Inequality as an Incentive

Jeffrey V. Butler
University of California, Merced

Eric Cardella
Texas Tech University *

This version: *April 13, 2021*

Abstract

We study how salient wage and income inequality affect firm-relevant behaviors. We propose a novel channel based on Just World Beliefs relating experienced inequality to subsequent behavior, use it to make broad predictions and then test those predictions in a laboratory experiment involving over 420 participants. Mirroring labor market trends, we randomly vary initial-task type (ability-intensive, effort-intensive) and associated wage levels (high, low). Next, we vary facets of initial-task inequality revealed to workers across treatments (wage inequality, income inequality, both or neither). All workers then perform a common set of subsequent tasks permitting a wide array of behaviors (pro-social, anti-social, pure effort). Consistent with our predictions, we find that: i) revealed inequality only affects behavior when associated with an ability-intensive task; ii) revealed wage inequality induces stronger behavioral responses than revealed income inequality, which primarily consist of *low-wage* workers *increasing* effort and *reducing* pro-sociality; and iii) revealing income inequality in addition to wage inequality produces countervailing effects. We find no effects of inequality on anti-social (lying) behavior. Supporting attitudinal evidence on Just World Beliefs is provided by a post-experiment survey. Our results may inform the literature on optimal compensation schemes and wage transparency.

JEL Classification: Z19, D91

Keywords: Inequality; Incentives; Just World Beliefs

*The authors thank George Akerlof, Yan Chen, Dan Friedman, Laura Giuliano, Einav Hart, Karla Hoff, Alessandro Santuoso, Alec Smith and Maggie Switek, as well as conference participants at the 2019 Bay Area Behavioral and Experimental Economics Workshop (BABEEW), ERINN 2019, the 2019 Texas Experimental Association Symposium, WEAI 2019 and seminar participants at George Mason University, Osaka University, the University of California, Santa Cruz, the University of Pennsylvania and the University of California, Merced for many helpful comments which have improved the paper.

1 Introduction

A slew of recent studies using a wide variety of methodologies document a common finding, that salient horizontal compensation inequality, i.e., inequality among similar workers, can produce substantial backlash (Bracha, Gneezy, and Loewenstein 2015; Breza, Kaur, and Shamdasani 2018; Breza et al. 2018; Card, Mas, Moretti, and Saez 2012; Dube, Giuliano, and Leonard 2019; Sseruyange and Bulte 2019). Even trivially small wage differences have been shown to reduce employee morale and productivity, increase turnover or increase anti-social behavior. Dube et al. (2019), for example, exploit an exogeneous (to the firm) policy shock increasing the wage gap among entry-level workers at a large US retailer by a mere 9 cents, from 1 cent to 10 cents, and estimate a large causal effect of negative peer-wage differentials on quits.

While inequality-induced backlash is rapidly becoming the accepted wisdom among economists, some research as well as everyday experience suggests that reactions to inequality may vary qualitatively across contexts. Far from always (weakly) inducing backlash, inequality may sometimes be celebrated and *demanded* by workers. Consider the classroom context, interpreting students as workers and grades as compensation. In that context, students themselves call for considerable inequality in the grade distribution and express support for (meritocratic, achievement-oriented) grading principles that naturally induce substantial dispersion (Gaultney and Cann 2001; Nisan 1985; Sabbagh, Faher-Aladeen, and Resh 2004).[1] In the context of athletics, the Olympic Games are a widely praised institution despite routinely conferring vastly unequal rewards across arbitrarily small – to the point of being essentially random – performance differences.[2] More broadly, large-scale surveys reveal that wealth and income inequality as well as principles naturally inducing such inequality, ostensibly even among similar individuals, are considered necessary and *ideal*. (John T. Jost and Sullivan 2003; Starmans, Sheskin, and Bloom 2017).

To the extent that *reactions* to experienced inequality are governed by fairness judgments, qualitatively different reactions across contexts may reflect contextual variation in the perceived basis of inequality (cf, Bracha et al. 2015; Breza et al. 2018). The latter examples above suggest that

[1] Gaultney and Cann (2001, p. 85) study college students in a large introductory Psychology class and find: "Participants indicated what they considered to be a fair grade distribution. They thought that instructors should give roughly equal numbers of As, Bs, and Cs, with fewer Ds, and even fewer Fs. A grade of B would be modal, and A, B, and C would each make up about 30% of the grade distribution."

[2] While some of Michael Phelps' close victories in the men's butterfly have entered the collective consciousness, examples of arbitrarily close Olympic victories are not hard to find. A cursory search for swimming results, for instance, yields that in the 2012 (London) Olympic Games a mere 0.01 seconds separated gold from silver in the 100-meter men's freestyle swimming finals, constituting a 0.02% difference in times. For comparison, the gold-to-gold comparison across the 2012 and 2016 Olympics in the same event was 0.08 seconds (0.1% difference), with the 2016 gold-medal winner being *slower* than the 2012 gold medalist. Thus, the time difference across these games of the fastest swimmers was approximately an order of magnitude greater than the difference between the first- and second-fastest swimmer in the more recent Games. This anecdote, one of many possible, highlights the possibility that random variation due to time and venue characteristics may swamp effort or ability differences separating medalists.

merit-based inequality, where rewards are allocated according to inputs of effort and ability, may be judged as inherently fair (see, e.g., Adams 1963; Hayek 2020). When relative inputs of effort and ability are only imperfectly observable, subjective fairness perceptions may depend crucially on subconscious cognitive and psychological processes which, in turn, may be context-dependent. For ease of exposition, we will broadly refer to the subconscious cognitive and psychological processes induced by experienced inequality as the *non-pecuniary incentives* of inequality in order to distinguish them from the pecuniary incentives that form the basis of traditional economic theory (wages, private benefits, etc).

In the present study, we propose a novel channel (Section 2) through which experienced inequality may interact with contextual factors and subconscious psychological processes to provide positive non-pecuniary effort incentives, improving subsequent firm-relevant behavior. The channel relies on experienced disadvantage *strengthening* Just World Beliefs which affects behavior in predictable ways. We test hypotheses derived from this channel in a laboratory experiment involving over 420 participants. Our laboratory experiment admits a wide array of potential responses to experienced inequality: pro-social and anti-social as well as responses in the quantity and quality of subsequent (real-effort) production. Allowing a wide array of responses is important because it is not well understood how these behaviors interrelate, yet they may all be important to a firm. A firm may care about pro-sociality to the extent that it facilitates cooperation, while it may also care about anti-sociality to the extent that it manifests itself as theft, fraud or other so-called "counterproductive workplace behaviors." Beyond our channel-derived hypotheses, allowing this wide array of responses may provide preliminary evidence on relationships among various behaviors.

Exploring how non-pecuniary inequality incentives affect firm-relevant behavior is a question of first-order importance because, ignoring these incentives, traditional economic theory *justifies* horizontal inequality. In the optimal contract of the standard principal-agent model, for example, identical workers making identical decisions generally receive unequal compensation.[3] Winter (2004) extends this result to multiple-agent settings and ex ante inequality, showing that identical agents must be offered different *contracts* in any optimal mechanism.[4] If, however, *non-pecuniary* inequality incentives affect firm-relevant behavior and are amenable to modeling, being predictable and non-trivial, then their omission from traditional wage-setting models is unwarranted. If, in

[3] In equilibrium, the optimal contract ties the agent's earnings to outcomes which are, in turn, only stochastically related to the agent's (effort) decisions. Consequently, identical employees exerting identical effort under this optimal contract will end up receiving different realized compensation levels.

[4] Tournament incentives may obviously also induce substantial inequality among similar workers (e.g., Gibbons 1987; Hopkins and Kornienko 2009, 2010; Lazear 2000; Stiglitz 1975). Among dissimilar workers, when types are unobservable inequality may be justified as allowing high-ability workers to credibly signal their ability (Spence 1973). Vertical inequality also enjoys a wide array of theoretical justifications, but this is not the focus of our study (Cullen and Perez-Truglia 2018; Prendergast 1999). For recent overviews of the literature on incentives in personnel economics and in economics more generally, see Lazear and Oyer (2007) and Lazear (2018).

addition, these (non-pecuniary) incentives vary predictably across contexts, then addressing their omission may improve our understanding of optimal compensation structures and serve as a point of departure into the broader endeavor of creating a theory of wage-setting where context matters.

Guided by our proposed channel, described in detail in Section 2, we conduct an experiment in which we randomly assign initial tasks and initial-task wage inequality. We then reveal or conceal aspects of initial-task inequality and observe how this experienced inequality affects *subsequent* behavior. Initial tasks vary in the extent to which performance depends primarily on ability ("ability-intensive") rather than primarily on effort ("effort-intensive"). This first contextual factor captures a decades-long and ongoing shift in the labor market toward occupations intensive in "nonroutine cognitive tasks" and away from occupations intensive in "routine noncognitive tasks" (Autor, Katz, and Krueger 1998; Autor, Levy, and Murnane 2003), respectively. A second contextual factor we vary is the type of initial-task inequality revealed to workers, wage inequality or income inequality. We incorporate this second factor to mirror a dichotomy in the economics literature on fairness. One influential class of models, distributional social preferences, is consequentialist in nature and inserts some aspect of the distribution of material outcomes into the utility function (Bolton and Ockenfels 2000; Charness and Rabin 2002; Fehr and Schmidt 1999). Other models are non-consequentialist and allow fairness judgments to depend on other aspects of the environment (Bohnet and Zeckhauser 2004; Butler and Miller 2018; Charness and Dufwenberg 2006; Dufwenberg and Dufwenberg 2018). Incorporating the distinction between consequentialist and non-consequentialist fairness by varying the type of inequality revealed may help separate the predictions of our proposed channel from the predictions of other influential classes of models.

As a preview of our results, we find that revealing initial-task *wage* inequality has substantial and context-dependent effects on subsequent behavior, and that these effects are concentrated on the disadvantaged (low-wage) employees. When the initial task is ability intensive, low-wage employees *increase* both the quantity and the quality of their subsequent production but *reduce* their pro-social behavior. When the initial task is effort-intensive, on the other hand, neither wage nor income inequality induce any discernible effect on the subsequent behavior of low-wage or high-wage workers. As we will argue later, the contextual variation we observe across the ability intensiveness of the initial task works against most existing explanations – distributional or otherwise – but is a prediction of the novel channel we propose. Attitudinal evidence from a post-experiment survey also provides support for our channel.

Overall, our results demonstrate that non-pecuniary inequality incentives are not always (weakly) negative and may sometimes improve firm-relevant behavior, such as effort-provision but are context-dependent. The primary contribution of our study is therefore to document predictable, qualitative and substantial contextual variation in non-pecuniary inequality incentives. A potentially important implication of our results is that optimal wage-secrecy policies may vary by industry and occupation or by the proportion of relatively low-wage workers in a firm on whom the effects we demonstrate are concentrated. In some contexts, revealed wage inequality may complement pecuniary incentives while in others wage secrecy may ameliorate the inequality-induced backlash documented in other studies. Empirically, contextual variation in non-pecuniary inequality incentives may help explain why there tends to be more wage dispersion net of human capital in occupations that rely more heavily on nonroutine and cognitive tasks (Van der Velde 2020; Visintin, Tijdens, Steinmetz, and de Pedraza 2015).

The remainder of the paper proceeds as follows. In Section 2 we outline our novel channel connecting experienced inequality to subsequent behavior and outline broad predictions. Next, we present our experimental design in detail. In section 4 we provide results and discuss various possible explanations. In the penultimate section, Section 5, we put our study in the context of closely related literature. In the final section we summarize and provide concluding remarks.

2 A Channel Relating Experienced Inequality to Subsequent Firm-Relevant Behavior

We take inspiration from research in economics and in (social) psychology to propose a novel channel relating experienced inequality to subsequent firm-relevant behavior. The conceptual heart of this channel is cognitive dissonance (Festinger 1957), the psychological distress associated with holding internally inconsistent beliefs. Avoiding cognitive dissonance induces often-subconscious cognitive and psychological processes aimed at reconciling inconsistencies through such stratagems as self-deception and reality denial. Economists have long recognized the potential economic consequences of cognitive dissonance, particularly for employee behavior (see, e.g., Akerlof and Dickens 1982).

Continuing in this tradition, we focus on cognitive dissonance arising from the inconsistency between a particularly relevant "motivated" belief and experienced compensation inequality. Motivated beliefs are beliefs that fulfill important psychological or functional needs, are protected through various self-deception mechanisms including reality denial, and may respond to the social and economic environment (Bénabou and Tirole 2016). The motivated beliefs we single out are "Just World Beliefs," hereafter JWBs, which refer to a well documented, universal and deep-seated *need* to believe that people generally get what they deserve (Benabou and Tirole 2006; Lerner 2013, 1965).

Highlighting the relevance of JWBs for incentive theory, Lerner (1965) asserts in his seminal JWBs study that it "... is functional as well as less dissonance arousing (Festinger, 1957) for a person to believe that his efforts will lead to a desired outcome. Few people would engage in extended activities if they believed that there were a random connection between what they did and the rewards they received." In light of this assertion, we conjecture that horizontal compensation inequality is inconsistent with JWBs except in particular contexts. One context where inequality is consistent with, and may be even demanded by, JWBs is the situation where awards are allocated according to merit. Consequently, in order to avoid cognitive dissonance there will be an automatic tendency to recast inequality as stemming from merit-based principles when possible. This reality-denial stratagem is facilitated by contextual cues, including contextual cues providing a superficial resemblance between the actual situation and (ideally) meritocratic situations. As reality denial can be thought of a relatively extreme way to reconcile JWBs and inequality, this tendency will be strongest where inequality would be the most distressing, i.e., among those disadvantaged by inequality. A perverse implication is that experienced inequality successfully reconciled with JWBs through such reality-denial then becomes additional evidence in favor of a just world, thus *strengthening* JWBs.

Behavioral implications of strengthened JWBs are suggested by (Benabou and Tirole 2006), which we adopt. There, the authors argue that strong JWBs give rise to an "American equilibrium," characterized by a strong work ethic and an emphasis on self-reliance, while a European equilibrium associated with weak JWBs is characterized by the opposite tendencies. Heuristically, if people generally get what they deserve and work deserves to be rewarded, then JWBs act as an intrinsic incentive to work. On the other hand, if others get what they deserve then the incentive for other-regard is weakened – others' outcomes are subjectively perceived to be less dependent on one's own decisions.

2.1 Three Broad Predictions

Based upon the novel channel just proposed, we can make three broad predictions relating our experimental factors to post-experienced-inequality behavior. For each of these predictions, we use the term "channel-consistent" as short-hand for the consequences of strengthened JWBs mentioned above – an increase in intrinsic work incentives and a decrease in pro-sociality.

