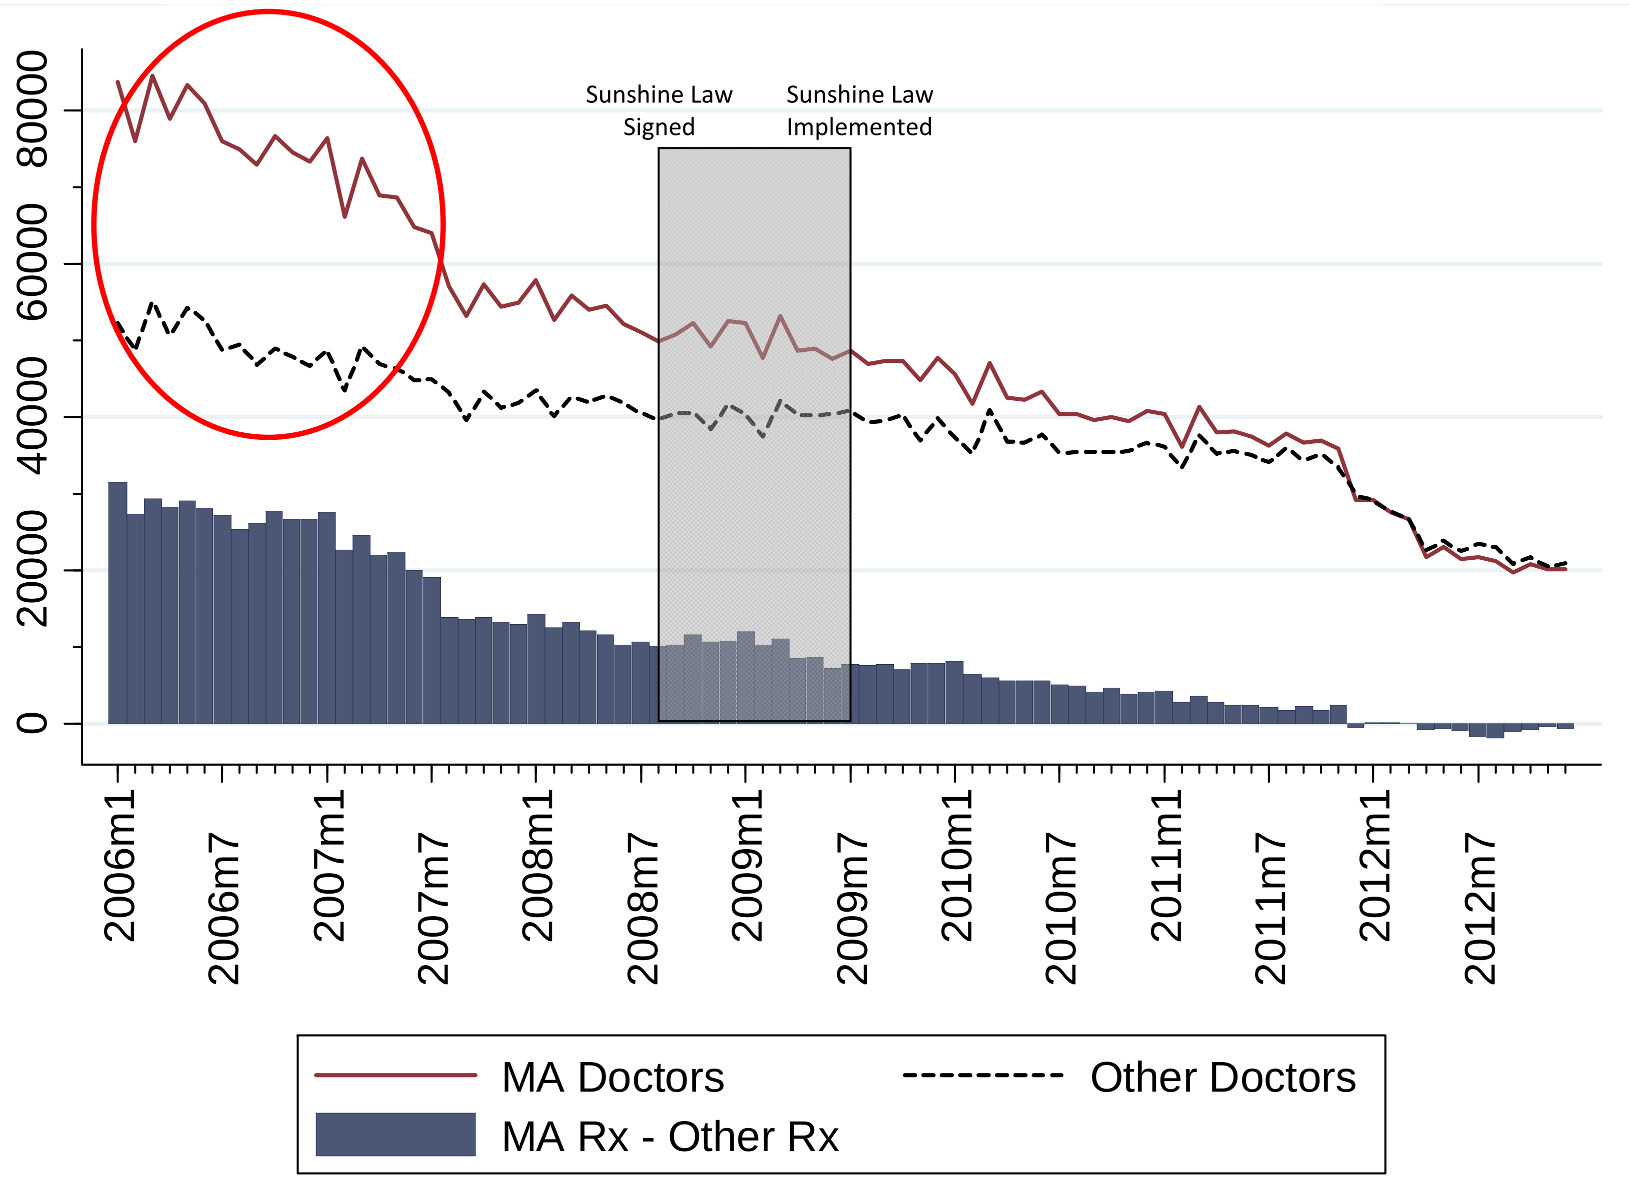

The “MA” group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The “Other” group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.
Note that shrinking distance between groups (the red circled region) is caused by psych-drugs coming off patent; there are more physicians that regularly prescribed psychotherapeutic drugs in Massachusetts than in the control states.

Figure 3B: Total Branded Rx (non-Psych Drugs only)

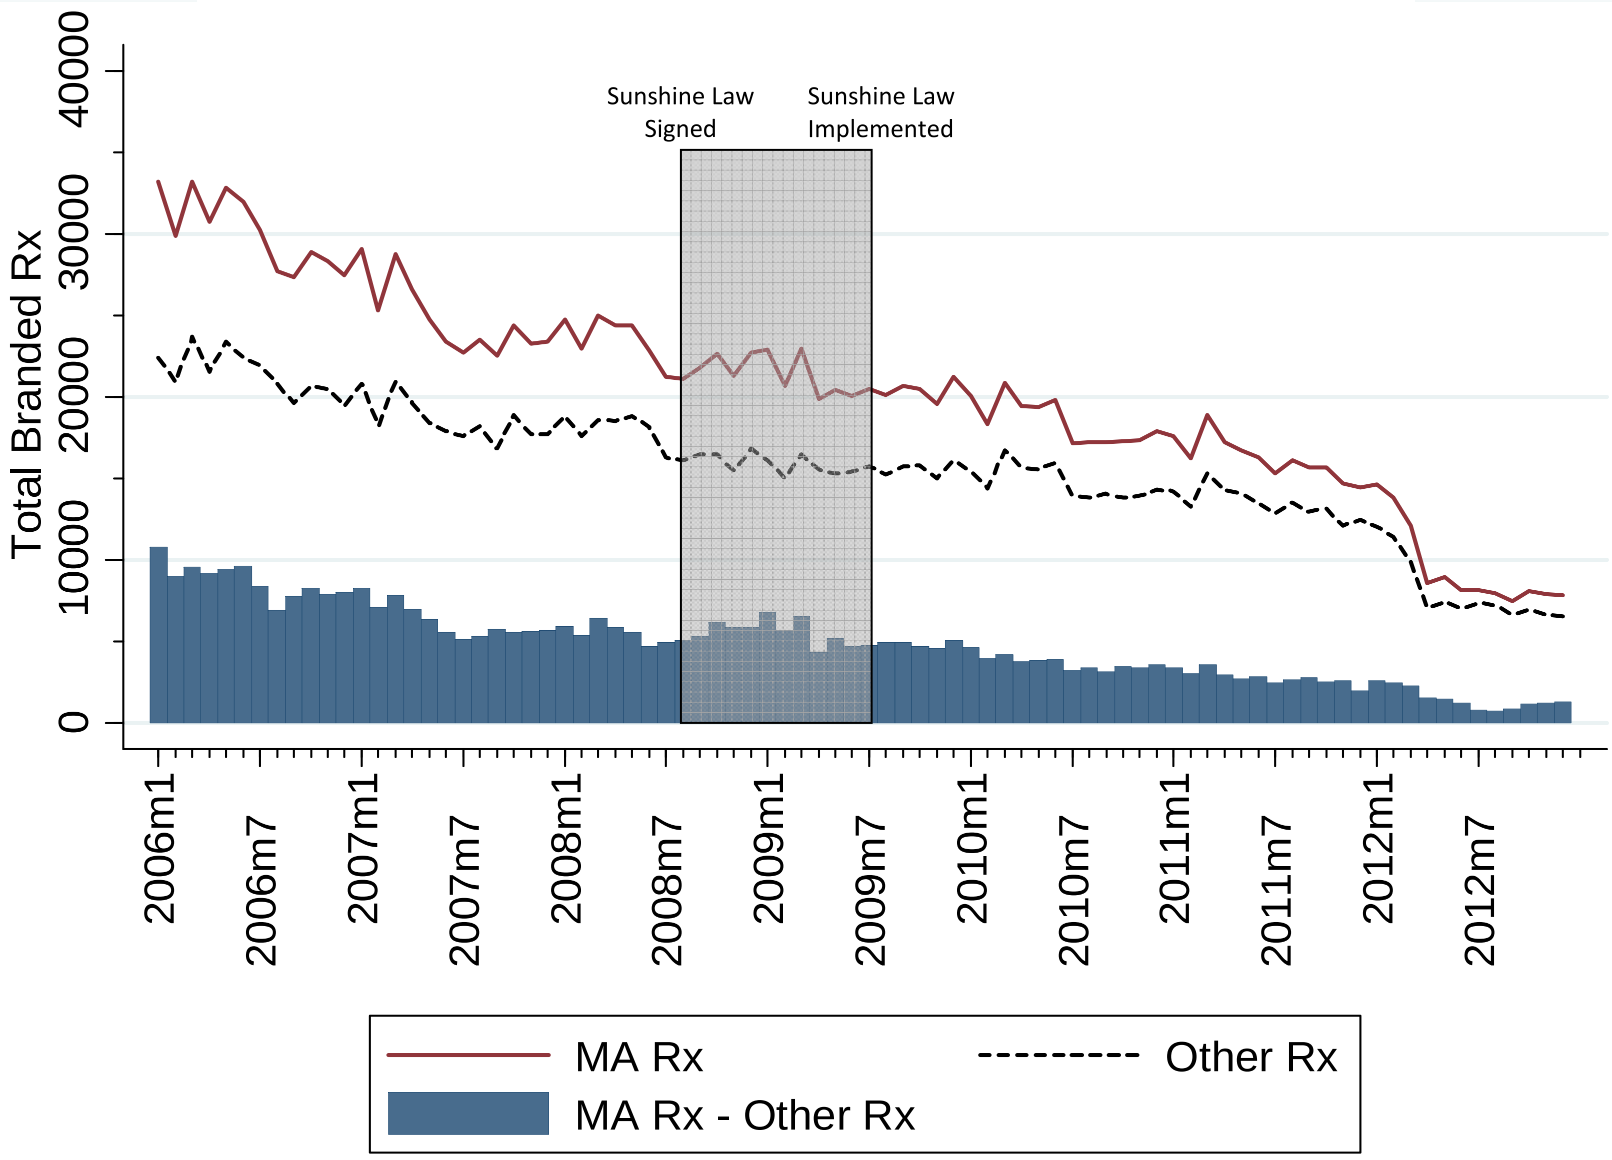

The “MA” group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The “Other” group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 4: Total Branded Rx (Always On Patent Only, Seven Years)

