Controls and the Asymmetric Stickiness of Norms

Scott Emett
Cornell University
sae67@cornell.edu

Ronald N. Guymon
Georgia State University
rguymon@gsu.edu

William B. Tayler
Brigham Young University
tayler@byu.edu

Donald Young
Georgia Institute of Technology
Donald.Young@scheller.gatech.edu

April 18, 2014

We appreciate helpful comments and suggestions from anonymous reviewers for the Management Accounting Section Midyear Meeting and the American Accounting Association Annual Meeting, and from Jeremy Bentley, Alisa Brinks, Willie Choi, Karim Jamal, Kathryn Kadous, Michael Maier, Erin Marshall, Greg Richins, Greg Rudin, Florin Sabac, Steve Smith, Bryan Stikeleather, and David Wood, as well as workshop participants at the Annual Meeting of the American Accounting Association, Brigham Young University's Accounting Research Symposium, Brigham Young University's Faculty Ethics Conference, Emory University's Experiment-Based-Research Brownbag, the Global Management Accounting Research Symposium, the Management Accounting Section Midyear Meeting, the University of Notre Dame's Excellence in Ethics Research Conference, and the University of Alberta. We are also grateful for programming assistance from Lisa Harris. This research benefited from a Dean's Research Grant from the Goizueta Business School at Emory University, and received the Excellence in Ethics Best Paper Award, presented by the Excellence in Ethics Research Conference at the University of Notre Dame.

Controls and the
Asymmetric Stickiness of Norms

ABSTRACT

This study highlights how formal controls and the behavior of peers can shape individual behavior in accounting contexts. We exogenously manipulate formal controls and visible peer behavior in a laboratory experiment, allowing us to precisely investigate the interaction of controls and social norms. In a setting where appropriate behavior is somewhat ambiguous, we demonstrate that individuals conform more to peer behavior when that behavior conflicts with individuals' own beliefs about appropriate behavior, which are shaped by formal controls. We further find that individuals preferentially attend and conform to the self-interested actions of peers when given mixed signals about peer behavior, causing these self-interested norms to be "stickier" than socially-interested norms for behavior. Our results suggest that the interaction of formal controls and normative influence can lead to a gradual movement toward noncompliance with management expectations or regulatory requirements in accounting contexts.

Key Words: Controls, Conformity, Norms, Social Norms, Ethics

JEL Classification: M40, M41, M49, C91

Data Availability: Contact the authors.

I. INTRODUCTION

Since the early "vision test" experiments of Solomon Asch (1951), studies on conformity to social norms have been a mainstay of psychology research. However, accounting researchers have only recently begun to investigate the importance of norms in accounting contexts. Within the last decade, a growing body of research has added to our understanding of the impact of norms on decision making in accounting settings. Among other findings, accounting studies have demonstrated that tax compliance increases when taxpayers are informed regarding high rates of compliance among their peers (Wenzel 2005a, 2005b), that firms located in geographic areas with high religious adherence have high accruals quality and low tax aggressiveness (Dyreng et al. 2012; McGuire et al. 2012), and that accounting control systems can influence conformity to the behavior of others (Tayler and Bloomfield 2011). Our study investigates how accounting control systems, peer behavior within control systems, and the interaction of these two constructs influence individual behavior. We also provide evidence on whether norms for self and social interest asymmetrically influence individual behavior in accounting settings.

Prior research has shown that environmental cues can influence individuals' beliefs about what behavior is appropriate in a given setting (Messick 1999; Tenbrunsel and Messick 1999; Gneezy and Rustichini 2000; Davidson and Stevens 2013). We examine how two critical environmental cues—control systems and peer behavior—interact to affect which types of peer behavior individuals attend to and comply with. While externally-imposed control systems are intended to motivate a desired behavior (Christ et al. 2014), they also influence people's beliefs about appropriate behavior in a given setting (e.g., Tenbrunsel and Messick 1999; Gneezy and Rustichini 2000). Individuals also look to their peers to infer what is appropriate behavior, as evidenced by prior work that shows that individuals are strongly influenced by the behavior of others (*descriptive norms*) (Asch 1951; Cialdini et al. 1990, 1991; Aronson 2002).

We examine how individuals respond to competing environmental cues from control systems and different types of peer behavior. Specifically, we first add empirical support for a prior finding that individuals who are initially subject to weak controls behave in a more socially-interested way and are more likely to conform to the behavior of others than individuals who are initially subject to strong controls (Tayler and Bloomfield 2011). Second, we investigate whether the influence of peer behavior depends on expectations determined by the initial strength of controls. Prior research on expectancy violation suggests that individuals will perceive peer behavior as more salient when the behavior conflicts with their prior beliefs. Thus, in settings with sufficient ambiguity about appropriate behavior, this increased salience will lead individuals to conform more to peer behavior that *conflicts* with their expectations for appropriate behavior. Third, we investigate how mixed signals regarding the behavior of others impacts social conformity. Research on motivated reasoning suggests that, when possible, individuals will seek out and conform to behavior that is consistent with their preferences more than to behavior that is inconsistent with their preferences (Kunda 1990). Accordingly, we expect an asymmetric response to peer behavior such that when individuals encounter both self-interested and socially-interested peer behavior, they will conform more to the self-interested behavior of others. In other words, while pressures from social norms can lead individuals to act in socially-interested ways, we predict that norms for self-interested behavior will be "stickier" than norms for socially-interested behaviors, resulting in a gradual movement toward noncompliance with management expectations or regulatory requirements in many accounting contexts.

To test our predictions, we run two experiments in which participants take part in a modified public-goods game. Participants are asked to contribute $5 to a public good fund in

each of 14 rounds. Participants can contribute any amount to the fund, in $0.05 increments, between $0 and $5. Participants face a known audit probability of either 90 percent or 0 percent (i.e. controls are either strong or weak) during the first two rounds to establish views about appropriate behavior in the setting. In all subsequent rounds, however, participants face a known audit probability of 45 percent. During each of these subsequent rounds, participants can maximize their payoffs by contributing $0 to the fund, regardless of their peers' contributions. However, we allow participants to see the contributions of three other peer participants. We manipulate which contributions from peers are displayed, showing either three high contributions in each round (socially-interested behavior) or three low contributions in each round (self-interested behavior). Additionally, after several rounds, we further manipulate, within subjects, the contributions of peers by reassigning peer groups and changing which peer contributions we display. In our first experiment, those who saw low contributions in early rounds see high contributions in late rounds, and vice versa. In our second experiment, we hold constant that all participants see high contributions in early rounds, and manipulate peer contributions (high vs. low) in late rounds.

Our primary dependent measure of interest is participants' conformity to the behavior of others. The word conformity can be used in multiple ways. The Merriam-Webster dictionary defines "conformity" as 1) "behavior that is the same as the behavior of most other people…in a group," and 2) "action in accordance with some specified standard" (Merriam-Webster 2005). For simplicity, our use of the word "conformity" in this manuscript refers to the first definition: the tendency to behave in a way similar to others. We use the word "compliance" to refer to the tendency to behave in accordance with a specified standard. Thus, an individual's contributions in our study (i.e., the individual's level of *compliance* with the request that they contribute to the

public good) is used, in conjunction with the manipulated compliance of a participant's peers, to calculate the individual's level of *conformity* in the study. In other words, conformity is a function of participants' contributions to the public good each round and the exogenously manipulated peer contributions.

Results provide support for our hypotheses. First, when initial controls are strong, individuals behave in a more self-interested manner and conform to peer behavior less than individuals in a setting where initial controls are weak. Second, we find that individuals conform more to peer behavior when that behavior conflicts with the expected behavior given the initial strength of controls. Specifically, when initial controls are strong, socially-interested norms are more compelling than self-interested norms; and when initial controls are weak, self-interested norms are more compelling than socially-interested norms. Finally, we find that overall social cooperation changes more when peer behavior shifts from being socially-interested to being self-interested than when peer behavior shifts from being self-interested to being socially-interested. This finding suggests that individuals preferentially attend to self-interested peer behavior when given mixed signals about peer behavior, thus causing norms for self-interest to be "stickier" than norms for social interest.

Our study contributes to accounting and psychology research in several ways. Prior work in accounting has relied on endogenously formed group behavior to examine conformity, whereas we manipulate these behaviors exogenously. By manipulating the descriptive norm, we are able to draw strong causal explanations, as opposed to correlational connections, regarding conformity in accounting contexts. Second, our results relating to asymmetry in the conformity to behavior of others have important implications for accounting models that examine the influence of norms (Davis et al. 2003; Fischer and Huddart 2008). Equilibrium outcomes in these

models are partly dependent on the response to both "good" and "bad" norms, which are assumed to be symmetric. Asymmetric stickiness of self-interested and socially-interested norms is likely to influence equilibrium predictions in these models. Third, our study provides additional evidence on how individuals use control systems as a social cue when inferring appropriate behavior in different settings, and on how control systems interact with other social cues (e.g., peer behavior) when influencing individual behavior.

Our research also has important implications for accounting practitioners. In addition to providing theory and evidence that accounting controls affect beliefs about appropriate behavior in unanticipated ways, our study provides insights into how individuals respond to mixed accounting signals in an open information environment. If non-compliant peers have a greater impact on individuals than compliant peers over time (as we find), regulators and organizations must continually manage a tendency of firms and employees, on average, to gravitate towards non-compliance. These efforts may include selective disclosure of peer behavior (Tafkov 2012; Hannan et al. 2012) or adjustments to accounting control systems. Importantly, though set in a management accounting context, this study has testable implications across a number of accounting domains where peer behavior is observable and formal controls vary in strength: on the factory floor, under different tax administrations, or in financial markets.

Our manuscript proceeds as follows. Section II reviews related literature and presents our predictions. Section III describes the experiment. Section IV presents statistical analyses and results. Section V concludes.

II. THEORY AND HYPOTHESES

Individuals tend to conform to the actions of those around them (Cialdini et al. 1990). Research in social psychology demonstrates that *descriptive norms*, which represent the behavior

of others, can have a powerful influence on individual behavior (Sherif and Murphy 1936; Asch 1951; Cialdini et al. 1990, 1991; Aronson et al. 2002). In perhaps the most famous demonstration of descriptive norm conformity, Asch (1951) shows that participants in his experiment agree with the wrong answers of other (confederate) participants regarding a simple comparison of the length of several lines, even though participants exhibit considerable discomfort in doing so.