Our first prediction is context-dependence associated with initial-task type. We expect channel-consistent behavioral patterns to be more likely when inequality is associated with an ability-intensive task than when it is associated with an effort-intensive task. We speculate that the superficial resemblance of an ability-intensive task to meritocratic environments where compensation inequality is demanded by fairness (cf, Starmans et al. 2017) will facilitate the operation of our proposed channel. This speculation is informed by existing research documenting the activation of the necessary subconscious beliefs distortion processes in ability-intensive environments (Butler 2016; Lerner 1965). Specifically, the task in the original Lerner (1965) study was described to third-party observers, who then exhibited beliefs distortion, as an "intellectual task," while Butler (2016) extended this finding to the context of *own*, i.e., first-party, wage disadvantage and showed the absence of beliefs distortion in a similar, but effort-intensive, environment.

Broad Prediction 1*: channel-consistent behavior will be more likely when inequality is associated with an ability-intensive initial task than when it is associated with an effort-intensive initial task.*

Our second broad prediction is context dependence associated with the type of inequality experienced. We consider two types of revealed inequality, wage inequality and the inequality embodied by the earnings distribution (income inequality). We mentioned above that we expect channel-consistent behavior to be more likely among those who experience personal *misfortune* as the inconsistency between JWBs and reality is the most acute for the disadvantaged. Between the two types of disadvantage we consider, we expect revealed wage disadvantage to be more likely to induce channel-consistent behavior than revealed income disadvantage for several reasons. First of all, income inequality will provide some positive comparisons for almost all workers and this countervailing "good news" should make the need to deny reality less acute.[5] Second, income inequality should more readily admit alternative distress-reducing justifications by making salient the fact that some portion of inequality is determined according to relative performance, a merit-based principle consistent with JWBs. A third reason for our prediction is suggested by analogy to victim-blaming. A robust result in the JWBs literature is that victim-blaming is more severe, embodying more beliefs distortion, when the (third-party) observer has less control over the victim's misfortune (see, e.g., Lerner and Miller 1978). In our study, the observer and the victim are one and the same, the worker, and the worker has less control over the misfortune of being assigned a low wage (none) than over the misfortune of receiving low initial-task income (some). Consequently, wage disadvantage should be more likely to induce the reality-denying beliefs distortion upon which our channel relies and hence more likely to induce channel-consistent behavior than income disadvantage.

Broad Prediction 2*: revealing wage inequality will be more likely to induce channel-consistent behavior than revealing income inequality; conditional on revealing wage inequality, channel-consistent behavior will be more likely among the disadvantaged (low-wage) workers than among the advantaged (high-wage) workers.*

[5] The initial-task earnings distribution will provide positive comparison opportunities for any worker not in the lowest category, a countervailing force to the distressing news of being less well paid than some.

Our third broad prediction relates to the prospect of revealing multiple types of inequality. As just discussed, because income inequality may provide countervailing good news for most workers, may be partially justified as being (partially) allocated according to a fair principle (performance) and is partially under the control of the "victim," revealing income inequality to the disadvantaged should lessen their psychological need for the kind of reality-denying beliefs distortion central to our proposed channel. As a consequence, we expect that revealing income inequality in addition to wage inequality will weaken the effects of wage inequality and reduce the likelihood of channel-consistent subsequent behavior.

Broad Prediction 3*: revealed income inequality will weaken the effects of revealed wage inequality on subsequent behavior.*

3 Experimental Design and Procedures

We conducted a laboratory experiment involving 423 student "employee" participants (44% female; average age was about 21). All experimental sessions were conducted at the Rawls College of Business at Texas Tech University in the Spring and Fall of 2018. Participants were recruited from a pre-existing college-maintained subject pool. In total, we conducted 31 experimental sessions.

Our experiment consisted of eight treatments – two task types and four inequality-information conditions – implemented in a between-subjects design. In addition, each of these eight treatments featured two randomly assigned piece-rate wage schemes. Some factors were manipulated only across sessions to minimize the chances that participants realized there were different experimental treatments. We make clear below which factors these were. Each session lasted approximately one hour and yielded an average compensation of $15.01. Participant instructions for all parts of the experiment are provided in an Instructions Appendix.

Each treatment was comprised of three decision-making phases plus a purely informational phase. Phases 1 and 2 involved real-effort tasks. In Phase 3, participants were provided an opportunity to donate to charity. Phase 1.5 was the purely informational phase, which we denote by assigning it a fractional phase designation. In Phase 1.5, all participants learned of their own Phase 1 performance and experimental income while some participants, depending on treatment, received additional information about relative wages or incomes. The experiment was multi-modal in order to plausibly introduce (temporal and psychological) separation between decision-making phases.[6] Phase 1 and 1.5 were programmed in oTree (Chen, Schonger, and Wickens 2016), Phases 2 and 3

[6] Specifically, each decision-making phase had to be explained as it arose, introducing time separation and taking focus off of the previously completed phase.

were implemented using pen-and-paper, while Phase 4 used Qualtrics, an on-line survey service. We describe each of these phases in detail below and summarize the timeline of the experiment in Figure 1.

3.1 Phase 1: Initial Real-Effort Task

To induce wage and income inequality, participants first completed a real-effort task for a specific piece-rate wage. As part of the experimental design, we consider two distinct tasks, which differ in the extent to which performance credibly depends on skill or ability versus effort alone. This factor was manipulated across sessions only: each session featured only one of the two tasks. At the outset of the experiment, we randomly determined which sessions would receive which task.

While we recognize that performance on any task takes some effort and some ability, we constructed our tasks to lie on opposite ends of the effort-ability spectrum. One of these two tasks, which we refer to as the "ability-intensive" task, consists of 48 Raven's Advanced Progressive Matrices. The other task, referred to as the "effort-intensive" task, asked participants to count the number of particularly easily spotted typographical errors, described below, in each of 48 sentences. Sentences were displayed one-at-a-time. Each participant performed only one of these tasks and did not know of the existence of the task they were not assigned.

Raven's matrices are multiple choice questions in which respondents select a figure that best completes a given pattern from among eight alternatives. Raven's matrices are designed to measure the test taker's reasoning ability, which is considered to be an important component of general (fluid) intelligence. We convey this interpretation to our participants in the instructions preceding the ability-intensive task and therefore consider it a plausible assumption that participants perceived performance on the task as *primarily*, although obviously not solely, depending on (cognitive) ability.

We constructed the effort-intensive task to be as parallel as possible to the ability-intensive task. Like the ability-intensive task, the effort-intensive task consisted of a sequence of 48 multiple choice questions, each with eight possible responses. Each of these questions asked the respondent to count the number of typos appearing in a particular sentence, with valid responses ranging from "none" to "7 or more." We deliberately constructed the sentences to contain only typos that were easily recognizable: duplicated words, missing words, or numbers that replaced letters. Our reasoning was that spotting and counting these kinds of errors requires little (reading) ability but a reasonable amount of effort. We conveyed this interpretation to our participants. Consequently, our maintained assumption is that participants perceived that performance on the effort-intensive task depended *primarily* – although obviously not solely – on effort.

Despite our efforts to make the ability-intensive and effort-intensive tasks as parallel as possible, they appear quite different on a computer screen. The ability-intensive task involves a grid of images, while the effort-intensive task features primarily text. Because of this, we decided that across-session randomization was necessary to minimize the possibility of participants realizing there were different tasks and hence different treatments.

Performance on each of the tasks was compensated through a piece-rate wage scheme. To induce explicit wage inequality we implemented two different piece-rate wages:

Low Wage: \$0.20 for each correct response; \$0.00 for each incorrect response.

High Wage: \$0.30 for each correct response; \$0.10 for each incorrect response.

Randomization into pay scheme occurred within-session. Each participant in each session was equally likely to be assigned each piece-rate wage. During Phase 1, participants only knew about their own piece-rate wage and did not know of the existence of the wage they were not assigned. This eliminates any scope for relative pay concerns to affect Phase 1 behavior. We also note that, conditional on responding, the two pay schemes feature the same expected marginal (monetary) incentives: $\$0.20 \times Prob(correct|respond)$.[7] Interpreting *performance* as the number of correct responses, this design choice mitigates the possibility of large performance differences due solely to monetary incentives which may otherwise introduce an unnecessary confound into the attribution of subsequent behavior to non-pecuniary inequality incentives.[8]

Prior to beginning their assigned task, participants were provided with a description of it and also informed of their own piece-rate wage. As part of this description they were informed that there would be a 10-minute time limit, with a timer appearing on-screen, to attempt as many of the 48 possible questions as they wished and that the phase would last the full 10 minutes regardless of their progress.

3.2 Phase 1.5: Information Revelation

After 10 minutes had elapsed, participants were informed of their own performance and their own income from the task. In addition to these basic pieces of information, we exogenously varied

[7] From an ex ante perspective, where we include the decision to respond at all as a choice variable, the marginal incentives would be $\$0.20 \times Prob(correct|respond)$ for low-wage workers and $\$0.20 \times Prob(correct|respond) + \$0.10$ for high-wage workers. Thus, there is a higher marginal incentive to *respond* for high-wage workers, but not a higher marginal incentive to *respond correctly* conditional on responding. We interpret responding correctly as "performance" and hence the latter as "performance incentives," which are identical across wage levels.

[8] As outlined below, a naïve application of Equity Theory (Adams 1963) could ostensibly justify a relatively low wage as being fair on the basis of relatively low performance even if relative performance differences were induced by differences in marginal incentives. Our attempt to neutralize this, and related, confounds by implementing wage schemes with identical marginal performance incentives separates us, to the best of our knowledge, from all of the existing experimental literature on wage inequality other than (Butler 2016).

the information participants received about the existence of the two possible pay schemes and the distribution of income in their session.

Within each session each participant was equally likely to be informed or to not be informed that there were two possible piece-rate wages. Conditional on being informed of these multiple possible wages, participants were told that each of the wage levels had been equally likely to be assigned. We think of revealing the existence of different piece-rate wages as implementing "Visible **Wage** Inequality" and concealing this information as implementing "**Invisible** Wage Inequality."

The other domain of information we manipulated was relative (experimental) income. Specifically, in the "Visible Income Inequality" condition, participants were provided with a binned relative frequency chart describing the distribution of earnings on the Phase 1 task in their specific session. In the "Invisible Income Inequality" condition, participants did not receive this chart. Through this manipulation we reveal comprehensive information about the income inequality within a session, allowing participants to locate themselves within the session's income distribution.

Because the Visible Income Inequality condition involves a large chart appearing on participants' screens, to minimize the possibility of participants discovering there were multiple treatments we varied this factor only across sessions. That is to say, either all participants in a particular session were assigned to the Visible Income Inequality condition or all participants in a particular session were assigned to the Invisible Income Inequality condition. Assignment of sessions to each of these two conditions was random and determined before any sessions were conducted.

Our experiment is identical for all participants after Phase 1.5, so that we we can summarize our experimental design as consisting of a 2 (wage inequality visibility) x 2 (income inequality visibility) x 2 (initial task type) full factorial design with eight distinct treatments (Table 1). Within each treatment there were also two randomly assigned wage levels, high or low. It is worth emphasizing that because information about wage and income inequality was provided to participants only after they *completed* the Phase 1 task, if at all, initial-task performance cannot be affected by these manipulations. In the absence of this design feature, one worry is that part of the reaction to inequality may have involved altering initial-task performance which would then have introduced an additional confound in the attribution of subsequent behavior. As one possible example, if low-wage workers had responded to their (known) wage disadvantage by withholding initial-task effort, then Equity Theory (Adams 1963) which relies on a ratio of inputs (effort) to rewards (compensation) might judge a low wage as fair *because* of low-wage workers' low effort.

Table 1: *Summary of Treatments*

Initial Task	Wage Inequality	Income Inequality	N
Ability-Intensive	Invisible	Invisible	56
Ability-Intensive	Visible	Invisible	58
Ability-Intensive	Invisible	Visible	49
Ability-Intensive	Visible	Visible	48
Effort-Intensive	Invisible	Invisible	54
Effort-Intensive	Visible	Invisible	53
Effort-Intensive	Invisible	Visible	52
Effort-Intensive	Visible	Visible	53

Notes: [**1**] As described in the text, the Ability-Intensive initial task involves 48 Raven's Matrices, commonly thought to be a culture- and language-free measure of (fluid) general intelligence. We use the Advanced Progressive Matrices which are appropriate for our college population rather than the Standard Progressive Matrices designed for a more general population. [**2**] The Effort-Intensive task involves 48 simplified proofreading questions. Each question involves one sentence in which we randomly inserted simple forms of typographical errors and asked participants to count the number we inserted. [**3**] Task format was constructed to be as similar as possible; each item of each task involved selecting the correct answer from among eight choices, and was approximately the same size and format on participants' computer screens.

3.3 Phase 2: Scantron Task

The Phase 2 task was another incentivized real-effort task designed to be ecologically valid for our student population. Participants were asked to prepare multiple versions of a Scantron answer key, corresponding to different versions of a statistics exam. They were each provided with a mock exam closely resembling an actual exam administered in a statistics class previously taught by one of this study's authors, as well as a master answer key for the original (mock) exam. Each participant also received a sheet listing the fifteen desired answer key versions (labeled generically as Version A – Version O), and fifteen Scantrons with which to create these versions by bubbling letters according to the version sheet.[9] The Phase 2 task was again timed: participants had 15 minutes to complete as many answer key versions from the list as they could. Before beginning, they were instructed that when the 15 minutes had expired they would report how many Scantrons they had completed and would be paid, *based solely on this self-report*, $1 for every completed Scantron. We only asked about the quantity produced, with no mention of quality.

By basing Phase 2 payments solely on self-reported productivity we make the *actual* quantity

[9]To facilitate statistical identification, all participants received the same exam copy, master answer key and version variations. The versions were created by taking random permutations of question orderings on the master key. To enhance ecological validity, through opacity in our instructions we created an environment where participants likely perceived the task as regular economically valuable work (Falk and Ichino 2006), as creating different versions of exams is a common practice in large introductory courses. At the same time, to avoid deception participants were not told anything about the intended use of the answer keys they would be preparing. They were simply informed that they were to make answer keys for different versions of the exam based on the exam master key provided.

and quality of Scantrons produced depend solely on intrinsic motivation, i.e., "work ethic." At the same time, tying compensation exclusively to self-reported quantity of Scantron production implements an ecologically valid form of a by-now standard paradigm for the investigation of anti-social behavior: lying about the privately observed outcomes (die rolls, coin flips) for financial gain (Fischbacher and Föllmi-Heusi 2013).

