Inequality as an Incentive

Jeffrey V. Butler
University of California, Merced

Eric Cardella
Texas Tech University *

This version: *September 6, 2020*

Abstract

We study the non-pecuniary incentives of salient inequality. In a laboratory experiment involving over 420 participants we randomly vary piece-rate wages for performance on real-effort tasks. Real-effort tasks are randomly assigned across treatments to be either ability- or effort-intensive. After completing their task, all participants learn their own earnings and may receive information about others' wages and incomes. This additional inequality information is randomly assigned across treatments to create four information conditions: participants learn about wage inequality, income inequality, both or neither. After the inequality-revelation stage, participants complete an additional real-effort task and have an opportunity to donate to charity. We find that revealing wage inequality on the ability-intensive task *increases* effort on the subsequent task but, at the same time, reduces charitable donations. We find generally weaker effects from revealing income inequality and no effects at all when the initial task is effort-intensive. We propose a novel mechanism that can explain the behaviors we observe and provide suggestive evidence supporting this mechanism from a post-experiment survey.

JEL Classification:

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

Wage and income inequality is a hotly debated topic, both in general discourse and in the academic literature. In some contexts, revealing essentially negligible horizontal inequality produces strong reactions, reducing morale and productivity or increasing turnover and anti-social behavior (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). In other contexts, even pronounced inequality elicits little backlash: students demand decidedly dispersed grading distributions; professional athletes or tech entrepreneurs who oppose the extreme compensation inequality within their occupations are the exception rather than the rule.

These seemingly stark differences in reactions to earnings inequality suggest there may be a lot to be gained from understanding their contextual determinants. A question of first-order importance is whether the reaction to horizontal inequality is qualitatively similar across contexts – being generally weakly negative as in the studies cited above – or, instead, whether salient inequality may sometimes serve as a positive (non-pecuniary) incentive, increasing effort over monetary incentives alone. This question is the focus of our study, which we view as a point of departure into the broader question of how non-pecuniary inequality incentives should be incorporated into a theory of wage-setting where context matters.

The importance of our central question is highlighted by the observation that horizontal wage inequality is a prescription of many theoretically optimal contracts. In principal-agent problems, for example, where effort is unobservable or unverifiable, standard analysis implies tying agents' earnings to outcomes to mitigate moral hazard. Outcomes have the benefit of being observable and contractable, but they are only stochastically related to agents' (effort) choices. If there is an idiosyncratic component to the effort-outcome relationship, then even identical employees exerting identical effort may end up receiving quite different compensation levels. Consequently, horizontal wage inequality would be a general feature of equilibrium.[1] Tying earnings to relative performance through tournament incentives is another commonly prescribed horizontal-inequality inducing practice justified by traditional economic theory (e.g., Gibbons 1987; Hopkins and Kornienko 2009, 2010; Lazear 2000; Stiglitz 1975).[2] These traditional prescriptions, however, only take into account the

[1] When employees are not identical, optimal contracts in this setting may entail ex ante wage inequality as well in order to account for variation in risk tolerance or effort costs. Employees' endogenous responses to such ex ante inequality may exacerbate ex post inequality.

[2] Other theoretical justifications in related settings include arguments from efficiency, where inequality permits high-ability workers to credibly signal their ability (Spence 1973). Vertical inequality within a firm across different job types may also be justified theoretically through its effects on beliefs about expected future returns to current effort (e.g., promotions) (Cullen and Perez-Truglia 2018). Yet another justification is due to Rosen (1982, p. 311): "Assigning persons of superior talent to top positions increases productivity by more than the increments of their abilities because greater talent filters through the entire firm by a recursive chain of command technology. These mul-

pecuniary incentives associated with wage contracts. If non-pecuniary inequality incentives vary systematically and substantially across contexts, then the optimality of the contracts prescribed by traditional economic theory may need to be reconsidered.

Our study examines experimentally whether and how the non-pecuniary incentives of wage and income inequality vary systematically and substantially across contexts. To make this question tractable, we single out a couple of contextual factors. One of these factors is the type of work being performed, whether productivity primarily depends on ability or on effort. This factor is potentially economically important to the extent that it captures the essential distinction between skilled and unskilled sectors of the economy, yet underexplored.[3] The other factor we single out is the type of inequality. We examine two types that are often conflated in existing research: wage inequality, which is associated with ex ante unequal treatment; and income inequality, which incorporates individuals' endogenous responses to ex ante contextual factors, and should thus be thought of as more of an ex post measure of inequality.

In a laboratory experiment we randomly assign our student "employees" one of two real-effort tasks and also randomly assign them one of two piece-rate wage levels. Productivity on one of these tasks depends primarily on ability, the "ability-intensive" task, while productivity on the alternative task depends primarily on effort.[4] We refer to this alternative task as the "effort-intensive" task. The two piece-rate wage levels implement the same marginal monetary performance incentives. While completing their assigned task, employees only know about the existence of their own wage level and their own task type.

Immediately after completing the task, participants learn of their own performance and earnings and may receive additional information about wage or income inequality in their session. We vary this additional information across treatments to create four inequality-information conditions: no inequality information revealed; revealed wage inequality only; revealed income inequality only; and revealed wage and income inequality. After this inequality-information revelation phase, participants complete an additional real-effort task for a common piece-rate wage and have the opportunity to donate some of their experimental earnings to charity. Both the subsequent real-effort task and the charitable donation opportunity are surprises – participants do not know about the second real-effort task when completing the initial real-effort task, nor do they know about the

tiplicative effects support enormous rewards for top level management in large organizations." For recent overviews of the literature on incentives in personnel economics and in economics more generally, see Lazear and Oyer (2007); Prendergast (1999) and Lazear (2018).

[3]To the best of our knowledge, the only other paper that explicitly studies this factor in the context of reactions to (wage) inequality is Butler (2016). Conversations with the authors suggest that other papers, particularly Hart and Piff (2018), may also admit this interpretation and thus could be considered as complementary evidence of the robustness of our findings.

[4]Of course, as with any task, productivity always depends on effort, ability and luck to varying degrees.

charitable donation opportunity when they are completing the second real-effort task. Examining behavior in both of these subsequent situations allows us to infer how the valence of reactions to inequality, i.e., whether reactions are positive or negative, depends on our contextual factors. We complement this behavioral data with a post-experiment survey that provides suggestive evidence about possible mechanism(s) relating inequality to subsequent behavior.

