

INEQUALITY AND RELATIVE ABILITY BELIEFS*

Jeffrey V. Butler

I present experimental evidence for a novel mechanism yielding inequality persistence. Just world beliefs research suggests that individuals believe they merit unequal treatment they experience. Merit depends on ability and effort so that disadvantage (advantage) may undermine (bolster) confidence in own relative ability. Because decisions determining economic success rely on such beliefs (e.g. competitiveness), inequality may self-perpetuate. In multiple experiments, I randomly assign unequal pay for an identical task where performance depends on cognitive ability. I find that pay level consistently and substantially affects beliefs but not performance. Finally, I show that among males high pay increases competitiveness by 33%.

What causes inequality to persist? This question is of perennial interest to economists both because of its implications for efficiency – underdeveloped or underutilised human capital may translate into forgone productivity – as well as on more general moral grounds. Given its long pedigree as a puzzle of interest, the literature on persistent inequality is large and offers many compelling theoretical explanations which can be grouped into a few broad categories (Piketty, 2000). One familiar category of explanations singles out the importance of external factors such as real or expected discrimination (Becker, 1971; Arrow, 1971) or access to credit markets upon which investment in human or physical capital depends (Loury, 1981; Banerjee and Newman, 1991, 1993; Galor and Zeira, 1993; Torvik, 1993; Gaviria, 2002). A second category offers as a primary explanation differences in the distribution of individual traits or preferences, for example, risk tolerance, a preference for competition or aspirations (inter alia, Hoff and Pandey, 2006; Ray, 2006; Niederle and Vesterlund, 2007; Mookherjee *et al.*, 2010). A third, nascent, category of explanations examines the ecology of inequality and suggests that social interactions, the social environment or culture shape traits, preferences or even cognitive ability in a way that reinforces existing inequality (Durlauf, 1996; Akerlof, 2002; Gneezy *et al.*, 2009; Bogliacino and Ortoleva, 2011; Afridi *et al.*, 2012; Leibbrandt *et al.*, 2013; Mani *et al.*, 2013; Bowles *et al.*, 2014).

In this article, I provide the first clean experimental evidence documenting a fourth, under-explored, class of explanations. This fourth class of explanations starts from the well-established notion that differences in the willingness to compete may account for a substantial fraction of economic inequality (Niederle and Vesterlund, 2007). Of

* Corresponding author: Jeffrey V. Butler, Einaudi Institute for Economics and Finance, via Sallustiana, 62, Rome 00187, Italy. Email: jeff.butler@eief.it.

I owe a great debt of gratitude to George A. Akerlof, Stefano DellaVigna, Luigi Guiso, Shachar Kariv, Frederic Vermeulen and Georg Weiszäcker as well as two anonymous referees for many helpful comments which greatly improved the manuscript. I also like to thank seminar participants at LUISS Guido Carli University, the European University Institute, Southern Methodist University, University of Arizona, University of Nebraska, Lincoln, University of Arkansas, the Federal Trade Commission and the University of Nevada, Las Vegas in addition to conference participants at the 2012 North-American ESA Conference, Tucson and Segment 7 of the Stanford Institute for Theoretical Economics (SITE) Summer 2013 Workshop for their attention and valuable feedback. I am also grateful to the National Science Foundation under Research Grant SES 04-17871 and UC Berkeley's X-lab for financial support.

particular importance may be ability-based competition, since this is the type of competition upon which many pathways to success depend – the quintessential example being matriculating at a highly selective college. Willingness to compete depends on at least two components: a preference to compete as well as beliefs about the likelihood of success. While the existing literature has largely focused on the preference component, in this study I focus on the beliefs component and investigate individuals' beliefs about their relative ability. Because individuals' subjective *ex ante* expected returns to ability-based competition depend positively on relative ability beliefs, those who believe they are relatively more able should, *ceteris paribus*, be more likely to enter into ability-based competition. At the same time, those who are less sanguine about their relative ability should be less likely to select into competitive environments. Consequently, to provide novel evidence for the existence of a link between persistent inequality and relative ability beliefs I build on the belief in a just world literature (Lerner, 1965; or, more recently Bénabou and Tirole, 2006) to investigate the ecology of beliefs formation, formulating and testing a specific hypothesis:

HYPOTHESIS 1. Experiencing unequal treatment changes individuals' beliefs about their relative ability in a way consistent with the unequal treatment.