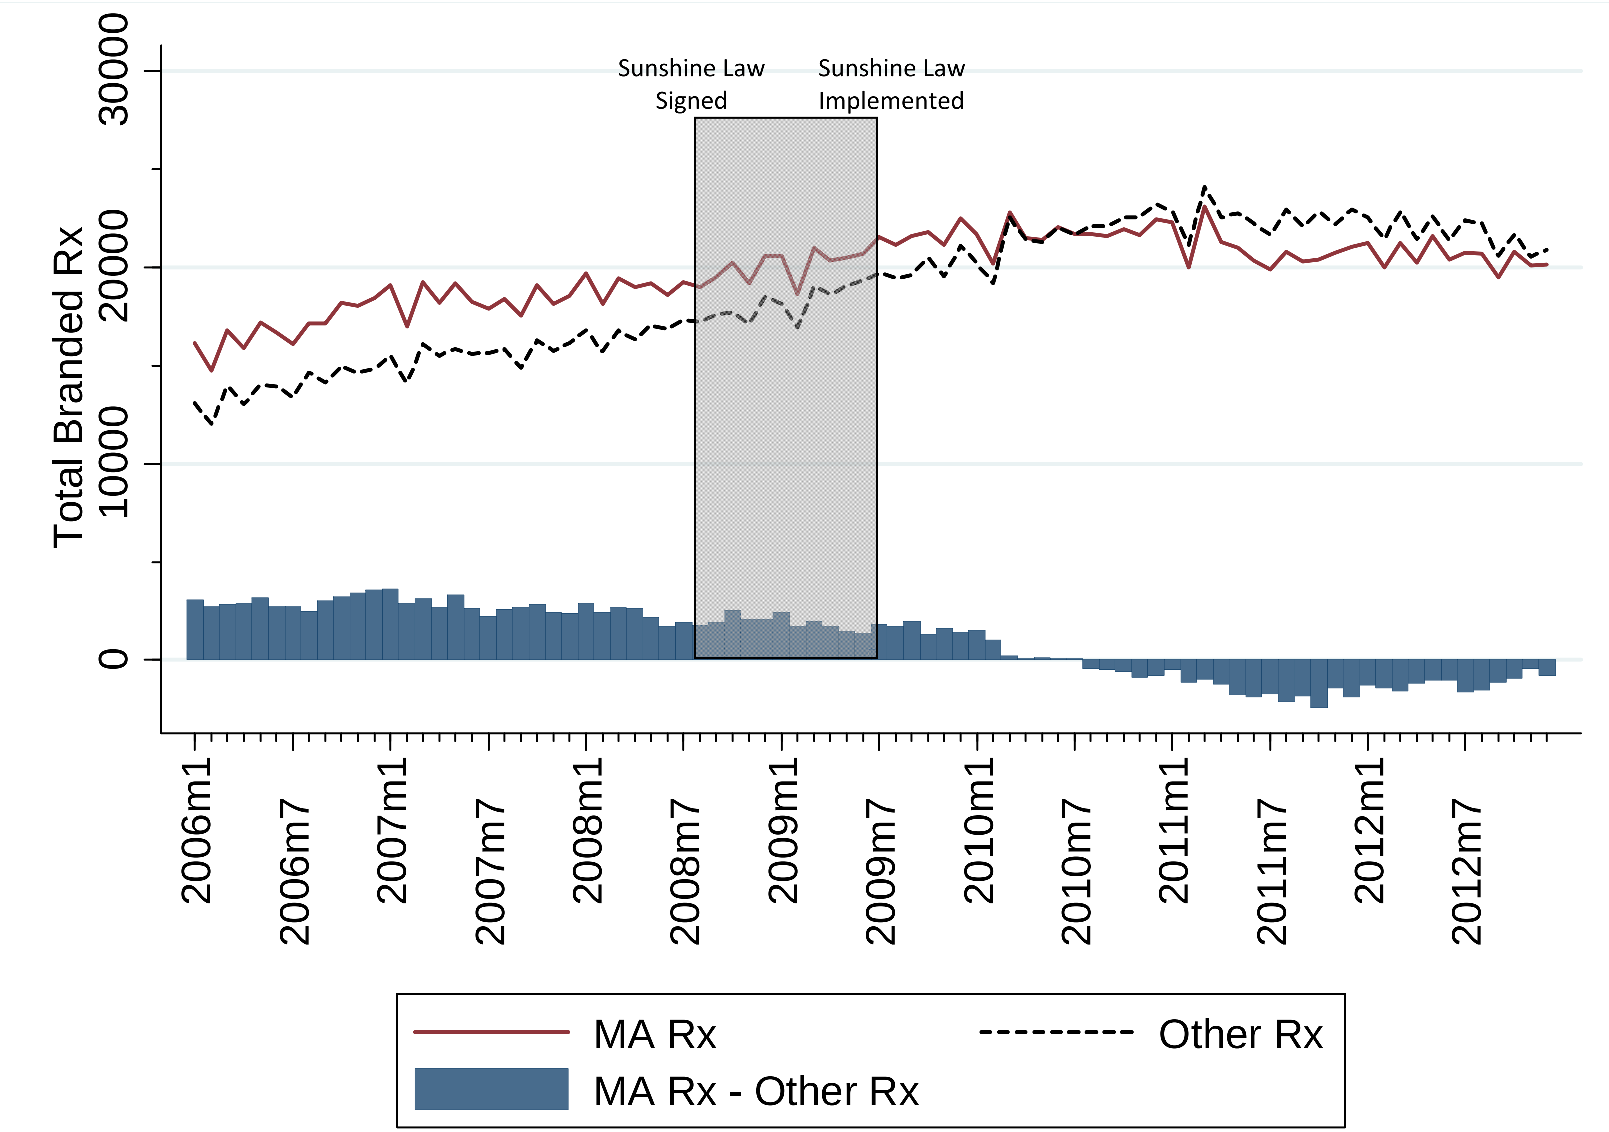

The "MA" group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The "Other" group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 4B: Total Branded Rx (Always On Patent Only)

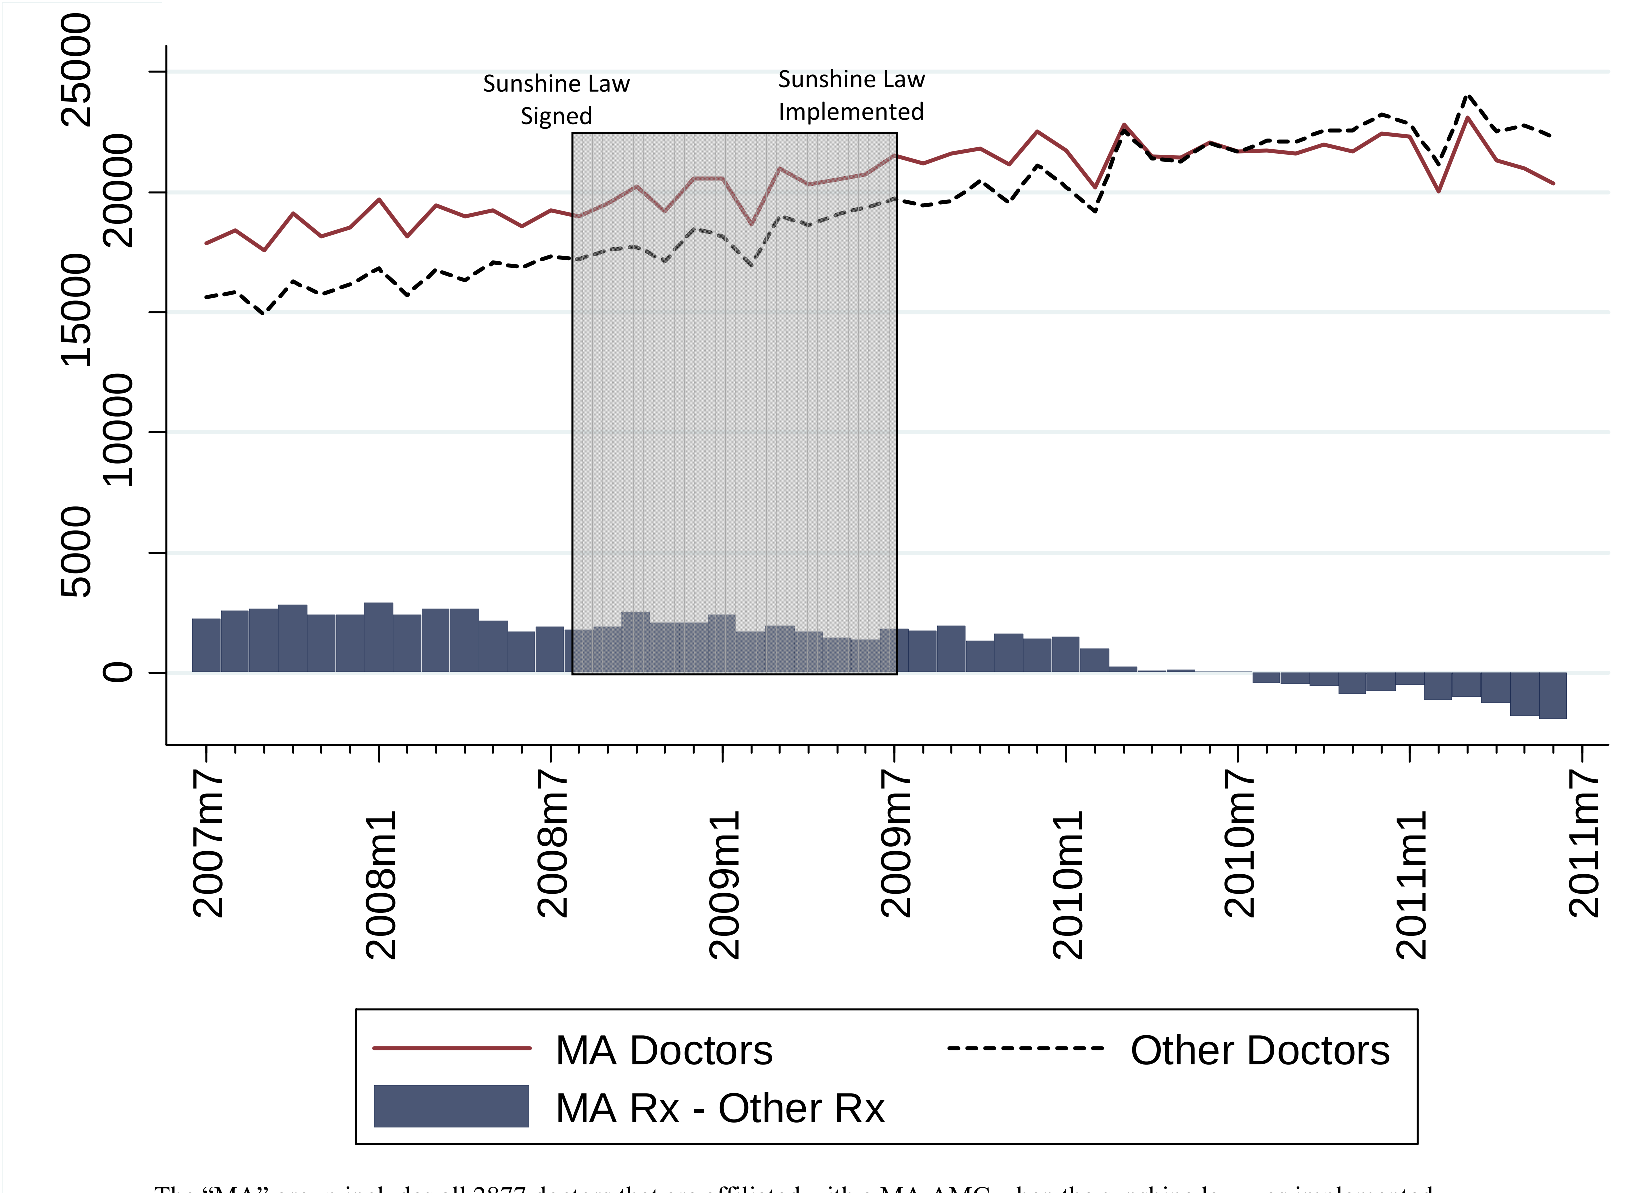

The "MA" group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The "Other" group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 5: Total Generic Rx (Four Years)

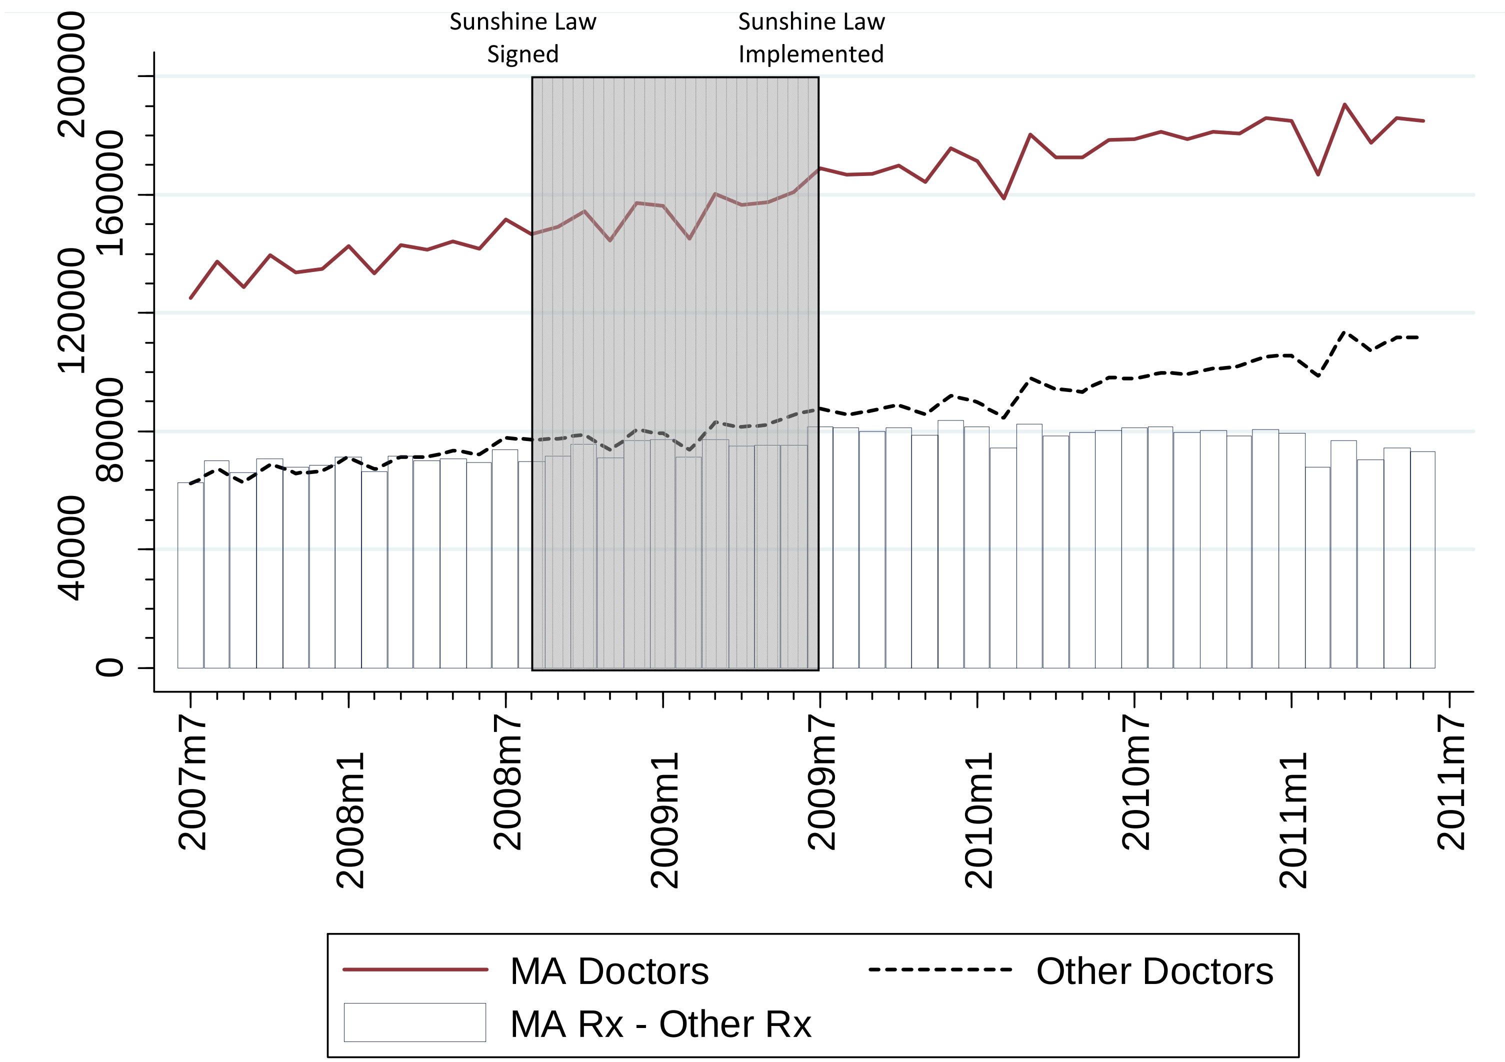

The “MA” group includes all 2877 doctors that are affiliated with a MA AMC when the sunshine law was implemented.
The “Other” group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Table 3: Sunshine Effects, OLS (Doctor-Month)

	(1)	**(2)**
DV	**Branded Rx**	**Generic Rx**
N (Doctor-Month)	**275,112**	**275,112**
Doctors	**5,755**	**5,755**
Sunshine law	-2.249**	2.497*
	(1.061)	(1.394)
Doctor, Month FE	YES	YES

*p<0.10; ** p<0.05; *** p<0.01
SEs are robust and clustered by AMC.

Table 4: Sunshine Effects, OLS (Doctor-Drug-Month)

	(3)	(3B)
DV	**Rx**	**Drug-Marketshare**
N (doctor-drug-month)	**7,541,535**	**7,237,375**
Doctors	**5,755**	**5,755**
Policies		
1. Sunshine	0.071	0.002***
	(0.063)	(0.000)
Interactions		
2. Sunshine*Branded	-0.635***	-0.006***
	(0.082)	(0.001)
Controls		
Branded Drug Indicator	YES	YES
Doctor*Drug, Month FE	YES	YES
Linear combinations of coefficients		
1 + 2	-0.564**	-0.005***
	(0.104)	(0.001)

*p<0.10; ** p<0.05; *** p<0.01
SEs are robust and clustered by AMC.