In order to make the latter analogy complete, we took great pains to ensure there was as much anonymity in our self-reported production design as in the standard die-rolling paradigm. In particular, immediately after the 15-minute time limit had expired, we had participants themselves take all of their Scantron forms, completed or not, to an unmonitored box at the back of the room. Because there was only one experimenter, who remained at the front of the room, it was apparent to participants that neither the quantity nor quality of their Scantrons was observable. Moreover, it was only after returning to their carrels that participants filled out an anonymized payment slip. On the slip, participants noted only their carrel number and the quantity of Scantrons they produced. It was also conveyed implicitly to participants, by the experimenter's position in the room and by the logistics of the experimental procedures, that actual Scantron-task productivity could only be assessed after sessions were completed and could never be attributed to any particular name, but rather only to participants' carrel numbers, and that, moreover, we could never match names to carrel numbers.[10] We paid using cash in envelopes marked with code numbers which could not be linked to participant names, since we never recorded who sat where. Participants retrieved their code-marked envelopes from a table at the beginning of the next Phase, rather than directly from the experimenter and, as an additional measure of anonymity which is uncommon, we also did not have participants sign receipts. We argue that this provides essentially the same level of anonymity as commonly used private die-rolling designs (Abeler, Nosenzo, and Raymond 2019; Dufwenberg and Dufwenberg 2018; Fischbacher and Föllmi-Heusi 2013; Gneezy, Kajackaite, and Sobel 2018; Mazar, Amir, and Ariely 2008), while still being familiar and appropriate for our subject pool.[11]

There are two primary advantages of our Phase 2 design over the more common designs which infer lying at the aggregate level. First, we can measure actual quantity and quality produced,

[10] With only one experimenter in the room, it would have been logistically impossible at the time to examine self-reports for accuracy, or Scantrons for quality, which would have been apparent to participants. In particular, the experimenter did not approach the back of the room, where the completed Scantrons were delivered, while the participants were still in the room.

[11] Our logic is as follows. In the typical dice-rolling experiment, lies can be identified at some aggregate level – say, the session level. Experimenters typically know who participated in which sessions, so that an experimenter conducting a dice-rolling experiment would know that John Smith certainly participated in a session where lying occurred. Consequently, the experimenter would know that with some probability John Smith told a lie. In our setting, we would know that somebody seated at carrel 33 certainly told a lie and that, with some probability, the person seated at carrel 33 was John Smith. Therefore, we would know that with some probability John Smith lied. Overall, therefore, both designs feature probabilistic personal identification of lying so that while there may be a difference in the probability with which liars are identifiable, this difference is a matter of degree rather than of kind.

which would be equivalent to observing the actual realized distribution of dice-rolls. Second, we can ex post link decisions, self-reports, demographics and attitudes at the (anonymized) individual level across all phases of the experiment. These advantages allow us to examine two types of anti-social behavior at the individual level. One type is outright lying about the number of Scantrons produced, which would be akin to misreporting a die-roll. The second type is more subtle and does not actually involve lying. Specifically, participants could "shirk" by reducing the accuracy of the Scantrons they produce. More of these lower quality Scantrons could be produced in the allotted time, allowing participants to increase their Phase 2 compensation *without* actually lying – we never asked them to report any quality measure.[12] To the extent that employers care about production quality, such shirking benefits the employee at the employers expense and may therefore be considered anti-social.

3.4 Phase 3: Charitable Giving Opportunity

At the beginning of Phase 3, participants received their cumulative earnings from the Phase 1 and Phase 2 tasks, in cash, in an envelope placed on a table. Upon retrieving their envelopes and returning to their carrells, they were then informed that they could anonymously donate to charity any amount they wished out of these earnings. To donate, they would simply leave some of their cash in their envelope along with a slip specifying how to divide the cash between two specific charities.[13] The two specific charities they could allocate money to were: South Plains Food Bank, a local charity; or the American Red Cross, a national charity.

For credibility, participants were informed that they would receive by email at the conclusion of the study a link where they could view a receipt for the total amount of money donated to each charity across all participants. For anonymity, participants were asked to leave their envelopes at their carrel when they left the experiment rather than hand them to the experimenter. As an added anonymity measure, even participants who chose to donate nothing were asked to place a completed donation slip in their envelope and leave the envelope at their carrel. Together, these procedures should ensure that donations were anonymous to the experimenters as well as to participants' peers.

We include this phase to provide an ecologically valid opportunity to behave pro-socially. Students are likely to face opportunities throughout the school year to donate to charities, particularly the well known and familiar charities we describe.

[12] Alternatively, they could produce a given number of Scantrons with less effort, improving their compensation-to-effort ratio.

[13] We did not specifically limit donations to the cash they received in the envelope. That is to say, they could have contributed cash they brought in with them, although we have no evidence that anybody did.

Figure 1: *Timeline*

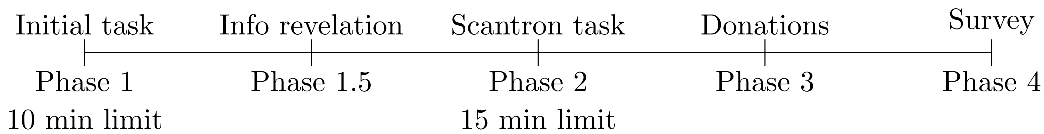

3.5 Phase 4: Post-Experiment Survey

After all decision-making phases were completed but before leaving their experimental session, all participants completed a post-experiment survey. The survey was anonymous and conducted using Qualtrics. We chose to use Qualtrics instead of programming the survey in oTree, as we did for the Phase 1 task, to introduce further mental separation between the various phases of the experiment. Completing the survey would draw focus to the computer screen and away from previous tasks. On the post-experiment survey we gathered general demographic data, various psychological measures and self-reported beliefs and attitudes on a variety of relevant topics. Importantly, we collected a measure of JWBs by asking respondents to report how much they agreed with the statement: "people generally get what they deserve." Responses were collected on a 7-point scale, ranging from 1="totally disagree" to 7="totally agree."

3.6 Matching measures across phases

Well after each session was completed, we matched outcome measures from all phases of the experiment.[14] Matching was accomplished by carrel number, which was automatically recorded for the computerized portions of the experiment (Phase 1 and the post-experiment survey) and, for the other phases of the experiment, noted on materials. It is worth mentioning that it is actually not possible for us, the experimenters, to match participants' names to carrel numbers or even payment amounts: we did not record who sat where, nor did we ask participants to fill out payment receipts.

[14] We should note that we incorporated two "distractor" tasks into the experimental design. The first of these tasks occurred between the end of the Phase 1 task and the revelation of inequality in Phase 1.5, and consisted of filler questions to ensure that participants who finished the Phase 1 task before time expired were occupied. The filler questions were uniformly randomly assigned and asked about the participant's political affiliation, gender or what service they used to watch television (e.g., cable, satellite, etc.). We do not analyze these questions, but instead simply control for the type of question participants were asked in our formal econometric estimates. The second distractor task was inserted between Phase 3 (charitable donation decision) and Phase 4 (post-experiment survey) and consisted of two questions eliciting intertemporal preferences. This (ostensibly) morally neutral distractor task introduces psychological distance between the morally charged charitable donation opportunity and the post-experiment survey. We, again, do not analyze responses to this second distractor task but note that preliminary analyses revealed no significant relationship between our experimental factors and the time-preference responses.

4 Results

We conducted chi-squared tests of independence across our treatments for each of the limited set of demographic measures we collected: age, gender, family income and GPA.[15] Only one of these tests rejected the null at anywhere near a conventional significance level, providing some reassurance that randomization into treatments was successful.[16] As a precaution, however, we control for available demographics in all of our econometric estimates. Before presenting those estimates, we start by examining the raw data.

In Table 2 we report means and standard errors for our primary outcome variables, by treatment and wage level. This exercise will allow us to introduce the outcome measures that our more formal estimates focus on. The first two columns report initial-task income as well as the number of correct initial-task answers (score), a measure of performance. These columns reveal that within-task performance was similar across wage levels and treatments. The substantial income differences across wage levels are therefore largely due to wage differences. The absence of within-task performance differences is consistent with performance being determined by, in our preferred interpretation, the identical marginal monetary incentives provided by our high and low piece-rate wage levels.[17] Average performance was consistently substantially higher (about 40%) on the effort-intensive initial task than on the ability-intensive task. One plausible interpretation is that (cognitive) ability more severely constrains performance on the ability-intensive task than on the effort-intensive task so that the same level of effort induced by the same marginal performance incentives (across tasks) produces lower performance on the former task.[18]

The other decision-making phases all occur after treatment-specific inequality revelation, so our experimental factors have the potential to affect behavior. For Phase 2, we report the average number of Scantrons actually produced (quantity), which hypothetically ranges from 0 to 15 and in the data ranges from 1 to 15. In the column labeled "quality" we report the average proportion of correctly bubbled items (out of 20) across all submitted non-blank Scantrons.[19] Finally, we

[15] Family income and GPA were both self-reported. From these self-reports we construct a three-category measure – high, medium and low – for income and GPA.

[16] The p-values associated with the chi-squared tests were: gender, $p = 0.702$; age, $p = 0.288$; family income, $p = 0.797$; GPA, $p = 0.085$.

[17] As discussed above, the marginal monetary incentives to respond at all, as opposed to marginal monetary *performance* incentives, are higher for high-wage workers since their piece-rate wage includes \$0.10 for *any* response. This results in high-wage workers responding to 1.74 more effort-intensive task questions and 3.98 additional ability-intensive task questions. Chi-squared tests suggest these are not significant increases (ability-intensive, $p = 0.339$; effort-intensive, $p = 0.778$), while some other tests yield significance for both initial tasks (simple t-tests, $p < 0.01$). In any event, the increase in the total number of responses without a concomitant increase in the total number of correct responses suggests that high-wage workers may have mindlessly clicked through later questions, or earlier more difficult questions, to collect the associated \$0.10 fee.

[18] Alternatively, the proofreading task could have been more enjoyable – although we find this interpretation unlikely.

[19] More precisely, we restrict attention to Scantrons on which a participant bubbled at least one item and compute

report a measure of anti-social behavior in the column labeled "lying." This measure is a simple indicator taking the value of one when self-reported quantity did not match actual quantity. An indicator function is appropriate because almost all misreports involved over-reporting production by *exactly one* Scantron so there was very little intensive-margin variation. At the same time, the overwhelming prevalence of *over*-reporting justifies our interpretation of the indicator as a measure of lying.[20] In the last column, we report the average proportion of participants' total Phase 1 and Phase 2 earnings donated to charity in Phase 3.

There are a few prominent patterns in the data. Production quality was surprisingly high despite the lack of any monetary quality incentives. At the same time, lying was surprisingly rare given the substantial monetary incentives and the availability of "small" lies (Gneezy 2005).[21] In light of our predictions, raw means already provide some initial positive evidence. Focusing on low-wage workers in Ability-Intensive treatments and comparing behavior when only wage inequality is revealed to behavior when no inequality is revealed – the first two treatments reported – suggests that, in line with our second prediction, experienced wage disadvantage increases subsequent effort (Phase 2 quantity and quality) and reduces pro-sociality (Phase 3 donations).[22]

4.1 Econometric estimates

While descriptive statistics are useful in providing an overview of the data, econometric models will allow us to control for potential nuisance factors, such as unintended demographic variation across treatments, and more clearly reveal the effects our experimental factors had on behavior. To account for spurious within-session correlations, in all of our estimates we cluster standard errors at the session level.[23] Our analysis will generally investigate the broad predictions we made in Section 2.1.

Our first broad prediction was that the context provided by an effort-intensive initial task was

the proportion of correct answers for each of these Scantrons. On such Scantrons, non-bubbled items are counted as incorrect. We then take the average of these proportions for each participant.

[20]That is to say, unintentional mistakes, a more generous alternative interpretation, would not be expected to produced the systematic upward bias we observe. An upward bias would also not be explained by workers incorrectly characterizing partially completed Scantrons as complete, as any Scantron with even one item bubbled is included in our quantity measure against which our lying indicator is defined.

[21]On average, participants could have doubled their Phase 2 earnings – the average self-report was 7.66 instead of the theoretical maximum of 15 – and increased their total experimental earnings by approximately 50% while, at the same time, saving themselves the considerable effort of filling out any Scantron answer keys at all by fully lying.

[22]Simple one-sided t-tests incorporating the directional hypotheses provided by our predictions further reveal some statistical support for our proposed channel even in the raw data without demographic controls. The tests yield p-values of $0.095, 0.064$ and 0.046 for quantity, quality and donations, respectively. The pattern with respect to pro-sociality is made more striking by the fact that in seven out of the eight treatments it is the low-wage workers that donate *more* to charity.

[23]In the Online Appendix we report in analyses using bootstrapped (session-)clustered standard errors to allay concerns about too few clusters (see, e.g., Cameron, Gelbach, and Miller 2008). None of our significance patterns were changed by these alternative standard errors.

less likely to give rise to the JWBs-strengthening processes we conjectured, and hence less likely to yield behavioral patterns consistent with that process (increased work ethic, decreased other-regard). To test this prediction, we split our data by (Phase 1) initial-task type and estimate our models of behavior as a function of experimental factors on each of these subsamples separately. Splitting the data is motivated by two considerations. First, the split will allow us to more conveniently report coefficients associated with all of our experimental factors and the interactions among them. Second, since randomization into initial-task type occurred only across sessions this split is conceptually clean: no worker performing the ability-intensive task could have affected any individual assigned to the effort-intensive task.[24]

In Table 3 we report OLS estimates restricted to observations from Effort-Intensive treatments and in Table 4 we report the same models restricted to Ability-Intensive treatments. As explanatory variables, we include indicators for each of our experimental factors – high wage, visible wage inequality, and visible income inequality – as well as all interactions among these indicators. All estimates also include demographic controls (omitted for readability).[25] Finally, to account for the possible reactions to one's own initial-task earnings, such as to falling short of an ex-ante subjective expected own-earnings target, we control for Phase 1 (initial-task) income where possible.[26]

The picture that emerges is easily summarized. In our Effort-Intensive sessions (Table 3) we find no evidence for any significant effect of any of our experimental factors on any subsequent behavior. In stark contrast, our Ability-Intensive sessions (Table 4) reveal substantial effects of at least some of our experimental factors on almost all subsequent behavior. We summarize this broad contextual variation as our first result.

Result 1: *consistent with our first broad prediction, we find little evidence that revealing inequality affects subsequent behavior when inequality is associated with an effort-intensive initial task, but substantial evidence when the initial task is ability-intensive.*

For the remainder of this section, we focus exclusively on the estimates from our Ability-Intensive sessions (Table 4). Our second broad prediction was context-dependence with respect to the type of inequality experienced. We argued that inequality which is the least explicable and which induces the most acute need for rationalization is the most likely to give rise to the JWBs-strengthening channel we propose. In terms of our experimental factors, this should imply

[24]For instance, the relative performance information for Ability-Intensive sessions never included any Effort-Intensive session component.

[25]Demographic controls were almost always non-significant, so we lose little from their omission.