A large body of research in social psychology starting with Lerner (1965) demonstrates that individuals have a deep-seated and fundamental need to believe the world is just and fair. Dozens of subsequent studies document that individuals often respond to observing others treated unequally by, perhaps subconsciously, forming beliefs about them that justify the inequality, for example, blaming victims for their own misfortune. For an overview of this expansive literature, see Lerner and Miller (1978), Furnham (2003) or Hafer and Bègue (2005). But what about the victims themselves? Do individuals' beliefs about their own relative merit respond similarly to unequal outcomes, so that individuals blame themselves for their own misfortunes? If so, this may provide a mechanism generating self-perpetuating inequality.

Specifically, suppose a just world is a meritocratic world, so that economic outcomes are ideally distributed based on merit. Merit is often thought of as being a combination of effort and ability. If just world beliefs cause individuals to align their beliefs about their own relative merit with their economic situation, disadvantage may cause the disadvantaged to believe they are relatively less able; advantage may cause the advantaged to believe they are relatively more able. Because beliefs about relative ability are central to many of the plans and paths generating heterogeneity in economic success, just world beliefs may lead the initially disadvantaged (advantaged) to make decisions which perpetuate the initial inequality. In this study, I shed light on one such decision, the decision to compete, providing evidence consistent with the idea that past inequality may affect future competitiveness.

HYPOTHESIS 2. Prior unequal treatment affects individuals' subsequent willingness to compete.

All together, the chain of reasoning leading inequality to self-perpetuate here involves several steps, none of which is obvious *a priori*. It is not certain that just world beliefs affect beliefs about one's own relative merit – particularly given the large

literature documenting various self-serving biases such as overconfidence (see the discussions in Moore and Healy, 2008; Grossman and Owens, 2012; Burks *et al.*, 2013). Even then, it is not obvious whether beliefs about the ability component of merit – a stable trait – will be affected. Finally, it is not guaranteed that the effect of inequality on beliefs will be sufficiently strong to change behaviour.

In this study, I present evidence from multiple experiments testing each step in this chain. While there are myriad ways to engender a feeling of being treated unequally in the laboratory, in the experiments reported herein I implement inequality in a familiar and economically relevant manner: unequal pay for equal performance. To keep theoretical performance incentives constant, I choose pay structures featuring identical marginal pay but different total compensation. Despite equal marginal incentives, one may still worry that the primary effect of inequality is on effort provision and related beliefs (Gill and Prowse, 2012; Hargreaves Heap *et al.*, 2013).¹ To shed light on this effort confound directly, in my second experiment I randomly vary whether performance depends on effort or cognitive ability. In all experiments I test the hypothesis that personally experienced unequal treatment colours individuals' beliefs about their own relative ability. In my second experiment I additionally test whether initial inequality affects individuals' subsequent willingness to compete.

In doing so, I make two main contributions. First of all, I present evidence for a specific and under-explored mechanism through which inequality may naturally persist. This mechanism is particularly important not because it supersedes previous explanations or explains all inequality persistence – inequality is a multifaceted phenomenon unlikely to be fully explained by one channel. Rather, its importance lies in providing minimal conditions under which initial inequality may become durable inequality: in environments where ability-based competition is particularly important for success or failure, experienced inequality itself is sufficient. Secondly, the study contributes to our understanding of the ecology of economically relevant beliefs formation. While the consequences of just world beliefs and, more generally, cognitive dissonance have been studied by economists before (Akerlof and Dickens, 1982; Bénabou and Tirole, 2006), their impact on beliefs about one's own economically relevant traits has so far eluded serious inquiry.

The remainder of the article is organised as follows. First, I discuss closely related literature. In Section 2, the design of Experiment 1 is presented in detail, results are presented and then discussed. In Section 3, the design of Experiment 2 is presented, followed by results and a discussion. In Section 4, I summarise, discuss and put in

¹ It must be noted that my focus on relative ability and effort provision in this study is motivated by the philosophical and historical connection between these concepts and the concept of merit. For persistence more generally, the important distinction is between attributions to dispositional/characterological factors and attributions to situational/behavioural factors (Kelley, 1967; Janoff-Bulman, 1979). The former factors are defined by their permanence and stability, so that cognitive ability may fit naturally into this class of factors. The defining characteristic of the latter class of factors is impermanence, so that effort fits less naturally into this class. Obviously, a reluctance to provide effort may be a stable trait (e.g. 'habitual laziness') or it may stem from short-lived contextual factors such as an individual's reaction to particular financial incentives. While I interpret effort as a transitory phenomenon and as a natural confounding factor for ability-based explanations in my study, in other contexts and for more general investigations of the relationship between effort and outcomes this may not be a valid interpretation.

context the findings from this study. Instructions for Experiments 1 and 2 appear in an online Appendix A. Data from all experiments reported herein are available online.