As a preview of our results, we find that revealing *wage* inequality related to the (intial) ability-intensive task has the strongest effects on subsequent behavior and that these effects are concentrated on the low-wage employees, *increasing* both the quantity and the quality of their subsequent production while *reducing* their charitable donations. Revealing income inequality has little effect on low-wage employees and a negative effect on the subsequent quantity produced by high-wage employees. In the alternative work environment, where the initial task is effort-intensive, we find no significant main effects of wage or income inequality revelation. We argue that these patterns are not consistent with many of the explanations inequality reactions offered in the existing literature and then go on to detail a novel mechanism, building on prior research in economics as well as in social psychology, that *can* explain the data. This mechanism relies on contextual variation in cognitive processes strengthening "Just World Beliefs" Lerner (2013, 1965) in reaction to intial-task inequality. Suggestive supporting evidence is provided by our post-experiment survey.

The primary contribution of our study is to provide evidence that non-pecuniary incentives do, in fact, vary systematically, substantially and predictably across contexts. Moreover, we show that these non-pecuniary incentives need not always be (weakly) negative. In a particularly important context, i.e., when work is ability-intensive, the non-pecuniary productivity incentives of wage inequality may actually be positive so that revealing wage inequality induces more effort than concealing it. One implication of this finding is that optimal wage-secrecy policies may vary by industry, by sector of the economy or by occupation. A secondary contribution is to offer a novel mechanism together with suggestive evidence supporting this mechanism that can potentially explain the patterns in behavior that we observe in our data.

The remainder of the paper proceeds as follows. First, we present our experimental design in detail. In section 3 we provide results pertaining to employee behavior. Next, we argue that the patterns in behavior we observe are inconsistent with several explanations common in the literature. We go on to lay out a novel mechanism that can explain the empirical (behavioral) results and provide attitudinal evidence from the post-experiment survey supporting this mechanism. In the penultimate section, we discuss closely related literature. We summarize and conclude in the final section.

2 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 an information revelation 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 a purely informational phase, which we denote by assigning it a fractional phase designation, in which all participants learned of their own Phase 1 performance and experimental income while some participants, depending on treatment, received additional information about relative wages and incomes. The experiment was multi-modal in order to introduce psychological separation between the phases. Phase 1 and 1.5 were programmed in oTree (Chen, Schonger, and Wickens 2016), Phases 2 and 3 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.

2.1 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.

One of these two tasks, which we refer to as the ability-intensive (AI) task, consists of 48 Raven's Advanced Progressive Matrices. The other task, referred to as the effort-intensive (EI) task, asked participants to count the number of typographical errors in each of 48 sentences, 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 intelligence. We convey this interpretation to our participants in the instructions preceding the AI task and therefore consider it a plausible assumption that participants perceived the task as being ability-intensive.

We constructed the EI task to be as parallel as possible to the AI task. Like the AI task, the EI 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 performance on the EI task depends primarily on effort and that, moreover, our workers perceived this to be the case.

Despite our efforts to make the AI and EI tasks as parallel as possible, they appear quite different on a computer screen. The AI task involves a grid of images, while the EI 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. Additionally, by keeping marginal (monetary) incentives for performance constant across the two pay schemes – \$0.20 per correct response – we mitigate the possibility of differences in performance due solely to monetary incentives (c.f., Butler 2016).

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 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. They were instructed that a timer would appear at the top of the computer screen during the task.

2.2 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 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 not, of the existence of the 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, which we denote by the acronym **VW**. Concealing this information implements invisible wage inequality, denoted **IW**.

The other domain of information we manipulated was relative (experimental) income. Specifically, in the visible income inequality condition, denoted (**VI**), 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, denoted (**II**), 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 VI 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, all participants in a particular session were assigned to either the VI condition or the II condition and which sessions received which condition was randomly 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 (task type) full factorial design with eight distinct treatments (Table 1). For ease of exposition, we refer to each treatment using the format [VW/IW]-[VI/II]-[AI/EI]. As an example, VW-VI-AI refers to the treatment with visible wage inequality, visible income inequality and an ability-intensive Phase 1 task.

Before describing the rest of the experiment it is worth noting that because information about wage and income inequality was provided to participants only after they completed the Phase 1 task,

Table 1: *Summary of Treatments*

Treatment Name	Task	Visible Wage Inequality	Visible Income Inequality	N
IW-II-AI	Raven's Matrices	No	No	56
VW-II-AI	Raven's Matrices	Yes	No	58
IW-VI-AI	Raven's Matrices	No	Yes	49
VW-VI-AI	Raven's Matrices	Yes	Yes	48
IW-II-EI	Proofreading	No	No	54
VW-II-EI	Proofreading	Yes	No	53
IW-VI-EI	Proofreading	No	Yes	52
VW-VI-EI	Proofreading	Yes	Yes	53

Notes: **[1]** Raven's matrices refer to Raven's Advanced Progessive matrices, commonly thought to be a culture- and language-free measure of general intelligence.measure. **[2]** For the proofreading task, we randomly inserted simple forms of typogaphical errors into sentences and participants had to count the number of typos appearing in each sentence. **[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.

performance and earnings on the Phase 1 cannot plausibly be affected by these manipulations. This design feature allows us to cleanly identify how knowledge of wage and income inequality impacted subsequent productivity and pro-sociality by observing behavior in the subsequent phases of the experiment.

2.3 Phase 2: Scantron Task

The Phase 2 task was another incentivized real-effort task designed to be ecological 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.[5] 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 completed and

[5]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.

would be paid, *based solely on this self-report*, $1 for every completed Scantron.[6]

Immediately after the 15 minutes had elapsed, participants placed all of their Scantrons, completed or not, in a box at the back of the room. The box was not monitored by the (lone) experimenter present in the room. Only after returning to their carrels did each participant fill out an anonymized payment slip, noting only their carrel number and the quantity of Scantrons they produced. Importantly, it was obvious that for logistical reasons neither the quantity nor quality of Scantrons produced were observable by the experimenter during the session. Great pains were taken to ensure, and to convey implicitly to participants, 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 the participant's carrel number, and that, moreover, we could never match names to carrel numbers.[7] We argue that this provides essentially the same level of anonymity as other commonly used experimental designs, particularly those where pay is based on self-reported behavior such as experiments in the dice-rolling paradigm (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.[8] The primary advantage of our Phase 2 design over the more common aggregate-decisions-only design is that we can: i) measure actual quantity and quality produced; and ii) link decisions, self-reports, demographics and attitudes at the (anonymized) individual level across all phases of the experiment.