Table 5: Sunshine + AMC Effects (OLS)

DV = Rx	**(4)**
N (doctor-drug-month)	**7,541,535**
Doctors	**5,755**
Policies	
1. Meals Ban	0.091*
	(0.052)
2. Consulting Restrictions	0.062
	(0.116)
3. Sunshine Law	0.099
	(0.076)
Interactions	
4. Meals*Branded	-0.249***
	(0.075)
5. Consulting*Branded	-0.069
	(0.095)
6. Sunshine*Branded	-0.540***
	(0.058)
7. Meals*Sunshine	-0.060
	(0.086)
8. Consulting*Sunshine	-0.079
	(0.121)
Double-Interactions	
9. Meals*Sunshine*Branded	-0.068
	(0.073)
10. Consulting*Sunshine*Branded	0.239
	(0.150)
Controls	
Branded Drug Indicator	YES
Doctor*Drug, Month FE	YES
Linear combinations of coefficients	
1 + 4	-0.157**
	(0.069)
2 + 5	-0.007
	(0.085)
3 + 6	-0.442***
	(0.107)
1 + 3 + 7	0.130*
	(0.076)
2 + 3 + 8	0.081
	(0.074)
1 + 3 + 4 + 6 + 7 + 9	-0.726***
	(0.087)
2 + 3 + 5 + 6 + 8 + 10	-0.289
	(0.210)

*p<0.10, **p<0.05, ***p<0.01.
SEs are robust and clustered by AMC.

Figure 6: Branded Rx Per Doctor, Meals

The "MA – Accepted Meals" group includes 430 doctors that accepted meals per the sunshine law data.
The "MA – No Meals" group includes 2447 MA doctors that didn't accept any meals per the sunshine law data.
The "Other" group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Figure 7: Branded Rx Per Doctor, Consulting

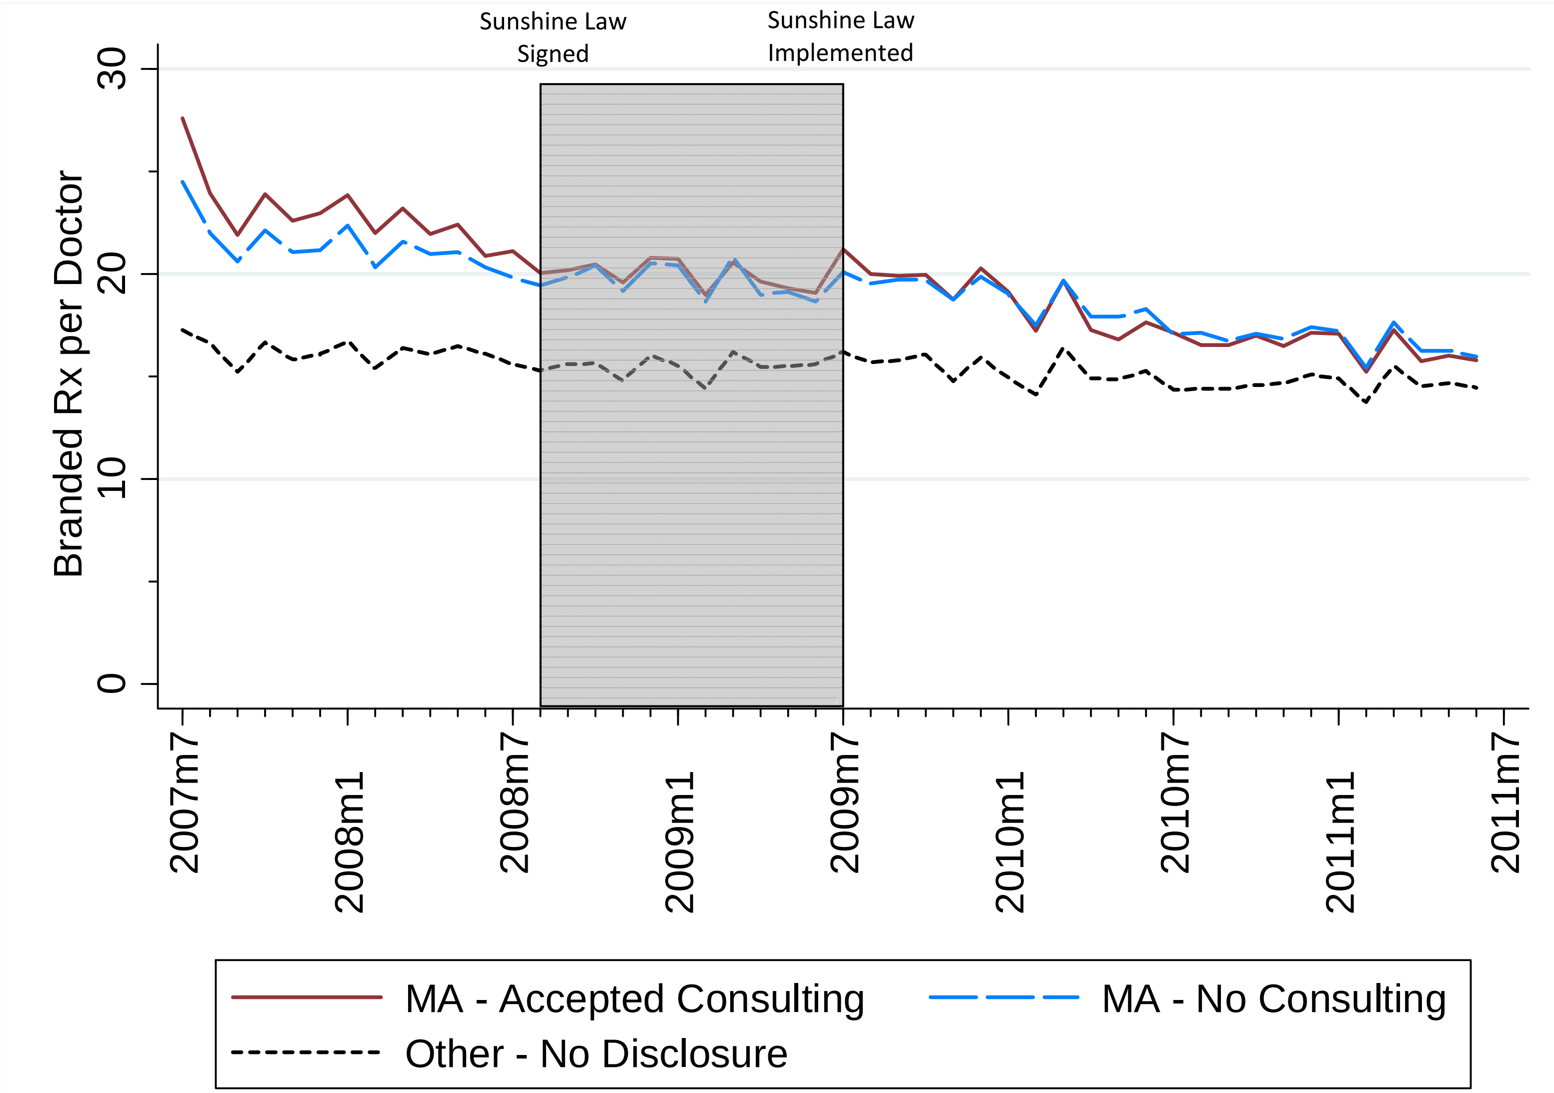

The "MA – Accepted Consulting" group includes 370 doctors that accepted consulting payments per the sunshine law data.
The "MA – No Consulting" group includes 2407 MA doctors that didn't accept any consulting payments per the sunshine law data.
The "Other" group includes the 2878 doctors that are not affiliated with a MA AMC when the sunshine law was implemented.

Table 6: Policy Effects (OLS) on Massachusetts Doctors, by Doctor Type

	(5)	(6)	(5B)	(6B)
DV	**Rx**	**Rx**	**Marketshare**	**Marketshare**
N (doctor-drug-month)	**7,541,535**	**7,541,535**	**7,237,375**	**7,237,375**
Doctors	**5,755**	**5,755**	**5,755**	**5,755**
Policies				
1. Sunshine	0.085	0.107	0.002***	0.001***
	(0.078)	(0.066)	(0.000)	(0.000)
Interactions				
2. Sunshine*Branded	-0.610***	-0.632***	-0.006***	-0.006***
	(0.090)	(0.079)	(0.001)	(0.001)
3. Sunshine*Meals-Doc	-0.005		-0.000	
	(0.231)		(0.000)	
4. Sunshine*Consult-Doc		-0.290**		-0.001
		(0.133)		(0.001)
5. Meals-Doc*Branded	0.904		-0.006***	
	(0.542)		(0.002)	
6.Consult-Doc*Branded		0.015		-0.006**
		(0.479)		(0.003)
Double-Interaction				
7. Sunshine*Branded*Meals-Doc	-0.202		0.000	
	(0.171)		(0.001)	
8. Sunshine*Branded*Consult-Doc		-0.027		-0.001
		(0.176)		(0.001)
Controls				
Branded Drug Indicator	YES	YES	YES	YES
Doctor*Drug, Month FE	YES	YES	YES	YES
Linear combinations of coeff				
1 + 2	-0.525***	-0.524***	-0.005***	-0.004***
	(0.108)	(0.092)	(0.001)	(0.001)
1 + 2 + 3 + 7	-0.732***		-0.004**	
	(0.106)		(0.002)	
1 + 2 + 4 + 8		-0.842**		-0.006***
		(0.222)		(0.001)

*p<0.10, **p<0.05, ***p<0.01
SEs are robust and clustered by AMC.

Table 7: Policy Effects (OLS) on High-Meals and High-Consulting Doctors

	(7)	(8)	(7B)	(8B)
DV	**Rx**	**Rx**	**Marketshare**	**Marketshare**
N (doctor-drug-month)	**7,541,535**	**7,541,535**	**7,237,375**	**7,237,375**
Doctors	**5,755**	**5,755**	**5,755**	**5,755**
Policies				
1. Sunshine	0.052	0.086	0.002***	0.002***
	(0.072)	(0.062)	(0.000)	(0.000)
Interactions				
2. Sunshine*Branded	-0.611***	-0.628***	-0.006***	-0.006***
	(0.079)	().077)	(0.001)	(0.001)
3. Sunshine*High-Meals	0.332		-0.001***	
	(0.334)		(0.000)	
4. Sunshine*High-Consulting		-0.220***		-0.000
		(0.042)		(0.000)
5. High-Meals*Branded	1.291		-0.007**	
	(0.834)		(0.003)	
6. High-Consulting*Branded		0.293		-0.007***
		(0.683)		(0.002)
Double-Interaction				
7. Sunshine*Branded*High-Meals	-0.375		0.001	
	(0.295)		(0.002)	
8. Sunshine*Branded*High-Consulting		-0.142		-0.003**
		(0.312)		(0.001)
Controls				
Branded Drug Indicator	YES	YES	YES	YES
Doctor*Drug, Month FE	YES	YES	YES	YES
Linear combinations of coeff				
1 + 2	-0.559***	-0.541***	-0.005***	-0.004***
	(0.102)	(0.096)	(0.001)	(0.001)
1 + 2 + 3 + 7	-0.602***		-0.005*	
	(0.175)		(0.002)	
1 + 2 + 4 + 8		-0.904**		-0.007***
		(0.373)		(0.002)

*p<0.10, **p<0.05, ***p<0.01
SEs are robust and clustered by AMC.

Table 8: Sunshine Law Effects on New Branded Drugs

	(9)	(9B)
DV	**Rx**	**Marketshare**
N (doctor-drug-month)	**7,541,535**	**7,237,375**
Doctors	**5,755**	**5,755**
Policies		
1. Sunshine	0.039	0.002***
	(0.063)	(0.000)
Drug Type		
2. New Branded Drug	-0.224***	-0.000
	(0.037)	(0.000)
Interactions		
3. Sunshine*Branded	-0.788***	-0.008***
	(0.097)	(0.001)
4. Sunshine*New Branded	0.463***	0.003***
	(0.050)	(0.001)
Controls		
Branded Drug Indicator	YES	YES
Doctor*Drug, Month FE	YES	YES
Linear combinations of coeff		
1 + 3	-0.748***	-0.006***
	(0.121)	(0.001)
1 + 3 + 4	-0.285***	-0.003***
	(0.089)	(0.001)

*p<0.10, **p<0.05, ***p<0.01
SEs are robust and clustered by AMC.

C H A P T E R   T H R E E

I. Introduction

Many non-profits offer incentives during fundraising drives to increase a prospective donor's motivation to give. Common incentives include matching donations or challenges, where third party benefactors make a contribution depending on whether certain goals are reached by other donors (Dove 2001). Many non-profits also use simpler and more direct economic incentives, such as offering small thank-you gifts to those that donate. However, while such thank-you gifts are recommended by many fundraising guidebooks (e.g., Dove 2001, Greenfield 2002), there is no clear empirical evidence that they are effective.

In fact, evidence from psychology suggests that economic incentives may interact with and reduce intrinsic incentives to donate (Deci 1971), making it possible for thank-you gifts to *decrease* donations. Indeed, some laboratory results suggest that thank-you gifts can decrease donations in certain fundraising contexts (Newman and Shen 2012). Similarly, there are economic models that suggest charitable donations are a self-signaling device for traits such as altruism; economic incentives can obscure the value of that self-signal and therefore decrease the likelihood of a donation (Benabou and Tirole 2006).

This study uses a large-scale field experimental design to measure how thank-you gifts can influence donation rates and amounts. In collaboration with a public radio station, I run a direct-mail field experiment that randomly varies whether prospective donors are offered a thank-you gift in exchange for a donation. The control group receives no thank-you gift offer, and the treatment group receives an offer of a thermos in exchange for a donation at or above a specific amount. An additional treatment group is instead offered a pro-social thank-you gift that only benefits others: the radio station will offer to donate meals to a food bank on the donor's behalf if they donate to the radio station. This *pro-social* gift can help identify whether crowding out is related to self-signaling, since this gift introduces no direct economic incentive to the donor to interfere with the self-signal value of the donation.

I find that donation rates decreased relative to controls when thank-you gifts were offered by direct-mail. This occurred with both types of thank-you gifts. 6.2% of donors in control treatments chose to donate, while only 4.3% of donors chose to donate in each of the gift treatments.[1] Furthermore, the percentage of donations made that were at or above the minimum required to qualify for the gift were unchanged from treatment to control; additionally, only 5 total donors in the treatment conditions (out of 66 that qualified for the gift) opted for the gift. Average donation amounts, conditional on donating, did not increase in treatment groups; if anything, average donation amounts were higher in the control group.

My results suggest that extrinsic incentives can interact with the intrinsic motives that are present in a non-profit fundraising context. Under certain conditions, such as those present in this study, this interaction can lead to motivation crowding out. This result contrasts with standard economic assumptions about relative price effects, which typically assume that marginal increases in extrinsic incentives should always weakly increase utility (Solow 1971; Arrow 1972).

These results suggest against a self-signaling mechanism for crowding out. First, the meals gift had a negative effect on donations, even though it provided no extrinsic incentive to the donor that would decrease the self-signaling value of the donation. Second, since gifts are always opt-in, donors concerned with self-signaling can simply refuse the gift and be at least as well off from a self-signaling standpoint as if no gift had been offered.

Instead, the results suggest that crowding out in this context may simply be due to attention and mindset. The direct-mail inserts called attention to both the gift and the donation amount required to be eligible for the gift. The inserts simultaneously detracted attention away from the intrinsic motives for donating, which were emphasized in the actual solicitation letter. This change in attention can cause donors to place more weight on the gift incentive (which is not a favorable incentive when compared to

[1] These figures are based on the first time each donor was eligible to receive a gift offer by mail. Donors that did not respond to this first solicitation were solicited by mail in subsequent months as well, and these data were not included in the analysis; however, including them will yield the exact same conclusions.

the required donation amount to be eligible), while also causing donors to place less weight on the intrinsic reasons to donate. Moreover, this change in attention can lead to a change in mindset (Gneezy, Li, Saccardo, and Samek 2014; Newman and Savary 2014), such as causing donors to view the donation in a cost-benefit mindset instead of a pro-social mindset. Since the cost-benefit terms for the gift are not favorable in this study (e.g., $180 for a thermos), this would also reduce donation rates.