1. Closely Related Literature

Several authors have proposed models in which social interaction and social comparison may perpetuate inequality by affecting individuals' preferences (e.g. aspirations) directly or, indirectly, by affecting the choices parents make for their children (Ray, 2006; Dalton *et al.*, 2010; Mookherjee *et al.*, 2010; Bogliacino and Ortoleva, 2011; Bowles *et al.*, 2014). Another preference-based story posits that individuals prefer to conform to existing stereotypes about their groups, undermining success among members of negatively stereotyped groups (Steele and Aronson, 1995). The current study differs from these by examining the effect of experienced inequality on individuals' beliefs rather than directly on their preferences as, for example, aspirations have typically been modelled. Although intuitively relative ability beliefs may affect aspirations, this link has not been made explicit in any model I am aware of, so that one view of the current study is that it provides some insight into one plausible determinant of aspirations-based inequality persistence. The distinction between preferences and beliefs is an important one, as each suggests a different remedy.

Another closely related strand of literature investigates how beliefs about relative ability are updated in response to informative signals about ego-relevant traits such as intelligence or beauty. Eil and Rao (2011) show that individuals tend to underweight negative signals relative to the benchmark provided by Bayes' rule, while Möbius *et al.* (2011) find that the way in which negative ego-relevant information is underweighted follows the spirit, if not the letter, of Bayes' rule. Grossman and Owens (2012) show that individuals tend to attribute poor past performance to bad luck rather than ability when noisy signals make this possible, so that overconfidence may be the result of overly strong prior beliefs rather than biased information processing. In Bénabou and Tirole (2002), the authors study why individuals might prefer to hold inflated beliefs about their own ability or efficacy, providing a theoretical justification for overconfidence. The primary difference between the current studies and this strand of the literature is that I focus on how an uninformative random shock to outcomes – randomly assigned inequality – affects ability beliefs (confidence). In doing so, I seek to understand the relationship between the social environment and confidence, more generally, rather than how biases in information processing can lead to overconfidence specifically.

In Bénabou and Tirole (2006), the authors study a macroeconomic implication of heterogeneity in the intensity of just world beliefs, showing how they can alter preferences over redistributive taxation by affecting beliefs about the relationship between success and effort. Supporting empirical evidence is provided by Fong (2001) and Di Tella *et al.* (2007). The former study uses a large-scale Gallup Poll conducted in 1998 to document that Americans who believe that outcomes are self-determined, for example, who believe the poor are poor because of their own laziness – are significantly less likely to support redistributive taxes. The latter study exploits a natural experiment in which property was essentially randomly allocated. The main result is that economic advantage can cause large shifts in some types of beliefs. However, some features of the natural experiment make it difficult to construct an unimpeachable control group with which to identify whether it is

advantage, disadvantage or both that moves beliefs. Moreover, the study detects no change in beliefs about the relationship between effort and success. Unlike these studies, I construct hypotheses about how the interaction between just world beliefs and personally experienced inequality may alter individuals' confidence in their own relative ability and test these hypotheses directly in a controlled environment.

On competitiveness and inequality, several studies document that historically disadvantaged groups perform worse in competitive environments. Gneezy *et al.* (2003) find that women perform worse than men under tournament incentives, particularly when members of both genders are present simultaneously.² One reason for this mixed-gender-session effect could be that the presence of both males and females makes gender a salient group boundary. Accordingly, Hoff and Pandey (2006) show that enhancing the salience of another unequal group boundary – Indian caste affiliation – leads members of low-caste groups to underperform. Afridi *et al.* (2012) find similar results in the context of the Chinese *hukou* system. Returning to gender, Niederle and Vesterlund (2007) document that women are half as likely as men to select tournament incentives over piece-rate pay and, although they perform no worse on average, report believing so. Gill and Prowse (2013) provide evidence that initial gender differences in competitiveness may be amplified over time: the negative effect on performance from losing a previous competition is stronger for women than for men. Finally, Gneezy *et al.* (2009) suggest that these gender biases in competitiveness may be culturally dependent. They conduct identical experiments in two societies – one quintessentially patriarchal (the Maasai in Tanzania); the other matrilineal (the Khasi in India) – and find that competitiveness patterns flip across these two societies. Maasai men are more likely to select tournament incentives than Maasai women, whereas Khasi women are more competitive than Khasi men. A potential explanation unifying these results is that making group affiliations salient evokes memories of unequal treatment, undermining (boosting) current beliefs about relative ability and performance among the disadvantaged (advantaged) which in turn affects the strength of tournament incentives and/or beliefs about the strength of tournament incentives differentially. This chain of reasoning involves several links, all of which are problematic. In particular, since the real-world people groups involved differ along many dimensions, it is unclear which dimension is being made salient by priming group affiliation.³ Alternatively, priming inequality may simply demoralise the disadvantaged group, reducing effort provision as in Hargreaves Heap *et al.* (2013) and the strength of incentives. The current study differs from almost all of these studies by randomly assigning unequal treatment – the exception being Hargreaves Heap *et al.* (2013) – so that theoretically groups differ only along this uninformative dimension. Moreover, differently from all of these studies I directly examine the beliefs link in the causal chain outlined above.