[6]In one interpretation, this allows participants to behave anti-socially, i.e., "cheat" their employer, in two distinct ways. One way would be to simply inflate their self-reported number of completed Scantrons and very deliberately lie for financial gain. The second way is more subtle. Participants could "shirk" by completing Scantrons inaccurately, which presumably requires less time and effort, while still reporting truthfully the number they produced. Shirking enables participants to earn a given amount of money with less effort, or to complete more Scantrons in the allotted time with a fixed level of effort, than if they filled out versions accurately. To the extent that employers care about accuracy, this also harms the employer by benefitting the employee.

[7]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. Moreover, 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. As an additional measure of anonymity which is uncommon, we also did not have participants sign receipts.

[8]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.

2.4 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. 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.[9] 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. As an added anonymity step, 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.

2.5 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. 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.

2.6 Matching measures across phases

Well after each session was completed, we matched outcome measures from all phases of the experiment.[10] Matching was accomplished by carrel number, which was automatically recorded for the

[9] 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.

[10] 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.

Figure 1: *Timeline*

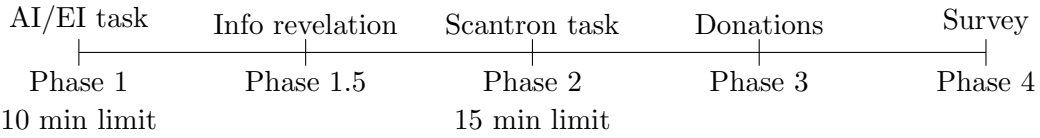

computerized portions of the experiment (Phase 1 and the post-experiment survey). For the other phases of the experiment, participants noted their carrel numbers on materials they submitted. It is worth noting that it is not possible for the experimenters to match participants' names to carrel numbers or even payment amounts as, e.g., no receipts were collected.

3 Results

As an initial step to provide reassurance that our randomization into treatments was successful, we conducted simple balance checks using the limited set of demographics we collected on the post-experiment survey. Demographics included age and gender as well as (self-reported, categorical) measures of family income and GPA. For each of these variables separately we conducted a Chi-squared test with the null hypothesis of independence across treatments. The only test that was significant at even near-conventional levels was the test associated with GPA ($p = 0.084$) suggesting that, overall, randomization into treatments was successful. As a precaution, however, in our formal econometric estimates below we control for available demographics.

Next, in Table 2, we report the raw means of our primary outcome variables for all treatments and pay schemes. For the initial (Phase 1) task, there were no apparent differences in performance within each task across treatments. This is consistent with task-specific performance being determined by marginal monetary incentives in the absence of inequality revelation. Across tasks, performance is consistently higher on the effort-intensive (EI) task which would be consistent with some additional factor, such as ability, constraining average performance on the AI task but not on the EI task.

For Phase 2 we report the average number of Scantrons actually submitted (quantity), which hypothetically ranges from 0 to 15, and ranges from 1 to 15 in the data. For our quality measure we report the average proportion of correctly bubbled items (out of 20) across all submitted non-blank Scantrons.[11] Finally, although it is not the focus of our study we also report the prevalence of "lying," i.e., misreporting the number of Scantrons produced.[12] There are no obvious patterns

[11] More precisely, we restrict attention to Scantrons on which a participant bubbled at least one item, compute the proportion of correct answers for each of these Scantrons, and then take the average of these proportions for each participant.

[12] We construct a simple indicator for the self-reported quantity not matching the participant's actual production

across treatments or pay schemes in our Phase 2 outcome measures. Overall, we note that quality is surprisingly high despite the lack of formal quality incentives and that lying is surprisingly rare given the presence of substantial financial incentives to misreport. 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% by lying. In the last column, we report the average proportion of participants' total Phase 1 and Phase 2 earnings donated to charity in Phase 3. A prominent pattern is that, with the lone exception of treatment VI-II-AI, it was typically low-wage employees who donated a larger proportion of their experimental earnings to charity.

3.1 Econometric estimates

While descriptive statistics are useful in providing an overview of the data, the large number of treatments involved complicates our ability to discern any patterns or effects. More formal econometric estimates 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. For ease of exposition we proceed by splitting the sample by Phase 1 task type and estimating our econometric models on each of the resulting subsamples separately. Since randomization into task type occurred only across sessions, this is a (relatively) clean split: no individual performing the AI task could have possibly affected any individual assigned to the EI task as, for instance, the relative performance information for AI-sessions never included any EI-session component. This route has the advantage of letting us feasibly report the coefficients associated with all of our experimental factors and their interactions, facilitating the interpretation of the estimated magnitudes.

3.2 Ability-intensive environment

We begin by analyzing the AI-session data and in Table 3 report OLS estimates of our primary behaviors of interest. As explanatory variables, we include indicators for high wage (HW), visible wage inequality (VW) and visible income inequality (VI) as well as all interactions among these variables. Each estimate also includes demographic controls that we omit for readability.[13] Finally, to account for the possible reactions to one's own initial-task earnings, such as to falling short of

quantity. Almost all misreports involved over-reporting one's production by exactly one Scantron. This suggests intentional misreporting – otherwise, we would expect some over-reports and some under-reports – and also implies that we lose little by using an indicator function instead of, for example, the difference between the number produced and the number reported.

[13]These demographic controls were almost always non-significant. We note exceptions to this pattern in either the text or in the table footnotes.

an (ex-ante) subjective expected earnings level, in columns 2-5 we control Phase 1 earnings (Phase 1 income).[14]

The estimate in column 1 serves as an important specification check. The absence of any significant effect of any of our exogenously manipulated experimental factors on initial-task performance is reassuring evidence that participants: i) did not anticipate the inequality in their sessions; and that ii) in the absence of non-pecuniary inequality incentives, marginal monetary incentives were decisive.