[26]To address the concern that Phase 1 income depends partially on the randomly assigned wage level, rather than including yet another interaction coefficient we conducted the same estimates using initial-task performance (score) instead of initial-task income which, as we saw, exhibited little variation with wage level. None of our findings, including patterns of statistical significance, were changed by this alternative specification.

two patterns: wage inequality should be more likely to induce behaviors consistent with JWBs-strengthening; and the effects of wage inequality should be more concentrated on low-wage workers, as their experienced disadvantage should produce a more acute need to rationalize inequality than high-wage workers' experienced advantage.

We examine the type of inequality first, and begin by focusing on low-wage workers. The main effects of revealing wage inequality (alone) to low-wage workers are captured by the coefficients associated with our *Visible Wage Inequality* indicator (row 2 of Table 4) while the main effects of revealing income inequality to low-wage workers are captured by the coefficients associated with *Visible Income Inequality* (row 3). The contrast in the findings clearly support our second broad prediction. Experienced *wage* disadvantage produces channel-consistent behavior, significantly *increasing* the subsequent quantity and quality of production and significantly *reducing* pro-sociality (charitable donations). Experienced *income* disadvantage, on the other hand, produces no significant effects.[27] The effects of revealed wage inequality are substantial as well as statistically significant. The one-unit average increase in Scantron quantity constitutes roughly one-half (48.9%) of a standard deviation relative to the pooled Ability-Intensive data. Controlling for quantity, the 7 percentage point increase in Scantron quality also represents roughly one-half (45.5%) of a standard deviation in the pooled Ability-Intensive data. The 12 percentage point reduction in the proportion of experimental earnings donated to charity (column 5) represents 42.5% of a standard deviation of the pooled Ability-Intensive data. We summarize these patterns as our next result.

Result 2a: *consistent with our second broad prediction, the main effects of revealed wage inequality on low-wage workers are significant, substantial and consistent with our proposed channel, while the main effects of revealing income inequality alone to low-wage workers are all non-significant.*

The second half of our second broad prediction concerns high-wage workers. We predicted that channel-consistent behavior would be less likely among high-wage workers than among low-wage workers. We consider first wage inequality and then income inequality. The main effects of revealed wage inequality on high-wage workers are captured by the sum of the coefficients on *Visible Wage Inequality* and their interactions with the high-wage indicator (row 2 plus row 4). The signs and magnitudes generally suggest a substantial weakening if not outright elimination of the effects of revealed wage inequality documented among low-wage workers: the point estimate of the effect of

[27] Outside of the context of our proposed channel, the contrast in the patterns between wage and income inequality are somewhat surprising, since the income distribution conveys relative (dis)advantage information similar to the information conveyed by wage inequality: the correlation between being assigned a low wage and being in the bottom half of the income distribution is nearly one and highly significant ($\rho = 0.868, p < 0.01$)). We examine behavior by position in the income distribution later, in the context of examining alternative explanations (Section 4.3).

revealed wage inequality on Scantron quantity is essentially zero (column 2, $1.00 - 1.17 = -0.17$); the effect on Scantron quality is less than half as large for high-wage workers as it is for low-wage workers, and non-substantial;[28] the point estimate for pro-sociality is approximately zero (column 5, $-0.12 + 0.11 = -0.01$). Turning next to the main effects of revealed income inequality, where we found little evidence for the operation of our channel among low-wage workers, we also find little evidence for our channel among high-wage workers. The only significant effect we find is with respect to subsequent Scantron quantity, where the interaction with the *High Wage* indicator is negative, highly significant and large enough in magnitude to produce a negative point estimate of the effect of revealed income inequality on Scantron quantity ($0.75 - 1.74 = -0.99$). Overall, we conclude that the patterns among high-wage workers are consistent with our second broad prediction, which we summarize as our next result.

Result 2b: *also consistent with our second broad prediction, we find channel consistent-behavior (increased effort, reduced pro-sociality) to be generally less likely and less pronounced among high-wage workers than among low-wage workers.*

Our final prediction was about the effects of revealing income inequality in addition to wage inequality. We argued that since this additional information may provide alternative ways to ex post rationalize inequality other than our relatively extreme proposed channel, it would mute the effects of revealed wage inequality alone. Because the evidence for our channel with regard to wage inequality is, as predicted, primarily present among low-wage workers, it is among these workers where we can expect to find that revealed income inequality *mutes* the effects. The additional effect of revealing income inequality can be found by considering the sum of the coefficients associated with *Visible Income Ineq* and the interaction term *(Visible Wage Ineq) x (Visible Income Ineq)*. The data reveal a pattern largely consistent with our predictions. The point estimates suggest that income inequality offsets the positive effect of wage inequality on low-wage workers' subsequent production quantity ($0.75 - 1.47 = -0.72$) and, although the interaction is not significant, quality ($0.05 - 0.08 = -0.03$) as well. We do not find an offsetting effect, however, on pro-sociality – the relevant point estimate is exactly zero ($-0.03 + 0.03 = 0.00$). We summarize this pattern as our final result.

Result 3: *partially consistent with our third broad prediction, among low-wage workers where*

[28]The point estimate of the effect on Scantron quality (column 3) is ($0.07 - 0.04 = 0.03$), which represents only 19.5% of a standard deviation in the pooled Ability-Intensive data.Some common guidelines using closely related effect-size measures such as Hedges-g suggest 0.25 standard deviations as a lower-bound on substantial effects. This is, for instance, the guidance provided by the US Department of Education's Institute of Education Sciences in their *What Works Clearinghouse: Procedures and Standards Handbook* available at https://ies.ed.gov/ncee/wwc/Handbooks.

revealing wage disadvantage increased the quantity and quality of production while reducing charitable donations, revealing income inequality in addition to wage inequality produced countervailing effects on Scantron quantity and quality but not on charitable donations.

Finally, we turn to anti-social behavior. We permitted lying for financial gain in order to provide workers with a wide array of possible responses to experienced inequality. While we made no predictions here, we can still use our data to shed light on some aspects of anti-sociality. In particular, an intuitively plausible conjecture is that pro- and anti-social behavior are two sides of the same coin, being opposite dispositions, and should therefore negatively co-vary. More pro-social people should be less anti-social and vice versa. To get a rough sense of this conjecture, we test whether our pro-social measure (*donations*) differs by our anti-social indicator (*lying*). We find no significant differences in the Ability-Intensive data, in the Effort-Intensive data, or in the pooled data (two-sided t-tests: $p = 0.756, p = 0.754, p = 0.750$, respectively).[29]

Summing up, our data support the idea that reactions to salient inequality exhibit substantial context dependence that is largely consistent with the novel channel we proposed. Specifically, we show that salient wage disadvantage associated with an ability-intensive initial task induces predictable variation in subsequent behavior that embodies a stronger work ethic and an increased emphasis on self-reliance. These effects on subsequent behavior are weakened by the revelation of an additional form of inequality, income inequality, as this finer measure may permit other justifications or rationalizations that are less extreme and hence more readily adopted than reality denial. In stark contrast, when inequality is associated with an effort-intensive task we find no evidence for any effect on subsequent behavior.

4.2 Examining Just World Beliefs

We next turn from behavior to attitudes to examine JWBs directly. If the patterns we document in our data are to be explained with our proposed channel, some necessary conditions would be that in our Ability-Intensive treatments: i) revealed *wage* inequality strengthens JWBs among low-wage workers; ii) revealed *income* inequality does not strengthen JWBs among low-wage workers; and iii) revealed income inequality provides countervailing effects to revealed wage inequality. We investigate these necessary conditions using a straightforward measure of JWBs collected on our post-experiment survey. Participants were asked to report how much they agreed with the statement "people generally get what they deserve." Responses were collected on a seven-point scale ranging

[29] We also conducted more formal OLS estimates by inserting *donations* as a control into our estimates of *lying* together with all other controls from our main analyses (e.g., as in column 4 of Table 3). We omit these estimates out of space concerns, but the results were the same: non-significant and small-in-magnitude (negative) coefficients on *donations* in estimates using the Ability-Intensive data, the Effort-Intensive data and the pooled data.

from 1="totally disagree" to 7="totally agree."

In Table 5 we provide OLS estimates of JWBs as a function of our experimental factors in our Ability-Intensive treatments. All estimates include demographic controls, omitted for readability. In column 1, we only add controls for our experimental factors and interactions among them. We find that the positive relationship between revealed wage inequality and JWBs among low-wage workers is highly statistically significant and substantial in magnitude, so that our first necessary condition finds support. The coefficient on *Visible Wage Inequality* of 0.53 represents 0.37 standard deviations of the JWBs measure in the Ability-Intensive session data. Consistent with our second necessary condition, we find that the main effect of revealed income inequality on low-wage workers JWBs is non-significant. Finally, the negative point estimate of the effect of revealing income inequality in addition to wage inequality on low-wage workers' JWBs provides some (limited) support for our third necessary condition concerning. The point estimate of the effect of revealing income inequality in addition to wage inequality is negative, albeit small, suggesting that salient income inequality partially offsets the salutary effects of salient wage inequality on JWBs.

In subsequent columns, we add controls for: initial-task earnings (column 2); Phase 2 behaviors (column 3); and Phase 3 behavior (column 4). We add these behavioral controls to account for the possibility that since we only measure JWBs at the end of the experiment, they could have been affected by actions taken during the experiment. While this timing complicates causal attribution of variation in JWBs to our experimental factors, investigating how JWBs vary across our treatments may still provide suggestive evidence. Overall, we find that patterns are mostly robust to these additional controls: the main effect of revealed wage inequality is always positive and significant, while the main effect of revealed income inequality is never significant. The point estimate of the effect of revealing income inequality in addition to wage inequality, however, turns from small and negative (columns 1 and 2) to small and positive (columns 3 and 4). On balance, we conclude that the relationship between JWBs and our experimental factors are generally supportive of our proposed channel.[30]

[30] Although we found no evidence for our proposed channel in our Effort-Intensive data and consequently have no a priori expectations about the relationship between our experimental factors and JWBs there, for completeness in the Online Appendix (Table A3) we repeat the exercise above using our Effort-Intensive session data. One could expect the patterns there to reflect how inequality is processed *in the absence* of our proposed channel, perhaps according to more well understood phenomena such as the Fundamental Attribution Error in psychology. One version of this phenomena asserts that positive outcomes tend to be subconsciously attributed to one's traits or actions, while negative outcomes are attributed to the environment. In line with this alternative inequality-perception process, the estimates in the Effort-Intensive session data suggest that low-wage workers attribute their wage disadvantage to an unfair world (lower JWBs) while high-wage workers infer no association between the fairness of the environment and their (random) advantage, leaving their JWBs unaffected. The point estimate of the effect of revealed wage inequality on the JWBs of high-wage workers is essentially zero in all estimates.

4.3 Inconsistency with many common explanations

So far we have argued that our data are consistent with predictions derived from our novel channel. Here, we argue that the data are not consistent with obvious alternative explanations. We start with loss aversion (Kahneman and Tversky 1979). While we do not know what the reference point against which a loss would be defined is, a plausible assumption is that participants came into the experiment with earnings expectations. Given this assumption, workers who earn little on the initial task are more at risk of ending up falling short of their reference point by the end of the experiment and thus have a stronger incentive to increase earnings in subsequent tasks. If we make the additional assumption that workers are averse to lying (Dufwenberg and Dufwenberg 2018; Gneezy 2005; Gneezy et al. 2018) so that they cannot increase Phase 2 earnings by simply inflating their self-reported Phase 2 productivity, then these stronger earnings incentives among initial-task low earners could manifest themselves as an increase in Scantron quantity and a reduction in charitable donations.

We examine this lying-and-loss aversion explanation by reporting in Table 6 estimates of Phase 2 and Phase 3 behavior as a function of the worker's place in the initial-task earnings distribution. Because the initial task type should not affect earnings expectations formed before the experiment, we pool the data across task types.[31] As task type could have some effect on subsequent behavior through, e.g., fatigue or exhaustion, we insert initial-task type as a control. To ensure adequate sample size and allow flexibility in the effect of earnings position, we construct mutually exclusive initial-task earnings quintile indicators and include them as explanatory variables. We find that: low initial-task earners produce marginally significant *fewer*, not more, Scantrons and they lie about production quantity statistically significantly more; they also donate (marginally significantly) more, not less, to charity; we find no relationship between Scantron quality and initial-task earnings.[32] Not only are the patterns in subsequent behaviors generally opposite to the patterns in our main analysis, the data suggest workers are not averse to lying as a strategy to increase earnings. Overall, therefore, this exercise provides little support for the notion that lying-and-loss aversion can explain our results.[33]

A second class of alternative explanations is distributional social preferences (Bolton and Ockenfels 2000; Charness and Rabin 2002; Fehr and Schmidt 1999). In their standard formulations,

[31] Likewise, we do not include our experimental factors. In the lying-and-loss aversion story their effects would be spurious as, e.g., revealing that there were different wage levels should not have affected own-earnings expectations.

[32] We note that lying aversion would have no implications for Scantron quality as we never workers anything about it. If anything, we would expect a negative effect stemming from the quantity-quality tradeoff in production if workers increase Scantron quantity solely to increase earnings without lying.

[33] In an Online Appendix, we repeat this exercise examining data from each initial-task type, separately, using task-type specific earnings quintiles (Tables A4 and A5). Again, we find no evidence that relative initial-task earnings can explain our findings.

these models are purely consequentialist and posit essentially that preferences are a function of a social reference point defined by features of the earnings distribution. The most widely adopted models, inequality aversion (Bolton and Ockenfels 2000; Fehr and Schmidt 1999), define this social reference point in a two-person setting as being equality of one own and one's co-player's earnings, with the utility function being steeper in own-earnings below this point.[34] As with lying-and-loss aversion, it could be argued that if workers are strongly lying averse and, at the same time, inequality averse, the concomitant increase in Scantron quantity and reduction in donations we observe among low-wage workers could be explained by the increased marginal utility of own-earnings for those workers falling below the (social) reference level of earnings.

We investigate the lying-and-inequality aversion explanation by, again, estimating Phase 2 and Phase 3 behaviors as functions of initial-task earnings quintiles (Table 7). In these new estimates, we include an indicator for our "Visible Income Inequality" factor as well as interactions with the earnings quintile indicators. Our reasoning for including this experimental factor is that one might expect a social reference point to be more salient, and its effects more pronounced, when workers are provided with information about the initial-task earnings distribution. Without such information workers can obviously still form beliefs about the initial-task earnings distribution, so we do not restrict attention to the Visible Income Inequality sessions alone. We find that low (initial-task) earners tend to *reduce*, not increase, their quantity and quality of Scantron production. We find no relationship between initial-task earnings and charitable donations. We find significantly more lying among workers in the first and third initial-task earnings quintiles (either below the 20th percentile or between the 40th and 60th percentiles) and, more generally, among workers assigned the ability-intensive initial task. We find no main or quintile-interaction effects of revealing the initial-task earnings distribution. Overall, we conclude that the lying-and-inequality aversion story does not provide a compelling explanation for the patterns we observe in the data with regard to our experimental factors.[35]

The final alternative explanation we consider is Equity Theory (Adams 1963), which has become a commonly used explanation for reactions to compensation inequality (see, e.g., Breza et al. 2018). Central to Equity Theory is the concept of an equity ratio, defined as the ratio of a worker's

[34]In a two-person setting, the effect on inequality of increasing own-earnings can be unambiguously calculated. When own-earnings are less than the co-player's earnings, the selfish motive of increasing own earnings and the social motive of reducing inequality are aligned. Above this equality reference point, the selfish and social motives are in opposition, reducing the marginal utility from increasing in own-earnings. Distributional preferences models typically leave the reference group unspecified. We will assume for simplicity that individuals compare themselves to a representative, i.e., average, other experimental participant to mimic the two-person setting. Without such an assumption the effects of increasing own-earnings on the inequality embodied by the earnings distribution would require some functional form assumptions on how individuals perceive inequality.