An alternative interpretation of the pay schemes used in my study is that they represent base pay plus a lump-sum bonus. In this light, the current study is most

² It is generally agreed that women have a long history of being disadvantaged both *de jure* and *de facto* in many countries.

³ In the last study, the direction of causation may even be an issue: it could be that societies whose women are more competitive – perhaps because of a competitiveness trait or gene – are more likely to be matrilineal or matriarchal.

directly related to Gill *et al.* (2013) where the authors investigate how exposure to bonus-based compensation affects productivity and cheating in the workplace. In their experiment, participants complete a real-effort task. In their main treatment (*T*), participants earn £2 plus a 50% chance of a £6 bonus for completing the task. In control conditions, participants are paid either £2 (*C1*) or £8 (*C2*). After completing the task, participants complete a second, similar, real-effort task for piece-rate pay and are also given the opportunity to cheat – to lie for personal monetary gain. The authors find that exposure to a bonus-based compensation scheme (*T*) does not affect productivity relative to the relevant control condition (*C1* or *C2*) on the second real-effort task but that it does increase cheating relative to the relevant control. While the study is silent on how inequality affects beliefs or competitiveness, the directly comparable results are consistent with mine: inequality does not affect actual effort provision. Moreover, the uniform effects across pay levels of exposure to inequality on an effort-intensive task are also consistent with my data. Assuming the similarity in effects carries over to (unmeasured) beliefs so that their participants' relative performance beliefs were uniformly bolstered by inequality, the increased cheating they document could be due to an enhanced sense of entitlement. Of course, exposure to inequality may also change behaviour by, for example, changing the perceived fairness of the situation – the interpretation Gill *et al.* themselves put forward. Whether and how the patterns documented in my study interact with more general concerns for morality and fairness are interesting open questions for future research on the (unintended) consequences of bonus-based compensation.

Outside economics, the current study builds on two closely veins of research in social psychology and sociology: belief in a just world (Lerner, 1965) and system justification theory (Jost *et al.*, 2003). A central hypothesis of both theories is that a basic human need for justice coupled with an aversion to cognitive dissonance causes individuals to rationalise an unequal status quo by choosing to believe existing inequality is deserved. Lerner (1965), in a seminal contribution to these literatures, provides non-incentivised experimental evidence that third-party observers rated workers who were randomly chosen to be paid as more deserving of pay than those workers who were randomly chosen not to be paid. Recognition of this type of phenomenon dates back to at least Allport (1954), while more recent social psychological research suggests that the phenomenon is fundamental by demonstrating its prevalence across multiple cultures and nations (Fiske *et al.*, 2007; Caprariello *et al.*, 2009; Cuddy *et al.*, 2009).⁴ Whether beliefs about oneself are affected by inequality in a manner similar to the way beliefs about others or other groups vary in just world/system justification research is an important open question whose answer is not obvious. The few existing studies provide mixed results (Hafer and Olson, 1998), however an important distinction is made in this literature between 'behavioural' and 'characterological' self-blame (Janoff-Bulman, 1979). In the former, individuals attribute bad outcomes to past behaviour so that similar outcomes may be avoidable in the future. The latter type of self-blame is more persistent, involving attributions to a relatively permanent trait (e.g. character). A related distinction is made between situational (impermanent) factors and dispositional

⁴ For a discussion of cognitive dissonance as well as other economically relevant implications (Akerlof and Dickens, 1982).