Descriptive norms influence compliance and conformity in a variety of accounting contexts. Tax research suggests that people conform in their taxpaying behavior to perceived compliance norms (e.g., Blanthorne and Kaplan 2008; Bobek et al. 2007; Davis et al. 2003; Wallschutzky 1984; Wenzel 2004, 2005a, 2005b). That is, people who believe their peers evade taxes have more accepting attitudes of tax avoidance and are more likely to cheat on their own taxes. In labor settings, descriptive norms of effort can develop within firms to establish an expectation among workers about appropriate levels of effort (Kandel and Lazear 1992); descriptive norms for reciprocity can generate noncompetitive wages and effort in a competitive market (Fehr et al. 1993, 1998); and descriptive norms of productivity influence the output of laborers working in the same team (Mas and Moretti 2009). Recent financial reporting studies find that religious descriptive norms are associated with fewer financial reporting irregularities (McGuire et al. 2012), less aggressive financial reporting choices (Dyreng et al. 2012), and more conservative going-concern opinions by auditors (Omer et al. 2013).

In recent years, accounting theorists have developed analytical models demonstrating how descriptive norms in accounting contexts likely influence behavior.[1] For example, Fischer

[1] As discussed by Cialdini et al. (1990), the term *norm* has multiple meanings, and academic disciplines differ in their definition and use of the term. For example, sociologists have traditionally construed norms narrowly as social prescriptions or "rules" that govern large societal groups (e.g., see Broom and Selznick 1968). Social psychologists embrace a broader definition of norms, recognizing that norms can form in small groups (e.g., work settings), manifest themselves as internalized beliefs (personal norms), and represent both prescribed behavior (explicit, injunctive norms) as well as typical behavior (descriptive norms) (e.g., see Cialdini et al. 1990, 1991). Our

and Huddart (2008) investigate the impact of norms in a principal-agent setting. They show, analytically, that a preference for conformity to norms can magnify the impact of incentives on actions desired by the principal (like hard work to increase output), while mitigating the unintended impact of incentives on actions the principal wants the agent to avoid (like ignoring quality to increase output). In another study, Davis et al. (2003) show that a preference for conformity to norms leads tax compliance under imperfect enforcement regimes to be history-dependent: with compliant populations, enforcement strength can be dropped substantially with no change in actual compliance, whereas with noncompliant populations, enforcement strength must be increased dramatically to raise compliance levels.

In summary, prior research suggests that when descriptive norms are salient to accounting decision makers, these norms influence compliance behavior because individuals tend to conform to the norm.

Conformity and Beliefs about Appropriate Behavior

Tayler and Bloomfield (2011) examine how controls affect beliefs about appropriate behavior, and how these beliefs influence conformity to descriptive norms. Consistent with prior work in psychology (Messick 1999; Tenbrussel and Messick 1999; Gneezy and Rustichini 2000), they find that when controls are initially very strong in a given setting, individuals tend to believe self-interested behavior is appropriate; whereas when initial controls are very weak, individuals tend to believe socially-interested behavior is appropriate. Further, they find that when individuals believe socially-interested behavior is appropriate, they attend more closely to the behavior of those with whom they interact (Bicchieri and Xiao 2009). Consistent with theory suggesting that adherence to norms increases as those norms become more salient (Cialdini et al.

---

discussion of norms in this manuscript is focused only on this latter manifestation of normative influence—descriptive norms.

1990, 1991), they find that this increased focus on the behavior of others leads individuals to conform more to the behavior of others. However, due to the nature of the various questions examined by Tayler and Bloomfield (2011), their study relied on endogenously determined descriptive norms. Though this design choice allows the authors to draw general inferences based on correlations, it is difficult to arrive at causal conclusions regarding the effect of norms on compliance behavior because of the possibility of endogenous selection or sorting.

In our study, we first examine whether initial control strength has causal implications for conformity to exogenously manipulated descriptive norms. Specifically, consistent with Tayler and Bloomfield (2011), we predict that when individuals are initially subject to strong controls, they will be less likely to conform to descriptive norms than if they are initially subject to weak controls. We state this hypothesis formally as follows:

> **H1:** Conformity with the descriptive norm will be higher in settings where initial controls were weak than when initial controls were strong.

Asymmetric Stickiness of Descriptive Norms

Prior work has demonstrated the importance of understanding the impact of norms on decision making in accounting settings. However, these studies either assume symmetry in the influence of "good" and "bad" norms, or they ignore the potential for asymmetric influence of norms on behavior. Our study focuses on norms in an accounting control setting and explicitly investigates the potential for asymmetry in the stickiness of norms for self-interested and socially-interested behavior.

Expectancy Violation and Norm Stickiness

While we predict greater conformity to descriptive norms when initial controls are weak than when they are strong, we do not expect that individuals who face strong initial controls will be immune to the influence of descriptive norms (nor do we predict this non-effect). However,

prior research does suggest that the pull of self- and socially-interested descriptive norms will depend, in part, on individuals' beliefs about appropriate behavior in a given setting.

Research on expectancy violation suggests that outcomes or actions that violate an individual's expectations are more salient to that individual than those that align with the individual's expectations (see Burgoon and Burgoon 2001; Clor-Proell 2009). As applied to descriptive norms, this implies that descriptive norms that conflict with an individual's beliefs about appropriate behavior will be more salient to the individual than descriptive norms that align with the individual's beliefs about appropriate behavior. For individuals who are uncertain about how to behave in a given setting, this increased salience will lead to increased adherence to the descriptive norm.

If the initial strength of controls influences beliefs about appropriate behavior in a given setting (our first hypothesis), expectancy violation theory suggests two complementary predictions. First, when individuals are subject to weak initial controls, they will likely conform more to self-interested descriptive norms than to socially-interested descriptive norms. Second, when individuals are subject to strong initial controls, they will likely conform more to self-interested descriptive norms than to socially-interested descriptive norms. We outline these predictions below by first stating the overall prediction and then stating the two complementary predictions that compose the overall prediction:

> **H2:** Individuals will conform more to a descriptive norm if the descriptive norm conflicts with the individual's beliefs about appropriate behavior in the setting than if the descriptive norm aligns with the individual's beliefs about appropriate behavior in the setting.
>
> **H2a**: Individuals who are initially subject to weak controls will conform more to the descriptive norm if the norm is self-interested than if it is socially-interested.

H2b: Individuals who are initially subject to strong controls will conform more to the descriptive norm if the norm is socially-interested than if it is self-interested.[2]

Motivated Reasoning and Norm Stickiness

In many settings, individuals have either a pecuniary or psychological incentive to act selfishly. Research in psychology on motivated reasoning suggests that individuals acquire, process, and evaluate data in ways consistent with their preferences (Kunda 1990; Gilovich 1991; Ditto and Lopez 1992; Ditto et al. 1998; Dawson et al. 2002; Ditto et al. 2003).[3] Gilovich (1991) proposes that the manner in which those succumbing to motivated reasoning approach evidence depends on how the evidence aligns with their preferences. With disagreeable evidence, people tend to ask "Must I believe this?" and attempt to disconfirm the validity of the bad news. This high standard for acceptance ("*Must* I?") increases the likelihood that an individual will reject the bad news. However, with more agreeable evidence, people ask "Can I believe this?" and pursue confirmatory information in an attempt to validate the good news. This low standard for acceptance ("*Can* I?") increases the likelihood that an individual will accept the good news. However, regardless of the valence of the news, rejection of preference-inconsistent information and acceptance of preference-consistent information is constrained by the reasonableness of the data (Kunda 1990). Thus, motivated reasoning does not imply that people blindly reject bad news and blindly accept good news, but rather, that people maintain an "illusion of objectivity"

[2] Note that support for H2a and H2b could occur mechanically due to ceiling or floor effects. For example, if individuals believe that socially-interested behavior is appropriate in a given setting and tend to comply at such high rates that they cannot move in the direction of a socially-interested descriptive norm, this could artificially appear to support H2a. Or, if individuals believe that self-interested behavior is appropriate tend to comply at such low rates that they cannot move in the direction of a self-interested descriptive norm, this could artificially appear to support H2b. As seen in Section IV, we do not observe such ceiling/floor effects in our study, allowing for clean tests of H2a and H2b.

[3] Evidence of motivated reasoning has been demonstrated in various accounting settings. Hales (2007), for example, shows that even with an incentive for accurate earnings forecasts, investors forecast earnings in a manner consistent with their directional preferences. Similarly, Thayer (2011) demonstrates that investors forgo a certain amount of information credibility in order to seek out preference-consistent information. In a management accounting context, Tayler (2010) shows that managers are swayed by preference-consistent information more than by preference-inconsistent information in strategy evaluation judgments.

(Pyszczynski and Greenberg 1987, 302) by "attempt[ing] to be rational and to construct a justification of their desired conclusion that would persuade a dispassionate observer" (Kunda 1990, 482-482).

In many settings, employees observe a mix of peer behavior, with some peers complying with management expectations and other peers ignoring management expectations in pursuit of personal gain. Motivated reasoning suggests that individuals with a pecuniary incentive to act selfishly will attend to information suggesting the presence of a self-interested descriptive norm (the self-interested behavior of peers) more than to information suggesting the presence of a socially-interested descriptive norm. When the descriptive norm is unambiguously socially-interested, even an individual succumbing to motivated reasoning will likely accept the preference-inconsistent information (Pyszczynski and Greenberg 1987). However, when the descriptive norm is mixed, with some individuals acting in the social interest and others acting in their own self-interest, we expect that motivational factors will cause individuals to attend to self-interested aspects of the descriptive norm more than to socially-interested aspects. The more an individual attends to a descriptive norm, the more susceptible the individual should be to that norm. Thus, we make the following prediction:

> **H3**: When a descriptive norm includes both self-interested and socially-interested behavior, individuals will conform more to the self-interested behavior of others than to the socially-interested behavior of others.

III. METHODOLOGY

To test our hypotheses, we conducted two laboratory experiments. One hundred seventy-four students from various fields of study took part in our first experiment. Similar to Bloomfield et al. (2009), given the abstract and relatively simple nature of our experimental task (described below), we did not limit participation in our study to accounting students. Fifty-three percent of

the participants were male; and participants' average age was 21. We administered the experiment using the experimental economics z-Tree computer program (Fischbacher 2007) in 11 separate sessions. Each session consisted of 15 or 16 participants and lasted approximately 75 minutes.

Design

The experiment employs a 2 x 2 x 14 design (initial-control strength x descriptive norm x 14 rounds). We manipulate initial-control strength between subjects, and manipulate the descriptive norm both between and within subjects. As described in detail below, the experiment consists of four phases. We manipulate initial-control strength during Phase 2 and the descriptive norm during Phases 3 and 4.