[35]In the Online Appendix we report results using task-type specific earnings quintiles (Table A6). We again find little support for the lying-and-inequality aversion explanation for our main results.

perceived total own-compensation from the firm (money earnings or other benefits) to his or her own-contributions to the firm (talent, productivity). If this equity ratio is sufficiently far away from a subjective reference ratio, a worker may act to restore equity. Reducing effort or (perhaps nefariously) increasing compensation are potential equity-restoring reactions when the equity ratio is perceived to be too low. Increasing effort or reducing compensation (e.g., donating more to charity) may restore equity when the equity ratio is too high.

As with the above frameworks, Equity Theory by itself cannot explain the dual pattern we document among low-wage workers: increased productivity devoid of compensation implications is a way to lower a too-high equity ratio; withholding donations would be a way to raise a too-low ratio. Thus, one's own equity ratio would seemingly need to be simultaneously too high and too low. Assuming strong lying aversion may once again reconcile the quantity response: if one cannot lie and the compensation from producing a higher quantity improves the equity ratio (an additional assumption), then increasing production quantity and reducing donations may both serve to raise a too-low equity ratio. The *increase in quality* we observe, however, remains puzzling. Quality has no explicit link to lying aversion, as we never asked workers about it nor did we observe it until well afterwards. Moreover, a quantity-quality tradeoff even among eminently honest workers should entail a quality reduction among workers producing a higher quantity in the absence of extra *effort*. If both the increased quantity and quality stem from extra *effort*, however, then the most plausible interpretation in an Equity Theory framework is one of increasing one's own contribution to the firm, a strategy to rectify an initially too-high equity ratio, bringing us back to our original contradiction – a ratio that is both too high, addressed by increased contribution, and too low, addressed by withholding donations. Finally, we note that Equity Theory does not predict some of the context-dependence we observe without additional assumptions. It is not clear, for instance, why information about the income distribution induces little or no response while wage inequality does or why an Ability-Intensive context induces a behavioral response while an Effort-Intensive context does not. Overall, then, we conclude that Equity Theory, even augmented by lying aversion, does not provide a reasonable or useful explanation for our findings.

Wrapping up, we argued in this section that frameworks commonly used to explain reactions to inequality cannot provide a unified explanation for the patterns in our data: increased quantity and quality of production together with reduced pro-sociality. Moreover, none of these models predict the context dependence that we observe across the type of work being performed or across the type of inequality revealed without additional strong assumptions.

5 Closely related literature

Our study contributes to several closely related literatures. Our study is most directly related to the growing literature documenting non-pecuniary incentives of salient inequality (Abeler, Altmann, Kube, and Wibral 2010; Akerlof and Yellen 1990; Angelova, Güth, and Kocher 2012; Bartling and Von Siemens 2011; Bolton and Werner 2016; Bracha et al. 2015; Breza et al. 2018; Butler 2016; Card et al. 2012; Charness, Cobo-Reyes, Lacomba, Lagos, and Pérez 2016; Charness and Kuhn 2007; Clark, Masclet, and Villeval 2010; Cohn, Fehr, Herrmann, and Schneider 2014; Dube et al. 2019; Gächter and Thöni 2010; Gill, Prowse, and Vlassopoulos 2013; Godechot and Senik 2015; Greiner, Ockenfels, and Werner 2011; Gross, Guo, and Charness 2015; Hennig-Schmidt, Sadrieh, and Rockenbach 2010; Nosenzo 2013; Sseruyange and Bulte 2019). Almost all of these studies document (weakly) negative reactions to salient wage disadvantage using methodologies that run the gambit from laboratory, natural and field experiments to theoretical studies motivated by observational data. None of the studies, to the best of our knowledge, examine the *type* of work being performed. Most of the studies consider only what we would consider effort-intensive work. Sseruyange and Bulte (2019), for example, study the effects of wage inequality on a bean-sorting task in Uganda while Dube et al. (2019) focus on entry-level positions at a large US retailer.

We contribute to this literature by being among the first to explicitly consider how the behavioral consequences of inequality depend on the type of work the inequality is associated with, i.e., whether work is ability-intensive or effort-intensive. The only other papers we are aware of are Butler (2016), which provides the conceptual underpinnings of the current study, and Hart and Piff (2018) which can be re-cast in our framework as varying the ability-intensiveness of work, but which involves deception.[36] Considering the type of work being performed is important and timely. Personal and societal economic success increasingly depend on ability-intensive environments as labor markets around the world continue a decades long trend toward occupations requiring higher levels of education and generalized problem-solving capabilities (Acemoglu and Autor 2011; Albanesi, Gregory, Patterson, and Şahin 2013; Autor 2014; Autor et al. 2003; Center 2016; David and Dorn 2013). A key contributing factor is continued technical innovation and incorporation of data analytics into many jobs (Autor et al. 1998, 2003; Muro, Liu, Whiton, and Kulkarni 2017).[37] If the consequences of inequality differ qualitatively and predictably across this dimension, from the firm's perspective the optimal degree of inequality revealed naturally through pay structures

[36] The design of Hart and Piff (2018) involves deception in that they tell subjects that inequality is based on either: i) performance on a preferences survey, and hence ostensibly random; or ii) performance on a "knowledge" quiz, which in our terminology could be viewed as an ability-intensive task. In reality, inequality is randomly assigned. They find (as we do) positive non-pecuniary incentives of salient inequality.

[37] Autor et al. (1998) asserts that "... changes in job task content – spurred by technological change – may plausibly be viewed as an underlying factor contributing to recent demand shifts favoring educated labor."

or concealed through formal or informal pay secrecy policies may also vary across this dimension.

A second closely related literature focuses on so-called "motivated beliefs," i.e., beliefs that fulfill important psychological or functional needs, are protected through various self-deception mechanisms including reality denial, and may respond to the social and economic environment (Bénabou and Tirole 2016). Economists have long understood that the avoidance of cognitive dissonance (Festinger 1957) may cause employees to subconsciously alter their beliefs about the features of their work environment, thereby effectively altering incentives (Akerlof and Dickens 1982). We focus on a specific instance of motivated beliefs identified by previous research as relevant for economic behavior (Benabou and Tirole 2006), Just World Beliefs (JWBs). In the seminal JWBs study, in fact, Lerner (1965) asserted "It is functional as well as less dissonance arousing (Festinger, 1957) for a person to believe that his efforts will lead to a desired outcome. Few people would engage in extended activities if they believed that there were a random connection between what they did and the rewards they received." Benabou and Tirole (2006) argue that stronger JWBs are more likely to yield an "American equilibrium" characterized by work ethic and an emphasis on self-reliance, while Frank, Wertenbroch, and Maddux (2015) show that priming JWBs induces a preference for tournament-style compensation over egalitarian compensation.[38] Beliefs that rationalize systemic inequality in cross-sectional survey data are a central finding in the "System Justification Theory" literature (John T. Jost and Sullivan 2003; Jost, Banaji, and Nosek 2004).[39]

While the potential economic importance of JWBs are clear, this potential has not been realized for several reasons. First of all, even though the original Lerner (1965) study looked at perceptions of wage differences among similar workers, the large follow-on literature quickly turned away from this readily interpretable and common setting to study *extreme* injustices such as sexual violence or substantial physical injury to document "victim-blaming." Since the vast majority of these studies involve perceptions of *others'* misfortune and results exhibit considerable context-dependence, it is not obvious whether less severe instances of *own*-misfortune such as own-wage disadvantage reliably induce similar processes. The very few existing JWBs studies examining own-misfortune and firm-relevant behavior (lying for financial gain) provide mixed results: Wenzel, Schindler, and Reinhard (2017) document a positive correlation between pre-measured JWBs and lying; Schindler, Wenzel, Dobiosch, and Reinhard (2019) manipulate JWBs through vignettes involving extreme own-misfortune and find that JWBs reduce lying for some tasks and vignettes but not for other

[38] This study, like Hart and Piff (2018), involves deception. In particular, participants are told they are choosing an actual compensation scheme related to a puzzle-solving task for a subsequent task but in reality they are just paid a flat fee for participation.

[39] John T. Jost and Sullivan (2003) shows that low-income respondents to large-scale US-representative surveys are more likely to believe that economic inequality is "legitimate and necessary" and to embrace meritocratic ideals than their more affluent counterparts.

tasks or vignettes; Butler (2016) ostensibly manipulates JWBs through own-wage disadvantage and directly documents beliefs distortions consistent with the original Lerner (1965) study extended to an own-misfortune context, i.e., self-blaming. We contribute to this literature by outlining a novel channel relating injustice to behavior through JWBs and then providing behavioral and attitudinal evidence for the operation of this channel using a ubiquitous and essentially trivial form of injustice (a small own-wage disadvantage) and firm-relevant subsequent behavior (productivity and pro-sociality). We also contribute by documenting important sources of contextual variation in our JWBs-based channel (type of work; type of inequality) that have been rarely (or never) studied before. Finally, we note that our study is a partial conceptual replication and extension of Butler (2016), which also studied randomly assigned wage inequality and exogenously varied the ability-intensiveness of work. Since that study used quite different participants – UC Berkeley students and Amazon Mechanical Turk workers – and found beliefs distortion consistent with a JWBs-protecting motive only for low-wage workers performing an ability-intensive task, the two studies provide complementary evidence of the external validity and robustness of each other's results.

Our study is also related to the large experimental literature in economics on the determinants of pro- and anti-social behavior. Our paper differs from much of this literature by focusing on reactions to revealed inequality on an initial task as a determinant of subsequent pro- and anti-sociality. We also differ from much of this existing literature methodologically. In contrast to the existing experimental literature typically featuring decontextualized and unfamiliar tasks such as (abstract) dictator games or dice-rolling for money (Abeler et al. 2019; Charness and Rabin 2002; Dufwenberg and Dufwenberg 2018; Fischbacher and Föllmi-Heusi 2013; Gneezy et al. 2018; Mazar et al. 2008), we construct all of our tasks to be ecologically valid, i.e., appropriate for, and familiar to, our student population: multiple-choice quizzes, Scantron bubbling and donating to well-known charities.[40] A second distinction between our study and almost all of the literature is that we provide separate opportunities for pro-social and anti-social behavior so that we may examine how the two interrelate and how they both vary with important contextual factors, such as experienced inequality. Our results suggest that rather than being a fixed trait or disposition, pro-social behavior varies substantially with context. We find no evidence for contextual variation of anti-social behavior, nor do we find an individual-level relationship between pro-social and anti-social behavior, suggesting they are driven by independent forces.

[40]Throughout the course of a typical day, students may proofread their own or others' work, take an in-class quiz, possibly involving bubbling Scantrons, or donate to well-known charities. Classroom quizzes often have an ability component, which we mimic by using questions from Raven's Advanced Progressive Matrices – a commonly used measure of general intelligence. Highlighting this connection, economist Greg Mankiw famously quipped "... math courses are one long IQ test" (Mankiw 2006). Many participants may have even taken a Raven's-like test, as many similar (unofficial) on-line tests purport to measure intelligence.

Lastly, our study is related to recent discussions regarding wage transparency within organizations, which remains a hot button issue in the popular press and has even been the target of recent policy reform (e.g., mandatory pay-ratio reporting under Dodd-Frank).[41] While some firms have embraced the idea of making salary information among its employees publicly available, the view on the merits of wage transparency remains polarized. Anecdotal evidence often supports the claim that pay transparency is a good thing, as it can result in higher productivity. However, some experts claim that wage transparency can result in lower productivity especially among the relatively underpaid subset of employees (see recent articles by Cooney (2018) and Conlin (2018)). Not surprisingly, there has been little formal empirical evidence on the causal impacts of wage transparency within firms on employee behaviors, given data limitations and identification challenges (c.f., Cullen and Perez-Truglia 2018). Importantly, the ramifications of policies aimed at increasing wage transparency hinge crucially on how employees might respond upon observing a (presumably) unequal pay distribution. Our results suggest the response to wage transparency might hinge crucially on the type of work employees are engaging in. Specifically, wage transparency among firms with (perceived) ability-intensive jobs (e.g., tech sector) can, indeed, be beneficial by boosting productivity and reducing shirking, especially among the subset of relatively low paid employees. This may help explain an empirical puzzle documented in two recent studies, that there is more wage dispersion net of human capital in more ability-intensive occupations (Van der Velde 2020; Visintin et al. 2015), perhaps because it is more readily accepted by employees in those occupations. As such, our study can contribute more broadly to the ongoing discussions of the merits of wage transparency.

6 Discussion and Concluding Remarks

In this paper we study the reactions to salient compensation inequality not captured by standard (pecuniary) economic incentives. A slew of recent studies demonstrate what is rapidly becoming the accepted wisdom among economists: that these *non-pecuniary* incentives are (weakly) negative, for instance reducing employee morale and productivity while increasing turnover. Against this backdrop, we propose a novel channel relating experienced inequality to subsequent firm-relevant behavior based on previous research in economics and social psychology suggesting that non-pecuniary inequality incentives are context-dependent and may sometimes be positive.

We use our proposed channel to make three broad predictions about how reactions to salient compensation inequality will vary across contexts and test these predictions in a laboratory experiment involving over 420 participants. The contextual factors we examined were the type of

[41] By way of example, several recent articles on the merits of pay transparency have appeared across a variety of reputable business publication outlets including (but not limited too): Time (Cooney 2018), Fortune (Fisher 2015), The Wall Street Journal (Shellenbarger 2016), and Business News Daily (Conlin 2018).

inequality – wage or income – and the type of work associated with inequality. We considered two types of work suggested by on-going trends in labor markets around the world. In one type, performance depends primarily on ability (ability-intensive) and in the other type, which we label "effort-intensive," performance is primarily a function of effort. We predicted that experienced *wage* disadvantage related to an *ability-intensive* task would increase subsequent effort provision but reduce pro-sociality by *strengthening* JWBs. We also predicted that experienced *income* inequality would induce a much weaker response alone and produce countervailing effects when revealed in addition to wage inequality, perhaps because it is perceived as being (partially) under the control of the workers themselves. Our results are broadly consistent with our predictions and, we argued, inconsistent with other commonly proposed explanations. Overall, the strong and predictable context-dependence of the reactions we document as well as the potential for a *positive* productivity response imply that the non-pecuniary incentives of revealed compensation inequality warrant further attention. In particular, when incorporated into considerations of optimal contracts (cf Koszegi 2014), non-pecuniary incentives may produce a theory of wage-setting where context matters.