From a practical perspective, if a non-profit chooses to give gifts, it must consider whether the economic benefits of the gift are enough to overcome possible negative crowding out effects. In many field experiments, the use of economic incentives actually increased pro-social actions in donors (Lacetera, Macis, and Slonim 2012, 2014); this may simply be because the economic incentives in these tests were desirable enough to overcome any crowding out effects. Alternatively, the saliency of the gift may also vary across these studies, also leading to variation in the degree of crowding out that might occur.

The rest of this paper is organized as follows. Section II discusses the relevant literature on both motivation crowding out and fundraising strategies. Section III analyzes the direct-mail membership renewal campaigns, including using donor attributes based on archival donation histories. Section IV discusses the cognitive mechanisms that may be driving the observed treatment effects. Section V examines a follow-up online experiment that tests the mechanisms directly. Section VI concludes.

II. Related Literature

Crowding Out

Motivation crowding out has been well-studied in cognitive psychology. Edward Deci (1971) first demonstrated that providing extrinsic motivation can reduce self-perceived intrinsic motivation on a task; in particular, Deci's subjects showed less intrinsic interest in a puzzle task if they had previously been compensated for working on the same task. Subsequent studies replicated this effect in other contexts (see Deci, Koestner, and Ryan 1999 for summaries), and some of these studies alternately termed this effect the "over-justification effect" (Lepper, Green and Nisbett 1973) and "cognitive evaluation theory" (Deci, Koestner, and Ryan 1999) among other names.

At around the same time as Deci's seminal study, the social scientist Richard Titmuss (1970) published his influential book *The Gift Relationship*, where he hypothesized that paying people to give blood can potentially decrease donation rates. Titmuss' hypothesis spawned a stream of field experiments, to mixed results (Ferrari et al. 1985; Reich et al. 2006; Mellstrom and Johannesson 2008; Goette and Stutzer 2008; Costa-Font, Jofre-Bonet and Yen 2012; Niza, Tung and Marteau 2013; Lacetera, Macis and Slonim 2012 and 2014).

Economic models have accounted for crowding out by allowing for interactions between extrinsic and intrinsic incentives (Frey and Jegen 2001). For instance, a field experiment showed that individuals offered low pay for a task were less likely to complete the task than if they were offered no pay at all (Gneezy and Rustichini 2000). Other studies suggested that financial incentives can crowd out willingness to volunteer (Frey and Goette 1999), to conform to a social norm (Fehr and Gachter 2000; Gneezy and Rustichini 2000), or to perform a civic duty (Frey and Oberholzer-Gee 1997). Non-monetary forms of extrinsic incentives, such as institutional rules and regulations, can also crowd out willingness to perform a civic duty (Frey 1997) or to adhere to a norm (Bohnet, Frey and Huck 2001).

Like these other studies, I use a field experiment to test for relationships between intrinsic and extrinsic incentives. However, instead of the contexts studied in previous papers, I examine whether extrinsic incentives can crowd out one's willingness to donate money to a non-profit. Importantly, this could cause extrinsic incentives (such as a conditional gift) to be viewed as akin to a purchase, which is not possible when the action in question is a service or duty (such as giving blood). A laboratory study has investigated crowding out in this fundraising context (Newman and Shen 2012); however, aside from being limited to a lab setting, their study does not use optional gifts or pro-social gifts, and thus it cannot test whether crowding out is driven by self-signaling or by alternative mechanisms.

Charitable Giving

Field experiments have been used to test other common fundraising strategies. These include donation matching (Karlan and List 2007; Meier 2007), donation seeding, (List and Lucking-Reiley 2002), and social pressure tactics (DellaVigna, List, and Malmendier 2012, 2013; Andreoni, Rao and Trachtman 2012), among others. Falk (2006) tests whether unconditional gifts from a non-profit can increase donations and shows that large gifts increased giving rates by 75%, while small gifts increased giving rates by 17%.[2] However, this strategy varies in a subtle but important way from thank-you gifts in fundraising: Falk invokes reciprocity by giving unconditional gifts that precede the solicitation request, which can cause the gift to be perceived differently than if they were conditional on donating.

Eckel, Herberich, and Meer (2014) run a similar test on gifts by collaborating with Texas A&M University. The authors use direct-mail to solicit Texas A&M alumni for donations, and they offer conditional opt-in or opt-out luggage tags in exchange for donations. They find no effects of conditional gifts on donations. However, their study was designed to measure possible positive effects of gifts on donations, and the test was not geared towards detecting crowding out. As a result, they offered the gift in exchange for any donation amount (even $1 or $5 donations); this favorable gift-to-donation value ratio makes crowding out more difficult to detect, since the high value of the gift (relative to donations) may cancel out any crowding out effects. In addition, their study targeted prospective first-time donors or lapsed donors, yielding a donation rate of less than 1%. This decreased statistical power and decreased the strength of the null effect that they found.

Like these other studies, I use a field setting to evaluate the efficacy of a commonly used fundraising strategy. Unlike Eckel, Herberich, and Meer (2014), I specifically implement an experiment designed to test whether conditional items can crowd out intrinsic motivation to give. I use conditions that are more conducive for detecting crowding out, such as requiring a significant minimum donation to obtain the gift and a gift-to-donation value ratio of 1:10, which is similar to many industry standards. My

[2] In addition, the gift items in Falk (2006) were postcards drawn by needy children who were supported by the non-profit. These gifts may invoke not only reciprocity but also empathy, so the results are not a pure reciprocity effect.

study also focuses on renewal candidates who have given in the past instead of prospective first-time donors, yielding donation rates much higher than 1% and generating sufficient statistical power to test for crowding out.

III. Direct-Mail Renewal Tests

Background

I collaborate with a non-profit public radio station to test for the effects of conditional thank-you gifts during the station's rolling direct-mail membership renewal fundraising campaign. The station defines a member as any individual that has contributed any amount, in any fundraising campaign, within the past twelve months. When a donor's membership is three months from expiration, the donor is solicited by mail to renew.[3] Unless the member renews, the renewal letters will continue to be mailed – one per month – up through the fifteenth month beyond their last donation; in other words, the member could receive donation letters each month when they are between nine and fifteen months removed from their last donation. (See Table 1 for the abbreviations that the station uses to refer to each of these months within the membership cycle, and which this paper will also adopt for brevity's sake).

The renewal letter is a short, 1-page solicitation message thanking the member for their past support and requesting a new donation to renew the membership. The solicitation letter also reminds donors of the many reasons why they donate, such as the station's original programming and the NPR content that they broadcast. The letter includes a detachable remit form (see Figure 1) that donors can mail back in order to renew their membership. These remit forms contain a unique identifier that allows the station to track exactly which mailing the donor is responding to. The suggested donation amounts listed on the remit form typically depended on the amount the donor has donated in past years.

[3]The only exceptions are the very few members who requested a "do-not-mail" designation (<1% of active donors).

Procedure

Testing was run by direct-mail from May–July 2014. The standard solicitation letter contains no reference to thank-you gifts, and this was used for the control group. For treatment groups, the same letter was sent, but a colored buck-slip insert was also included in the envelope that advertised one of two possible thank-you gifts (see Figure 1 for a black and white image of the slips). One treatment group received an offer of a *swag* gift item in exchange for a $180+ donation, which in this test was a tumbler (i.e., a thermos). A second treatment group received an offer of a *meals* gift item for a $180+ donation, which in this test meant the station would offer to donate 60 meals to the local food bank on the donor's behalf. Both gifts have a fair market value of $15, although this was not stated in the gift offer.

The remit forms attached to these letters were the same across all three conditions. For most donors, the suggested amounts listed on the remit form were $120, $180, and $240, although some donors saw suggested amounts of $25 and $89.[4] In treatment conditions, the remit form is simply altered to include a check-box for selecting the thank-you gift, but all other aspects of the remit form remained identical across conditions.[5]

In each month, all renewal members were randomly split into three conditions: control, swag, and meals. In the first month of testing, May 2014, all members scheduled to receive a mailer (that is, any member that was nine to fifteen months removed from their most recent donation) were randomly assigned to one of the three conditions. The parent company responsible for the randomization split approximately half of the subjects into the control group, and then split the remainder 50/50 between swag and meals treatments. In the second and third months of testing (June and July 2014), all new R1 members were randomly distributed across all three groups in similar fashion.[6]

[4] The station was testing a new ask string with the lower suggested amounts of $25 and $89 for donors that gave very low amounts in the past. Randomization of treatment assignment ensured that the ask strings were evenly and randomly distributed across treatment groups.
[5] To equate the gift conditions, donors were not asked to pay for shipping costs if they selected the thermos.
[6] Some control subjects who had been solicited in May were randomly assigned to new groups in June if they did not respond to the May mailer; the station did not require that they remain in the same treatment every month. They were still randomly assigned, and robustness checks show that all reported results also hold when running only on May data, before any treatment re-assignment.

Data

Over these three months, a total of 19,636 pieces of direct-mail were sent to 10,232 unique members about renewing their membership. Figure 2A breaks down the number of mailers sent by month and condition. Although more donors were assigned to the control group than the other groups, the process was random (the station randomized approximately 50% of all subjects to the control group). Figure 2B breaks down the mailers according to the renewal cycle. R1 refers to donors who are nine months removed from their last donation and receiving a mailer for the first time in 2014. Similarly, R2 donors are those receiving a second mailer this year; note that R2 donors in May would have received a mailer in April, but would not have been eligible for a gift offer until May, when the experiment began.

I limit the data to observations where a donor was eligible to receive a gift offer for the first time. This includes all donors solicited in May, as well as the R1 donors in June and July. These observations are outlined in red in Figure 2B. This avoids confounds that would otherwise arise from including more than one observation for each donor. For instance, donors essentially self-select into receiving multiple mailers by choosing to ignore the first mailer, and thus these donors would be weighted too heavily in any analysis that treats every mailer as an independent observation. In addition, the station would sometimes re-assign donors to a different condition if they did not respond to the first mailer; while the re-assignment was still random amongst the remaining conditions, this could also introduce bias if not treated carefully. I avoid these confounds by simply using the restricted data specified, although *all* results in this paper are robust to using the entire dataset (see appendix).

Figure 3 demonstrates that the randomization process was generally successful. In aggregate, subjects showed similar distributions in demographics, including race, gender, and age.[7] Recent donation histories also seem fairly similar across conditions, although this historical data was rather noisy and thus hard to parse.[8] The average number of pledges and the average pledge amount from the past year were

---

[7] Gender and race data were inferred using U.S. national Census data from 1990 and 2000 on surname-race statistics and name-gender statistics.

[8] In particular, a single pledge would sometimes be split into multiple observations in the archival data, making both total-pledges and average-donations-per-pledge difficult to determine. In addition, both pledge and payments data

generally similar across conditions. However, compared to the treatment subjects, control subjects were more likely to have donated by phone or mail than by web. This is likely because older donors are less likely to give by web, and due to randomness the control subjects happened to be very slightly older on average than those in the treatment groups. Nevertheless, controlling for these slight differences in groups will still yield similar results, as will be shown in the Results section.

Main Results

Figure 4 displays summary statistics by condition and difference-in-means t-tests between conditions. 6.2% of donors in the control group chose to donate by mail, while 4.2% and 4.3% of donors in the meals and swag conditions chose to donate. These differences-in-means are significant at the $p<0.01$ level, suggesting that donors were less likely to donate (by almost 33%) when offered either type of gift. Since I have no data on whether donors actually opened the envelope and saw the buck-slip that advertised the gift, these regressions represent intend-to-treat effects. Actual treatment effects on the treated are likely even larger than observed.

Table 2 reports standard OLS estimates of how thank-you gifts affected renewal rates. The dependent variable is a binary indicator representing whether a donor chose to donate in response to the mailer. In Model (1), the regressor of interest is an indicator for whether any thank-you gift was offered; in Model (2), the regressors of interest are binary indicators for whether a swag item or meals item was offered. I also include month fixed effects to account for seasonality trends. Both models suggest that thank-you gifts significantly decreased donation rates. Although the coefficient sizes are small, they suggest a decrease in giving rate similar to what is observed in Figure 4.

Since the dependent variable is binary, Table 2 also reports logistic regressions as a robustness check. I present OLS as my main results because they do not yield out-of-sample predictions. Nevertheless, the standard logistic regression in Model (3) yields similar results as before. In addition,

---

were recorded as separate observations within the dataset, yielding multiple observations for each donation; while pledges and payments were flagged separately, the data did not always add up correctly for all donors.

Model (4) uses a penalized maximum likelihood logistic model that accounts for the fact that donation rates are relatively rare events; this specification is typically reserved for samples with very few occurrences of a particular outcome, e.g. 200 or fewer of a particular outcome (Allison 2012). This is not particularly applicable in this instance, since I have over 500 decisions to donate. Nevertheless, results are also robust to this alternate specification, and yield very close results to the standard logit in Model (3). In future regressions, I will report only OLS for ease of interpretation, and because they never yield out-of-sample predictions.

Figure 5 examines whether average pledge amounts differ between treatments, conditional on choosing to donate. A total of 509 mailers resulted in donations, with average donations of \$107 for donors in the control condition, \$96 for donors in the swag treatment, and \$94 for donors in the meals treatment. This suggests that, if anything, gift offers led to lower donations on average. Consistent with this, 4.3% of donors in each of the gift groups donated at least \$180 (the minimum required to be eligible for a gift in the gift treatments), compared to 5.8% of donors in the control group, suggesting that the \$180 threshold did not cause donors to give at the threshold more frequently. Even conditional on donating, donors in the gift groups did not give at the threshold more often; 9.9% of those that donated in the swag group and 8.3% of those that donated in the meals group gave at or above the threshold, compared to 9.3% of those that donated in the control group. Not surprisingly, this is in part because donors were not particularly tempted by the gifts. Of those that gave enough to qualify for the gift, only four out of ten accepted the meals gift, and only one out of thirteen accepted the swag gift.

Table 3 reports regression results with conditional pledge amount as the dependent variable, and also includes month fixed effects. Models (5) and (6) only include observations where pledges were made. Model (5) uses OLS to demonstrate that average donation amounts were not different between treatment and control. Model (6) uses OLS to show that the likelihood of donating at least \$180 was not different between groups. Together, Tables 2 and 3 suggest that in this three month test period, the gift offers decreased total donations by approximately $12000, which equates to approximately 15% of the total amount (~$80k) raised during this three-month mailer campaign.[9]

Since the mail renewal campaign is ongoing year-round, the first month of testing (May 2014) included donors from all parts of the renewal cycle. That is, it included some donors that had already received mail solicitation letters in April 2014 or before, but are still receiving mailers in May because they did not respond to the previous mailers. Since donors self-select into receiving these repeat mailers, it makes sense to control for the number of mailers each donor had received when the test began. Figure 6 shows donation rates by treatment and renewal-month status; R1 stands for the first time a donor is receiving a mailer this year, R2 is the second time, and so forth. The figure shows that, while there is significant variation in donation rates by Renewal Month, R1 donors are driving most of the observed crowding out effects.[10] The red stars above some of the treatment group bars represent the p-values for two-tailed difference-in-means t-tests between that group and the relevant control group.

Table 4A displays OLS regressions that control for renewal-month status or interact it with treatment indicators. I use binary indicators for R2 through R7 (R1 is the omitted indicator). Model (9) shows that gift effects are still significant when controlling for renewal-month status, without interactions.[11] Model (10) adds interactions between renewal-month status and treatment. While many of these interaction terms are significant, they only denote the additional incremental effect of treatment on a donor in that particular Renewal Month; the net effect of treatment is actually represented by the linear combination of the main effect of treatment plus the relevant interaction term. Table 4B displays the relevant linear combinations of effects for Model (10), and these linear combinations demonstrate that gift effects are only present in R1-meals, R1-swag, and R6-swag.[12]

---

[9] There were 5,874 treatment mailers. Average pledges were $101 across all conditions, and Table 3 cannot reject the null that average donations were the same across treatments. Table 2 suggests that donation rates dropped 2% in response to gift offers. This leads to an estimated decrease in donations of about $12000.