We purposefully employ an experimental task that abstracts from "real-world" accounting contexts, because contextual cues can exert significant influence on individuals' beliefs regarding appropriate behavior in a task (e.g., see Liberman et al. 2004 and Choi et al. 2012, 2013). By abstracting our experimental task from familiar settings, we isolate the effects of controls and peer behavior and enhance our ability to examine the influence of these factors on individual behavior. Though we avoid "mundane realism" (Berkowitz and Donnerstein 1982; Swieringa and Weick 1982), our setting (described below) is intended to map into real-world accounting settings where imperfect accounting controls influence the behavior of individuals who can see the reported behavior of their peers, and where individuals make decisions in a social dilemma in which the interest of others are at odds with each individual's self-interest (Kerr 1983). Thus, our setting speaks to numerous tax, financial reporting, and managerial

accounting contexts in which controls influence behavior, and where decision makers must weigh the well-being of others against their own incentives.[4]

Task

Participants complete fourteen rounds of a modified public-goods game in an experimental laboratory. In each of the 14 rounds in the study, we endow participants with 10 laboratory dollars.[5] Participants are asked to contribute $5 to a public good fund each round, but are free to choose an amount to contribute between $0 and $5. However, participants face a known probability (*p)* that their contributions will be audited and penalized if they contribute below the $5 request. If incurred, the penalty is two times the difference between $5 and the participant's contribution in that round. We manipulate *p* between conditions during the first two rounds of the experiment as described below, and hold *p* constant at 45 percent in all subsequent rounds for all participants.

Participants are informed that their total contributions to the fund will be doubled and given to an anonymous "matched participant." Participants are also informed that they will receive double the total contributions their matched participant gives to a separate fund.[6] Participants remain matched with their same matched participant throughout a given phase of the experiment. We inform participants that approximately one week after the experiment session,

[4] This would include individual tax compliance settings such as those modeled by Davis et al. (2003), financial reporting environments such as those investigated by Graham et al. (2005), and relative performance evaluation settings such as those examined by Tafkov (2012) and Hannan et al. (2012).
[5] Hereafter, we will abbreviate laboratory dollars simply as $.
[6] This task shares features of the first step in a gift exchange game (Akerlof 1982; Hannan 2005; Kuang and Moser 2009) as well as a trust game (Berg et al. 1995; Charness et al. 2004). However, unlike these games, contribution decisions are made simultaneously (rather than sequentially), and contribution enforcement is determined exogenously. The game also shares features of a typical public goods game, but is "impure" (Davis & Holt 1993) because participants are excluded from consuming the good they contribute to, and there is only one contributor to each "public good" (see Ledyard (1995) for a review of studies utilizing public-goods games).

laboratory dollars will be converted to US dollars using a predetermined rate, and that expected payoffs are US $15, with a minimum payoff of US $5.[7]

In our experiment, participants can control the portion of their payoff that is determined by their own contributions, but cannot control the portion of their payoff that is determined by the matched participant. Accordingly, participants can maximize their payoff by maximizing the following equation which represents the controllable portion of their payoff:

$10 – Contribution – 2*p* ($5 – Contribution),

where *p* is the audit probability and "Contribution" is the amount (between $0 and $5) that the participant contributes to the fund. When *p* is less than 0.50 (as is the case for all participants in rounds 3-14), participants maximize their payoff by contributing $0 to the common fund. When *p* is greater than 0.50, participants maximize their payoff by contributing the full $5 to the common fund.

We manipulate the *initial* strength of controls participants face in the study. When initial controls are strong, participants are subject to an audit probability of 90 percent during Rounds 1 and 2. When initial controls are weak, participants are informed that contributions in Rounds 1 and 2 will not be audited.[8] In Rounds 3-14, all participants face the same likelihood of an audit—45 percent. Whether or not an audit occurs in Rounds 1-14 is randomly predetermined based on the stated probability of an audit for that round.[9]

[7] We use a predetermined rate as opposed to a relative performance rate to discourage participants from trying to minimize other participants' payoffs in an effort to boost their own payoff. Participants are explicitly told that "You will not learn the exact exchange rate at this time; however, the rate has been predetermined and will not change; thus it is possible for all participants to make more than the expected $15 payment."

[8] We define a control as any mechanism imposed by management to encourage (or discourage) a specific behavior (Christ et al. 2014). Thus, we are careful to label the setting where the audit probability is 0 as one in which controls are weak, though not entirely "absent." That is, even mentioning an "audit probability"—albeit 0—is likely to influence behavior.

[9] We generated two random, predetermined audit-outcome sequences, allowing us to manipulate the control-enforcement outcomes between participants who faced the same audit probabilities. This manipulation allows us to control for an effect of audit outcomes on contribution decisions when testing our hypotheses. Enforcement outcome

During some of the rounds, participants see the contributions of other participants (the descriptive norm). We manipulate whether the descriptive norm is socially- or self-interested by displaying either high or low peer contributions.[10] As described in more detail below, we also manipulate the descriptive norm within subjects, such that by the end of the study all participants have seen peers who give high contributions as well as peers who give low contributions. We do not reveal to participants the identity of their peer contributors; and we inform participants that the peer contributions they see do not affect their own payoffs in any way.[11] Further, at no point in the study do we reveal to participants the contributions of their *matched* participant (whose contributions *do* impact their payoff). To prevent participants from trying to infer their matched participant's contributions, we also do not reveal participants' payoff at any point in the study. Ensuring that peer contributions have no effect on a participant's payoff and that matched participant contributions remain confidential is essential to investigate the effects of norms on behavior without the confounding influence of other well-known behavioral tendencies (reciprocal cooperation (Fehr et al. 1998), threatening to punish bad behavior with "trigger strategies" (Rasmusen 2001), trying to control which equilibrium outcome is attained (King and Wallin 1995), building a reputation (Schwartz and Young 2002), signaling intentions to comply (Alm et al. 1999; Tyran and Feld 2002), etc.).

---

does not significantly affect participant contributions, and controlling for this factor does not change statistical inferences for any of our analyses, so we do not discuss this factor further (and do not manipulate this factor in our second experiment, discussed below).

[10] Peer contributions were generated as part of a pilot study. The observed peer contributions in the pilot study were determined real-time by other (anonymous) participants' decisions. In the current study, determining in advance which contributions would be displayed (rather than limiting these contributions to those made in a given session) allows us to hold constant the exact contributions seen within experimental conditions. While participants in the study were not told anything that was untrue, we did not inform participants that data for peer contributions was pre-selected or that peers took part in a different experiment session. Participants were simply told, "…you will be allowed to see the contributions chosen by three other peer participants (with the same audit probability as you)."

[11] Specifically, we tell participants during the instructions phase, "The contributions of the three peer participants you see will not affect your payoff in any way, nor will your contributions affect these peer participants' payoffs."

In the next subsections, we provide additional detail on a typical round of play, the four phases of the experiment, and our experimental manipulations. Panel A of Figure 1 summarizes the phases of Experiment 1 and the experimental manipulations within each phase.

[Insert Figure 1 here]

Structure of a Typical Round

Figure 2 shows two sample screen shots from a single round of the experiment. In a typical round, participants first learn the probability of an audit and choose how much to contribute to the fund (see Panel A of Figure 2). Before locking in their contribution amount for a given round, participants are able to view what their actual payoff will be if they are or are not audited, as well as their expected payoff for different contribution levels (calculated as a weighted average of payoffs in the event of an audit and in the event of no audit). Participants also see that their final payoff depends on the contributions of their matched participant; but they cannot see how much their matched participant contributes. After locking in their contribution, participants learn whether or not they were audited in that round, as well as their resulting payoff (see Panel B of Figure 2). As described below, in most rounds participants also learn how much three of their peers contributed during the round.

[Insert Figure 2 here]

Phase 1: Rounds 1 and 2

We begin the experiment session by providing participants written instructions about the task. As part of the instructions, we explain each of the elements of the task: contributions to a fund that benefits another participant, audits and penalties, the audit probability, peer contributions, the sequence of events in each round, expected payoffs, etc. After answering any questions participants have about the task, we allow participants to familiarize themselves with

the computer interface. Once all participants have read the instructions and reviewed the computer interface, participants begin the experiment task.

As explained above, during the first two rounds of the experiment, we manipulate (between subjects) the audit probability that participants face. Specifically, half of the participants face an audit probability of 90 percent and half face an audit probability of 0 percent. Thus, initial controls are either strong (at 90 percent audit probability) or weak (at 0 percent audit probability).[12] We do not allow participants to see peer contributions during Rounds 1 and 2.

Phase 2: Round 3

In Round 3, we set the probability of an audit to 45 percent for all participants, across all conditions.[13] Participants begin this round with no knowledge of peer participant contributions, though they do learn peer contributions for Round 3 *after* locking in their own contributions. Concealing peer contributions until the end of Round 3 allows us to measure the effect of initial controls on contributions, absent any effect of the behavior of others.

Phase 3: Rounds 4-8

During each of Rounds 4-8 we hold constant the probability of an audit at 45 percent across all conditions. We manipulate descriptive norms between subjects by manipulating which peer contributions are displayed at the end of each of each of these rounds.[14] Specifically, half of the participants see that their peers contributed low amounts during each round (low-contribution descriptive norm) and half see that their peers contributed high amounts during each round (high-contribution descriptive norm). For the low-contribution descriptive norm, peer contributions are

[12] Based on a predetermined random audit outcome tied to the stated audit probability, all participants who faced an audit probability of 90 percent in Rounds 1 and 2 were audited in each of these rounds, while those who faced an audit probability of 0 percent were not audited.

[13] Using a 45 percent audit probability ensures that the absolute value of the change in audit probability from Round 2 to Round 3 is the same for participants who start in Rounds 1 and 2 with an audit probability of 90 percent and for participants who start in Rounds 1 and 2 with an audit probability of 0 percent.

[14] Because peer contributions for a given round are revealed at the end of the round, participants know Round 3 peer contributions before selecting their Round 4 contributions, and so on.

$0.00, $0.05, or $0.10 for Round 3, and all three are $0.00 for Rounds 4-8. For the high-contribution descriptive norm, peer contributions are $5.00, $4.95, and $4.90 for Round 3, and all three are $5.00 for Rounds 4-8.

Phase 4: Rounds 10-14

Prior to beginning Round 9, the program alerts participants that during the next six rounds of the experiment they will see the contributions of a different set of peers and have a new matched participant. During each of Rounds 9-14, we hold constant the probability of an audit at 45 percent across all conditions. We manipulate the descriptive norm participants encounter in these rounds, with participants who saw high peer contributions for Rounds 3-8 seeing low peer contributions for Rounds 9-14, and vice versa. Specifically, participants who saw high contributions for Rounds 3-8 see peer contributions of $0.00, $0.05, or $0.10 for Round 9 and peer contributions of $0.00 for Rounds 10-14. Participants who saw low contributions for Rounds 3-8 see peer contributions of $5.00, $4.95, and $4.90 for Round 9 and peer contributions of $5.00 for Rounds 10-14.

As before, the peer contributions are displayed after participants make their individual contribution decision for each round. Accordingly, participants in Round 9 have not yet experienced the within-subjects manipulation of descriptive norms associated with Phase 4 of our experiment. Thus, we exclude Round 9 contributions from our analyses (discussed in the next section).

Following Round 14, the program displays several debriefing and demographic questions. At the end of the experiment we remind participants that they will be paid within about one week.