Taken at face value, our study suggests many avenues for future research. The first avenue relates to the recent debate that has sprung up about the optimal degree of wage transparency within organizations. Namely, a question that has been posed is: should employee-level wage and earnings data be publicly shared among employees? Proponents tout its ability to potentially boost productivity. The results from our study support this assertion and suggest one possible mechanism relying on subconscious rationalization of inequality as fair in particular settings. We note two important caveats. For comparability, we follow much of the (laboratory and natural) experimental literature in randomly assigning inequality. In the real world, there is often a manager or firm that is perceived to assign compensation inequality and this more personal injustice may induce different reactions than those we document. A second caveat is that in our experimental setting we were able to carefully control both the type of work and the type of inequality revealed. Such careful controls may not be possible in real-world situations as, for instance, employees may have more latitude in their subjective perceptions of the ability-intensiveness of their work than they had in our experiment.

This brings up a second avenue for future research. Because performance on any task or in any profession always involves some combination of effort, ability and luck, a plausible conjecture is that employees' perceptions of the ability-intensiveness of their environment may be affected by emphasizing these aspects to differing degrees in job descriptions, mission statements or through other aspects of corporate culture. If perceptions of ability-intensiveness matter for employees' reactions to inequality, the characterization of work may be an important policy lever in designing a firm's wage transparency policies and corporate culture more broadly, particularly where moral hazard is unavoidable. By characterizing work as ability-intensive, firms may be able to induce higher, and higher quality, productivity from transparent wage policies than by characterizing work as depending primarily on effort.

A third avenue for future research relates to cooperation. Cooperation is held to be central to firm success, so the short- and longer-run implications of the reduced pro-sociality we observe for overall firm productivity warrant more careful consideration. Anecdotally, our findings appear to be consistent with tech firms and internet start-ups having reported success in moving to more pay transparency, which likely reveals a lot of pay inequality.[42] This suggests that the productivity effects of revealed inequality dominate the effects on cooperation, but confirming this appearance would obviously require careful empirical study.

Finally, more research into the mechanism(s) at work is necessary. We conjecture a novel channel relating experienced inequality to subsequent behavior and provide experimental evidence on behavior as well supporting survey evidence on JWBs. However, as noted, our supporting attitudinal evidence is only suggestive. As a modicum of additional reassurance, it is worth mentioning that some crucial features of our data and of the channel we conjecture to explain it have been documented in a wide array of settings. Evidence for the underlying cognitive processes we conjecture, or very closely related processes, has appeared across multiple disciplines, across a long time horizon and across various settings: in laboratory experiments using student subject pools; in on-line experiments using working adults; in field experiments; and also in observational data (inter alia, Benabou and Tirole 2006; Bénabou and Tirole 2016; Butler 2016; Hart and Piff 2018; Lerner 2013, 1965). Still, more direct causal evidence on the channels through which contextual factors affect how inequality is perceived, processed, rationalized and reacted to is warranted.

References

Abeler, J., S. Altmann, S. Kube, and M. Wibral (2010): "Gift exchange and workers' fairness concerns: when equality is unfair," *Journal of the European Economic Association*, 8, 1299–1324.

Abeler, J., D. Nosenzo, and C. Raymond (2019): "Preferences for truth-telling," *Econometrica*, 87, 1115–1153.

[42]For example, see the recent popular press articles by (Burkus 2016; McLaren 2019; Weller 2017).

Acemoglu, D. and D. Autor (2011): "Skills, tasks and technologies: Implications for employment and earnings," in *Handbook of labor economics*, Elsevier, vol. 4, 1043–1171.

Adams, J. S. (1963): "Towards an understanding of inequity." *The Journal of Abnormal and Social Psychology*, 67, 422.

Akerlof, G. A. and W. T. Dickens (1982): "The Economic Consequences of Cognitive Dissonance," *The American Economic Review*, 72, 307–319.

Akerlof, G. A. and J. L. Yellen (1990): "The fair wage-effort hypothesis and unemployment," *The Quarterly Journal of Economics*, 105, 255–283.

Albanesi, S., V. Gregory, C. Patterson, and A. Şahin (2013): "Is job polarization holding back the labor market," *Lib. Street Econ.(March)*.

Angelova, V., W. Güth, and M. G. Kocher (2012): "Co-employment of permanently and temporarily employed agents," *Labour Economics*, 19, 48–58.

Autor, D. H. (2014): "Skills, education, and the rise of earnings inequality among the "other 99 percent"," *Science*, 344, 843–851.

Autor, D. H., L. F. Katz, and A. B. Krueger (1998): "Computing inequality: have computers changed the labor market?" *The Quarterly journal of economics*, 113, 1169–1213.

Autor, D. H., F. Levy, and R. J. Murnane (2003): "The skill content of recent technological change: An empirical exploration," *The Quarterly journal of economics*, 118, 1279–1333.

Bartling, B. and F. A. Von Siemens (2011): "Wage inequality and team production: An experimental analysis," *Journal of Economic Psychology*, 32, 1–16.

Benabou, R. and J. Tirole (2006): "Belief in a just world and redistributive politics," *The Quarterly journal of economics*, 121, 699–746.

Bénabou, R. and J. Tirole (2016): "Mindful economics: The production, consumption, and value of beliefs," *Journal of Economic Perspectives*, 30, 141–64.

Bohnet, I. and R. Zeckhauser (2004): "Trust, risk and betrayal," *Journal of Economic Behavior & Organization*, 55, 467–484.

Bolton, G. and P. Werner (2016): "The influence of potential on wages and effort," *Experimental Economics*, 19, 535–561.

Bolton, G. E. and A. Ockenfels (2000): "ERC: A theory of equity, reciprocity, and competition," *American economic review*, 90, 166–193.

Bracha, A., U. Gneezy, and G. Loewenstein (2015): "Relative pay and labor supply," *Journal of labor economics*, 33, 297–315.

Breza, E., S. Kaur, and Y. Shamdasani (2018): "The morale effects of pay inequality," *The Quarterly Journal of Economics*, 133, 611–663.

Burkus, D. (2016): "Why do we keep salaries secret?" *Forbes*.

Butler, J. V. (2016): "Inequality and relative ability beliefs," *The Economic Journal*, 126, 907–948.

Butler, J. V. and J. B. Miller (2018): "Social risk and the dimensionality of intentions," *Management Science*, 64, 2787–2796.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 90, 414–427.

Card, D., A. Mas, E. Moretti, and E. Saez (2012): "Inequality at work: The effect of peer salaries on job satisfaction," *American Economic Review*, 102, 2981–3003.

Center, P. R. (2016): "The state of American jobs," .

Charness, G., R. Cobo-Reyes, J. A. Lacomba, F. Lagos, and J. M. Pérez (2016): "Social comparisons in wage delegation: Experimental evidence," *Experimental Economics*, 19, 433–459.

Charness, G. and M. Dufwenberg (2006): "Promises and partnership," *Econometrica*, 74, 1579–1601.

Charness, G. and P. Kuhn (2007): "Does pay inequality affect worker effort? Experimental evidence," *Journal of labor economics*, 25, 693–723.

Charness, G. and M. Rabin (2002): "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 117, 817–869.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree–An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Clark, A. E., D. Masclet, and M. C. Villeval (2010): "Effort and comparison income: Experimental and survey evidence," *ILR Review*, 63, 407–426.

Cohn, A., E. Fehr, B. Herrmann, and F. Schneider (2014): "Social comparison and effort provision: Evidence from a field experiment," *Journal of the European Economic Association*, 12, 877–898.

Conlin, B. (2018): "The Pros and Cons of Salary Transparency," *Business News Daily*.

Cooney, S. (2018): "Should you Share your Salary with Co-Workers? Here's what Experts Say," *Time Magazine*.

Cullen, Z. and R. Perez-Truglia (2018): "How much does your boss make? the effects of salary comparisons," *NBER Working Paper no. 24841*.

David, H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 103, 1553–97.

Dube, A., L. Giuliano, and J. Leonard (2019): "Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior," *American Economic Review*, 109, 620–663.

Dufwenberg, M. and M. A. Dufwenberg (2018): "Lies in disguise–A theoretical analysis of cheating," *Journal of Economic Theory*, 175, 248–264.

Falk, A. and A. Ichino (2006): "Clean evidence on peer effects," *Journal of labor economics*, 24, 39–57.

Fehr, E. and K. M. Schmidt (1999): "A theory of fairness, competition, and cooperation," *The quarterly journal of economics*, 114, 817–868.

Festinger, L. (1957): *A theory of cognitive dissonance*, vol. 2, Stanford university press.

Fischbacher, U. and F. Föllmi-Heusi (2013): "Lies in disguise–an experimental study on cheating," *Journal of the European Economic Association*, 11, 525–547.

Fisher, A. (2015): "How Pay Transparency can Keep People from Quitting," *Fortune*.

Frank, D. H., K. Wertenbroch, and W. W. Maddux (2015): "Performance pay or redistribution? Cultural differences in just-world beliefs and preferences for wage inequality," *Organizational Behavior and Human Decision Processes*, 130, 160–170.

Gächter, S. and C. Thöni (2010): "Social comparison and performance: Experimental evidence on the fair wage–effort hypothesis," *Journal of Economic Behavior & Organization*, 76, 531–543.

Gaultney, J. F. and A. Cann (2001): "Grade expectations," *Teaching of Psychology*, 28, 84–87.

Gibbons, R. (1987): "Piece-rate incentive schemes," *Journal of Labor Economics*, 5, 413–429.

Gill, D., V. Prowse, and M. Vlassopoulos (2013): "Cheating in the workplace: An experimental study of the impact of bonuses and productivity," *Journal of Economic Behavior & Organization*, 96, 120–134.

Gneezy, U. (2005): "Deception: The role of consequences," *American Economic Review*, 95, 384–394.

Gneezy, U., A. Kajackaite, and J. Sobel (2018): "Lying Aversion and the Size of the Lie," *American Economic Review*, 108, 419–53.

Godechot, O. and C. Senik (2015): "Wage comparisons in and out of the firm. Evidence from a matched employer–employee French database," *Journal of Economic Behavior & Organization*, 117, 395–410.

Greiner, B., A. Ockenfels, and P. Werner (2011): "Wage transparency and performance: A real-effort experiment," *Economics Letters*, 111, 236–238.

Gross, T., C. Guo, and G. Charness (2015): "Merit pay and wage compression with productivity differences and uncertainty," *Journal of Economic Behavior & Organization*, 117, 233–247.

Hart, E. and P. K. Piff (2018): "When Inequality is Good for Groups: The Moderating Role of Fairness," *SSRN Working Paper*.

Hayek, F. A. (2020): *The constitution of liberty: The definitive edition*, Routledge.

Hennig-Schmidt, H., A. Sadrieh, and B. Rockenbach (2010): "In search of workers' real effort reciprocity–a field and a laboratory experiment," *Journal of the European Economic Association*, 8, 817–837.

Hopkins, E. and T. Kornienko (2009): "Status, Affluence, and inequality: Rank-based comparisons in games of status," *Games and Economic Behavior*, 67, 552–568.

——— (2010): "Which inequality? The inequality of endowments versus the inequality of rewards," *American Economic Journal: Microeconomics*, 2, 106–37.

John T. Jost, Brett W. Pelham, O. S. and B. N. Sullivan (2003): "Social inequality and the reduction of ideological dissonance on behalf of the system: evidence of enhanced system justification among the disadvantaged," *European Journal of Social Psychology*, 33, 13–36.

Jost, J. T., M. R. Banaji, and B. A. Nosek (2004): "A decade of system justification theory: Accumulated evidence of conscious and unconscious bolstering of the status quo," *Political psychology*, 25, 881–919.

Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–292.

Koszegi, B. (2014): "Behavioral contract theory," *Journal of Economic Literature*, 52, 1075–1118.

Lazear, E. P. (2000): "The power of incentives," *American Economic Review*, 90, 410–414.

——— (2018): "Compensation and Incentives in the Workplace," *Journal of Economic Perspectives*, 32, 195–214.

Lazear, E. P. and P. Oyer (2007): "Personnel economics," Tech. rep., National Bureau of economic research.

Lerner, M. (2013): *The Belief in a Just World: A Fundamental Delusion*, Springer Science & Business Media.

Lerner, M. J. (1965): "Evaluation of performance as a function of performer's reward and attractiveness." *Journal of Personality and Social Psychology*, 1, 355.

Lerner, M. J. and D. T. Miller (1978): "Just world research and the attribution process: Looking back and ahead." *Psychological bulletin*, 85, 1030.

Mankiw, G. (2006): "Why Aspiring Economists Need Math," http://gregmankiw.blogspot.com/2006/09/why-aspiring-economists-need-math.html, accessed: 2019-07-26.

Mazar, N., O. Amir, and D. Ariely (2008): "The dishonesty of honest people: A theory of self-concept maintenance," *Journal of marketing research*, 45, 633–644.

McLaren, S. (2019): "Why These 3 Companies Are Sharing How Much Their Employees Make," *LinkedIn Talent Blog*.

Muro, M., S. Liu, J. Whiton, and S. Kulkarni (2017): "Digitalization and the American workforce," .

Nisan, M. (1985): "Justice in the classroom: Teachers' and students' perceptions of principles that guide grading distribution," *School and education*, 141–155.

Nosenzo, D. (2013): "Pay secrecy and effort provision," *Economic inquiry*, 51, 1779–1794.

Prendergast, C. (1999): "The provision of incentives in firms," *Journal of economic literature*, 37, 7–63.

Sabbagh, C., R. Faher-Aladeen, and N. Resh (2004): "Evaluation of grade distributions: a comparison of Jewish and Druze students in Israel," *Social Psychology of Education*, 7, 313–337.

Schindler, S., K. Wenzel, S. Dobiosch, and M.-A. Reinhard (2019): "The role of belief in a just world for (dis) honest behavior," *Personality and Individual Differences*, 142, 72–78.

Shellenbarger, S. (2016): "Open Salaries: The Good, the Bad, and the Awkward," *Wall Street Journal*.

Spence, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics*, 87, 355–374.

Sseruyange, J. and E. Bulte (2019): "Wage Differentials and Workers' Effort: Experimental Evidence from Uganda," *Oxford Bulletin of Economics and Statistics*.