[10] Note that if a donor gives during their R1 month, their donation renews their membership for 12 months from the date their current membership expires (and not 12 months from their new donation date).

[11] These results are even stronger when pooling both treatments into a single indicator representing any gift offered.

[12] These are linear combinations (i.e., joint significance tests) of the Swag and Meals indicators with the relevant Renewal-month*Treatment interaction coefficient.

These latter results suggest that once a donor has chosen to ignore the first mailer they receive, the gift offer no longer matters in subsequent mailers. This may be because these donors have chosen not to donate regardless; they may also disproportionately represent individuals who are not opening the mailers. This group also includes individuals who choose to wait until R4 to donate (the twelve month mark, when their membership officially expires). However, note that if you donate in a month prior to R4, your membership is extended by the full 12 months from the month your membership expires, and not the month that you donated.

Results by Sub-Groups

One concern, as illustrated in Figure 3, is the fact that some donor-level traits seemed to vary slightly across groups, despite the randomization. In particular, control groups seem slightly older, slightly more Caucasian, and slightly more likely to give by mail or phone (and slightly less likely to give by web). The most important difference is in the rate of donating by mail, since this is the factor most likely to influence base donation rates across groups. This difference in donation method is likely related to the slight difference in age across groups, as older donors tend to prefer giving by mail or phone more-so than younger donors. This difference resulted in double- and triple-checking of the randomization process, which nevertheless seems valid; donors scheduled to receive a mailer (based simply on their last donation made) were randomized into groups at the individual level via a random-number style algorithm.

Table 5 displays OLS regressions that control for these factors. Model (11) uses an indicator to represent whether they gave by mail in 2013, and also interacts this indicator with each treatment. Month and Renewal-Month fixed effects are included.[13] The main effects of Swag and Meals show that those that did not give by mail in the past year *did* experience crowding out; they were 1.3% and 1.4% less likely to donate if offered the Meals and Swag items, respectively. Not surprisingly, any donor that gave by mail in the past year was much more likely to donate by mail again this year (9.7% more likely, regardless of condition). The linear combination terms suggest that those that gave by mail in the past

[13] Results are robust to excluding these.

year also experienced crowding out when offered the meals gift; the swag gift also showed near marginal significance for crowding out ($p < 0.12$). Thus, both those that gave by mail in 2013 and those that did not' seem to show some crowding out in response to gift offers.

Model (12) splits the donors by race, since the control group showed a potentially slightly higher ratio of Caucasians. I include a variable for Caucasian that corresponds to a value between 0 and 1, representing the likelihood that the given donor's last name represents a Caucasian ethnicity. I interact this with each treatment variable. The results suggest that Caucasian donors drive the crowding out results in both gift treatments. However, this may simply be driven by sample size, since approximately three-fourths of all donors in the dataset (in each condition) are Caucasian. Nonetheless, it is useful to note that the treatment groups had slightly lower rates of Caucasians, and yet still demonstrated crowding out in that group, suggesting that any bias from the small difference in race distribution across groups would bias against finding a crowding out effect.

Model (13) includes a variable to capture donor age. The radio station obtained age data by purchasing a consumer behavior dataset from a third party vendor, which matched donor names and addresses to publicly available data sources on age. Only approximately 75% of donors in the dataset were matched, so regressions including age drop approximately 25% of the data. Model (13) includes a measure of age that is normalized between 0 and 1. This Model shows main effects of treatment that are still negative, but not near significance; likewise, the linear combinations suggest that the effects are essentially zero for the oldest in the data (the reported linear combinations represent treatment effects on the hypothetical 100-year old donor). There may be several reasons for the lack of statistical significance in these results; first, as mentioned, the regression drops 25% o the data for having no age data. Second, there is likely significant noise in the age data; the oldest matches in the data were above 110 years of age, which is unlikely to be accurate (indeed, although I called the linear combination effects a representation of the hypothetical 100-year old donor, there are donors in the dataset with a calculated age of 100). Although all ages of 100+ were dropped, there may still be mismatches elsewhere in the data that may be confounding results.

Results Based on Past Thank-You Gift Behavior

Importantly, the station has offered thank-you gifts in the past – just never by mail. The station has offered thank-you gifts through on-air advertisements and on their online donations page since at least 2007. Thus, their archival data includes information on whether donors have accepted similar thank-you gifts in the past, including the meals gift. Importantly, since they have never been offered by mail before, donors would not have forestalled the gift offer that was included in the direct-mail experiment. I use this information on past thank-you gifts to estimate crowding out effects among different donor types. Those that have accepted gifts in the past could conceivably respond positively to gifts, while those that have never accepted a gift could be the ones driving the observed crowding out effects.

Table 6 re-analyzes the direct-mail field experiment using additional information from the archival data. I run the same base linear models as before, but I run them separately on those that have ever accepted a thank-you gift in the past, and those that have never done so.[14] These regressions are Models (14) and (15), respectively. These results show that crowding out is driven primarily by those that did *not* accept a thank-you gift when they first joined as a member; for these donors, swag and meals gifts caused decreases of 2.0% and 2.3% in donation rates, respectively. Those that have accepted gifts in the past showed no crowding out effects, although the coefficients still skewed negative. Importantly, those that have accepted gifts before did not respond positively to gift offers, even when the gift was an item that they would likely still want even if they had selected the item in prior years (i.e., the meals gift).

Donors that have accepted thank-you gifts in the past are more likely to have given via the web in the past (since traditionally that is the only type of donation that offered gifts, although there are nonetheless plenty of exceptions in the station's history). Models (16) and (17) control for this by including an indicator for whether a donor has ever given via web, and interacting this indicator with treatment. The results show that the results from Models (14) and (15) are robust to controlling for whether a donor has given by web in the past, although crowding out is driven primarily by those that

[14] Results are the same when running the specification on all donors and using main effect and interaction terms to separate donor types. I present the results as separate regressions for ease of interpretation.

have never given by web before (this is partly a sample size effect, since those that have never given by web represent a majority of the donors in the dataset).

IV. Cognitive Mechanisms

The observed negative effects run counter to standard economic concepts. Unlike standard marginal price effects, adding economic incentives appears to decrease donor response. There are several possible theories that can explain parts of the results in this paper, but only one theory that can convincingly explain every result.

Benabou and Tirole (2006) model motivation to donate as a self-signaling mechanism. In their words, "it [is] rational to define oneself partly through one's past choices: '*I am the kind of person who behaves in this way*.'" This implies that donors give in part so that they can then tell themselves that "*I am the kind of person who donates money to benefit a public radio station*." A reward offer will dilute the value of this self-signal, since donors may instead view the signal as "*I am the kind of person who donates money in order to receive a thermos*." Thus, for the gift to help donation rates, it must present greater utility to the individual than the decrease in utility caused by a dilution of the self-signaling value of the donation.

There are two reasons why this self-signaling interpretation cannot explain these results. First, it is not clear why the meals gift would dilute the self-signaling value of a donation, since it does not offer direct economic utility to the donor. Second, the fact that the reward is opt-in should theoretically provide donors the opportunity to avoid diluting the value of the self-signal. If the donor simply chooses not to opt-in to the gift, then they will remember their actions exactly as if no gift was offered, and the self-signaling value should not be any different. For this second reason, Benabou and Tirole's model would predict that neither type of gift should cause crowding out, especially if donors refuse the item (which almost all do).

It could be that the gift offer provides information to the donor. For instance, it could imply that the non-profit is not really in need of donations. This seems unlikely simply because most donors should understand that the gift is meant to induce donations, and not because the non-profit is flush with resources. In addition, it seems unlikely to assume that many donors give based purely according to which organization they deem to require the funds the most; most give because they enjoy the programming and would like to support it. Finally, most of these donors listen to the station regularly and almost certainly have heard the on-air advertisements for thank-you gifts. Thus, they should not be surprised by the gift offer in the mailer, and the offer should not be providing any new information to the donor regarding the degree to which the station needs funds.

Perhaps the effect is instead driven by including a glossy, colored insert in the mailer, and the content of the insert does not matter. For instance, if donors saw the insert and assumed it was an advertisement, they may be more likely to throw the mailer away without reading any further, leading to the observed effects. However, the envelope contains the radio station's logo on the outside, so donors should have some knowledge about the likely contents of the mailer, especially since these are donors who have given in the past and have received these mailers previously. In addition, the negative gift effects occur even among those that have donated by mail in the past, which represent the donors most likely to recognize the envelopes and mailers for what they are.

Some could posit a reference effect to explain the observed behavior. Donors in gift treatments may view $180 as an important reference point, since that is the amount required to qualify for the gift. However, the most likely reference points in all treatments would be the suggested donation amounts, which are the same across treatments. In addition, those that choose to donate are no more likely to give at or above the threshold in treatment vs. control conditions; this suggests that, conditional on choosing to donate, the threshold is not affecting how much donors choose to donate. In addition, the vast majority of those that give at or above the threshold are not even opting for the gift, suggesting that the gift is not likely to cause donors to commit to the threshold as an important reference point. These all suggest against a reference effect.

Instead, the negative effect of gifts is likely due to an attention and saliency mechanism. The colored insert mentions the gift item and the minimum required donation to receive it, but makes no mention of any intrinsic reason to donate (such as supporting your favorite programs). This causes people to shift their attention away from their intrinsic motivation for donating and toward the economic incentive being advertised. This can reduce the amount of weight that donors put on the intrinsic incentive when choosing whether to donate, while simultaneously increasing the weight placed on the economic incentive. In addition, this shift in attention can also cause donors to use a less pro-social and more cost-benefit mindset (Gneezy et al. 2014; Savary and Newman 2014); a cost-benefit approach would reduce donations in this context, since neither gift represents good economic value at the $180 threshold.

In some respects, this is the reverse to fundraising strategies that ask you to purchase an item for a high price, but where proceeds of the sale go to charity. Such a strategy explicitly calls attention to the intrinsic incentive to give, so that despite the economic incentive involved (the item being purchased), donation rates increase significantly relative to direct solicitation requests (Koppel and Shulze 2009; Newman and Savary 2014). In the next section, I test this mechanism more directly.

V. Mechanical Turk Experiment

Procedure

To directly test the role of attention and mindset in crowding out, I designed a 30-minute survey for Mechanical Turk workers to fill out. At the end of the survey I gave workers the opportunity to donate a portion of their earnings to charity. Subjects were paid $1.50 for the survey, and they were able to earn an additional $1.00 to $1.40 bonus from playing various social preference games within the survey. They were then asked whether they would like to donate $1 of that bonus to charity. Some subjects would be offered a conditional thank-you gift contingent on donating; of these, some were offered the gift up front and in a cost-benefit frame, while others were offered the gift in a more pro-social and less salient way. The exact language used is displayed in Figure 7.

The survey required workers to finish several tasks before choosing whether to donate. These tasks included asking subjects to play a short series of fairness, trust, and reciprocity games. These are games derived from Charness and Rabin (2002). Subjects earned from $1.00 to $1.40 on these games, depending on their choices and those of a randomly selected partner (the partner being another survey taker, chosen at random after the survey is submitted). These tasks measured whether attitudes toward thank-you gifts are correlated with adherence to any of these particular norms. Importantly, these games also provided a source of income that the worker can choose to donate to charity.[15]

After subjects finished these games, the survey asked subjects if they would like to donate $1 of their bonus to a charity. Subjects were informed that they earned at least $1 in bonus, but they were not told the exact amount they earned (thus all subjects were operating off of the same information about their pay). The charity options were the Red Cross, the Boys and Girls Clubs of America, the United Way, and the American Cancer Society. In the control treatment, no incentives were offered for donating. The swag treatment offered donors a $0.10 Amazon promotional code if they donated $1, while the meals treatment offered a $0.10 donation to Feeding America (a national food bank organization) for the same $1 donation. I chose $0.10 because it sets the gift-to-donation ratio at 10%, which is similar to the ratio offered by the radio station.

Results

Figure 8 displays the summary statistics for when the less salient and pro-social framing was used. There was no difference in donation rates between control, swag, and meals conditions. There was also no difference in response to gifts across any measure of donor heterogeneity, including donor responses to the social games portion of the survey.

Figure 9 displays the same statistics for when gifts are framed more prominently, as in the mailer solicitation. I find that the meals condition reduced donation rates 33% relative to controls, suggesting

[15] Due to the Mechanical Turk setup, I cannot ask subjects to donate any portion of their $1.50 fee for taking the survey. Thus it has to be donated from their worker bonus, which is earned through these games.

crowding out. However, this is still only significant at the $p<0.067$ threshold. The results improve slightly when excluding subjects that took the survey too fast or too slow (<5 minutes or >90 minutes), but results are still only marginally significant. Nonetheless, this suggests that attention and mindset can change whether crowding out occurs.

It is unclear why this re-framing matters for the meals condition but not the swag condition. Most likely, the value of the swag simply is of equal or greater magnitude than any decrease in utility that might result from crowding out. Also note that donors request the gift much more frequently than they did in the radio station campaigns; this may be because the gift is provided by the experimenter, and not at the expense of the non-profit they are donating to.

VI. Discussion

This study demonstrates that intrinsic and extrinsic incentives can interact in a fundraising setting. Introducing extrinsic incentives can crowd out a prospective donor's intrinsic motives for supporting a non-profit. In the direct-mail tests, conditional gifts decreased giving rates by over 30% relative to controls, leading to a potential 15% decrease in funds. This is despite the fact these gifts are always opt-in only, and that donors almost never opted for the gift anyways, even if they were eligible for them.

This evidence is not consistent with the self-image or self-signaling interpretation of crowding out in Benabou and Tirole (2006). In this field experiment, donors could maintain the self-signaling value of their donation by simply refusing the gift. They could then still interpret their act of donating as a charitable act not confounded by any extrinsic incentives. In addition, the meals gift is pro-social in nature and also should not decrease the self-signal value of a donation. For these reasons, self-image is unlikely to be the driving force behind the observed negative gift effects.

Instead, crowding out likely occurred because the direct-mail inserts caused prospective donors to shift attention away from their intrinsic motives and toward the utility value of the gift. This shift can also cause donors to change their mindset toward their donation; a conditional gift offer may cause donors to adopt a cost-benefit or exchange mindset toward the donation, instead of a pro-social mindset. Such an interpretation is also consistent with the Mechanical Turk results; changes in framing that increased gift saliency and cost-benefit mindsets marginally increased crowding out. Stronger changes, such as emphasizing the item more prominently, could result in even more crowding out.

Additionally, these effects were driven primarily by a specific type of donor. Only those that had never accepted a gift in the past were negatively affected by the gift offer. Just as importantly, those that had accepted gifts in the past did not respond positively to the gift offer. This lack of positive response is not because donors already have the item, since donors that accepted the meals gift in the past should still be interested in providing another set of meals to the food bank. Altogether, these results suggest that gifts for these particular fundraising campaigns have not benefited the non-profit.

Importantly, this does not suggest that extrinsic incentives should never be used in a fundraising context. Instead, it suggests that any extrinsic incentive that is introduced must have sufficient value and desirability to overcome a possible endogenous reduction in intrinsic incentives. In some cases, the extrinsic incentive may be enough to overcome any crowding out. This may be what occurred in field experiments that offered pay for blood donations (Lacetera, Macis, and Slonim 2014).

More generally, these results emphasize that non-profit organizations must carefully consider the psychological incentives that their donors face. While conditional incentives might make sense in a standard for-profit marketing perspective, they carry much different implications in a non-profit setting. Non-profit organizations can better optimize their fundraising strategies by being more cognizant of the psychological incentives that are particular to the contexts in which they operate.