3.2.1 Revealing wage inequality

In columns 2-5 we consider behavior after our inequality-information phase. We focus first on the main effects of revealing wage inequality (VW and HWxVW). Among low-wage workers we find that visible wage inequality increased the subsequent quantity of Scantrons produced (column 2) and, controlling for quantity, also the subsequent quality (column 3) of production. Relative to the respective pooled AI-session measures, the one-unit increase in quantity implied by the coefficient on VW in column 2 constitutes about one-half (48.9%) of a standard deviation, while the 7 percentage point increase in quality also represents roughly one-half (45.5%) of a standard deviation. For high-wage workers, in stark contrast, we find that revealing wage inequality has little effect on the subsequent quantity or quality of production. For quantity, the sum of coefficients on VW and HWxVW in column 2 is essentially zero or, if anything, slightly negative. At the same time, while the reduction in quality among high-wage workers is not by itself statistically significant, the point estimate of the effect of visible wage inequality on high-wage workers' production quality – $0.07 - 0.04 = 0.03$ – is not substantial in magnitude, representing less than one-fifth of a standard deviation of our AI-session quality measure.[15] We summarize these patterns as our first result:

Result 1: *Revealing wage inequality associated with an ability-intensive Phase 1 task significantly and substantially increases the subsequent quantity and quality of (Scantron) production among low-wage employees. In stark contrast, revealing wage inequality has no effect on the subsequent quantity of production among high-wage workers and a non-substantial effect on their quality of production.*

Because Phase 2 earnings were based on self-reports rather than actual production, one possible way to rectify Phase 1 wage inequality would be to inflate these self-reports. In column 4 we

[14] To address the concern that Phase 1 income depends partially on the randomly assigned wage level, we conducted the same estimates using initial-task performance instead of initial-task income. None of our findings, including patterns of statistical significance, were changed by this substitution.

[15] The point estimate of $0.07 - 0.04 = 0.03$ represents 19.5% of a standard deviation of the pooled AI-quality measure. In practice, a threshold of one-quarter of a standard deviation in magnitude, as captured by commonly used formal measures like Hedges-g, is commonly used to separate negligible effects from substantial effects.

estimate a linear probability model of "lying" as a function of our experimental factors.[16] Despite substantial financial incentive to lie, only about 13% of AI-session participants overall lied and, when they did lie they lied minimally. More importantly, we find no significant variation in lying across our experimental factors.

In the last column, column 5, we estimate an OLS model of the proportion of experimental income donated to charity in Phase 3. Among low-wage workers, we find a significant and substantial *decrease* in pro-sociality or other-regard as measured by charitable donations. The 12 percentage point reduction in the proportion of experimental earnings donated to charity constitutes more than 40 percent (42.5%) of a standard deviation of the pooled AI-session donation measure; it also constitutes over one-half (56.6%) of the pooled AI-session donation-proportion average. In stark contrast, the point estimate of the effect of revealing wage inequality to high-wage workers is essentially zero ($-0.12 + 0.11 = 0.01$) even though the coefficient on HWxVW is not by itself significant. We summarize these patterns in our second result.[17]

Result 2: *Revealing wage inequlity to low-wage employees significantly and substantially* reduces *subsequent other-regard as measured by the Phase 3 charitable donations. In contrast, the point estimate of the magnitude of this effect on high-wage employees is essentially zero.*

3.2.2 Revealing income inequality

We turn next to the effect of revealing *income inequality* in the same, ability-intensive, environment. The direct effects, as captured by the coefficients on VI and HWxVI, suggest that revealing income inequality by itself generally has no significant effect on subsequent behavior. The lone exception to this pattern is that learning about income inequality significantly and substantially reduces the subsequent quantity of Scantrons produced by high-wage employees. The point estimate of the reduction, 0.99 units, amounts to about one-half (48.4%) of a standard deviation of the pooled AI-session quantity measure and 13.5% of the AI-session mean.

Considering next interaction effects, we also generally find a lack of any significant effects on subsequent behavior with the exception of the quantity of Scantrons produced in Phase 2. The negative and significant interaction effect of revealing income inequality in addition to wage inequality to low-wage employees (VWxVI) is large enough in magnitude to more than completely

[16] Recall, we constructed an indicator for whether a participant's self-reported quantity of production matched his or her actual quantity produced. Because in the vast majority of cases, an inaccurate report was off by only one Scantron, we lose little by using a dummy variable. Additionally, in the vast majority of cases where an inaccurate report occurred, it was an over-report, so characterizing the misreporting as lying rather than simply a mistake seems warranted.

[17] We find the same results using the donation totals rather than donation proportions, so the patterns are not driven by low-wage workers simply having less to donate.

overturn the positive effect from revealing wage inequality alone. The implied point estimate is a reduction in the quantity of subsequent Scantron production of 0.47 units, or slightly less than one-quarter (23%) of a standard deviation of pooled AI-session production. Among high-wage workers, revealing income inequality in addition to wage inequality mutes the negative subsequent quantity response associated with revealing income inequality alone. The sum of the coefficients on VWxVI and HWxVWxVI is positive and implies an interaction effect point estimate of 0.77, or 37.7% of a standard deviation of (pooled AI-session) quantity. We summarize these patterns as our third result.

Result 3: *Among low-wage workers, revealing income inequality alone has no significant effect on subsequent behavior but, when revealed in addition to wage inequality, eliminates the positive effect of wage inequality on subsequent production quantity. We find qualitatively different patterns among high-wage workers. For this group, revealed income inequality alone reduces subsequent prodution quantity but exhibits a positive interaction with wage inequality, muting wage inequality's negative effects on subsequent quantity produced.*

3.3 Effort-intensive environment

We next turn to the effort-intensive environment and re-estimate all of our econometric specifications (Table 4). The picture that emerges is easily summarized. We find no significant effect of any of our experimental factors or their interactions on subsequent behavior. Among non-experimental factors, the lone significant relationship we find is in column 5. We find a negative and significant association between the number of Scantrons produced (Phase 2 quantity) and the proportion of total experimental income donated to charity. To put this in context, donation proportions in the EI-session data have a standard deviation of 0.275 (mean = 0.183) while the actual number of Scantrons produced ranges from 1 to 15. Taking the point estimate of -0.02 at face value, the proportion of experimental earnings donated to charity by those who produce the maximum number of Scantrons is about one standard deviation lower than those who produce the minimum number of Scantrons.