(permanent) factors in attribution theory (Kelley, 1967). In this study, I try to separate these two types of blame/attributions, showing that unequal treatment gives rise to the latter – blaming one’s ability, a dispositional factor – rather than the former – blaming one’s effort, which I interpret here as a behavioural/situational factor. However, given the large literature documenting overconfidence, one might *a priori* expect myriad psychological mechanisms protecting self-image and self-esteem to provide an overwhelmingly countervailing force (Bénabou and Tirole, 2002) against both types of self-blame. Unlike these veins of research the current study uses experiments with monetary incentives and an incentive compatible belief elicitation mechanism to investigate whether beliefs about one’s own relative ability are affected by randomly assigned initial inequality and whether the effect, if any, persists to affect future decisions.

2. Experiment 1

2.1. *Design and Procedures*

Experiment 1 was conducted in the XLab facilities at the University of California, Berkeley using the software z-tree (Fischbacher, 2007). Participants were recruited among students and staff of the university. Three separate treatments were conducted. Each participant took part in just one of the three treatments. All three treatments shared a common structure:

- (i) participants were informed about the task they would be performing;
- (ii) they then learned about the pay structure and, if applicable, were randomly assigned one of two pay levels;
- (iii) participants subsequently completed ten rounds of the analytical ability-intensive task that had been described to them initially; and
- (iv) after all ten rounds of the task were completed, participants’ beliefs about how well they performed relative to other participants were elicited in an incentive compatible manner.

Importantly, participants learned about the belief-elicitation stage only after having completed the analytical ability task so that belief elicitation should not have affected behaviour.

2.1.1. *Pay structure: varying inequality across treatments*

To investigate the impact of inequality on relative performance and related beliefs, I varied the pay structure across the three treatments comprising the experiment. In the payoff inequality (PI) treatment, half of the participants in each session were randomly assigned to the high-pay group (HP) while the other half were assigned to a group earning low pay (LP).⁵ Participants assigned to HP earned \$4 for each correct answer on the analytical ability-intensive task (described below) and \$2 for each incorrect answer. Participants assigned to LP were paid \$2 for each correct answer and \$0 for each incorrect answer. An individual’s pay group persisted over all ten rounds of the

⁵ Before assigning pay groups, participants were informed that pay group assignment would be random and that each participant was as likely to be assigned to the HP group as to the LP group.

task. When submitting answers, participants were informed of the pay group of each preceding guesser.

The two remaining treatments can be viewed as control sessions in the sense that they implement each of the two pay structures involved in the PI treatment separately. In control-high pay (C-HP) all participants earned \$4 for each correct answer on the analytical ability-intensive task and \$2 for each incorrect answer. Meanwhile, in control-low pay (C-LP) participants earned \$2 for each correct answer and \$0 for each incorrect answer.

The essential feature of these two pay structures – HP and LP – is that marginal performance incentives are identical across pay groups. Each correct answer always pays two dollars more than each incorrect answer. This implies that task performance should not vary across pay groups which will prove essential to cleanly identifying the effect of inequality on beliefs. At the same time, the two pay structures combined implement salient inequality: the most a participant assigned to the LP group could earn was the least an HP participant could earn from performance on the (same) task.

2.1.2. *The analytical ability task in experiment 1*

The analytical ability task used is an urn-guessing game common in the social learning literature (Anderson and Holt, 1997). The experimenter starts by choosing one of two biased urns: ‘Urn A’ contains two red balls and one white ball; ‘Urn B’ contains two white balls and one red ball. The game involves n individuals each of whom, in random sequence order, submits one guess about which of two urns the experimenter has chosen. An individual wins $\$X$ for a correct guess and $\$Y < \X for an incorrect guess. At the time of his or her guess, each individual privately observes one draw (with replacement) from the chosen urn and all previously submitted guesses, if any.⁶ To complete the description, it is common knowledge that the experimenter chooses among the urns with equal probability, so that individuals’ common prior beliefs about the likelihood of Urn A should be 1/2.

This task was chosen for three main reasons. First and foremost, it provides each player with a lot of information about his or her own performance as well as pecuniary incentives to pay attention to, and form correct beliefs about, how well others perform without making such comparisons the main focus of the task. Second, the task provides an objective performance measure: number of correct guesses.⁷ And third, the urn-guessing task is plausibly ability-intensive: a lot of effort with no analytical ability will

⁶ For example, an individual guessing third would observe one private draw from the urn (with replacement) plus the guesses – but not the private draws – of the first and second-guesser. The individual chosen to guess first in the sequence would observe only his or her private draw from the urn.