References

Allison P. Logistic regression for rare events. *Statistical Horizons Blog; February 13, 2012. Accessed May 28, 2015 at statisticalhorizons.com/logistic-regression-for-rare-events.*

Andreoni J. Giving with impure altruism: Applications to charity and Ricardian equivalence. *Journal of Political Economy 1989; 97(6): 1447-1458.*

Andreoni J. Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal 1990; 100(401): 464-477.*

Andreoni J, Rao J, Trachtman H. Avoiding the ask: A field experiment on altruism, empathy, and charitable giving. Working Paper, 2012.

Arrow K. Gifts and exchanges. *Philosophy and Public Affairs* 1972; 1, Summer, 343-362.

Benabou R, Tirole J. Incentives and prosocial behavior. *American Economic Review* 2006; 96(5): 1652-1678.

Berinsky A, Margolis M, Sances M. Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. Working Paper, 2013.

Bohnet I, Huck S, Frey B. More order with less law: On contract enforcement, trust, and crowding. *American Political Science Review* 2001; 95: 131-134.

Charity Navigator. Tax benefits of giving. 2014. Available at www.charitynavigator.org. Retrieved on September 27, 2014.

Cialdini R. Influence: The Psychology of Persuasion. HarperBooks: USA, 2006.

Deci E. Effects of externally mediated rewards on intrinsic motivation. *Journal of Personality and Social Psychology* 1971; 18: 105-115.

Deci E, Koestner R, Ryan R. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin* 1999; 125(3): 627-668.

DellaVigna S, List J, Malmendier U. Testing for altruism and social pressure in charitable giving. *Quarterly Journal of Economics* 2012; Vol 127(1): 1-56

DellaVigna S, List J, Malmendier U, Rao G. The importance of being marginal: Gender differences in generosity. *American Economic Review Papers and Proceedings* 2013; 103(3): 586-590.

Dove K, Lindauer J, Madvig C. Conducting a Successful Annual Giving Program. Jossey-Bass: USA, 2001.

Eckel C, Herberich D, Meer J. It's not the thought that counts: A field experiment on gift exchange and giving at a public university. Work in progress.

Falk A. Gift-exchange in the field. *Econometrica* 2007;75(5): 1501-1511.

Fehr E, Gachter S. Do incentive contracts crowd out voluntary cooperation? Working Paper No. 34, Institute for Empirical Research in Economics, Working Paper Series, Zurich University.

Frey B. A constitution for knaves crowds out civic virtues. *Economic Journal* 1997; 107(443): 1043-1053.

Frey B, Goette L. Does pay motivate volunteers? Unpublished Manuscript. Institute for Empirical Research in Economics, University of Zurich.

Frey B, Jegen R. Motivation crowding theory. *Journal of Economic Surveys* 2001; 15(5): 589-611.

Frey B, Oberholzer-Gee F. The cost of price incentives: an empirical analysis of motivation crowding out. *American Economic Review* 1997; 87(4): 746-755.

Giving USA: The Annual Report on Philanthropy for the year 2011 (2012). Chicago: Giving USA Foundation.

Gneezy A, Li C, Saccardo S, Samek A. Changing mindset: Using different appeals increase payments under consumer elective pricing. *Working Paper, 2014.*

Gneezy U, Rustichini A. Pay enough or don't pay at all. *Quarterly Journal of Economics* 2000; 115(3): 791-810.

Gneezy U, Rustichini A. A fine is a price. *Journal of Legal Studies* 2000; 29: 1-18.

Greenfield J. Fundraising Fundamentals: A Guide to Annual Giving for Professionals and Volunteers. Wiley: USA, 2001.

Karlan D, List J. Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *American Economic Review* 2007; 97(5): 1774-1793.

Koppel H, Schulze G. On the channels of pro-social behavior: Evidence from a natural field experiment. Working Paper, 2009.

Lacetera N, Macis M, Slonim R. Will there be blood? Incentives and displacement effects in pro-social behavior. *American Economic Journal of Economic Policy* 2012; 4: 186-223.

Lacetera N, Macis M, Slonim R. Rewarding volunteers: A field experiment. *Management Science* 2014; 5: 1107-1129.

Lepper M, Green D. The Hidden Costs of Reward: New Perspectives on Psychology of Human Motivation. 1978; Hillsdale, NY: Erlbaum.

Lepper M, Greene D, Nisbett R. Undermining children's interest with extrinsic rewards: A test of the overjustification hypothesis. *Journal of Personality and Social Psychology* 1973; 28(1): 129-37.

List J, Lucking-Reiley D. The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign. *Journal of Political Economy* 2002; 110(1): 215-233.

Meier S. Do subsidies increase charitable giving in the long run? Matching donations in a field experiment. *Journal of the European Economic Association* 2007; 5(6):1203-1222.

Mellstrom C, Johannesson M. Crowding out in blood donation: Was Titmuss right? *Journal of the European Economic Association* 2008; 6: 845-863.

Newman G, Shen J. The counterintuitive effects of thank-you gifts on charitable giving. *Journal of Economic* Psychology 2012; 33(5): 973-983.

Newman G, Savary J. When do incentives help and when do they hurt? The effects of relational norms on charitable giving. *Working Paper, 2014.*

Roeger KL, Blackwood AS, Pettijohn SL. The Nonprofit Almanac 2012.Urban Institute Press: USA, 2012.

Shang J, Croson R. The impact of downward social information on contribution decisions. *Experimental Economics* 2008; 11(3): 221-233.

Solow R. Blood and thunder. *Yale Law Journal* 1971; 80: 170-183.

Titmuss, R. The Gift Relationship. 1970; London: Allen and Unwin.

Table 1: Classification of Renewal Call Types

Categorization Name	Months until Renewal	Months since Last Donation
R1	3	9
R2	2	10
R3	1	11
R4	0	12
R5	Overdue (by 1 month)	13
R6	Overdue (by 2 months)	14
R7	Overdue (by 3 months)	15

Figure 1: Buck-slips for Direct-Mail Renewal Testing (Glossy and Colored)

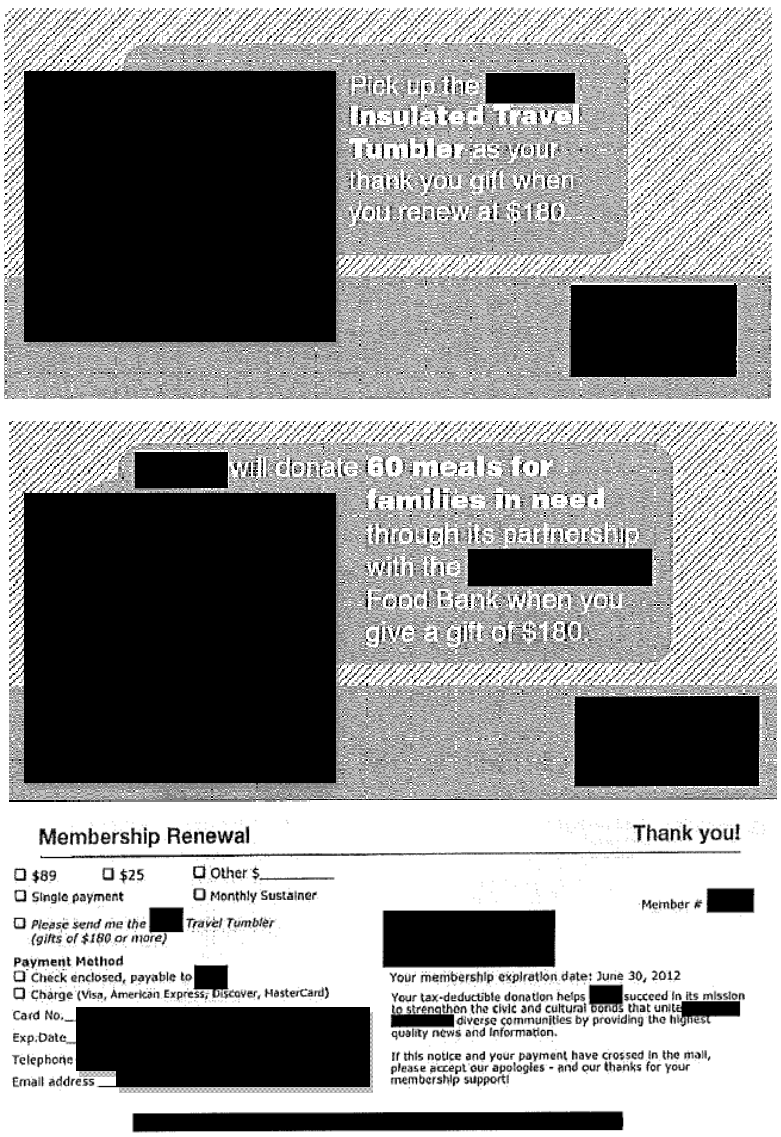

Note: Name and images redacted.

Figure 2A: Solicitations by Month and Condition

	Pieces of Mail		
Month	**Control**	**Swag**	**Meals**
May-14	3,080	1,942	1,778
Jun-14	2,746	1,851	1,972
Jul-14	2,337	1,992	1,938
Total	8,163	5,785	5,688

Figure 2B: Solicitations by Month and Renewal-Status

Renewal Status	May	June	July	Total
R1 – 1st notice	426	616	2,597	**3,641**
R2	847	383	578	**1,808**
R3	1,513	806	344	**2,663**
R4	528	1,302	685	**2,515**
R5	799	454	1,096	**2,349**
R6	2,407	701	351	**3,459**
R7	278	2,307	616	**3,201**
Total	**6,800**	**6,569**	**6,267**	**19,636**

Figure 3: Subject Traits, by Treatment Assignment

Variable	Total Obs^ (out of 10013)	Control	Meals	Swag
Demographics				
Male	8807	47%	47%	47%
Caucasian	8194	74%	72%	72%
Age	7381	58 (15)	56 (15)	54 (16)
Pledge History From 2013				
Months from Last Pledge (from Jul-14)	-	13.0 (2.7)	12.3 (2.8)	12.5 (2.9)
Pledges per Donor in 2013	-	0.83	0.83	0.86
Total Pledged per Donor in 2013	-	$ 65.6 (81.6)	$ 59.1 (73.6)	$ 59.2 (76.2)
Donation Methods Used in 2013				
Mail	-	4.8%	2.9%	2.7%
Web	-	27.4%	34.1%	35.4%
Telemarketing	-	13.0%	10.8%	8.3%

^This column represents the number of successful merges (based on name) with the dataset providing that particular demographic variable. 8,807 of the 10,013 donor names matched the U.S. Census' database on first-names and gender probabilities; 8,194 matched the U.S. Census data on last-names and race probabilities. 7,381 matched the purchased third-party consumer behavior data providing ages (matched donors based on name and address).

Figure 4: Donation Rates (May donors and/or R1 donors)

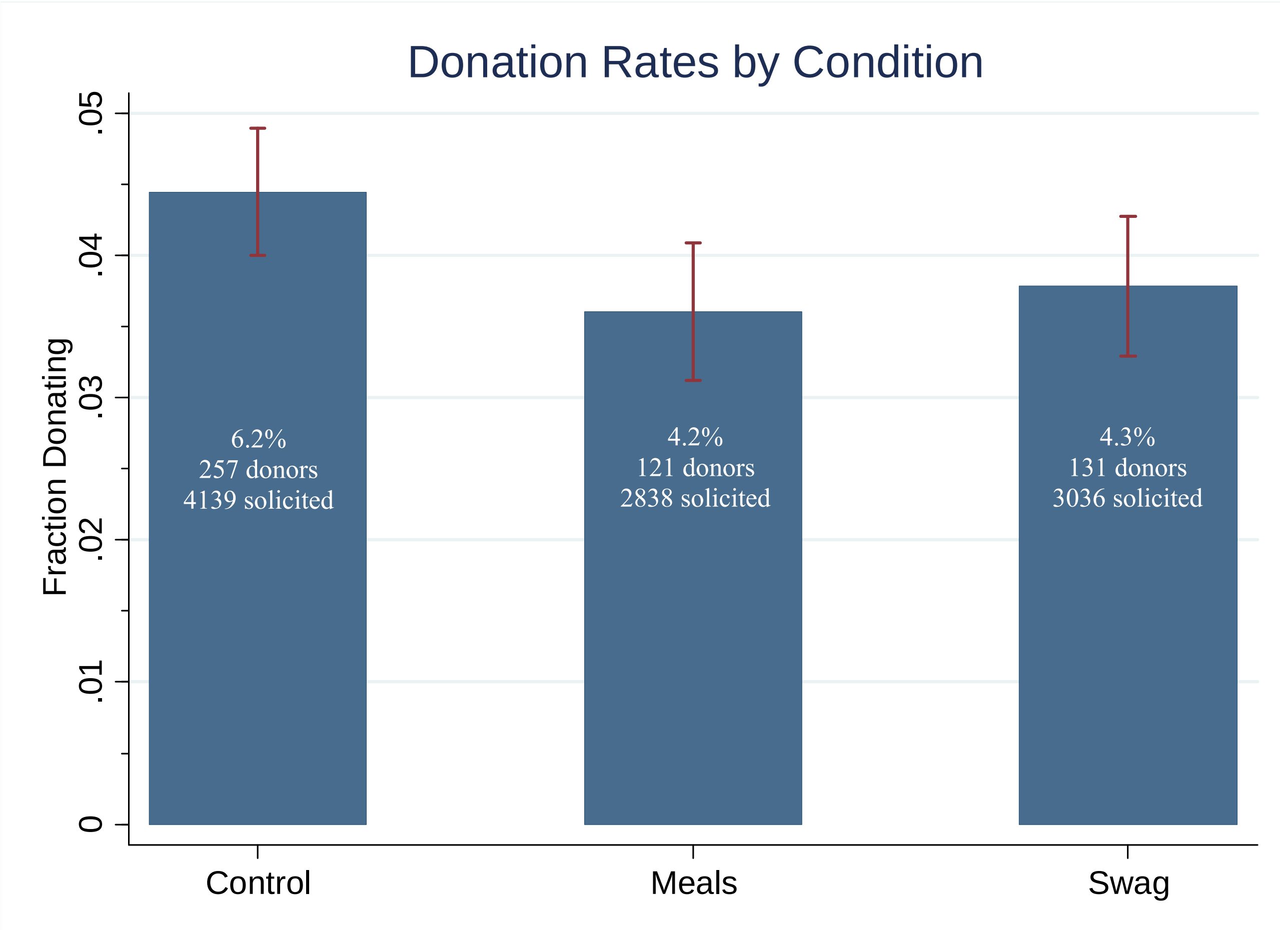

Condition 1	Condition 2	T-statistic	P-value (2-tail)
Control	Swag	3.508	0.001***
Control	Meals	3.529	0.000***
Meals	Swag	-0.097	0.923

Table 2: Likelihood of Donating

	(1)	(2)	(3)	(4)
Specification	OLS	OLS	Logit	Penalized Likelihood Logit
Dep. Var.	Chose to Donate	Chose to Donate	Chose to Donate	Chose to Donate
N	10,013	10,013	10,013	10,013
R^2 / Pseudo-R^2	0.005	0.005	0.005	0.005
Any Gift Offered	-0.017***			
	(0.005)			
Swag Offered		-0.017***	-0.339***	-0.338***
		(0.005)	(0.111)	(0.110)
Meals Offered		-0.017***	-0.349***	-0.347***
		(0.005)	(0.114)	(0.113)
Constant	0.066***	0.066***	-2.641***	-2.639***
	(0.004)	(0.004)	(0.067)	(0.068)
Month FE	YES	YES	YES	YES