Result 4: *In an effort-intensive environment, we find no significant effects of revealing wage or income inequality on subseqeunt behavior.*

Summing up, reactions to inequality exhibit substantial context dependence. When inequality is associated with an effort-intensive task, we find no evidence for any effect on subsequent behavior. When the inequality is associated with an ability-intensive (initial) task, however, we find evidence

for strong reactions that differ qualitatively depending on relative wage level and the type of inequality revealed. Revealing wage inequality has positive effects on the subsequent productivity of the disadvantaged, i.e., low-wage workers, but negative effects on their pro-sociality. For high-wage workers, revealed wage inequality negatively affects subsequent productivity without affecting pro-sociality. The effects of revealing income inequality are concentrated on subsequent quantity of production only and, again, vary qualitatively by relative wage level: revealed income inequality overturns the positive quantity effects associated with revealed wage inequality among low-wage workers, while this interaction is positive for high-wage workers.

4 Potential mechanisms

4.1 Inconsistency with many common explanations

An oft-conjectured mechanism for explaining reactions to wage or income inequality is Equity Theory (Adams 1963). Central to this theory is the concept of an equity ratio, defined as the ratio of a worker's own perceived total compensation from the firm (money earnings or other benefits) to his or her contribution to the firm (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. It should be apparent, then, that the dual pattern we document among low-wage workers is difficult to reconcile in an Equity Theory framework. Essentially, a straightforward application of Equity Theory would explain increased productivity as a way to lower a too-high equity ratio, while withholding donations, and hence increasing compensation, is 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.

Stepping outside of Equity Theory to assume strong lying aversion may reconcile the quantity response: if one cannot lie, then increasing production quantity is one way to increase compensation so that, just like reducing donations, it could raise a too-low equity ratio. However, this does not rescue Equity Theory as an explanation. There would still be the issue of the increase in *quality* that we observe. There are neither direct monetary incentives nor obvious psychic costs associated with the quality dimension of production – participants never report, nor are they compensated for, quality. To the contrary, if low-wage workers are increasing quantity primarily to increase earnings, then a standard quantity-quality tradeoff argument should imply, if anything, *lower* production quality from this group.[18] Overall, Equity Theory, even augmented by lying aversion, cannot

[18] Another way Equity Theory could work as an explanation is if reactions are to the immediately previous task.

provide a unified explanation for our findings.

Another obvious obvious explanation for the patterns in our data could be loss aversion (Kahneman and Tversky 1979). A plausible assumption for the reference point is ex ante expected earnings. In this mechanism, low-wage participants would be more likely to come up short of their reference point at the end of the experiment, which would increase their marginal utility for money during the experiment. If, as above, lying is not feasible because of high psychic costs, then the increased production quantity and reduced charitable donations among low-wage workers could both be explained by their higher marginal utility for money. As straightforward as this explanation seems, it has several issues. First and foremost, there is no obvious channel for the revelation of wage inequality to matter: learning others received a relatively high wage does not affect whether you will fall short of your own reference earnings level. The second drawback is similar to the objection above: even if we assume strong lying aversion, it is still difficult to explain the increased *quality* of production that we observe. Quality has no direct positive monetary incentives attached, may have indirect negative monetary incentives (quantity-quality tradeoff) and does not obviously involve lying.

A third class of potential explanations is distributional social preferences (Bolton and Ockenfels 2000; Charness and Rabin 2002; Fehr and Schmidt 1999). In their standard formulations, these models would predict that behavioral reactions to inequality would be driven by features of the distribution of money earnings alone. Because of this feature, there are several issues with positing distributional social preferences as the primary explanation for our findings. The first is familiar, so we will not restate it in detail. We simply note that the puzzle of increased production quality remains even if we make assumptions that would explain increased production quantity as a vehicle to increase earnings. A second issue is that a primary concern about one's place in the distribution of money earnings, as these models posit, leaves little room for revealed *wage* inequality to matter once the *income* distribution has been revealed. While revealed wage inequality *could* play a role in these models by affecting beliefs about the (unknown) income distribution, once the income distribution is revealed wage inequality should clearly not affect beliefs about the (known) income distribution and, in turn, should not affect behavior.[19] Our data appears to contradict this prediction. For

So, for instance, low-wage workers could perceive their Phase 1 compensation as *too high* when completing the Phase 2 task and increase effort to restore equity. Then, having put in a lot of effort relative to their compensation in Phase 2, when they reach Phase 3 they reduce donations to increase their overall compensation from the experiment in order to, once again, restore equity. We cannot rule this out, but the story would require some explanation as to why learning one's wage level is relatively low leads one to perceive one's compensation-to-contribution (equity) ratio as relatively high.

[19]These models typically leave unspecified what the proper reference population is. We make an assumption in our argument that an obvious choice would be the participants in one's own session. If the reference population is wider (e.g., all participants in the experiment), then there could be scope for wage inequality to matter but this would take us outside of the realm of the standard analyses of distributional social preferences models.

instance, in Table 3, among low-wage workers the point estimate of the effect of revealing wage inequality in addition to income inequality on the proportion of experimental earnings donated to charity ($-0.12 + 0.03 = -0.09$) is substantial, at 31% of a standard deviation of the AI-session measure, while the effect on subsequent production quantity (-0.47) is nearly substantial, at 23% of a standard deviation of the AI-session measure.

4.2 An alternative mechanism: inequality-enhanced JWBs

As an alternative to the most widely accepted models for reactions to inequality highlighted above, we conjecture a novel alternative mechanism that builds on previous research in economics and in social psychology. The primary forces underlying this novel mechanism are cognitive dissonance, as in Akerlof and Dickens (1982), and "Just World Beliefs" (JWBs) (Lerner 2013, 1965), which refer to a deep-seated, universal and fundamental need to believe that people generally "get what they deserve."

A common and intuitive example of JWBs in action is the automatic tendency to "blame the victim," but the original experimental evidence involved showing that individuals subconsciously rationalize third-party pay differences as being *deserved*, despite being explicitly informed that pay had been randomly assigned (Lerner 1965). Butler (2016) extended this original finding to personal (own vs. other) wage disadvantage and documented context-dependence: ex post rationalization of one's own low wage was more likely to occur in an ability-intensive environment than in an effort-intensive environment.[20] The broad economic implications of JWBs are laid out theoretically and empirically in Benabou and Tirole (2006). The authors highlight the role of relatively strong JWBs in producing an "American/*laissez faire*" equilibrium, characterized by a strong work ethic and an emphasis on self-reliance, in contrast to a "European" equilibrium produced by weak JWBs and characterized by the opposite preferences. Heuristically, JWBs provide intrinsic effort incentives and ameliorate moral hazard. If people generally get what they deserve, then subjectively expected compensation accrues irrespective of the employer observing, or contracting on, effort. JWBs may reduce other-regarding behavior as well, since if other people generally get what they deserve, then the expected consequences for others of one's own failure to act are lessened.