⁷ It is true that this measure certainly includes some noise due to being correct through luck as well as skill – as would many tests or tasks I could have chosen. In particular, one might be concerned that those who guess later are more likely to be part of an informational cascade and that this fact undermines the interpretation of total correct guesses as a performance measure. There are two factors which partially ameliorate this concern:

- (i) cascades in this game are not irrational – entering into a cascade (or not) requires analytical ability; and
- (ii) sequence orders and groups were randomly re-determined across rounds so that, whatever effect there is from moving late in a sequence, it is unlikely that an individual is consistently assigned a late sequence order.

not produce performance better than random chance in this game, while a lot of analytical ability coupled with a little effort should produce superior performance.

To get a feel for the type of reasoning required to perform well in this game, briefly consider individuals' optimal guesses.⁸ To simplify matters, assume that an individual's guess equals his or her private draw whenever the individual believes both urns are equally likely.⁹ The first guesser should guess Urn *A* (*B*) upon observing a red (white) draw since the posterior probability of Urn *A* (*B*) is now 2/3. The second guesser, irrespective of the first guesser's guess in period 1, should also guess Urn *A* conditional on privately observing a red ball drawn and Urn *B* otherwise.¹⁰ Thus, the first two guesses are perfectly informative of the first two draws. For guessers $n > 2$, simply note that:

- (i) if the history of guesses in periods 1, ..., $n-1$ is such that the number of Urn *A* (*B*) guesses exceeds the number of Urn *B* (*A*) guesses by at least 2, then guesser n and all subsequent guessers optimally submit a guess of Urn *A* (*B*) regardless of their privately observed draws; and
- (ii) for all other histories, guesser n 's optimal guess is Urn *A* (*B*) conditional on a private draw of red (white).

Thus, guessing correctly in this game requires a bit of analytical ability but not an unreasonably large facility. Performance also obviously depends on a bit of luck (guessing sequence order) and effort (paying attention) but a plausible assertion is that performance is more sensitive to ability than effort.¹¹

All participants in Experiment 1 played ten separate rounds of the analytical ability task. Before each round, the experimenter randomly chose either Urn *A* or Urn *B* with equal probability. Participants were then randomly (re-)divided into groups of 8.¹² Within each group, the sequence in which participants would guess was randomly re-determined with each participant being assigned a sequence order number

⁸ For brevity's sake, I only sketch the logic here. For more detail, see the classical references in the huge literature on social learning using this game: Bikhchandani *et al.* (1992); Anderson and Holt (1997).

⁹ That is to say, they guess Urn *A* (*B*) conditional on a private draw of red (white) whenever their posterior belief – incorporating all available information including their private draw – suggests that Urn *A* and Urn *B* are equally likely. Relaxing this assumption changes how unbalanced the history of guesses needs to be in order to trigger a cascade, but does not add much intuition.

¹⁰ To see this note that if 2's private draw is red (white) and 1's guess was Urn *B* (*A*), then 2's posterior is 1/2 since it can be inferred that 1's private draw was white (red). In this case, by assumption, 2 goes with his private information. On the other hand, if 2's private draw is red (white) and 1's guess was Urn *A* (*B*), then 2 has essentially observed two red (white) balls being drawn and should guess Urn *A* (*B*). In all cases, 2 guesses Urn *A* (*B*) whenever 2's private draw is red (white).

¹¹ A potential confound to this assertion is the use of a counting heuristic – simply counting the number of previous Urn *A* and Urn *B* guesses agreeing with the majority opinion – which also performs well. It is not clear how one arrives at this heuristic from the set of all possible heuristics but it seems reasonable that some might. Partly to address this concern, additional experiments (Experiment 2, detailed in later subsections, and a pilot for Experiment 2 (omitted but with similar results)) were conducted using more traditional measures of ability: questions from the logical reasoning section of the LSAT (pilot for Experiment 2) and Raven's progressive matrices (Experiment 2). The latter is a widely used culture-free measure of general intelligence.

¹² To avoid the need to have the number of participants in each session divisible by 8 – a quite severe restriction which would have resulted in sending many participants home with just a show-up fee – as many 8-person groups of participants as possible were formed with the remaining participants forming a group of less than eight. Each session consequently featured at most one group with less than eight participants.