All SEs are Huber-White robust except Model (4).
*p<0.10; ** p<0.05; *** p<0.01

Figure 5: Average Conditional Donations

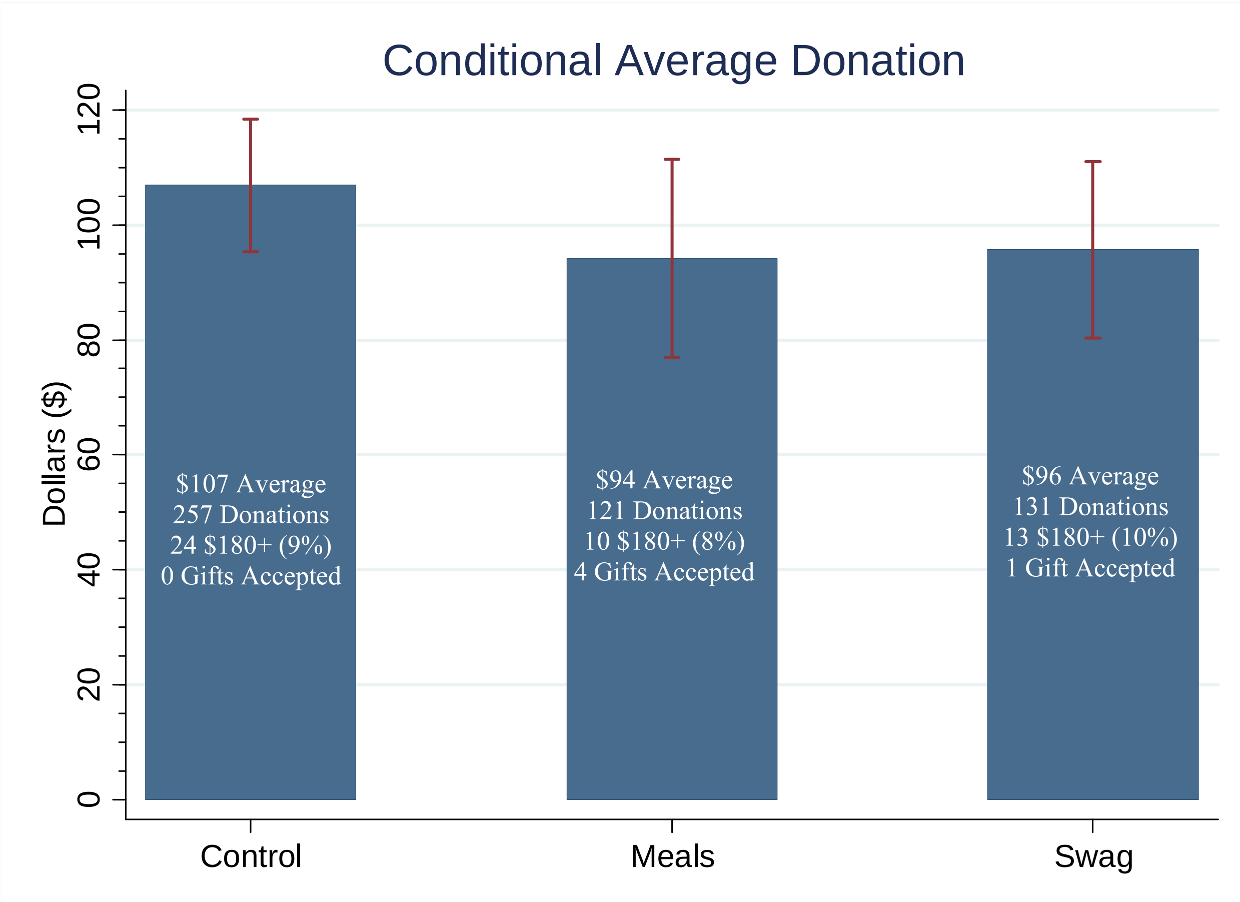

Condition 1	Condition 2	T-statistic	P-value (2-tail)
Control	Swag	1.131	0.259
Control	Meals	1.220	0.223
Meals	Swag	-0.130	0.897

Table 3: Pledge Amounts (Conditional on Donating)

	(5)	(6)
DV	Conditional Pledge Amount	{0,1}: 1 = Pledge ≥ $180
Specification	OLS	OLS
N	509	509
R^2	0.006	0.001
Swag Offered	-10.478	0.029
	(9.713)	(0.038)
Meals Offered	-13.063	-0.008
	(10.357)	(0.036)
Constant	106.902***	0.130***
	(6.540)	(0.023)
Month Fixed Effects	YES	YES

All SEs are Huber-White robust.
*p<0.10; ** p<0.05; *** p<0.01

Figure 6: Donation Rates by Treatment and Renewal-Month Status

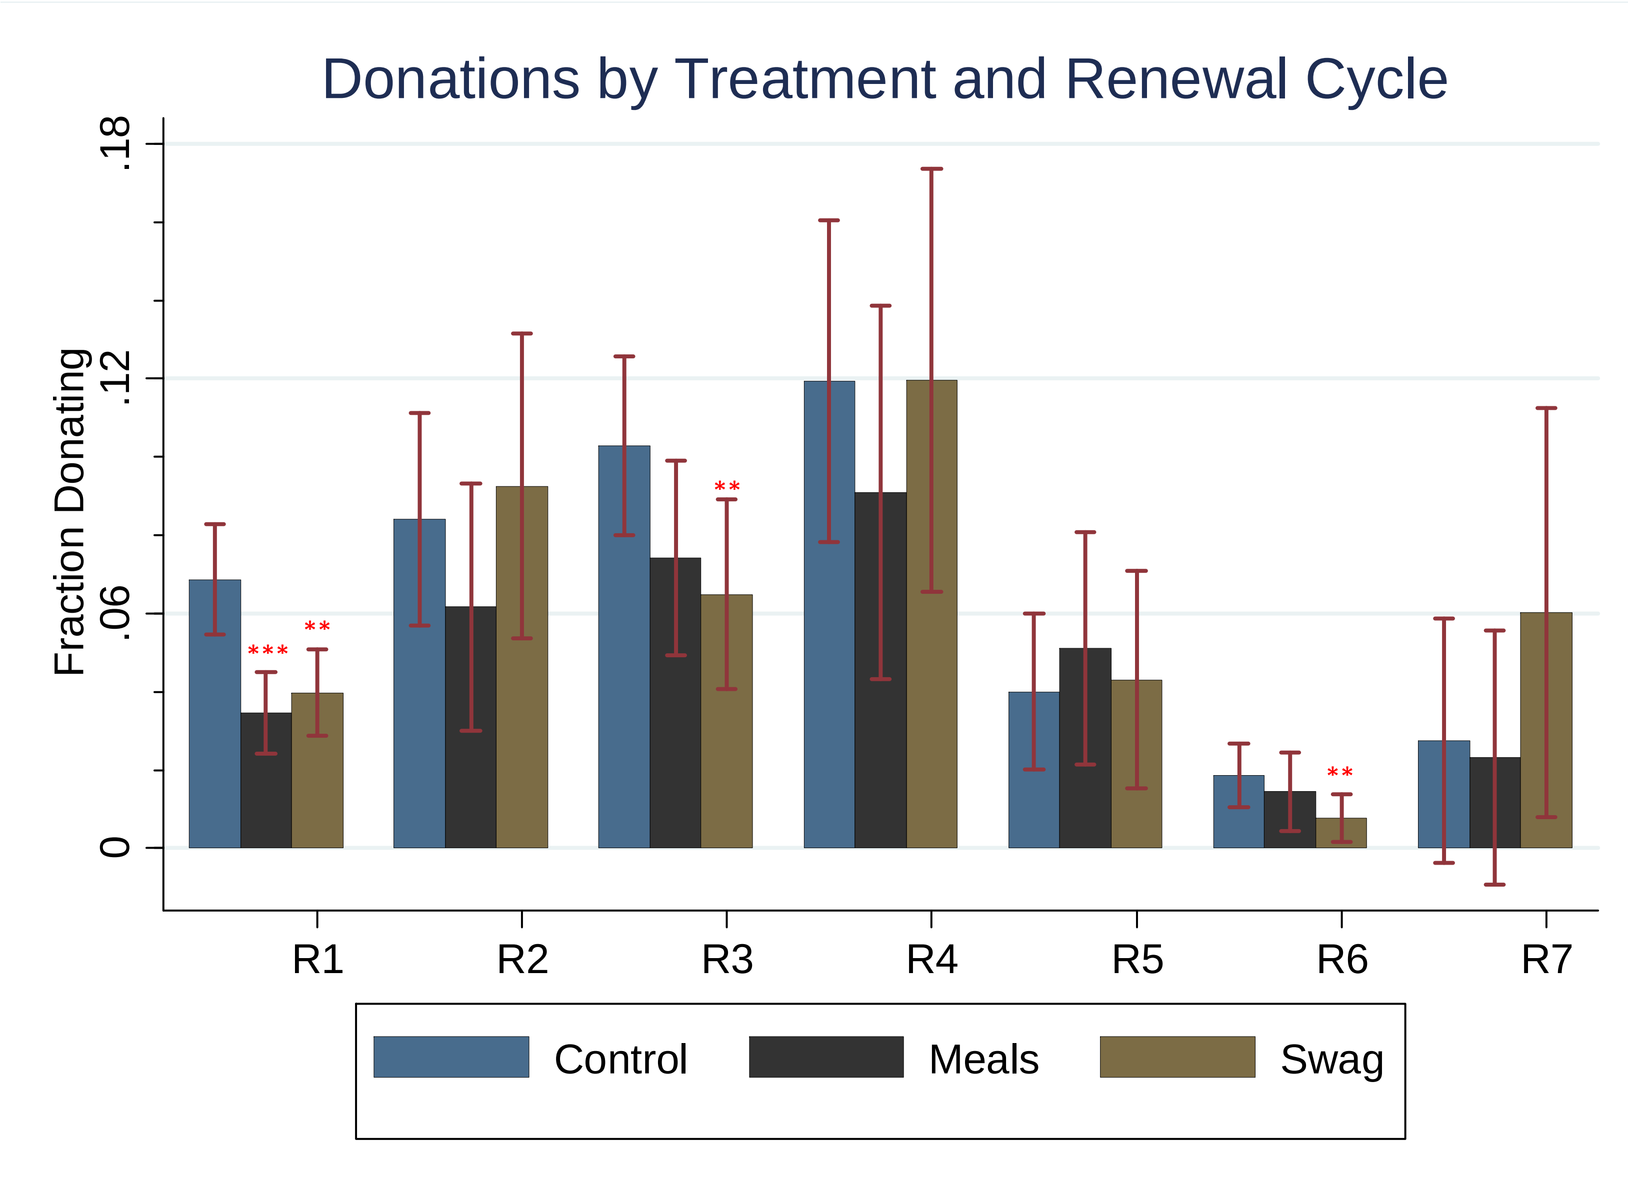

Table 4A: Likelihood of Donating by Renewal-Cycle

OLS	(7)	(8)	(9)	(10)
Renewal Months	All	All	R1	R2-R7
N	19,636	19,636	3,641	15,995
R^2	0.015	0.013	0.023	0.004
Main Effects				
Swag Offered	-0.006*	-0.027***	-0.023**	-0.010
	(0.003)	(0.009)	(0.009)	(0.006)
Meals Offered	-0.008**	-0.032***	-0.029***	-0.013*
	(0.003)	(0.009)	(0.009)	(0.007)
Interactions#				
Swag*R2		0.020		
		(0.017)		
Swag*R3		0.017		
		(0.015)		
Swag*R4		0.032**		
		(0.014)		
Swag*R5		0.032***		
		(0.012)		
Swag*R6		0.015		
		(0.010)		
Swag*R7		0.035***		
		(0.011)		
Meals*R2		0.034**		
		(0.017)		
Meals*R3		0.016		
		(0.015)		
Meals*R4		0.033**		
		(0.013)		
Meals*R5		0.040***		
		(0.012)		
Meals*R6		0.030***		
		(0.011)		
Meals*R7		0.027***		
		(0.010)		
Constant	0.082***	0.097***	0.131***	0.086***
	(0.006)	(0.008)	(0.016)	(0.010)
Renewal-Month Status FE	YES	YES	NO	YES
Month FE	YES	YES	YES	NO^

All SEs are Huber-White robust. *p<0.10; ** p<0.05; *** p<0.01
^Since individuals who declined to donate in May were not included in the June and July data, all R2-R7 donors in the sample are from May. Results are similar if all data, including R2-R7 donors in June and July, are included.
#The interactions represent the incremental effect of each treatment – renewal month pair beyond the main treatment effect. The net effect of treatment for each treatment-renewal month pair is the linear combination each interaction term with the relevant main effect (as seen below in Table 4B).

Table 4B: Joint Significance Tests for Model (10)

	R1	R2	R3	R4	R5	R6	R7
Swag	-0.027***	-0.006	-0.009	0.006	0.006	-0.011**	0.008
	(0.009)	(0.014)	(0.012)	(0.010)	(0.008)	(0.005)	(0.007)
Meals	-0.032***	0.002	-0.016	0.001	0.007	-0.002	-0.005
	(0.009)	(0.015)	(0.012)	(0.010)	(0.009)	(0.006)	(0.005)

Table 5: Likelihood of Donating, by Past Donation Method (OLS)

DV = Chose to Donate	(11)	(12)	(13)
N	10013	8194	7476
R^2	0.027	0.027	0.036
Treatment			
Swag Offered	-0.011**	0.006	-0.005
	(0.005)	(0.011)	(0.018)
Meals Offered	-0.015***	0.010	-0.002
	(0.005)	(0.012)	(0.019)
Relevant Donor-Level Traits			
Gave by Mail in 2013	0.086***		
	(0.026)		
White^		0.057***	
		(0.011)	
Age(Normalized)			0.165***
			(0.026)
Interactions			
Swag * Mail-in-2013	-0.052		
	(0.042)		
Meals * Mail-in-2013	-0.080**		
	(0.037)		
Swag * White		-0.033**	
		(0.015)	
Meals * White		-0.038**	
		(0.016)	
Swag * Age(Normalized)			0.014
			(0.037)
Meals * Age(Normalized)			0.003
			(0.038)
Constant	0.120***	0.086***	-0.011
	(0.016)	(0.019)	(0.020)
Renewal-Month Fixed Effects	YES	YES	YES
Month Fixed Effects	YES	YES	YES
Linear Combinations of Coefficients			
Swag + Swag*Mail	-0.064		
	(0.042)		
Meals + Meals*Mail	-0.095***		
	(0.037)		
Swag + Swag*White		-0.028***	
		(0.008)	
Meals + Meals*White		-0.028***	
		(0.008)	
Swag + Swag* Age(Normalized)			0.009
			(0.020)
Meals + Meals* Age(Normalized)			0.001
			(0.020)

All SEs are Huber-White robust. *p<0.10; ** p<0.05; *** p<0.01
^White is a continuous variable between 0 and 1.
#Age(Normalized) is simply the donor's age, normalized to a scale between 0 and 1.

Table 6: Direct-Mail Renewals with Archival Data

DV = {0,1}: 1 = donated	(14)	(15)	(16)	(17)
Specification	OLS	OLS	OLS	OLS
Population	Accepted Gift Previously	Never Accepted Gift	Accepted Gift Previously	Never Accepted Gift
N	4554	5459	4554	5459
R^2	0.021	0.029	0.032	0.032
Treatment				
Swag Offered	-0.005	-0.020***	-0.024*	-0.030***
	(0.009)	(0.007)	(0.013)	(0.011)
Meals Offered	-0.013	-0.023***	-0.020	-0.027**
	(0.008)	(0.007)	(0.013)	(0.011)
Past Donation Methods				
Gave by Web in the Past			-0.060***	-0.028***
			(0.010)	(0.011)
Interactions				
Swag * Gave by Web in Past			0.037**	0.023*
			(0.017)	(0.014)
Meals * Gave by Web in Past			0.007	0.013
			(0.016)	(0.014)
Constant	0.114***	0.143***	0.114***	0.143***
	(0.020)	(0.027)	(0.020)	(0.027)
Renewal-Month Fixed Effects	YES	YES	YES	YES
Month Fixed Effects	YES	YES	YES	YES
Linear Combinations of Coefficients				
Swag + Swag*Gave by Web			-0.013	-0.007
			(0.011)	(0.009)
Meals + Meals*Gave by Web			-0.013	-0.014
			(0.008)	(0.009)

All SEs are Huber-White robust.
Results are similar when using logit specifications.
*p<0.10; ** p<0.05; *** p<0.01

Figure 7: Mechanical Turk Experiment, Donation Requests

Less Salient and more Prosocial Frame, Swag:

Would you like to donate $1 of your earnings to one of the national charities below? If you do, we will deduct this from your earnings and make the donation on your behalf. We will also send you a confirmation message through mTurk messaging. In addition, as a thank-you for your donation, we will send you a $0.10 Amazon promotional code through mTurk messaging.