Starmans, C., M. Sheskin, and P. Bloom (2017): "Why people prefer unequal societies," *Nature Human Behaviour*, 1, 1–7.

Stiglitz, J. E. (1975): "Incentives, risk, and information: notes towards a theory of hierarchy," *The Bell Journal of Economics*, 552–579.

Van der Velde, L. (2020): "Within occupation wage dispersion and the task content of jobs," *Oxford Bulletin of Economics and Statistics*, 82, 1161–1197.

Visintin, S., K. Tijdens, S. Steinmetz, and P. de Pedraza (2015): "Task implementation heterogeneity and wage dispersion," *IZA Journal of Labor Economics*, 4, 1–24.

Weller, C. (2017): "A CEO who makes everyone's salary transparent says people are more productive than ever," *Business Insider*.

Wenzel, K., S. Schindler, and M.-A. Reinhard (2017): "General belief in a just world is positively associated with dishonest behavior," *Frontiers in psychology*, 8, 1770.

Winter, E. (2004): "Incentives and discrimination," *American Economic Review*, 94, 764–773.

Tables and Figures

Tables

Table 2: *Descriptive Statistics*

Treatment	Wage	N	Phase 1		Phase 2			Phase 3
			income	score	quantity	quality	lying	donations
Invisible Wage Ineq	Low	32	4.18	20.91	6.91	0.88	0.16	0.30
Invisible Income Ineq			(0.202)	(1.012)	(0.476)	(0.041)	(0.065)	(0.061)
Ability-Intensive Task	High	24	7.29	19.38	7.75	0.92	0.13	0.24
			(0.374)	(1.259)	(0.405)	(0.037)	(0.069)	(0.065)
Visible Wage Ineq	Low	24	4.03	20.17	7.75	0.96	0.25	0.16
Invisible Income Ineq			(0.236)	(1.181)	(0.367)	(0.018)	(0.090)	(0.048)
Ability-Intensive Task	High	34	7.70	19.03	7.74	0.96	0.09	0.22
			(0.224)	(0.881)	(0.320)	(0.015)	(0.049)	(0.045)
Invisible Wage Ineq	Low	28	4.01	20.04	7.64	0.94	0.14	0.25
Visible Income Ineq			(0.176)	(0.878)	(0.368)	(0.023)	(0.067)	(0.071)
Ability-Intensive Task	High	21	7.52	20.00	6.76	0.95	0.10	0.17
			(0.305)	(1.223)	(0.377)	(0.034)	(0.066)	(0.043)
Visible Wage Ineq	Low	20	3.85	19.25	6.85	0.92	0.10	0.19
Visible Income Ineq			(0.150)	(0.075)	(0.472)	(0.045)	(0.069)	(0.058)
Ability-Intensive Task	High	28	7.63	19.79	7.25	0.96	0.11	0.13
			(0.298)	(1.043)	(0.351)	(0.019)	(0.060)	(0.046)
Invisible Wage Ineq	Low	27	5.76	28.81	7.67	0.99	0.04	0.17
Invisible Income Ineq			(0.258)	(1.292)	(0.430)	(0.003)	(0.037)	(0.053)
Effort-Intensive Task	High	27	10.06	28.07	7.67	0.98	0.04	0.14
			(0.304)	(1.160)	(0.377)	(0.006)	(0.037)	(0.039)
Visible Wage Ineq	Low	26	6.00	30.00	7.88	0.96	0.08	0.23
Invisible Income Ineq			(0.251)	(1.253)	(0.405)	(0.029)	(0.053)	(0.059)
Effort-Intensive Task	High	27	10.60	29.89	8.26	0.96	0.11	0.15
			(0.189)	(0.897)	(0.402)	(0.025)	(0.062)	(0.046)
Invisible Wage Ineq	Low	22	5.74	28.68	7.59	0.97	0.09	0.28
Visible Income Ineq			(0.289)	(1.447)	(0.477)	(0.009)	(0.063)	(0.078)
Effort-Intensive Task	High	30	10.30	29.07	7.93	0.95	0.03	0.12
			(0.248)	(1.000)	(0.426)	(0.026)	(0.033)	(0.043)
Visible Wage Ineq	Low	30	5.53	27.67	6.73	0.96	0.10	0.26
Visible Income Ineq			(0.214)	(1.072)	(0.380)	(0.023)	(0.056)	(0.064)
Effort-Intensive Task	High	23	10.66	31.30	8.09	0.94	0.04	0.13
			(0.260)	(0.828)	(0.510)	(0.035)	(0.043)	(0.032)

Notes: [**1**] Standard errors appear in parentheses. [**2**] "Phase 1 income" refers to participants' earnings from the phase 1 task only; "Phase 1 score" refers to the number of correct responses (out of a maximum of 48) on the phase 1 task; "Phase 2 quantity" refers to the actual number of Scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted Scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the proportion of total experimental income earned in Phases 1 and 2 combined that was donated to either of the two possible charities (American Red Cross or the South Plains Food Bank.

Table 3: *OLS Estimates Effort-Intensive*

	(1)	(2)	(3)	(4)	(5)
	Phase 1		Phase 2		Phase 3
	score	quantity	quality	lying	donations
High Wage	-1.09	-1.17	-0.06	0.10	0.03
	(0.888)	(0.701)	(0.037)	(0.083)	(0.081)
Visible Wage Inequality	1.30	0.16	-0.03	0.05	0.08
	(1.214)	(0.613)	(0.033)	(0.046)	(0.069)
Visible Income Inequality	-0.05	-0.09	-0.02	0.07	0.12
	(1.429)	(0.440)	(0.014)	(0.071)	(0.084)
(High Wage) x (Visible Wage Ineq)	0.44	0.31	0.01	0.02	-0.04
	(1.989)	(0.675)	(0.019)	(0.089)	(0.053)
(High Wage) x (Visible Income Ineq)	1.52	0.50	-0.01	-0.05	-0.10
	(1.796)	(0.755)	(0.019)	(0.061)	(0.104)
(Visible Wage Ineq) x (Visible Income Ineq)	-2.50	-0.96	0.01	-0.07	-0.14
	(1.796)	(0.909)	(0.046)	(0.116)	(0.116)
(HighWage) x (VisWageIneq) x (VisIncIneq)	2.37	0.53	-0.01	-0.01	0.13
	(2.614)	(1.020)	(0.045)	(0.133)	(0.126)
Phase 1 income		0.23	0.01	-0.02	-0.02
		(0.156)	(0.009)	(0.015)	(0.014)
Phase 2 quantity			-0.01	-0.00	-0.02***
			(0.008)	(0.006)	(0.006)
Constant	31.90***	4.23**	0.98***	0.33**	0.18
	(2.823)	(1.442)	(0.099)	(0.141)	(0.222)
Observations	212	212	212	212	212
R-squared	0.089	0.118	0.056	0.063	0.143

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Effort-Intensive Phase 1 task with the dependent variable labeled in the column heading. "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Experimental controls are as follows: *HW* is an indicator variable for the High-Wage pay scheme, *VW* is an indicator for visible pay inequality, *VI* is an indicator variable visible (experimental) income inequality. [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: *OLS Estimates Ability-Intensive*

	(1)	(2)	(3)	(4)	(5)
	Phase 1		Phase 2		Phase 3
	score	quantity	quality	lying	donations
High Wage	-1.23	0.71	-0.01	0.07	0.02
	(1.582)	(0.582)	(0.060)	(0.106)	(0.110)
Visible Wage Inequality	-0.51	1.00**	0.07**	0.12	-0.12**
	(1.088)	(0.417)	(0.033)	(0.117)	(0.051)
Visible Income Inequality	-0.62	0.75	0.05	-0.01	-0.03
	(1.382)	(0.513)	(0.037)	(0.076)	(0.096)
(High Wage) x (Visible Wage Ineq)	0.35	-1.17***	-0.04	-0.16	0.11
	(2.360)	(0.340)	(0.042)	(0.113)	(0.082)
(High Wage) x (Visible Income Ineq)	1.04	-1.74**	-0.01	-0.04	-0.10
	(2.236)	(0.634)	(0.044)	(0.113)	(0.131)
(Visible Wage Ineq) x (Visible Income Ineq)	0.21	-1.47**	-0.08	-0.13	0.03
	(1.553)	(0.613)	(0.056)	(0.149)	(0.122)
(HighWage) x (VisWageIneq) x (VisIncIneq)	-0.56	2.24**	0.05	0.19	-0.01
	(2.915)	(0.948)	(0.087)	(0.172)	(0.142)
Phase 1 income		0.07	0.01	-0.02	-0.01
		(0.081)	(0.008)	(0.022)	(0.020)
Phase 2 quantity			0.03**	-0.01	-0.03**
			(0.011)	(0.014)	(0.012)
Constant	32.36***	7.17***	0.71***	0.33	0.27
	(4.051)	(1.156)	(0.161)	(0.262)	(0.243)
Observations	207	207	207	207	207
R-squared	0.122	0.162	0.190	0.074	0.123

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Ability-Intensive Phase 1 task with the dependent variable labeled in the column heading. "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Experimental controls are as follows: *HW* is an indicator variable for the High-Wage pay scheme, *VW* is an indicator for visible pay inequality, *VI* is an indicator variable visible (experimental) income inequality. [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: *JWBs Estimates (OLS) Ability-Intensive*

	Just World Beliefs			
	(1)	(2)	(3)	(4)
High Wage	0.22	-0.06	-0.15	-0.14
	(0.217)	(0.262)	(0.244)	(0.277)
Visible Wage Inequality	0.53***	0.53***	0.48**	0.42**
	(0.157)	(0.156)	(0.184)	(0.188)
Visible Income Inequality	0.36	0.38	0.34	0.33
	(0.218)	(0.223)	(0.230)	(0.237)
(High Wage) x (Visible Wage Ineq)	0.11	0.07	0.17	0.22
	(0.474)	(0.482)	(0.491)	(0.494)
(High Wage) x (Visible Income Ineq)	-0.04	-0.06	0.14	0.09
	(0.303)	(0.292)	(0.291)	(0.293)
(Visible Wage Ineq) x (Visible Income Ineq)	-0.39	-0.40	-0.30	-0.29
	(0.384)	(0.382)	(0.401)	(0.410)
(HighWage) x (VisWageIneq) x (VisIncIneq)	-0.60	-0.57	-0.80	-0.80
	(0.739)	(0.736)	(0.709)	(0.718)
Phase 1 income		0.09	0.09	0.08
		(0.057)	(0.053)	(0.056)
Phase 2 quantity			0.15***	0.13**
			(0.043)	(0.049)
Phase 2 quality			-0.95	-1.06*
			(0.591)	(0.575)
Phase 2 lying			0.02	-0.00
			(0.352)	(0.356)
Phase 3 donations				-0.56
				(0.322)
Constant	1.86	1.23	1.04	1.27
	(1.145)	(1.293)	(1.261)	(1.341)
Observations	207	207	207	207
R-squared	0.087	0.092	0.125	0.136

Notes: **[1]** Each column reports an OLS regression restricted to Ability-Intensive session data. The dependent variable in each regression is our measure of *Just World Beliefs*, which is increasing in participants' degree of agreement with the statement "People generally get what they deserve." Valid responses range from 1="totally disagree" to 7="totally agree." **[2]** The columns differ in the additional controls included. Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. **[3]** The additional controls are: "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of Scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted Scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the proportion of total phase 1 and phase 2 earnings donated to either of the possible charities in Phase 3 (American Red Cross and the South Plains Food Bank). **[4]** Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: *Phase 2 and 3 behavior by Initial-Task Earnings*

Initial-task Earnings Percentile	(1)	(2)	(3)	(4)
	Phase 2			Phase 3
	quantity	quality	lying	donations
20th	-0.80*	-0.02	0.14**	0.10*
	(0.426)	(0.019)	(0.054)	(0.058)
40th	-0.79*	-0.01	0.09**	0.07*
	(0.414)	(0.022)	(0.037)	(0.040)
60th	-0.16	-0.01	0.09**	0.06
	(0.391)	(0.021)	(0.040)	(0.040)
80th	-0.40	-0.00	0.08	0.03
	(0.421)	(0.021)	(0.049)	(0.041)
Ability-Intensive Initial Task	0.05	-0.02	0.03	0.00
	(0.267)	(0.016)	(0.036)	(0.034)
Phase 2 quantity		0.01	-0.00	-0.03***
		(0.007)	(0.007)	(0.006)
Constant	6.98***	0.96***	0.13	0.06
	(0.815)	(0.091)	(0.104)	(0.121)
Observations	419	419	419	419
R-squared	0.095	0.044	0.054	0.094

Notes: [**1**] Each column reports an OLS regression pooling data across both initial tasks with the dependent variable as labeled in the column heading. These are: "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Control variables include a set of indicators for the workers' initial-task earnings quintile. The earnings quintile indicators are exclusive, so that, "40^{th}" is an indicator taking the value of one if the workers' initial-task earnings fell below the 40^{th} percentile threshold but above the 20^{th} percentile threshold. The excluded category is "above the 80^{th} percentile." Other control variables are: and indicator for whether the initial task was ability-intensive, as opposed to effort-intensive and Phase 2 quantity (where appropriate). [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: *Phase 2 and 3 behavior by Initial-Task Earnings and Visible Income Inequality*

	(1)	(2)	(3)	(4)
Initial-task Earnings Percentile		Phase 2		Phase 3
	quantity	quality	lying	donations
20th	-0.73	-0.05*	0.20**	0.12
	(0.539)	(0.026)	(0.089)	(0.076)
40th	-0.86*	-0.03	0.09*	0.03
	(0.441)	(0.031)	(0.048)	(0.042)
60th	0.28	-0.00	0.09	0.06
	(0.529)	(0.014)	(0.052)	(0.050)
80th	-0.44	-0.01	0.10	0.00
	(0.568)	(0.025)	(0.059)	(0.043)
Visible Income Ineq	-0.14	-0.02	0.01	-0.04
	(0.438)	(0.018)	(0.036)	(0.046)
(Visible Income Ineq) x (20th)	-0.08	0.05	-0.12	-0.04
	(0.652)	(0.040)	(0.104)	(0.088)
(Visible Income Ineq) x (40th)	0.19	0.05	-0.02	0.08
	(0.700)	(0.037)	(0.076)	(0.083)
(Visible Income Ineq) x (60th)	-0.85	-0.00	0.00	0.00
	(0.651)	(0.037)	(0.081)	(0.068)
(Visible Income Ineq) x (80th)	0.10	0.01	-0.04	0.05
	(0.638)	(0.031)	(0.074)	(0.072)
Ability-intensive initial task	0.02	-0.02	0.03	0.00
	(0.249)	(0.016)	(0.036)	(0.034)
Phase 2 quantity		0.01	-0.00	-0.03***
		(0.008)	(0.007)	(0.006)
Constant	7.04***	0.97***	0.11	0.08
	(0.955)	(0.099)	(0.114)	(0.128)
Observations	419	419	419	419
R-squared	0.106	0.051	0.060	0.100