$n = 1, \dots, 8$. Consequently, neither the specific individuals in each group nor the sequence order any specific individual was assigned persisted across rounds. At the end of each round, after all members of all groups had finished submitting their guesses, the chosen urn was revealed. After all ten rounds of this task were completed, one round was chosen at random to count towards participants' earnings.

2.1.3. Eliciting relative performance beliefs

After all ten rounds of urn-guessing were completed, participants learned there would be a belief elicitation phase. Individuals were presented with the question: 'compared to others, how accurate were your urn predictions?' To respond, participants selected between one of two possible answers which appeared in random order: 'top 50%'; or 'bottom 50%'. A correct answer yielded one additional dollar, while an incorrect answer paid no additional money.

Eliciting beliefs in this way maintains incentive compatibility under weak assumptions on preferences while maintaining simplicity. Whereas many methods widely used to elicit beliefs are either quite complex, confusing or require strong assumptions such as risk neutrality (Schlag and Van der Weele, 2013), a sufficient assumption for the simple procedure used here to provide proper incentives for truthful revelation of beliefs is that preferences obey the axioms of expected utility.¹³ Intuitively, participants are faced with a choice between two assets which pay one dollar in different states of the world. Normalising the utility of no additional earnings to $u(w + 0) = 0$, the expected utility of choosing 'Top 50%' is simply $\text{Prob}(\text{Performance above median}) \times u(w + 1)$. Similarly, the expected utility of choosing 'Bottom 50%' is $\text{Prob}(\text{Performance below median}) \times u(w + 1)$. Since $u(w + 1)$ is a factor common to both of these expected utilities, the preferred choice is the asset whose payoff-relevant state the decision-maker (subjectively) considers most likely.¹⁴ In the light of the coarse hypotheses of the current inquiry, I found the potential benefits of using a mechanism capable of eliciting beliefs more precisely did not justify the added complexity or stronger assumptions required.

¹³ While expected utility is a sufficient assumption, and eases exposition, the mechanism is incentive compatible under weaker assumptions. For example, denote by $(x_1, p_1; x_2, p_2)$ the lottery yielding x_i with probability p_i , $i = 1, 2$. Let \bar{m} (\underline{m}) denote the message 'I performed above (below) the median'. Assume utility takes the form $\mathbb{U}(x_1, p_1; x_2, p_2) = \sum_{i=1}^2 f(p_i)u(x_i)$. Assume that:

- (i) $u(\$1) = 1$ and $u(\$0) = 0$; and
- (ii) that $f(p) \geq f(p') \Leftrightarrow p \geq p'$.

Suppose without loss of generality that an individual's true subjective belief about the probability that s/he performed above (below) the median is $p > 0.5$ ($1 - p < 0.5$). For the mechanism to be incentive compatible, it must be that there is no strictly profitable deviation from a truthful report, in this case \bar{m} . The utility of the lottery induced by reporting \bar{m} is $\mathbb{U}(1, p; 0, (1 - p)) = f(p)$, while deviating and reporting \underline{m} yields a lottery with utility $f(1 - p)$. Since $f(p) \geq f(p') \Leftrightarrow p \geq p'$ reporting \bar{m} always yields weakly more utility than reporting \underline{m} , so that no strictly profitable deviation from the truthful report exists. Notice that these assumptions allow for $f(p)$ to take the form of the type of distorted probability perception/weighting functions central to the original formulation of prospect theory (Kahneman and Tversky, 1979). They also allow for incomplete preferences, implying they are strictly weaker than expected utility.

¹⁴ If the individual believes these two events are equally likely, for example, because s/he is certain his/her performance is exactly median – then the two options yield the same expected utility, making the individual indifferent between the two options. This latter case should add noise, but not a bias, to the elicited beliefs.

A more specific concern one may have about how beliefs are elicited here is the lack of an ‘exactly median’ performance belief option. There are three reasons I omitted this option. First of all, an expected utility maximising individual will never strictly prefer to report an exactly median belief: because an exactly median performance is a subset of both of the events ‘top 50%’ and ‘bottom 50%’, choosing either of these latter two options always yields weakly higher expected utility. Second, given the wealth of information participants receive on their own and others’ performances during the course of the experiment, holding an exactly median belief is unlikely *a priori*. Third, previous research suggests such ‘middle’ options may be focal and chosen in spite of proper incentives for truthful revelation of beliefs or preferences (Harrison and Rütstrom, 2008). Balancing these concerns suggested leaving out a middle option. I realise, however, that omitting an exactly-median option comes at the cost of additional noise: people who are certain they performed exactly at the median will be indifferent between the two available options as they are simultaneously in both top and bottom 50% of performers.¹⁵