Experiment 2

We conducted a second experiment to rule out possible alternative explanations for our findings (explained subsequently). Forty-eight students from various fields of study took part in our second experiment. Forty-four percent of the participants were male; and participants' average age was 20. We administered the experiment in three separate sessions, with each session consisting of 16 participants.

Experiment 2 is identical to Experiment 1, with the following exceptions. Experiment 2 employs a 1 x 2 x 14 design (initial-control strength x descriptive norm x 14 rounds). We held constant that initial controls were weak and that the descriptive norm during Rounds 3-8 (Phase 3) was high for all participants. However, we manipulate the descriptive norms participants see in Rounds 9-14 (Phase 4), with half of the participants seeing low peer contributions ($0.00) for Rounds 9-14, and half seeing high contributions ($5.00) for these rounds. Panel B of Figure 1 summarizes the phases and experimental manipulation of Experiment 2.

IV. RESULTS

We begin with a discussion of results from Experiment 1. We measure participants' contributions during each of the 14 rounds of the experiment. Average contributions by initial control strength, descriptive norm, and phase of the experiment are plotted in Figure 3.

[Insert Figure 3 here]

To test our predictions, we convert raw contributions to a measure of conformity for each period (*i)* as follows: [15]

$$Conformity_i = 1 - \frac{|contribution_i - descriptive\ norm_{4-8}|}{5},$$

[15] Our results are inferentially identical when conducting hypothesis tests with raw contributions. We test our hypotheses with *conformity*, as our hypotheses are stated in terms of conformity to the descriptive norm.

where $contribution_i$ is equal to participants' contribution during Round *i* of the experiment, $descriptive\ norm_{4\text{-}8}$ is equal to 5 for participants who saw high peer contributions during Rounds 4-8 of the experiment and 0 for participants who saw low peer contributions during Rounds 4-8 of the experiment. The denominator (5) represents the maximum possible difference between participants' contributions and the descriptive norm. Thus, *conformity* equals 1 when participants completely conform to the descriptive norm they see in Rounds 4-8, and equals 0 when participants do not conform at all to the descriptive norm they see in Rounds 4-8. In testing H1 and H2, this measure allows us to test the extent to which participants conform to the first descriptive norm they see in the experiment. In testing H3, this measure allows us to test the extent to which participants continue to conform to the initial descriptive norm (Rounds 4-8) when exposed to a conflicting descriptive norm in a later phase of the experiment (Rounds 10-14).

Though not explicitly hypothesized, based on prior research we expect that initial control strength will influence perceptions of appropriate behavior in the setting, and thus affect participant contributions (Tayler and Bloomfield 2011). Specifically, strong initial controls should lead to less pro-social behavior than weak initial controls. We test the effectiveness of our initial control manipulation by examining data only from Round 3 because participants in this round all face the same audit probability (45 percent), have been exposed to the manipulated initial-control strength (in Rounds 1 and 2), but have yet to see their peers' contributions. This allows us to isolate the effect of initial-control strength from that of peer contributions.

Panel A of Table 1 reports cell sizes, means, and standard deviations of participants' contributions in Round 3 for the strong initial controls and weak initial controls conditions. We collapse across the descriptive-norm conditions in this analysis since participants in the same

initial-control-strength conditions but different descriptive-norm conditions have received identical information up to this point in the study.

[Insert Table 1 here]

To test whether the initial strength of controls induced participants to perceive self-interested behavior to be appropriate in the setting, we conduct an analysis of variance (ANOVA) with initial control strength as a two-level, between-subjects factor, and participants' contributions in Round 3 as the dependent variable. As seen in Table 1, Panel A, participants contributed more in Round 3 when initial controls were strong (mean = 2.93) than when initial controls were weak (mean = 3.53). As shown in Table 1, Panel B, the main effect of initial-control strength is significant ($F = 5.34$, $p = .011$).[16] These results are consistent with prior research (Kachelmeier and Shehata 1997; Christ et al. 2008; Tayler and Bloomfield 2011) and indicate a successful manipulation of perceptions of appropriate behavior in the setting via initial control strength. We next investigate the joint effect of initial control strength and descriptive norms on *conformity*, as well as their persistence over time.[17]

Tests of Hypothesis 1

Our first hypothesis predicts that individuals who were subject to weak controls initially will conform to the descriptive norm more than individuals who were subject to strong controls initially. In testing H1, we only examine conformity to the descriptive norm in Rounds 4-8 because participants in these rounds have seen their peers' prior-round contributions (the descriptive norm) before making their own current-round contribution decisions, and because the nature of the descriptive norm (high contributions or low contributions) does not change during

[16] Unless otherwise stated, all reported p-values associated with directional predictions are one-sided.

[17] Another possible manipulation check would have been to ask participants, following Rounds 1 and 2, "What behavior do you believe is appropriate in this setting"? However, asking this question would likely bias subsequent decisions and limit the impact of peer contributions on behavior by overemphasizing personal beliefs. Thus, we rely on differences in individual behavior between conditions to proxy for average beliefs about behavior in the setting.

these rounds within each descriptive-norm condition. Further, all participants in these rounds are now subject to the same audit probability (45%).

Panel A of Table 2 reports descriptive statistics for participants' conformity to the descriptive norm during Rounds 4-8 of the experiment. We test H1 with a mixed-model ANOVA that controls for the dependency in participants' responses across rounds of the experiment.[18] Panel B of Table 2 reports the results of this analysis with initial-control strength and descriptive norm as two-level factors, and conformity as the dependent variable. As seen in Panel A of Table 2, participants conformed more to the descriptive norm in Rounds 4-8 when subject to weak initial controls (mean conformity = 0.59) than when subject to strong initial controls (mean conformity = 0.52), consistent with H1. As seen in Table 2, Panel B, this reduction in conformity to the descriptive norm is significant ($F = 2.57$, $p = 0.057$). These results suggests that individuals are more swayed by the behavior of others if they are initially subject to weak controls than if they are initially subject to strong controls, consistent with H1.

[Insert Table 2 here]

Tests of Hypothesis 2

Our second hypothesis predicts that the extent to which individuals conform to the descriptive norm will be greatest when the descriptive norm is most salient, determined in part by the extent to which the descriptive norm violates expectations for behavior in the setting. Thus, we expect that when individuals are subject to weak initial controls, they will conform more to self-interested descriptive norms than to socially-interested descriptive norms (H2a). We also expect that when individuals are subject to strong initial controls, they will conform more to socially-interested descriptive norms than to self-interested descriptive norms (H2b). In testing

[18] In particular, the mixed-model employs a repeated-measures design, and we assume a compound symmetric covariance structure for participants' responses across periods.

H2, we examine participant contributions in Round 3 (prior to the disclosure of peer contributions) and Rounds 4-8 (after the disclosure of peer contributions). Panel A of Table 3 reports descriptive statistics for participants' conformity to the descriptive norm during Round 3 and Rounds 4-8 of the experiment.

[Insert Table 3 here]

We test H2 with a mixed-model ANOVA that controls for the dependency in participants' responses across rounds of the experiment. Panel B of Table 3, reports the overall ANOVA results with initial-control strength, descriptive norm, and experiment phase (Round 3 vs. Rounds 4-8) as two-level factors, and conformity as the dependent variable. As seen in Panel B of Table 3, we find support for the predicted three-way interaction between initial-control strength, descriptive norm, and experiment phase ($F = 7.41$, $p=0.004$), consistent with H2. Below, we decompose this interaction, focusing separately on participants who were subject to weak initial controls (H2a) and subjects who were subject to strong initial controls (H2b).

We test H2a by analyzing the two-way interaction between descriptive norm and experiment phase for participants who were subject to weak initial controls. We summarize our tests of H2a in Panel C of Table 3. H2a predicts that participants subject to weak initial controls will conform more to self-interested descriptive norms than to socially-interested descriptive norms. For participants subject to weak initial controls, we find that mean conformity increases from 0.32 when participants do not see a descriptive norm in Round 3 to 0.41 when participants are exposed to low peer contributions in Rounds 4-8, and this difference is significant ($t = 2.03$, $p = 0.02$). In contrast, mean conformity increases only slightly from 0.73 when participants do not see a descriptive norm in Round 3 to 0.76 when participants are exposed to high peer contributions in Rounds 4-8, but this difference is not significant ($p > 0.10$). Although these

simple effects are consistent with H2a, the predicted interaction between experiment phase and descriptive norm is not significant ($p = 0.170$).

We test H2b by analyzing the two-way interaction between descriptive norm and experiment phase for participants who were subject to strong initial controls. We summarize our tests of H2b in Panel D of Table 3. H2b predicts that participants subject to strong initial controls will conform more to socially-interested descriptive norms than to self-interested descriptive norms. For participants subject to strong initial controls, we find that mean conformity increases from 0.55 when participants do not see a descriptive norm in Round 3 to 0.70 when participants are exposed to high peer contributions in Rounds 4-8, and this difference is significant ($t = 3.30$, $p < 0.001$). In contrast, mean conformity decreases slightly from 0.37 when participants do not see a descriptive norm in Round 3 to 0.34 when participants are exposed to low peer contributions in Rounds 4-8 (this difference is not significant ($p > 0.10$)). In addition, the predicted interaction between experiment phase and descriptive norm is significant ($t = 2.88$, $p = 0.002$), supporting H2b.

The combined results of H2, H2a, and H2b suggest that individuals conform more to peer behavior when that behavior conflicts with participants' expectations, which are shaped by the initial strength of controls. Specifically, when initial controls are strong, socially-interested norms are more compelling than self-interested norms; and when initial controls are weak, self-interested norms are more compelling than socially-interested norms.

Tests of Hypothesis 3

Our final hypothesis predicts that when individuals are exposed to a mixed descriptive norm, with some individuals acting in the social interest and others acting in their own self-interest, individuals will conform more to the self-interested behavior of others than to the

socially-interested behavior of others. Because participants see unambiguous high contributions or unambiguous low contributions in Rounds 4-8 (depending on the descriptive-norm manipulation), we operationalize a mixed descriptive norm by exposing participants who saw high contributions in Rounds 4-8 to low contributions in Rounds 9-14, and exposing participants who saw low contributions in Rounds 4-8 to high contributions in Rounds 9-14.[19] Given this operationalization, H3 predicts that participants who initially see their peers making *low* contributions in Rounds 4-8 will continue conforming to this selfish descriptive norm, even when they see peers making *high* contributions during Rounds 10-14. In contrast, H3 predicts that participants who initially see their peers making *high* contributions in Rounds 4-8 will decrease their conformity to this socially-interested descriptive norm when they see peers making *low* contributions in Rounds 10-14. In testing H3, we examine how participants' conformity to the descriptive norm in Rounds 4-8 changes when the descriptive norm flips in Rounds 10-14.[20]

[Insert Table 4 here]

Panel A of Table 4 reports descriptive statistics for participants' conformity to the descriptive norm during Rounds 4-8 and Rounds 10-14 of the experiment for the two descriptive-norm conditions.