Notes: [**1**] Each column reports an OLS regression pooling data across both initial tasks with dependent variable as labeled in the column heading. These are: “Phase 2 quantity” refers to the number of scantrons submitted; “Phase 2 quality” refers to the average accuracy of submitted scantrons; “Phase 2 lying” is the proportion of participants who misreported the number of Scantrons they submitted; “Phase 3 donations” refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Control variables include a set of indicators for the workers’ initial-task earnings quintile, an indicator for the sessions in which we provided participants with information about the initial-task earnings distribution (“Visible Income Ineq”) and interactions between these two sets of indicators. The earnings quintile indicators are exclusive, so that, “40^{th}” is an indicator taking the value of one if the workers’ initial-task earnings fell below the 40^{th} percentile threshold but above the 20^{th} percentile threshold. The excluded category is “above the 80^{th} percentile.” Other control variables are: and indicator for whether the initial task was ability-intensive, as opposed to effort-intensive, and Phase 2 quantity (where appropriate). [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Online Appendix
Inequality as an Incentive

Jeffrey V. Butler
University of California, Merced

Eric Cardella
Texas Tech University

This version: *April 7, 2021*

Table A1: *OLS Estimates Effort-Intensive, Bootstrapped Standard Errors*

	(1)	(2)	(3)	(4)	(5)	(6)
	Phase 1	Phase 2			Phase 3	JWBs
	score	quantity	quality	lying	donations	
High Wage	-1.09	-1.17	-0.06	0.10	0.03	-0.09
	(1.122)	(0.737)	(0.037)	(0.094)	(0.087)	(0.475)
Visible Wage Inequality	1.30	0.16	-0.03	0.05	0.08	-0.97**
	(1.269)	(0.649)	(0.034)	(0.055)	(0.076)	(0.470)
Visible Income Inequality	-0.05	-0.09	-0.02	0.07	0.12	-0.58
	(1.449)	(0.541)	(0.019)	(0.078)	(0.088)	(0.625)
(HighWage) x (VisWageIneq)	0.44	0.31	0.01	0.02	-0.04	1.02**
	(2.159)	(0.705)	(0.025)	(0.093)	(0.067)	(0.399)
(HighWage) x (VisIncIneq)	1.52	0.50	-0.01	-0.05	-0.10	0.51
	(2.005)	(0.792)	(0.022)	(0.075)	(0.106)	(0.609)
(VisWageIneq) x (VisIncIneq)	-2.50	-0.96	0.01	-0.07	-0.14	1.18*
	(1.893)	(0.986)	(0.050)	(0.130)	(0.117)	(0.706)
(HighWage) x (VisWageIneq)	2.37	0.53	-0.01	-0.01	0.13	-1.32
x(VisIncIneq)	(2.684)	(1.035)	(0.049)	(0.150)	(0.122)	(0.882)
Phase 1 income		0.23*	0.01	-0.02	-0.02	-0.17**
		(0.138)	(0.009)	(0.017)	(0.013)	(0.084)
Phase 2 quantity			-0.01	-0.00	-0.02***	-0.03
			(0.007)	(0.006)	(0.007)	(0.034)
Constant	31.90***	4.23***	0.98***	0.33**	0.18	5.52***
	(3.544)	(1.549)	(0.110)	(0.166)	(0.215)	(1.118)
Observations	212	212	212	212	212	209
R-squared	0.089	0.118	0.056	0.063	0.143	0.161

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Effort-Intensive Phase 1 task with the dependent variable labeled in the column heading. "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Experimental controls are as follows: *HW* is an indicator variable for the High-Wage pay scheme, *VW* is an indicator for visible pay inequality, *VI* is an indicator variable visible (experimental) income inequality. [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Bootstrapped standard errors, clustered by session appear, in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A2: *OLS Estimates Ability-Intensive, Bootstrapped Standard Errors*

	(1)	(2)	(3)	(4)	(5)	(6)
	Phase 1	Phase 2			Phase 3	JWBs
	score	quantity	quality	lying	donations	
High Wage	-1.23	0.71	-0.01	0.07	0.02	-0.14
	(1.367)	(0.569)	(0.060)	(0.109)	(0.090)	(0.293)
Visible Wage Inequality	-0.51	1.00**	0.07**	0.12	-0.12**	0.41**
	(1.038)	(0.498)	(0.033)	(0.118)	(0.053)	(0.188)
Visible Income Inequality	-0.62	0.75	0.05	-0.01	-0.03	0.29
	(1.466)	(0.590)	(0.038)	(0.085)	(0.113)	(0.284)
(HighWage) x (VisWageIneq)	0.35	-1.17***	-0.04	-0.16	0.11	0.21
	(2.123)	(0.380)	(0.041)	(0.110)	(0.091)	(0.496)
(HighWage) x (VisIncIneq)	1.04	-1.74**	-0.01	-0.04	-0.10	0.15
	(2.225)	(0.675)	(0.046)	(0.132)	(0.162)	(0.447)
(VisWageIneq) x (VisIncIneq)	0.21	-1.47**	-0.08	-0.13	0.03	-0.22
	(1.528)	(0.704)	(0.073)	(0.172)	(0.145)	(0.482)
(HighWage) x (VisWageIneq)	-0.56	2.24**	0.05	0.19	-0.01	-0.84
x (VisIncIneq)	(2.899)	(1.014)	(0.107)	(0.165)	(0.175)	(0.770)
Phase 1 income		0.07	0.01	-0.02	-0.01	0.08
		(0.079)	(0.009)	(0.022)	(0.019)	(0.049)
Phase 2 quantity			0.03***	-0.01	-0.03**	0.12***
			(0.010)	(0.014)	(0.012)	(0.045)
Constant	32.36***	7.17***	0.71***	0.33	0.27	0.37
	(4.138)	(1.531)	(0.195)	(0.298)	(0.277)	(1.584)
Observations	207	207	207	207	207	207
R-squared	0.122	0.162	0.190	0.074	0.123	0.117

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Ability-Intensive Phase 1 task with the dependent variable labeled in the column heading. "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Experimental controls are as follows: *HW* is an indicator variable for the High-Wage pay scheme, *VW* is an indicator for visible pay inequality, *VI* is an indicator variable visible (experimental) income inequality. [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Bootstrapped standard errors, clustered by session appear, in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A3: *JWBs Estimates (OLS) Effort-Intensive Sessions*

	Just World Beliefs			
	(1)	(2)	(3)	(4)
High Wage	-0.84***	-0.05	-0.08	-0.07
	(0.265)	(0.414)	(0.441)	(0.446)
Visible Wage Inequality	-1.03**	-0.98**	-0.97**	-0.95**
	(0.396)	(0.397)	(0.400)	(0.406)
Visible Income Inequality	-0.58	-0.58	-0.58	-0.55
	(0.653)	(0.630)	(0.631)	(0.607)
(High Wage) x (Visible Wage Ineq)	1.00***	1.02***	1.02***	1.01***
	(0.287)	(0.300)	(0.299)	(0.300)
(High Wage) x (Visible Income Ineq)	0.46	0.50	0.51	0.48
	(0.640)	(0.629)	(0.631)	(0.609)
(Visible Wage Ineq) x (Visible Income Ineq)	1.30*	1.21*	1.18	1.13
	(0.696)	(0.675)	(0.688)	(0.666)
(HighWage) x (VisWageIneq) x (VisIncIneq)	-1.41	-1.34	-1.32	-1.28
	(0.884)	(0.870)	(0.872)	(0.837)
Phase 1 income		-0.18**	-0.17**	-0.18**
		(0.079)	(0.081)	(0.083)
Phase 2 quantity			-0.03	-0.04
			(0.030)	(0.031)
Phase 2 quality			0.09	0.11
			(0.561)	(0.564)
Phase 2 lying			0.01	-0.01
			(0.226)	(0.232)
Phase 3 donations				-0.30
				(0.434)
Constant	4.25***	5.38***	5.42***	5.47***
	(0.730)	(0.833)	(1.013)	(1.066)
Observations	209	209	209	209
R-squared	0.131	0.159	0.161	0.164

Notes: [**1**] Each column reports an OLS regression restricted to Effort-Intensive session data. The dependent variable in each regression is our measure of *Just World Beliefs*, which is increasing in participants' degree of agreement with the statement "People generally get what they deserve." Valid responses range from 1="totally disagree" to 7="totally agree." [**2**] The columns differ in the additional controls included. Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**3**] The additional controls are: "Phase 1 income" is participants' earnings from the the Phase 1 task alone; "Phase 2 quantity" refers to the number of Scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted Scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the proportion of total phase 1 and phase 2 earnings donated to either of the possible charities in Phase 3 (American Red Cross and the South Plains Food Bank). [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A4: *Behavior by Initial-Task Earnings Quintiles, Effort-Intensive Sessions*

Initial-task Earnings Percentile	(1) Phase 2 quantity	(2) Phase 2 quality	(3) Phase 2 lying	(4) Phase 3 donations
20th	-1.40**	-0.03	0.07	0.10
	(0.484)	(0.030)	(0.054)	(0.058)
40th	0.33	0.00	0.10	0.04
	(0.486)	(0.009)	(0.083)	(0.078)
60th	-0.08	-0.03	0.10*	0.06
	(0.618)	(0.026)	(0.049)	(0.056)
80th	0.16	-0.03	-0.00	-0.05
	(0.495)	(0.028)	(0.036)	(0.059)
Phase 2 quantity		-0.01	-0.00	-0.02***
		(0.008)	(0.008)	(0.007)
Constant	6.11***	1.04***	0.18	0.05
	(1.058)	(0.108)	(0.115)	(0.154)
Observations	212	212	212	212
R-squared	0.134	0.046	0.065	0.131

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Effort-Intensive Phase 1 task with the dependent variable labeled in the column heading. These are: "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Control variables include a set of indicators for the workers' initial-task earnings quintile. These are exclusive, so that, "40^{th}" is an indicator taking the value of one if the workers' initial-task earnings fell below the 40^{th} percentile threshold but above the 20^{th} percentile threshold. The excluded category is "above the 80^{th} percentile." [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A5: *Behavior by Initial-Task Earnings Quintiles, Ability-Intensive Sessions*

	(1)	(2)	(3)	(4)
Initial-task Earnings Percentile	Phase 2 quantity	Phase 2 quality	Phase 2 lying	Phase 3 donations
20th	-0.03	-0.06*	0.09	0.03
	(0.408)	(0.027)	(0.078)	(0.067)
40th	-0.10	0.02	0.06	0.01
	(0.349)	(0.026)	(0.078)	(0.050)
60th	-0.04	-0.01	0.12	0.08
	(0.555)	(0.027)	(0.092)	(0.068)
80th	0.36	0.01	-0.00	-0.02
	(0.345)	(0.021)	(0.064)	(0.056)
Phase 2 quantity		0.03**	-0.01	-0.03**
		(0.011)	(0.011)	(0.012)
Constant	8.26***	0.77***	0.17	0.13
	(1.003)	(0.177)	(0.245)	(0.220)
Observations	207	207	207	207
R-squared	0.132	0.186	0.068	0.094

Notes: [**1**] Each column reports an OLS regression restricted to observations from sessions featuring the Ability-Intensive Phase 1 task with the dependent variable labeled in the column heading. These are: "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Control variables include a set of indicators for the workers' initial-task earnings quintile. These are exclusive, so that, "40th" is an indicator taking the value of one if the workers' initial-task earnings fell below the 40th percentile threshold but above the 20th percentile threshold. The excluded category is "above the 80th percentile." [**3**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**4**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A6: *Behavior by Initial-Task Earnings Quintiles and Revealed Income Inequality*

	(1)	(2)	(3)	(4)
Initial-task Earnings Percentile		Phase 2		Phase 3
	quantity	quality	lying	donations
20th	-0.76*	-0.05**	0.13*	0.07
	(0.425)	(0.025)	(0.073)	(0.063)
40th	0.19	-0.01	0.01	-0.02
	(0.489)	(0.023)	(0.066)	(0.052)
60th	-0.03	-0.02	0.15**	0.01
	(0.499)	(0.015)	(0.073)	(0.045)
80th	0.20	-0.02	-0.03	-0.01
	(0.413)	(0.020)	(0.044)	(0.048)
Visible Income Ineq	-0.22	-0.01	-0.00	-0.05
	(0.386)	(0.018)	(0.044)	(0.056)
(Visible Income Ineq) x (20th)	-0.00	0.04	-0.13	-0.02
	(0.658)	(0.034)	(0.082)	(0.087)
(Visible Income Ineq) x (40th)	-0.09	0.02	0.13	0.09
	(0.612)	(0.029)	(0.109)	(0.086)
(Visible Income Ineq) x (60th)	-0.10	-0.01	-0.11	0.14
	(0.803)	(0.041)	(0.097)	(0.083)
(Visible Income Ineq) x (80th)	0.01	0.01	0.06	-0.07
	(0.580)	(0.035)	(0.066)	(0.074)
Ability-intensive initial task	-0.24	-0.03**	0.07**	0.03
	(0.189)	(0.014)	(0.032)	(0.024)
Phase 2 quantity		0.01	-0.01	-0.03***
		(0.007)	(0.007)	(0.006)
Constant	6.92***	0.98***	0.10	0.09
	(0.906)	(0.092)	(0.114)	(0.136)
Observations	419	419	419	419
R-squared	0.104	0.054	0.088	0.116

Notes: [**1**] Each column reports an OLS regression using all observations, i.e., pooling the Ability-Intensive and Effort-Intensive treatments. Dependent variables are labeled in the column headings. These are: "Phase 2 quantity" refers to the number of scantrons submitted; "Phase 2 quality" refers to the average accuracy of submitted scantrons; "Phase 2 lying" is the proportion of participants who misreported the number of Scantrons they submitted; "Phase 3 donations" refers to the total amount donated across both charity options (American Red Cross and the South Plains Food Bank). [**2**] Control variables include: a set of indicators for the workers' initial-task earnings quintile; an indicator for our 'Visible Income Inequality" treatments, where workers were provided a binned relative frequency chart of own-session initial-task earnings; interactions between earnings quintile and our 'Visible Income Inequality" treatments. [**3**] The initial-task earnings quintile indicators are mutually exclusive and computed for each task type, separately, so that "40th" is an indicator taking the value of one if the workers' initial-task earnings fell below the 40th percentile threshold but above the 20th percentile threshold for the specific initial task type they performed. The excluded category is "above the 80th percentile." [**4**] Included in each estimate, but not reported for readability, are additional controls: gender, age, high- and low-GPA dummies, high- and low-family-income dummies, as well as dummies for the type of filler questions inserted between stages 1 and 1.5, which were randomly assigned and asked about gender, political affiliation or a purchasing decision. [**5**] Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.