Less Salient and more Prosocial Frame, Meals:

Would you like to donate $1 of your earnings to one of the national charities below? If you do, we will deduct this from your earnings and make the donation on your behalf. We will also send you a confirmation message through mTurk messaging. In addition, as a thank-you for your donation, we will make a $0.10 contribution to Feeding America, a national food bank organization.

More Salient and Cost-Benefit Frame, Swag:

Earn an additional $0.10 Amazon promotional code by donating $1 of your earnings to one of the national charities below. If you do, we will deduct this from your earnings and make the donation on your behalf. We will also send you a confirmation message through mTurk messaging that contains the promotional code.

More Salient and Cost-Benefit Frame, Meals:

In addition, you can earn a $0.10 donation for Feeding America, a national food bank organization, if you donate $1 of your earnings to one of the national charities below. If you do, we will deduct this from your earnings and make the donation on your behalf. We will also send you a confirmation message through mTurk messaging.

Figure 8: Mechanical Turk Donations, Telemarketing Language

Treatment	Subjects	Number of Donors	% Donating	Gifts Given
Control	69	11	15.9%	N/A
Swag	69	13	18.8%	10
Meals	70	13	18.6%	9

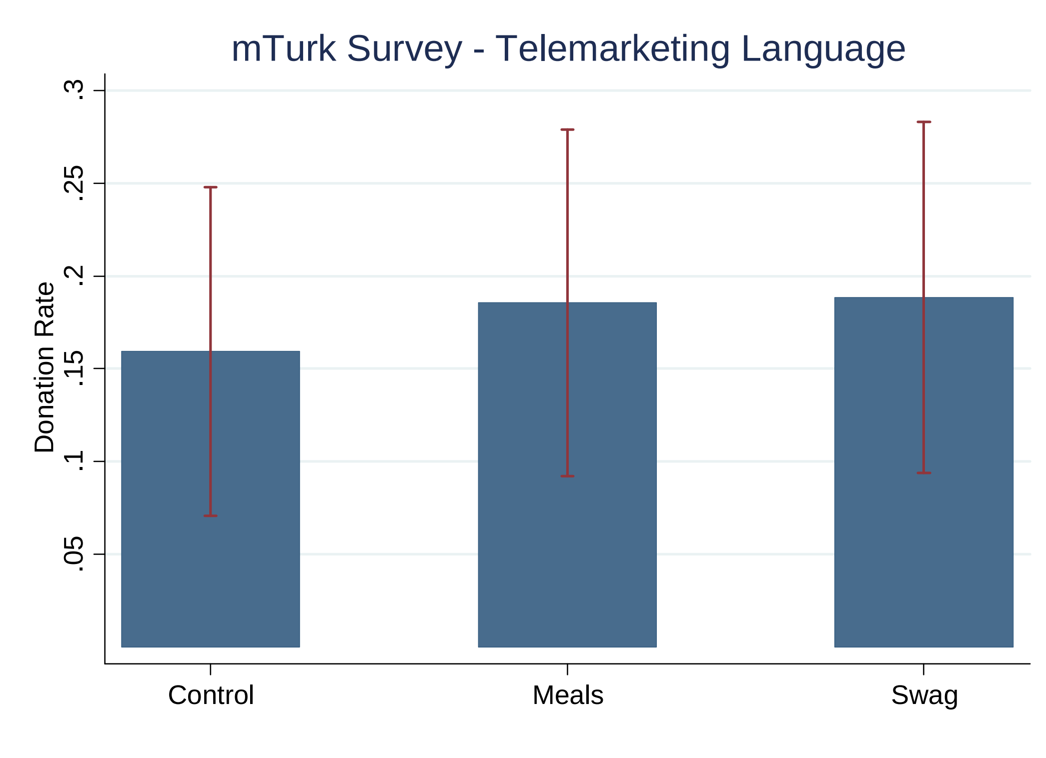

Condition 1	**Condition 2**	**T-statistic**	**P-value (2-tail)**	**P-value (1-tail)**
Control	Swag	-0.446	0.656	0.672
Control	Meals	-0.407	0.684	0.658
Meals	Swag	0.040	0.968	0.516

1-tail test conforms with ex-ante hypotheses that Pr (Condition 2 < Condition 1)

Figure 9: Mechanical Turk Donations, Mailer Language

Treatment	Subjects	Number of Donors	% Donating	Gifts Given
Control	140	30	21.4%	N/A
Swag	122	27	22.1%	17
Meals	138	20	14.5%	18

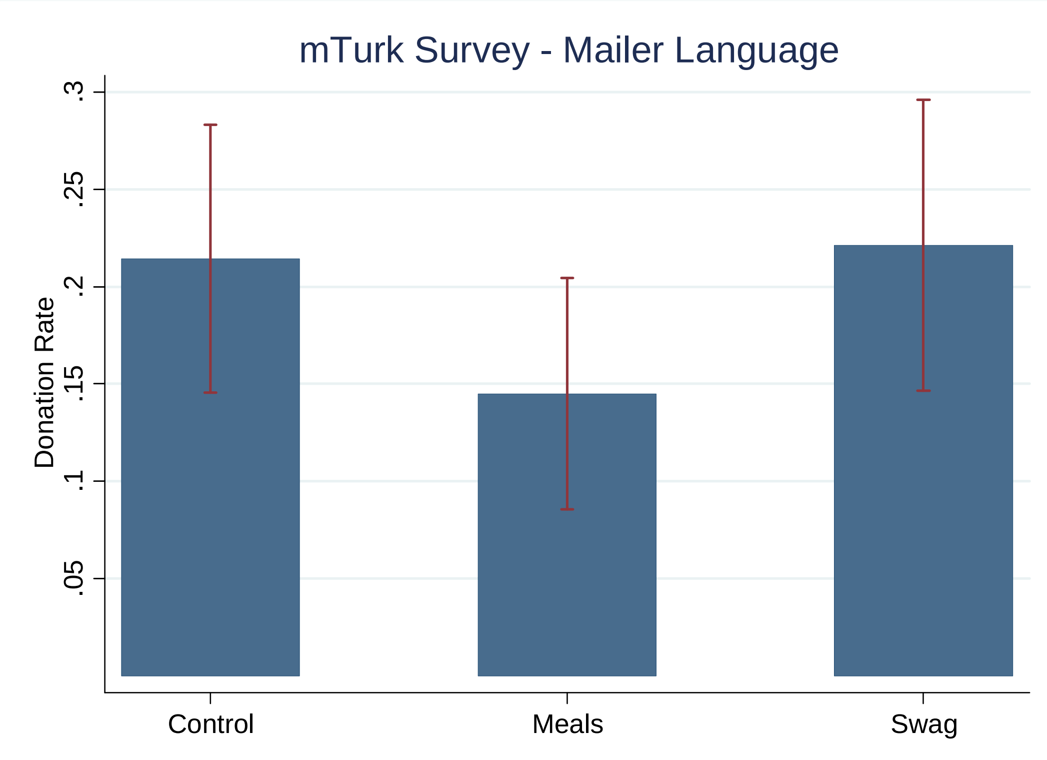

Condition 1	Condition 2	T-statistic	P-value (2-tail)	P-value (1-tail)
Control	Swag	-0.137	0.891	0.554
Control	Meals	1.506	0.133	0.067*
Swag	Meals	-1.599	0.111	0.933

1-tail test conforms with ex-ante hypotheses that Pr (Condition 2 < Condition 1)

5 minutes < Completion Time < 90 minutes only

Condition 1	Condition 2	T-statistic	P-value (2-tail)	P-value (1-tail)
Control	Swag	0.075	0.934	0.470
Control	Meals	1.506	0.106	0.053*
Swag	Meals	-1.599	0.134	0.933

1-tail test conforms with ex-ante hypotheses that Pr (Condition 2 < Condition 1)

A P P E N D I X A

INSTRUCTIONS TO CHAPTER I

[Instructions shown are for Standard sessions. Self-image sessions used identical instructions, but also added the *italicized* text shown at the end of some pages.]

Introduction

Welcome and thank you for participating. At this point, please don't talk to any other participants and please turn off all cell phones and electronics. If you haven't yet, please remove all personal materials from your desk (you may place them underneath your seat).

If at any point you have any questions, please raise your hand and I will come to your seat to answer your question in private. This will avoid delaying the experiment or disrupting others' during a task.

Today's session is a study on individual judgment and decision-making. Over the course of the experiment, you will make choices in a series of tasks, and these choices may impact your earnings as well as the earnings of others. In this particular study, you may also earn various items. If earned, these items will be handed out, along with your monetary earnings, at the end of the experiment by the lab manager.

For today, the items you can earn in this experiment are: a duffel bag, a pen, and/or an additional $2 in cash.

There is no deception or dishonesty in this experiment. It is an anonymous experiment; your name will not be used in any way by the researchers conducting this study.

You have a set of handouts in front of you. Please *do not* flip through these until you are instructed to do so. In addition, please do not mark up your handouts, and do not take them with you when you leave. They will be re-used for future sessions.

General Instructions: Section 1

In Section 1 of this experiment, the computer will randomly assign you to one of two roles: **role 0** or **role 1**. You will retain this designation throughout Section 1.

The computer will randomly pair you with one other participant. If you are a role 0, you will be paired with a role 1. If you are a role 1, you will be paired with a role 0. You will remain anonymous to one another. Each role will have different responsibilities. These responsibilities will vary as we progress through the experiment, although **role 0** will always be designated the primary decision-maker. As the experiment progresses, you will also be re-matched with a new, random partner for each task; that random partner will always be of a different role than you.

Please click "OK" on your screen to learn your role assignment.

Section 1 Continued

In this part of Section 1, for each pair of people, the computer will give either person an item with probability ¾. With probability ¼, neither person will receive an item. If you receive an item, you will see a picture of the item on your screen. If you received an item, it means your partner did not.

If you did not receive an item, the computer *may or may not* announce whether your partner received an item and what that item is. If you received an item, the computer will tell you whether your partner is aware that you received an item.

Following this, the role 0 person will be designated the primary decision-maker in the pair. He or she will be given $10. Role 0 must then choose how much of the $10 to give to their role 1 partner. The amount that **role 0** keeps will be his or her earnings for this part of the experiment (plus any item that they received). The amount that role 0 gives to role 1 will be **role 1's** earnings for this part of the experiment (plus any item that they received).

You will repeat this task for 8 rounds; each round will be with a new, randomly matched person of the proper role type. At the end of the experiment, one of these rounds will be selected for payment. Items and cash earned in the selected round will be paid to you at the end of the experiment.

[Self-Image Sessions Only]

At the end of the entire experiment, if you are a role 0, we will tell you how your individual choices compare to past participants that were role 0 in previous sessions. For instance, you will see the average amount you gave to your partner in rounds where you did not receive an item from the computer, accompanied with a percentile rank that compares what you gave to what past participants gave in the same scenario.

Section 1 Continued

In this part of Section 1, for each pair of people, the computer will give **role 1** people an item with probability ¾. With probability ¼, neither person will receive an item. If you receive an item, you will see a picture of the item on your screen. Your partner will not know whether you received an item.

If an item was given to **Role 1,** he or she may choose to keep this item or instead give it to their role 0 partner by clicking the "Give" button. If this is chosen, the picture of the item will disappear from role 1's screen and will show up on role 0's screen.

Following this, the role 0 person will be designated the primary decision-maker in the pair. He or she will be given $10. Role 0 must then choose how much of the $10 to give to their role 1 partner. The amount that **role 0** keeps will be his or her earnings for this part of the experiment (plus any item that they received). The amount that role 0 gives to role 1 will be **role 1's** earnings for this part of the experiment (plus any item that they kept).

Note to **role 0s**: If a role 0 does not receive a gift, it may be because role 1 did not receive an item from the computer this round. In addition, role 1s receive at most only one item per round, and they can only give the item that they received that round.

You will repeat this task for 8 rounds; each round will be with a new, randomly matched person of the proper role type. At the end of the experiment, one of these rounds will be selected for payment. Items and cash earned in the selected round will be paid to you at the end of the experiment.

[Self-Image Sessions Only]

At the end of the entire experiment, if you are a role 0, we will tell you how your individual choices compare to past participants that were role 0 in previous sessions. For instance, you will see the average amount you gave to your partner in rounds where you did not receive an item from your partner, accompanied with a percentile rank that compares what you gave to what past participants gave in the same scenario. Likewise, we will also report average amounts you gave in response to other actions by your partner, accompanied by a similar percentile rank.

Section 1 Continued

In this part of Section 1, for each pair of people, the computer will give **role 1** people an item with probability ¾. With probability ¼, neither person will receive an item. If you receive an item, you will see a picture of the item on your screen. Your partner will not know whether you received an item.

If an item was given to **Role 1,** he or she may choose to keep this item or instead gift it to their role 0 partner by clicking the "Give" button. However, if "Give" is chosen, the computer will override this choice with probability ½ and force role 0 to keep the item. If the computer overrides the choice, it will be announced to both players that a gift was attempted but overridden by the computer. If the computer instead allows the item to be given, then the picture of the item will disappear from role 0's screen and appear on role 1's screen.

Following this, the role 0 person will be designated the primary decision-maker in the pair. He or she will be given $10. Role 0 must then choose how much of the $10 to give to their role 1 partner. The amount that **role 0** keeps will be his or her earnings for this part of the experiment (plus any item that they received). The amount that role 0 gives to role 1 will be **role 1's** earnings for this part of the experiment (plus any item that they kept).

Note to **role 0s**: If a role 0 does not receive a gift, it may be because role 1 did not receive an item from the computer this round. In addition, role 1s receive at most only one item per round, and they can only give the item that they received that round.

You will repeat this task for 8 rounds; each round will be with a new, randomly matched person of the proper role type. At the end of the experiment, one of these rounds will be selected for payment. Items and cash earned in the selected round will be paid to you at the end of the experiment.

[Self-Image Sessions Only]

At the end of the entire experiment, if you are a role 0, we will tell you how your individual choices compare to past participants that were role 0 in previous sessions. For instance, you will see the average amount you gave to your partner in rounds where you did not receive an item from your partner, accompanied with a percentile rank that compares what you gave to what past participants gave in the same scenario. Likewise, we will also report average amounts you gave in response to other actions by your partner, accompanied by a similar percentile rank.

Section 2: General Instructions

In Section 2, you will answer a series of questions about yourself. Most of these will be multiple choice questions, but several will also require open text responses. Please respond carefully and truthfully. When you finish the questionnaire, the computer will randomly select three rounds from Section 1 to determine your pay for the experiment. You must finish the questionnaire to be paid in full.

Note: When typing your responses to the open-ended text questions, please make sure to hit the "ENTER" key when you are done typing. Otherwise, the computer will not record your response. Upon hitting "ENTER," the text that you entered should disappear; this confirms that your response was recorded by the computer.

[Self-Image Sessions Only]
At the end of the experiment, all role 0's will also see three successive screens that report their average allocations to their partner in different scenarios, as well as how their individual choices compare with role 0 choices from past sessions.

When you are finished with the questionnaire, please wait patiently at your seat. When all subjects are finished, I will dismiss you. To get paid, please bring the slip of paper that lists your computer number; it is included in your handouts. Your lab manager will pay you based upon the ID number of your computer.