We test H3 with a mixed-model ANOVA that controls for the dependency in participants' responses across rounds of the experiment. Panel B of Table 4 reports ANOVA

[19] Another possible operationalization of a mixed descriptive norm would be to show participants peer contributions with an average close to the midpoint of the contribution range (e.g., $5.00, $2.50, and $0.00; or $2.50, $2.50, and $2.50). We avoided this possibility for multiple reasons. First, one aim of the current study is to draw causal conclusions based on an exogenously manipulated, unambiguous descriptive norm. Our ability to test H1and H2 would have been severely hindered by use of a mixed descriptive norm in early rounds. Second, using one of the mixed peer groups suggested above only in Rounds 9-14 (so as not to affect tests of H2 and H3) would not have led to an average peer contribution of $2.50 if viewed in light of *all* peers in the study (participants in the high contribution descriptive-norm condition would have a higher peer-contribution average, across all rounds, than participants in the low contribution descriptive-norm condition). Finally, the strength of our manipulation may have been diminished by using a "middle-ground" peer contribution (e.g., $2.50), rather than the unambiguously high and unambiguously low peer contributions.

[20] Round 9 contributions are excluded from the analysis because participants do not see any contributions of their new group of peers for Rounds 9-14 until the end of Round 9.

results with initial-control strength, descriptive norm, and experimental phase (Rounds 4-8 vs. Rounds 10-14) as two-level factors, and conformity to the initial descriptive norm (i.e., the descriptive norm in Rounds 4-8) as the dependent variable. As seen in Table 4, Panel A, participants who initially see their peers contributing high amounts in Rounds 4-8 decrease their conformity to this socially-interested descriptive norm in Rounds 10-14 when they see their peers contributing low amounts. Mean conformity to the initial descriptive norm decreases from 0.73 in Rounds 4-8 to 0.61 in Rounds 10-14. In contrast, participants who initially see their peers contributing low amounts in Rounds 4-8 continue conforming to this self-interested descriptive norm when they see their peers contributing high amounts in Rounds 10-14. Mean conformity to the initial descriptive norm remains almost constant across phases of the experiment, with conformity of 0.37 in Rounds 4-8 and 0.36 in Rounds 10-14. As seen in Table 4, Panel B, the predicted interaction of descriptive norm x experiment phase is significant ($F = 13.64$, $p = 0.001$). These results suggest that when individuals are exposed to a mixed descriptive norm, with some individuals acting in the social interest and others acting in their own self-interest, individuals will conform more to the self-interested behavior of others than to the socially-interested behavior of others. This asymmetric response to descriptive norms suggests that, over time, norms for self-interested behavior will tend to be more compelling than norms for socially-interested behavior.

Additional Analysis: Experiment 2

One possible alternative explanation for our results relating to H3 is that participants in our study simply contribute less over time. A decrease in contributions could be driven by altruism fatigue (e.g., see Andreoni 1988 for examples where subjects play a repeated game and the provision of the public good ‘decays’ toward the free riding level with each repetition), by

learning, or by some other process causing a gradual movement to the Nash equilibrium of low (indeed, zero) contributions. To address this possibility, we conducted a second study where all participants are subject to weak initial controls and high peer contributions in Rounds 4-8, but where contributions in Rounds 10-14 are either high or low. Focusing on these two conditions allows us to compare the tendency for participants to conform to the high contributions of their peers even after several rounds of giving (when peer contributions are high in Rounds 4-8 and Rounds 10-14) to the conformity of participants who see high contributions in Rounds 4-8 and low contributions in Rounds 10-14.

Consistent with Experiment 1, participants who initially see their peers contributing high amounts in Rounds 4-8 decrease their conformity to this socially-interested descriptive norm in Rounds 10-14 when they see their peers contributing low amounts. Mean conformity to the initial descriptive norm decreases from 0.65 in Rounds 4-8 to 0.61 in Rounds 10-14. In contrast, participants who see their peers contributing high amounts in Rounds 4-8 and Rounds 10-14 increase their conformity to this socially-interested descriptive norm in Rounds 10-14. Mean conformity to the initial descriptive norm increases from 0.72 in Rounds 4-8 to 0.82 in Rounds 10-14. The interaction of descriptive norm x experiment phase is significant ($p = 0.003$), suggesting that results from H3 are not driven simply by a gradual decline in contributions over time, but rather by the exposure of participants in Experiment 1 to a mixed descriptive norm that includes low peer contributions.

V. CONCLUSIONS

Our study demonstrates how formal controls and the behavior of peers shape individual behavior in accounting contexts. First, we show that individuals who are initially subject to weak controls conform to the behavior of others more than people who are initially subject to strong

controls. Additionally, we show that individuals conform more to the behavior of others when that behavior conflicts with individuals' expectations about appropriate behavior. Finally, we show that individuals exposed to a mix of socially- and self-interested peer behavior will tend to conform more to the self-interested behavior of others than to the socially-interested behavior of others, suggesting that self-interested norms are "stickier" than socially-interested norms.

These results imply that decision makers, on average, will gravitate toward deviant behavior in organizations because of two interaction effects: (1) the asymmetric "pull" of descriptive norms on the behavior of individuals who have different beliefs about appropriate behavior, given the control environment, and (2) the asymmetric impact of conflicting descriptive norms. These findings have broad implications for the extent to which organizations and regulators must manage a tendency for individuals to act in their own self-interest.

Our study builds on prominent models of conformity, such as Bicchieri's (2006) model of social norm activation. Bicchieri argues that people conform to a known rule when they (a) believe that a sufficiently large subset of the population conforms to the rule (i.e. "empirical expectations") and (b) believe that a sufficiently large subset of the population expects them to conform to the rule (i.e., "normative expectations"). Our study suggests that people allow contextual features of a task (e.g., the strength of controls) to shape their "normative expectations." In particular, we provide evidence that the presence of a strong (weak) initial control ironically leads people to believe that others expect them to act selfishly (unselfishly). Furthermore, our study suggests that people develop "empirical expectations" in a biased manner. That is, we provide evidence that people attend and conform more to selfish behavior than to unselfish behavior and to behavior that conflicts with their "normative expectations" than

to behavior that aligns with their “normative expectations.” Thus, our study clarifies how empirical and normative expectations arise and develop in many work settings.

We note that this study has limitations which provide avenues for future research. For example, we use an abstract experiment to investigate the relationship between accounting controls and social norms. That is, we employ a modified public-goods game to make inferences about the constructs of interest. While contextual features of actual work settings may weaken or strengthen the effects we observe, they are unlikely to change the directional impact that these constructs have on employees in real work settings. Nevertheless, future researchers may wish to examine these constructs in a more contextualized setting, such as through field studies or experiments involving accounting experts and contextual labels. Further, we do not attempt to investigate methods of overcoming the asymmetric effect of norms on behavior in the current study. Future research may investigate how controls may be designed to influence, or even reverse, this asymmetry.

FIGURE 1
Experiment Design

Panel A: Experiment 1 Timeline and Manipulations

Panel A of Figure 1 depicts a timeline for Experiment 1and includes descriptions of how and when experimental manipulations occurred. As seen below, the experiment consists of four main phases. We manipulate initial control strength in Phase 1 by setting the probability of an audit equal to 0% or 90% in those rounds. In Phase 2, we set audit probability to 45% for all participants. In Phase 3, we manipulate descriptive norms by altering which peer contributions are seen by participants. In Phase 4, we manipulate descriptive norms again by altering (within subjects) which peer contributions are seen by participants.

Manipulation	**Phase 1** **Rounds 1-2**	**Phase 2** **Round 3**	**Phase 3** **Rounds 4-8**	**Phase 4** **Rounds 10-14**
Initial-Control Strength	Audit Probability is manipulated between subjects. Half of the participants face a 90% audit probability and half face a 0% audit probability.	Audit probability is held constant across all conditions at 45%.	Audit probability is held constant across all conditions at 45%.	Audit probability is held constant across all conditions at 45%.
Descriptive Norm	Participants cannot see peers' contributions	Participants cannot see peers' contributions	Half of the participants see their peers contribute high amounts, and half see their peers contribute low amounts.	Participants whose peers contributed high amounts in Rounds 3-8 now see peers contribute low amounts, and vice versa.

Panel B: Experiment 2 Timeline and Manipulations

Panel B of Figure 1 depicts a timeline for Experiment 2 and includes descriptions of how and when experimental manipulations occurred. As seen below, the experiment consists of four main phases. We hold constant initial control strength in Phase 1, setting the probability of an audit equal to 0% in those rounds, and set audit probability to 45% for all participants in all subsequent Phases. We hold constant descriptive norms in Phase 3, showing all participants high peer contributions. In Phase 4, we manipulate descriptive norms by altering which peer contributions are seen by participants (half of the participants see their peers contribute high amounts, and half see their peers contribute low amounts).

Manipulation	**Phase 1** **Rounds 1-2**	**Phase 2** **Round 3**	**Phase 3** **Rounds 4-8**	**Phase 4** **Rounds 10-14**
Descriptive Norm	Participants cannot see peers' contributions	Participants cannot see peers' contributions	All participants see their peers contribute high amounts.	Half of the participants see their peers contribute high amounts, and half see their peers contribute low amounts.

FIGURE 2
Screen Shots from a Typical Round of the Experiment

Panel A: Screen Shot Before Locking in a Contribution Amount

As seen in the screenshot below, before locking in their contribution amount for a given round, participants learn the probability of an audit that round and are able to view their expected payoff for different contribution levels, as well as what their actual payoff will be if they are or are not audited.