[20] On context-dependence, there are many possible explanations but Butler (2016) speculates, without direct evidence, that an ability-intensive environment is reminiscent enough of typical meritocratic settings to bring to mind a narrative in which inequality is not only consistent with, but actually required by, a just and fair world. A quintessential example would be the finals of an objectively scored Olympic sport, such as a track event. It would be considered grossly unfair for the second-place finisher to receive a better medal than, or even the same medal as, the first-place finisher. The contextual activation of a meritocracy narrative permits a relatively easy (cognitive) path to avoid the cognitive dissonance that would otherwise result from simultaneously holding strong JWBs and observing unjustifiable inequality. In a phenomenon reminiscent of "victim blaming," individuals (subconsciously) distort their beliefs about the basis of inequality to subjectively perceive it as being deserved. The primary difference between victim-blaming and this process is, of course, that the "victim" is oneself.

We build on all of these bodies of evidence and extend them to conjecture that once experienced inequality has been subconsciously ex post rationalized as warranted, deserved and, ultimately, fair, this experienced inequality becomes "evidence" in favor of a just world. As a (perverse) consequence, experienced inequality actually strengthens JWBs. This mechanism yields several predictions. First of all, because previous research shows wage inequality is more likely to be rationalized in an ability-intensive environment, we should be more likely to find the behavioral residue of strengthened JWBs – enhanced work ethic and emphasis on self-reliance – in an ability-intensive environment than in an effort-intensive environment. Secondly, because the threat of cognitive dissonance is the most severe for low-wage workers, it is among this group that we would expect subconscious inequality rationalization to to be most likely to occur.[21] It is therefore low-wage workers who should be most likely to exhibit the dual (behavioral) pattern. Thirdly, because our conjectured mechanism is relatively extreme, it is less likely to be activated when there are other, less extreme, alternatives. Prior research has documented such alternative rationalization channels for income inequality specifically.[22] Moreover, revealing income inequality may make the threat of cognitive dissonance less severe since the income distribution conveys *some* positive information about relative earnings for almost everyone, i.e., for anybody who is not in the lowest earnings bin. Therefore, we would expect the revelation of income inequality to make the dual pattern – increased work ethic and decreased other-regard – less likely. Overall, our novel mechanism provides a unified explanation for many of the patterns we find in the data: i) the context-dependence of the effects of inequality revelation; ii) the relative strength of the effects of the revelation of wage inequality compared to income inequality revelation; and iii) the puzzling dual pattern of increased productivity but reduced pro-sociality among low-wage workers.

Beyond mere consistency with the data, it would be ideal to have some supporting evidence of the mechanism. Since the mechanism ultimately operates through strengthened JWBs, reassuring evidence of the mechanism itself would be variation in JWBs consistent with our sketch of its operation above. Most importantly, this would entail showing that revealing wage inequality to low-wage workers in an ability-intensive environment strengthens their JWBs since this phenomenon would be the most difficult to reconcile with other explanations.

To provide such evidence, we collected a straightforward measure of JWBs on our post-experiment survey.[23] Participants were asked how much they agreed with the statement: "people generally get

[21] Recall, the goal of the subconscious rationalization is the avoidance of cognitive dissonance.

[22] Income inequality incorporates individuals' endogenous responses to pecuniary and non-pecuniary incentives. For example, Breza et al. (2018) argues that observing others who work harder eliminates the backlash from earnings inequality ostensibly because people can attribute others' higher earnings to their harder work.

[23] We note an important caveat. We collected our measure of JWBs only at the end of the experiment, on the post-experiment survey. Because participants made many decisions between the inequality-information revelation phase and the post-experiment survey phase, we can not make explicit causal claims about the relationship between

what they deserve." Responses were collected on a 7-point Likert scale ranging from 1="totally disagree" to 7="totally agree." Using this measure of JWBs as our dependent variable, in Table 5 we estimate the same models as in our primary analyses of behavior.[24] For ease of exposition, we again split the data by the type of Phase 1 task. The first two columns restrict attention to AI sessions, while the remaining columns use only observations from sessions with an effort-intensive initial task.

Examining the first two columns, consistent with our novel mechanism we find that revealing AI-task wage inequality to low-wage participants substantially and significantly *increased* their JWBs. The coefficient on VW implies a 0.53 point increase in JWBs, which represents over one-third (36.8%) of a standard deviation of AI-session JWBs. Revealing income inequality has little effect on JWBs, either alone (VI) or in terms of its interaction with revealed wage inequality (VWxVI). The weaker effect of income inequality on JWBs on low-wage workers is also consistent with our conjecture. None of our other experimental factors has a significant main or interaction effect on JWBs.

Turning to the effort-intensive environment (columns 3-4), we find a qualitatively different picture that is broadly consistent with what might be expected in the absence of the subconscious ex post rationalization of inequality. Low-wage workers responded to learning about (inexplicable) wage inequality with weakened JWBs, suggesting they perceived this disadvantage as unfair. Revealed income inequality had no significant main effect on their JWBs but a marginally significant countervailing interaction with wage inequality, perhaps reflecting the "some good news" character of the income distribution for almost everybody. High-wage workers were apparently not fazed by wage inequality in their favor. The coefficient on HWxVW is highly significant, equal in magnitude but opposite in sign to the coefficient on VW. They did perceive the high-wage scheme, which involves positive payment for wrong answers, as unfair. Finally, there is a negative relationship between initial-task performance and JWBs on average. This could reflect disappointment with poor performance being redirected toward JWBs or, of course, it could be evidence of the intrinsic motivational effect of JWBs and reflect the pattern of less motivated participants exerting less

inequality revelation and JWBs: the intervening decisions may obviously have their own effects on JWBs. Patterns in JWBs may nonetheless provide suggestive evidence about the plausibility of our conjectured mechanism.