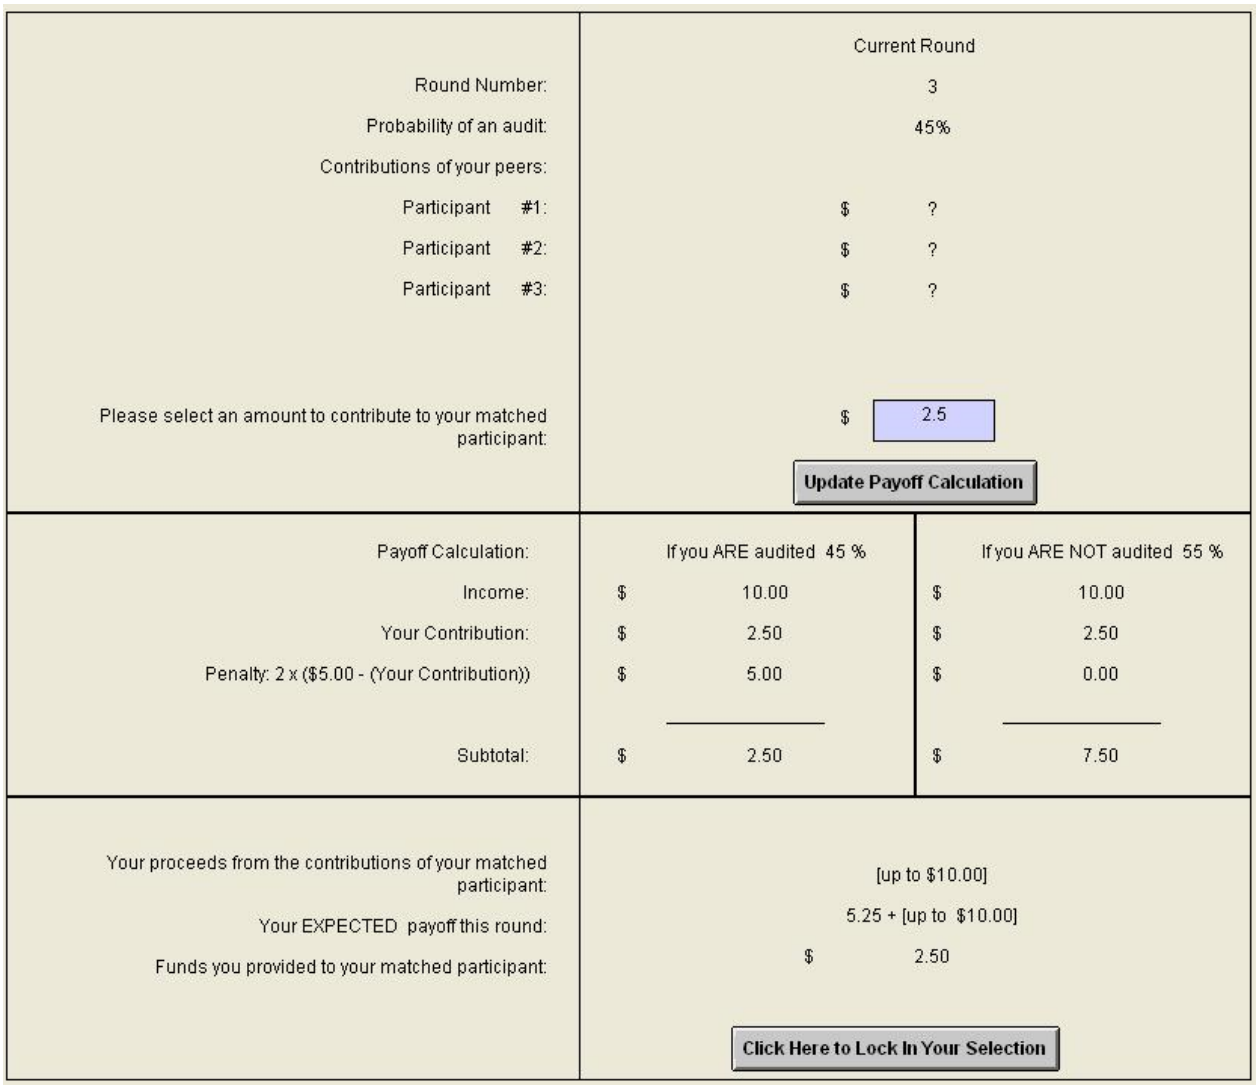

Panel B: Screen Shot After Locking in Contribution Amount

As seen in the screenshot below, after locking in their contribution amount for a given round, participants learn whether or not they were audited in that round, as well as their resulting payoff. In Rounds 3-14, participants also learn how much three of their peers contributed that round.

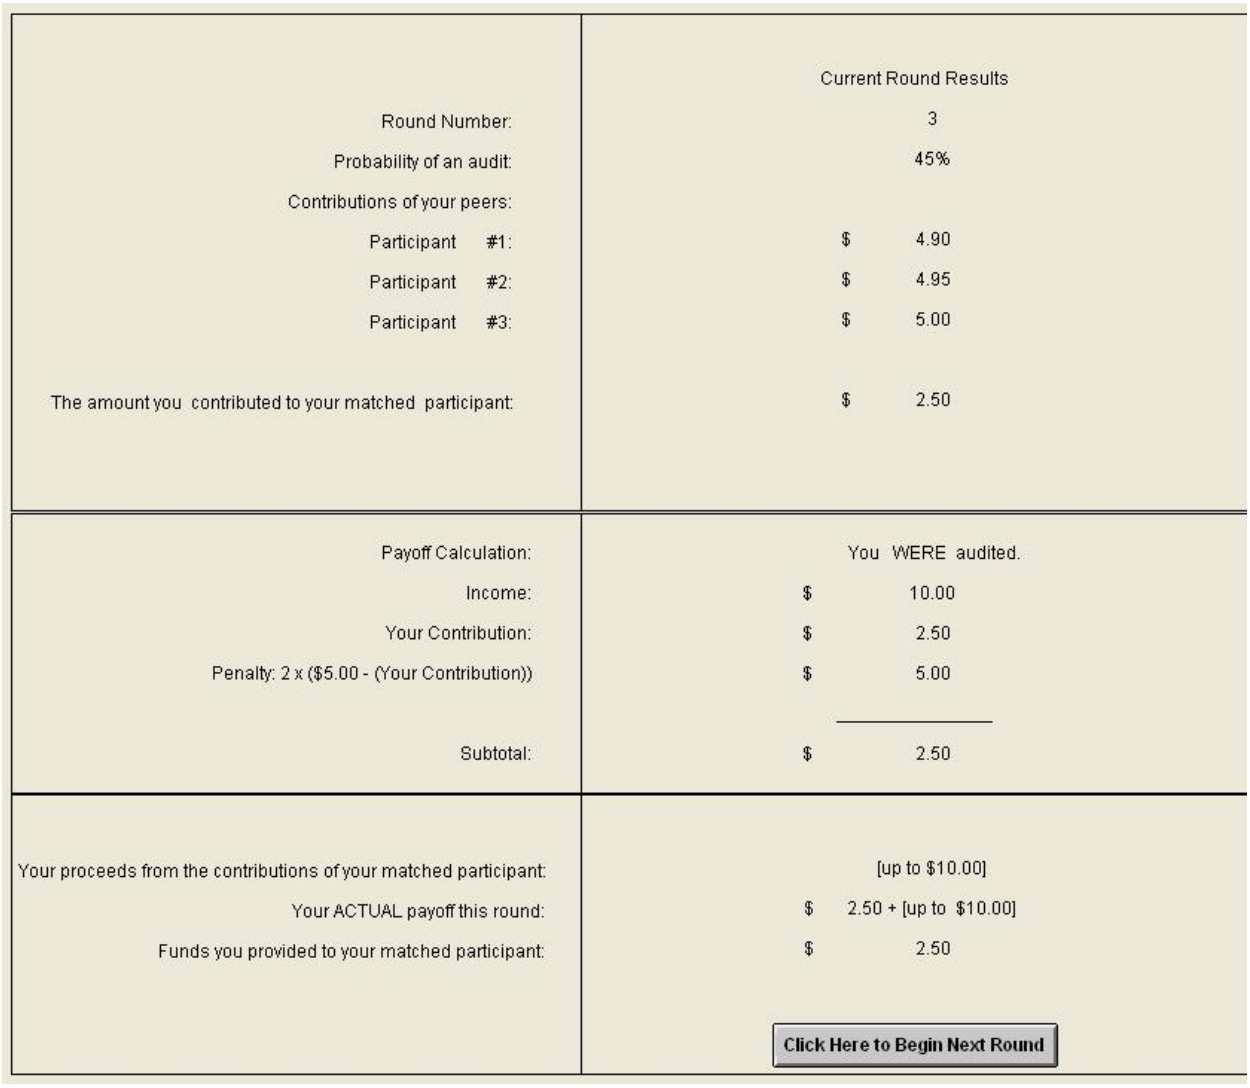

	Current Round Results
Round Number:	3
Probability of an audit:	45%
Contributions of your peers:	
Participant #1:	$ 4.90
Participant #2:	$ 4.95
Participant #3:	$ 5.00
The amount you contributed to your matched participant:	$ 2.50

Payoff Calculation:	You WERE audited.
Income:	$ 10.00
Your Contribution:	$ 2.50
Penalty: 2 x ($5.00 - (Your Contribution))	$ 5.00
Subtotal:	$ 2.50

Your proceeds from the contributions of your matched participant:	[up to $10.00]
Your ACTUAL payoff this round:	$ 2.50 + [up to $10.00]
Funds you provided to your matched participant:	$ 2.50

Click Here to Begin Next Round

FIGURE 3
Average Contributions in Experiment 1

Average contributions by initial-control strength (strong vs. weak), descriptive norm (high contributions in Rounds 4-8 then low contributions in Rounds 9-14 vs. low contributions in Rounds 4-8 then high contributions in Rounds 9-14), and experiment phase (Rounds 1-2, Round 3, Rounds 4-8, and Rounds 10-14) are plotted below. Round 9 is excluded because participants in this round have not yet seen the contributions of their new set of peers.

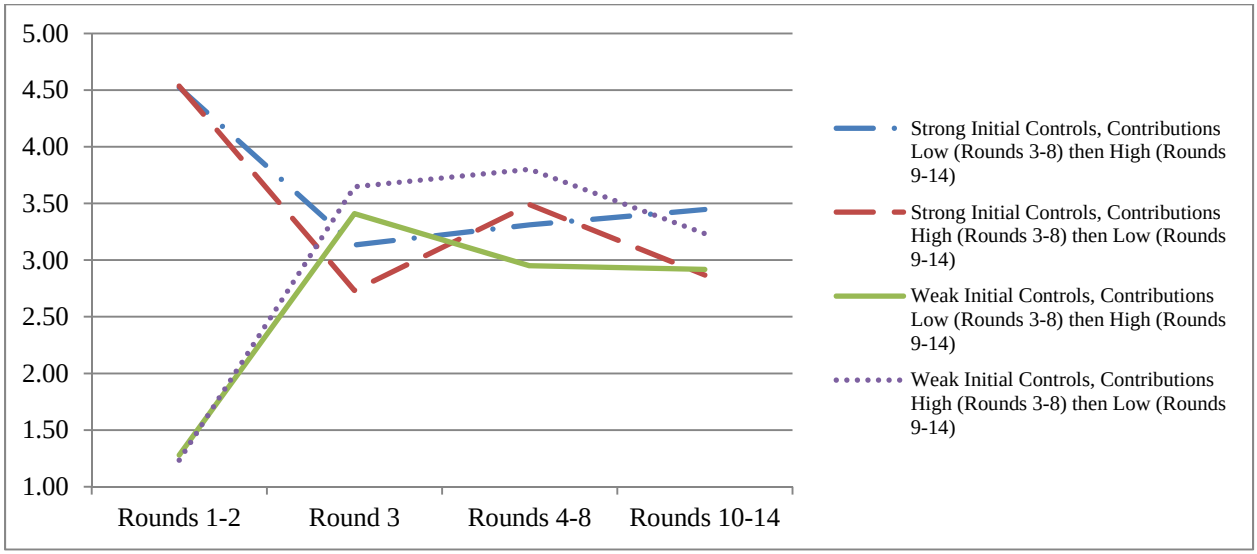

TABLE 1
Effect of Initial-Control Strength on Contributions in Round 3

Panel A: Mean Contributions (Standard Deviation)

	Round 3
Strong Initial-Controls	2.93 (1.85) n = 86
Weak Initial-Controls	3.53 (1.55) n = 88
Overall	3.23 (1.72) n = 174

Panel B: Analysis of Variance

Factor	df	F-ratio	p-value
Initial Control Strength	1	5.34	0.011
Error	172		

Panel A summarizes contributions of participants for the two initial-control-strength conditions. Initial control strength was manipulated by altering the probability that participants' contributions would be audited during Rounds 1 and 2. Half of the participants faced a 90% probability of being audited in Rounds 1 and 2 (strong initial controls) and half faced a 0% probability of being audited (weak initial controls). Contributions of experimental currency ranged from $0 to $5. Panel B presents a statistical test of the main effect of initial control strength in Round 3. We use one-sided equivalent p-values when testing hypotheses.

TABLE 2
Tests of Hypothesis 1

Panel A: Descriptive Statistics of Mean Conformity in Rounds 4-8 (Std. Dev.)

	Descriptive Norm		
	Low Contribution	High Contribution	Overall
Strong Initial Controls	0.34 (0.40) n=43	0.70 (0.37) n=43	0.52 (0.43) n = 86
Weak Initial Controls	0.41 (0.40) n=44	0.76 (0.31) n=44	0.59 (0.40) n = 88
Overall	0.37 (0.40) n=87	0.73 (0.34) n=87	0.55 (0.42) n = 174

Panel B: Analysis of variance

Factor	df	F-ratio	p-value
Initial-Control Strength	**1**	**2.57**	**0.057**
Descriptive Norm	1	70.83	<0.01
Initial-Control Strength x Descriptive Norm	1	0.01	0.906

Panel A summarizes conformity for each of the four conditions generated by crossing our initial-control-strength and descriptive-norm manipulations in Rounds 4-8. Conformity is calculated as follows: $Conformity_i = 1 - \frac{|contribution_i - descriptive\ norm_{4-8}|}{5}$, where $contribution_i$ is equal to participants' contribution during Round i of the experiment, descriptive $norm_{4\text{-}8}$ is equal to 5 for participants who saw high peer contributions during Rounds 4-8 of the experiment and 0 for participants who saw low peer contributions during Rounds 4-8 of the experiment. The denominator (5) represents the maximum possible difference between participants' contributions and the descriptive norm. Initial-control strength was manipulated by altering the probability that participants' contributions would be audited during Rounds 1 and 2. Half of the participants faced a 90% probability of being audited in Rounds 1 and 2 (strong initial controls) and half faced a 0% probability of being audited (weak initial controls). Descriptive norms were manipulated by showing participants either high or low prior-round contributions from a group of three peers before participants made contribution decisions in Rounds 4-8. Panel B reports the results of statistical tests of the main and interaction effects of initial-control strength and descriptive norms in Rounds 4-8. We use one-sided equivalent p-values when testing hypotheses. Tests of hypotheses are bolded.

TABLE 3
Tests of Hypotheses H2, H2a and H2b

Panel A: Descriptive Statistics of Mean Conformity (Std. Dev.)

		Experiment Phase	
		Round 3	**Rounds 4-8**
Strong Initial Controls	Low-Contribution Descriptive Norm	0.37 (0.37) n=43	0.34 (0.40) n=43
	High-Contribution Descriptive Norm	0.55 (0.37) n=43	0.70 (0.37) n=43
Weak Initial Controls	Low-Contribution Descriptive Norm	0.32 (0.31) n=44	0.41 (0.40) n=44
	High-Contribution Descriptive Norm	0.73 (0.31) n=44	0.76 (0.31) n=44

Panel B: Overall Analysis of Variance

Factor	df	F-Ratio	p-value
Initial-Control Strength	1	2.28	0.067
Descriptive Norm	1	55.58	<0.001
Experiment Phase	1	6.89	0.005
Initial-Control Strength x Descriptive Norm	1	1.73	0.190
Initial-Control Strength x Experiment Phase	1	0.01	0.940
Descriptive Norm x Experiment Phase	1	1.89	0.171
Initial-Control Strength x Descriptive Norm x Experiment Phase	**1**	**7.41**	**0.004**

Panel C: Weak Initial Controls Contrasts (H2a)

Contrast	df	t-ratio	p-value
Descriptive Norm x Experiment Phase	**1**	**0.96**	**0.170**
Simple effect of Expt. Phase for Low-Contribution Descriptive Norm	**1**	**2.03**	**0.022**
Simple effect of Expt. Phase for High-Contribution Descriptive Norm	1	0.68	0.501

Panel D: Strong Initial Controls Contrasts (H2b)

Contrast	df	t-ratio	p-value
Descriptive Norm x Experiment Phase	**1**	**2.88**	**0.002**
Simple effect of Expt. Phase for Low-Contribution Descriptive Norm	1	-0.78	0.439
Simple effect of Expt. Phase for High-Contribution Descriptive Norm	**1**	**3.30**	**<0.001**

Panel A summarizes conformity for each of the eight conditions generated by crossing our initial-control-strength, descriptive-norm, and experiment-phase manipulations. Conformity is calculated as follows: $Conformity_i = 1 -$

$\frac{|contribution_i - descriptive\ norm_{4-8}|}{5}$, where $contribution_i$ is equal to participants' contribution during Round i of the experiment, descriptive $norm_{4\text{-}8}$ is equal to 5 for participants who saw high peer contributions during Rounds 4-8 of the experiment and 0 for participants who saw low peer contributions during Rounds 4-8 of the experiment. The denominator (5) represents the maximum possible difference between participants' contributions and the descriptive norm. Initial-control strength was manipulated by altering the probability that participants' contributions would be audited during Rounds 1 and 2. Half of the participants faced a 90% probability of being audited in Rounds 1 and 2 (strong initial controls) and half faced a 0% probability of being audited (weak initial controls). Descriptive norms were manipulated by showing participants either high or low prior-round contributions from a group of three peers before participants made contribution decisions. Experiment phase was manipulated by not displaying peer contributions before participants made contribution decisions in Round 3, but displaying peer contributions before participants made contribution decisions in Rounds 4-8. Panel B reports the results of statistical tests of the main and interaction effects of initial-control strength, descriptive norms, and experiment phase (Round 3 vs. Rounds 4-8). Panel C and Panel D decompose the three-way interaction of initial-control strength, descriptive norms, and experiment phase into two-way interactions for participants who faced weak initial controls (Panel C) and for participants who faced strong initial controls (Panel D). We use one-sided equivalent p-values when testing hypotheses. Tests of hypotheses are bolded.

TABLE 4
Tests of Hypothesis 3

Panel A: Descriptive Statistics of Mean Conformity to Descriptive Norm in Rounds 4-8 (Std. Dev.)

		Experiment Phase	
		Rounds 4-8	**Rounds 10-14**
Descriptive Norm	High Contributions in Rounds 4-8; Low Contributions in Rounds 10-14	0.73 (0.34) n=87	0.61 (0.40) n = 87
	Low Contributions in Rounds 4-8; High Contributions in Rounds 10-14	0.37 (0.40) n=87	0.36 (0.40) n=87

Panel B: Analysis of variance

Factor	df	F-ratio	p-value
Initial-Control Strength	1	4.1	0.022
Descriptive Norm	1	60.46	<0.001
Experiment Phase	1	19.25	<0.001
Descriptive Norm x Experiment Phase	**1**	**13.64**	**<0.001**
Initial-Control Strength x Descriptive Norm	1	0.08	0.784
Initial-Control Strength x Experiment Phase	1	0.57	0.450
Initial-Control Strength x Descriptive Norm x Phase	1	0.15	0.702

Panel A summarizes conformity for each of the four conditions generated by crossing our descriptive-norm and experiment phase manipulations. Conformity is calculated as follows: $Conformity_i = 1 - \frac{|contribution_i - descriptive\ norm_{4-8}|}{5}$, where $contribution_i$ is equal to participants' contribution during Round i of the experiment, descriptive $norm_{4\text{-}8}$ is equal to 5 for participants who saw high peer contributions during Rounds 4-8 of the experiment and 0 for participants who saw low peer contributions during Rounds 4-8 of the experiment. The denominator (5) represents the maximum possible difference between participants' contributions and the descriptive norm. Initial-control strength was manipulated by altering the probability that participants' contributions would be audited during Rounds 1 and 2. Half of the participants faced a 90% probability of being audited in Rounds 1 and 2 (strong initial controls) and half faced a 0% probability of being audited (weak initial controls). Descriptive norms were manipulated by showing participants either high or low prior-round contributions from a group of three peers before participants made contribution decisions. Experiment phase was manipulated by switching the descriptive norm between phases of the experiment. Participants who saw high peer contributions during Rounds 4-8 of the experiment saw low peer contributions during Rounds 10-14 of the experiment, and vice versa for participants who saw low peer contributions during Rounds 4-8 of the experiment. Panel B reports the results of statistical tests of the main and interaction effects of initial-control strength, descriptive norms, and experiment phase (Round 3 vs. Rounds 4-8). We use one-sided equivalent p-values when testing hypotheses. Tests of hypotheses are bolded.

REFERENCES

Aarts, H., and A. Dijksterhuis. 2003. The Silence of the Library: Environment, Situational Norm, And Social Behavior. *Journal of Personality and Social Psychology* 84: 18-28.

Akerlof, G. A. 1982. Labor contracts as a partial gift exchange. *Quarterly Journal of Economics* 97: 543-569.

Alm, J.; G. H. McClelland, and W. D. Schuize. 1999. Changing the Social Norm of Tax Compliance and by Voting. *Kyklos* 52: 141-71.

Andreoni, J. 1988. Why free ride?: Strategies and learning in public goods experiments. *Journal of public Economics* 37(3): 291-304.

Aronson, E.; T. D. Wilson; and R. M. Akert. 2002. *Social Psychology, Fourth edition.* Upper Saddle River, NJ: Prentice Hall.

Asch, S. E. 1951. Effects Of Group Pressure Upon The Modification and Distortion Of Judgment. In *Groups, Leadership and Men*, edited by H. S. Guetzkow, 177-190. Pittsburgh, PA: Carnegie Press.

Berg, J., J. Dickhaut, and K. McCabe. 1995. Trust, reciprocity, and social history. *Games and Economic Behavior* 10: 122-142.