[24] One exception is that we use initial-task performance here instead of initial-task earnings as in previous tables. Using initial-task earnings instead would change nothing. Our motivation for controlling for initial-task performance here is to account for the idea that individuals who chose different effort levels resulting in different performance levels may have very different reactions to the revelation of inequality. For example, some low-wage participants may have reduced initial-task performance for some reason, perhaps because they simply perceived their assigned piece-rate wage to be low in absolute terms. If that were the case, then reaction to the revelation of wage-disadvantage as measured by JWBs may have been offset by having already (endogenously) responded to the (absolutely) low wage by underperforming. By controlling for initial-task performance we get closer to the ideal comparison of identical employees making identical decisions but with different knowledge about the presence of inequality.

effort.[25]

One puzzle that remains is why the large variation in EI-session JWBs did not translate into substantial variation in behavior. There are several possible explanations. It could be that this simply reflects the confounds we mentioned earlier that complicate any causal interpretation of these measures. More intriguingly, it could be that the initial task brought to mind narratives that colored the interpretation of appropriate behavior in subsequent tasks. It could be, for example, that both a narrative of meritocracy and strengthened JWBs are necessary to produce the non-pecuniary incentives we have inferred. We leave this question for future research.

5 Closely related literature

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). Our study is set apart by being among the first to consider how the behavioral consequences of inequality may depend on the content of work, i.e., whether work is ability-intensive or effort-intensive.[26] This is an important distinction, as personal and societal economic success increasingly depends on skill- or ability-intensive employment sectors. 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 or concealed through formal or informal pay secrecy policies may also vary across this dimension. Moreover, much of the literature in this domain focuses on how ex-ante revelation of pay inequality impacts effort provision in the designated task, whereas we focus on how ex-post revelation of pay inequality impacts behavior on

[25] Performance on the effort-intensive task theoretically ranges from 0 to 48, and ranges from 11 to 40 in the data, so the coefficient here of -0.04 represents a non-negligible effect on JWBs.

[26] The only other papers we are aware of are Butler (2016) and Hart and Piff (2018). Conversations with one of the authors of the latter paper suggest that their study can be interpreted as involving an ability-intensive task even if that was not their explicit interpretation, so that the two papers may be viewed as providing complementary evidence on the robustness of each other's findings. Specifically, in a design that involves deception Hart and Piff (2018) tell subjects that experimental earnings 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. They find (as we do) positive non-pecuniary incentives of salient inequality.

subsequent tasks.[27]

Our study is also related to the nascent literature on motivated beliefs in economics. Researchers have long understood that in order to avoid cognitive dissonance (Festinger 1957) individuals may subconsciously color their beliefs and this process may alter economic incentives (Akerlof and Dickens 1982; Bénabou and Tirole 2016). However, evidence on the determinants and consequences of motivated beliefs is scant. On the determinants, a primary exception is Di Tella, Galiant, and Schargrodsky (2007), who exploit a natural experiment to show that randomly assigned land wealth shifts attitudes towards "market beliefs," such as the belief that individuals can succeed through their own efforts. On the consequences, Frank, Wertenbroch, and Maddux (2015) show that priming JWBs increases makes experimental participants more likely to prefer tournament-style compensation over egalitarian compensation for performance on a puzzle-solving task.[28] Motivated beliefs that ostensibly rationalize personal disadvantage are also characteristic of the "System Justification Theory" literature in social psychology. John T. Jost and Sullivan (2003), for example, 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. We contribute to this literature by outlining a novel mechanism through which the economic environment can interact with subconscious beliefs-formation processes and readily accessible culturally determined narratives (Collier 2016; Shiller 2017) to affect motivated beliefs and, in turn, behavior. We provide direct evidence on behavior and, through post-experiment surveys, suggestive indirect evidence on the mechanism.

Our study is also related to the large experimental literature in economics investigating 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 features 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.[29]

[27] One recent exception is Sseruyange and Bulte (2019), who conduct a field experiment in Uganda and find that piece-rate wage inequality on an initial bean-sorting task reduces effort provision in a subsequent non-incentivized bean-sorting task.

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

[29] 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

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).[30] 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. As such, our study can contribute more broadly to the ongoing discussions of the merits of wage transparency.

6 Discussion and Concluding Remarks

Summing up, our study documents that the behavioral responses to salient inequality vary qualitatively across contexts. We focused on two primary contextual factors: the type of work involved and the type of inequality revealed. Our primary finding was a puzzling dual pattern in behavior in response to wage inequality when work is ability-intensive, where the disadvantaged increased their subsequent quality and quantity of production but decreased pro-sociality. We found weaker effects associated with revealing income inequality, generally, and no behavioral effects at all when inequality occurred in an effort-intensive environment. We outlined a novel mechanism that relies on contextual variation in cognitive processes strengthening Just World Beliefs in reaction to salient inequality. We argued that this novel mechanism is capable of providing a unified explanation for many of the behavioral patterns in our data and then went on to provide suggestive supporting

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.

[30] 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).

attitudinal evidence from our post-experiment survey. The strong context-dependence of the reactions we document, as well as the potential for a *positive* productivity response imply that the non-pecuniary incentives associated with inequality warrant further attention and, in particular, should be properly incorporated into our considerations of optimal contracts (cf Koszegi 2014).

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. However, it must be noted that we carefully controlled two factors – 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 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.[31] This suggests that the productivity effects of revealed inequality dominate the effects on cooperation, but this obviously would require careful empirical studies to confirm.

Finally, more research into the mechanism(s) at work is necessary. We conjecture a novel mechanism and provide supporting evidence. However, as noted, our supporting evidence is only

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

suggestive. As a modicum of additional reassurance, it is worth mentioning that some crucial features of our data and of the mechanism we conjecture to explain it have been documented in a wide array of settings. Evidence for the underlying cognitive processes we conjecture, in particular, 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.

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.

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

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.

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.

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.

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 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.

Collier, P. (2016): "The cultural foundations of economic failure: A conceptual toolkit," *Journal of Economic Behavior & Organization*, 126, 5–24.

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*.

Di Tella, R., S. Galiant, and E. Schargrodsky (2007): "The formation of beliefs: evidence from the allocation of land titles to squatters," *The Quarterly Journal of Economics*, 122, 209–241.

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.

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., 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*.

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.

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.

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*.

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.

Rosen, S. (1982): "Authority, control, and the distribution of earnings," *The Bell Journal of Economics*, 311–323.