Berkowitz, L., and E. Donnerstein. 1982. External Validity Is More Than Skin Deep - Some Answers to Criticisms of Laboratory Experiments. *American Psychologist* 37 (3):245-257.

Bicchieri, C., and E. Xiao. 2009. Do the Right Thing: But Only if Others Do So. *Journal of Behavioral Decision Making* 22: 191–208.

Blanthorne, C., and S. Kaplan. 2008. An Egocentric Model of the Relations Among The Opportunity to Underreport, Social Norms, Ethical Beliefs, and Underreporting Behavior. *Accounting, Organizations and Society* 33: 684–703.

Bloomfield, R. J., W. B. Tayler, and F. Zhou. 2009. Momentum, Reversal, and Uninformed Traders in Laboratory Markets. *Journal of Finance* 64: 2535-2558.

Bobek, D. D., R. W. Roberts, and J. T. Sweeney. 2007. The Social Norms of Tax Compliance: Evidence From Australia, Singapore, And The United States. *Journal of Business Ethics* 74: 49–64.

Broom, L., and P. Selznick. 1968. *Sociology; a text with adapted readings*. New York, NY: Harper and Row.

Burgoon, J. K., and M. Burgoon. 2001. Expectancy Theories. In *The New Handbook of Language and Social Psychology*, edited by W. P. Robinson and H. Giles, 79–99. London, U.K.: John Wiley & Sons, Ltd.

Charness, G., G. R. Frechette, and J. H. Kagel. 2004. How robust is laboratory gift exchange? *Experimental Economics* 7: 189-205.

Christ, M. H., K. L. Sedatole, K. L. Towry, and M. A. Thomas. 2008. When Formal Controls Undermine Trust and Cooperation. *Strategic Finance* January: 38-44.

Christ, M. H., S. A. Emett, W. B. Tayler, and D. A. Wood. 2014. Compensation or Feedback: Motivating Performance in Multidimensional Tasks. Working Paper, University of Georgia, Cornell University, and Brigham Young University.

Choi, J., G. Hecht, and W. B. Tayler. 2012. Lost in Translation: The Effects of Incentive Compensation on Strategy Surrogation. *The Accounting Review* 87: 1135-1164.

Choi, J., G. Hecht, and W. B. Tayler. 2013. Strategy Selection, Surrogation, and Strategic Performance Measurement Systems. *Journal of Accounting Research* 51 (1): 105-133.

Cialdini, R. B., C. A. Kallgren, and R. R. Reno. 1991. A Focus Theory of Normative Conduct: A Theoretical Refinement and Reevaluation of the Role of Norms in Human Behavior. *Advances in Experimental Social Psychology* 21: 201–34.

Cialdini, R. B., R. R. Reno, and C. A. Kallgren. 1990. A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology* 58: 1015–26.

Clor-Proell, S.M. 2009. The Effects of Expected and Actual Accounting Choices on Judgments and Decisions. *The Accounting Review* 84: 1465-1494.

Davidson, B. I. and D. E. Stevens. 2013. Can a Code of Ethics Improve Manager Behavior and Investor Confidence? An Experimental Study. *The Accounting Review* 88 (1): 51-74.

Davis, D. D., and C. A. Holt. 1993. *Experimental Economics*. Princeton, N.J.: Princeton University Press.

Davis, J. S., G. Hecht, and J. D. Perkins. 2003. Social Behaviors, Enforcement, and Tax Compliance Dynamics. *The Accounting Review* 78: 39-69.

Dawson, E., T. Gilovich, and D. T. Regan. 2002. Motivated Reasoning and Performance on the Wason Selection Task. *Personality and Social Psychology Bulletin* 28: 1379–1387.

Ditto, P. H., and D. F. Lopez. 1992. Motivated Skepticism- Use of Differential Decision Criteria for Preferred and Nonpreferred Conclusions. *Journal of Personality and Social Psychology* 63: 568-584.

Ditto, P.H., G. D. Munro, A. M. Apanovitch, J. A. Scepansky, and L. K. Lockhart. 2003. Spontaneous Skepticism: The Interplay of Motivation and Expectation in Responses to Favorable and Unfavorable Medical Diagnoses. *Personality and Social Psychology Bulletin* 29: 1120–1132.

Ditto, P.H., J. A. Scepansky, G. D. Munro, A. M. Apanovitch, and L. K. Lockhart. 1998. Motivated Sensitivity To Preference-Inconsistent Information. *Journal of Personality and Social Psychology* 75: 53–69.

Dyreng, S. D., W. J. Mayew, and C. D. Williams. 2012. Religious Social Norms and Corporate Financial Reporting. *Journal of Business Finance & Accounting* 39: 845-875.

Fehr, E., G. Kirchsteiger, and A. Riedl. 1993. Does Fairness Prevent Market Clearing? An Experimental Investigation. *Quarterly Journal of Economics* 108: 437–460.

Fehr, E., E. Kirchler, A. Weichbold, and S. Gachter. 1998. When Social Norms Overpower Competition-Gift Exchange in Experimental Labor Markets. *Journal of Labor Economics* 16: 325–351.

Fischbacher, U. 2007. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 10: 171-178.

Fischer, P., and S. Huddart. 2008. Optimal Contracting with Endogenous Social Norms. *American Economic Review* 98: 1459-1475.

Gilovich, T. 1991. *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*. New York, NY: Free Press.

Gneezy, U., and A. Rustichini. 2000. A Fine Is a Price. *Journal of Legal Studies* 29: 1–17.

Graham, J. R., C. R. Harvey, and S. Rajgopal. 2005. The economic implications of corporate financial reporting. *Journal of Accounting & Economics* 40 (1-3):3-73.

Hales, J. 2007. Directional preferences, information processing, and investors' forecasts of earnings. *Journal of Accounting Research* 45: 607-628.

Hannan, R. L. 2005. The combined effect of wages and firm profit on employee effort. *The Accounting Review* 80: 167-188.

Hannan, R. L., G. P. McPhee, A. H. Newman, and I. D. Tafkov. 2012. The effect of relative performance information on performance and effort allocation in a multi-task environment. *The Accounting Review* 88: 553-575.

Kachelmeier, S. J., and M. Shehata. 1997. Internal Auditing and Voluntary Cooperation in Firms: A Cross-Cultural Experiment. *The Accounting Review* 72: 407-431.

Kandel, E., and E. Lazear.1992. “Peer Pressure and Partnerships, *Journal of Political Economy* 100: 801–17.

Kerr, N. L. 1983. Motivation Losses in Small Groups: A Social Dilemma Analysis. *Journal of Personality and Social Psychology* 45: 819-828.

King, R. R., and D. E. Wallin. 1995. Experimental Tests of Disclosure with an Opponent. *Journal of Accounting & Economics* 19: 139-67.

Kuang, X., and D. V. Moser. 2009. Reciprocity and the Effectiveness of Optimal Agency Contracts. *The Accounting Review* 84: 1671-1694.

Kunda, Z. 1990. The Case for Motivated Reasoning. *Psychological Bulletin* 108: 480-498.

Ledyard, J. O. 1995. Public Goods: A Survey of Experimental Research. In *The Handbook of Experimental Economics*, edited by J. H. Kagel and A. E. Roth. Princeton, N.J.: Princeton University Press.

Liberman, V., S. M. Samuels, and L. Ross. 2004. The Name of the Game: Predictive Power of Reputations versus Situational Labels in Determining Prisoners’ Dilemma Game Moves. *Personality and Social Psychology Bulletin* 30: 1175-1184.

Mas, A., and E. Moretti. 2009. Peers at Work. *American Economic Review* 99: 112–145.

McGuire, S. T., T. C. Omer, and N. Y. Sharp. 2012. The Impact of Religion on Financial Reporting Irregularities. *The Accounting Review* 87: 645-673.

Merriam-Webster Inc. 2005. *The Merriam-Webster dictionary*. Springfield, Mass.: Merriam-Webster.

Messick, D. M. 1999. Alternative Logics for Decision Making in Social Settings. *Journal of Economic Behavior & Organization* 39: 11–28.

Omer, T. C., N. Y. Sharp, and D. Wang. 2013. *The Impact of Religion on the Going Concern Reporting Decisions of Local Audit Practice* Offices. Working paper, Texas A&M University.

Pyszcynski, T., and J. Greenberg. 1987. Toward an integration of cognitive and motivational perspectives on social inference—A biased hypothesis-testing model. *Advances in Experimental Social Psychology* 20: 297-340.

Rasmusen, E. 2001. *Games and Information: An Introduction to Game Theory, Third edition*. Malden, MA: Blackwell.

Schwartz, S. H., and J. A. Fleishman. 1978. Personal Norms and the Mediation of Legitimacy Effects on Helping. *Social Psychology* 41: 306-315.

Schwartz, S. T., and R. A. Young. 2002. A Laboratory Investigation of Verification and Reputation Formation in a Repeated Joint Investment Setting. *Contemporary Accounting Research* 19: 311-342.

Sherif, M., and G. Murphy. 1936. *The Psychology of Social Norms.* New York, NY and London, U.K.: Harper & Brothers.

Swieringa, R. J., and K. E. Weick. 1982. An Assessment of Laboratory Experiments in Accounting. *Journal of Accounting Research* 20: 56-101.

Tafkov, I. D. 2012. Private and public relative performance information under different compensation contracts. *The Accounting Review* 88: 327-350.

Tayler, W. B. 2010. The Balanced Scorecard as a Strategy-Evaluation Tool: The Effects of Implementation Involvement and a Causal-Chain Focus. *The Accounting Review* 85: 1095-1117.

Tayler, W. B., and R. J. Bloomfield.2011. Norms, Conformity, and Controls. *Journal of Accounting Research* 49: 753-790.

Tenbrunsel, A. E., and D. M. Messick. 1999. Sanctioning Systems, Decision Frames, and Cooperation. *Administrative Science Quarterly* 44: 684-707.

Thayer, J. 2011. Determinants of Investors' Information Acquisition: Credibility and Confirmation. *The Accounting Review* 86: 1-22.

Tyran, J. R., and L. P. Feld. 2002. Tax Evasion and Voting: An Experimental Analysis. *Kyklos* 55: 197-221.

Wallschutzky, I. G. 1984. Possible Causes of Tax Evasion. *Journal of Economic Psychology* 5: 371–384.

Wenzel, M. 2004. An Analysis of Norm Processes in Tax Compliance. *Journal of Economic Psychology* 25: 213–228.

Wenzel, M. 2005a. Motivation or rationalization? Causal Relations between Ethics, Norms, and Tax Compliance. *Journal of Economic Psychology* 26: 491–508.

Wenzel, M. 2005b. Misperceptions of Social Norms about Tax Compliance: From Theory to Intervention. *Journal of Economic Psychology* 26: 862-883.