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

Shiller, R. J. (2017): "Narrative economics," *American Economic Review*, 107, 967–1004.

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*.

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

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

Tables and Figures

Tables

Table 2: *Descriptive Statistics*

Treatment	Pay scheme	N	Phase 1		Phase 2			Phase 3
			income	score	quantity	quality	lying	donations
IW-II-AI	Low Wage	32	4.18	20.91	6.91	0.88	0.16	0.30
	High Wage	24	7.29	19.38	7.75	0.92	0.13	0.24
VW-II-AI	Low Wage	24	4.03	20.17	7.75	0.96	0.25	0.16
	High Wage	34	7.70	19.03	7.74	0.96	0.09	0.22
IW-VI-AI	Low Wage	28	4.01	20.04	7.64	0.94	0.14	0.25
	High Wage	21	7.52	20.00	6.76	0.95	0.10	0.17
VW-VI-AI	Low Wage	20	3.85	19.25	6.85	0.92	0.10	0.19
	High Wage	28	7.63	19.79	7.25	0.96	0.11	0.13
IW-II-EI	Low Wage	27	5.76	28.81	7.67	0.99	0.04	0.17
	High Wage	27	10.06	28.07	7.67	0.98	0.04	0.14
VW-II-EI	Low Wage	26	6.00	30.00	7.88	0.96	0.08	0.23
	High Wage	27	10.60	29.89	8.26	0.96	0.11	0.15
IW-VI-EI	Low Wage	22	5.74	28.68	7.59	0.97	0.09	0.28
	High Wage	30	10.30	29.07	7.93	0.95	0.03	0.12
VW-VI-EI	Low Wage	30	5.53	27.67	6.73	0.96	0.10	0.26
	High Wage	23	10.66	31.30	8.09	0.94	0.04	0.13

Notes: **[1]** Treatments are denoted using a triple (V/I)W-(V/I)I-(E/A)I, where the first component refers to whether W(age) inequality was visible or not, the second component refers to the visibility of I(ncome) inequality and the last component refers to the task type – A(bility) or E(ffort) intensive. **[2]** Columns 4-8 present averages of the variables listed in the column headings. "Phase 1 income" refers to participants' earnings from the phase 1 task only; "Phase 1 score" refers to the number of correct questions (out of a maximum of 48) on the phase 1 task; "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 experimental income donated to either of the charity options (American Red Cross and the South Plains Food Bank).

Table 3: *OLS Estimates AI only*

	(1)	(2)	(3)	(4)	(5)
	Phase 1	Phase 2			Phase 3
	score	quantity	quality	lying	donations
HW	-1.23	0.71	-0.01	0.07	0.02
	(1.582)	(0.582)	(0.060)	(0.106)	(0.110)
VW	-0.51	1.00**	0.07**	0.12	-0.12**
	(1.088)	(0.417)	(0.033)	(0.117)	(0.051)
VI	-0.62	0.75	0.05	-0.01	-0.03
	(1.382)	(0.513)	(0.037)	(0.076)	(0.096)
HWxVW	0.35	-1.17***	-0.04	-0.16	0.11
	(2.360)	(0.340)	(0.042)	(0.113)	(0.082)
HWxVI	1.04	-1.74**	-0.01	-0.04	-0.10
	(2.236)	(0.634)	(0.044)	(0.113)	(0.131)
VWxVI	0.21	-1.47**	-0.08	-0.13	0.03
	(1.553)	(0.613)	(0.056)	(0.149)	(0.122)
HWxVWxVI	-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 AI 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); "JWB" is the participant's response to the Just World Beliefs question. **[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 EI Only*

	(1)	(2)	(3)	(4)	(5)
	Phase 1	Phase 2			Phase 3
	score	quantity	quality	lying	donations
HW	-1.09	-1.17	-0.06	0.10	0.03
	(0.888)	(0.701)	(0.037)	(0.083)	(0.081)
VW	1.30	0.16	-0.03	0.05	0.08
	(1.214)	(0.613)	(0.033)	(0.046)	(0.069)
VI	-0.05	-0.09	-0.02	0.07	0.12
	(1.429)	(0.440)	(0.014)	(0.071)	(0.084)
HWxVW	0.44	0.31	0.01	0.02	-0.04
	(1.989)	(0.675)	(0.019)	(0.089)	(0.053)
HWxVI	1.52	0.50	-0.01	-0.05	-0.10
	(1.796)	(0.755)	(0.019)	(0.061)	(0.104)
VWxVI	-2.50	-0.96	0.01	-0.07	-0.14
	(1.796)	(0.909)	(0.046)	(0.116)	(0.116)
HWxVWxVI	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 EI 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); "JWB" is the participant's response to the Just World Beliefs question. [**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*

	(1)	(2)	(3)	(4)
	AI only		EI only	
HW	0.22	0.24	-0.84***	-0.87***
	(0.217)	(0.217)	(0.265)	(0.291)
VW	0.53***	0.53***	-1.03**	-0.97**
	(0.157)	(0.159)	(0.396)	(0.397)
VI	0.36	0.37	-0.58	-0.58
	(0.218)	(0.222)	(0.653)	(0.623)
HWxVW	0.11	0.10	1.00***	1.00***
	(0.474)	(0.472)	(0.287)	(0.308)
HWxVI	-0.04	-0.05	0.46	0.50
	(0.303)	(0.299)	(0.640)	(0.616)
VWxVI	-0.39	-0.39	1.30*	1.20*
	(0.384)	(0.384)	(0.696)	(0.665)
HWxVWxVI	-0.60	-0.60	-1.41	-1.29
	(0.739)	(0.741)	(0.884)	(0.848)
Phase 1 performance		0.01		-0.04**
		(0.017)		(0.018)
Constant	1.86	1.41	4.25***	5.57***
	(1.145)	(1.501)	(0.730)	(0.806)
Observations	207	207	209	209
R-squared	0.087	0.089	0.131	0.161

Notes: **[1]** Each column reports an OLS regression restricted to observations from either AI or EI sessions using as its dependent variable our measure of *Just World Beliefs*, i.e., 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]** 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; "Phase 1 performance" is the number of correct responses on the Phase 1 task, which theoretically ranges from 0 to 48; **[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. **[4]** Robust standard errors clustered by session appear in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.