Another minor point which needs addressing is hedging. There is typically a concern when paying for both beliefs and performance that individuals might try to use the belief elicitation incentives to hedge against a bad outcome, i.e., one might report doing poorly to move money into this bad state, reducing the expected variation in their earnings. However, since each individual knew his or her own performance – and hence earnings – at the time they answered the performance belief question, there is no scope for literal hedging here. In fact, ruling out hedging is one justification for asking about retrospective performance, as I do here, rather than eliciting performance beliefs on a future task – which may at first glance seem to be a more natural measure of relative ability beliefs.¹⁶

2.2. Experiment 1 Results

All together, eight sessions were conducted and 152 individuals participated. The PI treatment constituted four of these sessions, while two sessions each of C-LP and C-HP were conducted. In total, there were 83 participants in the PI treatment, 36 in C-LP and another 33 participants in C-HP. This information is summarised in Table 1.

Reassuringly, task performance did not vary substantially by pay group suggesting that marginal incentives are what matter for actual performance. Across all treatments and pay groups, participants guessed the correct urn about two-thirds of the time on average (Table 2). If anything, the low-pay group performed slightly better. Beyond simple means, Figure 1 reports histograms of participants’ performances across experimental conditions. The histograms reveal substantial heterogeneity in the number of urn guesses individuals made correctly while, at the same time,

¹⁵ For those unconvinced about the prudence of omitting an exactly-median option, Experiment 2 (detailed in later subsections) provides participants with such a middle option. It is included there because participants learn nothing about their relative performance, making it much more likely *a priori* that participants truly hold an exactly-median relative performance belief.

¹⁶ Specifically, the concern is that this would introduce an incentive for an individual to report believing they will perform below the median and then purposely perform poorly in order ensure this outcome obtains and thereby guarantee at least the earnings associated with a correct belief.

Table 1
Summary of Experiment 1 Design

	PI treatment		Control	
	High-pay group	Low-pay group	C-HP	C-LP
Pay for correct guess	\$4	\$2	\$4	\$2
Pay for incorrect guess	\$2	\$0	\$2	\$0
Participants	41	42	33	36

Table 2
Actual Performance, Experiment 1

	PI treatment		Control	
	High-pay group	Low-pay group	C-HP	C-LP
Number of correct guesses	6.54	6.69	6.63	6.81
SE	(0.25)	(0.25)	(0.31)	(0.19)
Observations	41	42	33	36

Notes. The Table reports the average number of correct guesses across all 10 rounds of the urn-guessing task. The maximum possible value would be 10. Standard errors appear in parentheses.

documenting that even the distributions of performance are strikingly similar across conditions.

RESULT 1. *Salient inequality did not affect task performance.*

Participants' beliefs were a different story, however. Consider first only raw means, without inserting any controls. PI participants assigned to the high-pay group were significantly more likely to report believing their performance ranked in the top half (Figure 2). Fully 85% of high-pay PI participants believed they were in the top half, while only 60% of low-pay participants put themselves in this top category ($p = 0.004$, one-tailed t-test).¹⁷ On the other hand, pay levels had essentially no impact on performance beliefs outside of the context of inequality: 78% of C-LP participants and 76% of C-HP participants believed they were in the top half of performers, which is in line with previous studies on overconfidence in general.

To get a sense of whether high relative pay enhanced participants' performance beliefs, low pay undermined beliefs or both, I pool observations across C-HP and C-LP – where beliefs do not vary by pay level – and compare beliefs in these pooled control sessions to beliefs for each pay level in the PI treatment, separately. The proportion of low-pay PI participants believing their performance ranked in the top half was significantly lower in the PI treatment than in the pooled control treatments (60%

¹⁷ All reported significance levels refer to one-tailed t-tests unless otherwise specified.




Fig. 1. Actual Performance, Experiment 1

Notes. The Figure presents histograms of the proportion of participants reporting believing the low-pay group (LP) or high-pay group (HP) performed better, as a whole, in terms of number of questions answered correctly. Observations are restricted to ability-intensive task participants. The urn-guessing task in Experiment 1 (maximum possible = 10), by pay group and treatment. The x-axis of each histogram refers to the number of correct guesses, while the y-axis record the proportion of participants answering a specific number of urn guesses correctly.

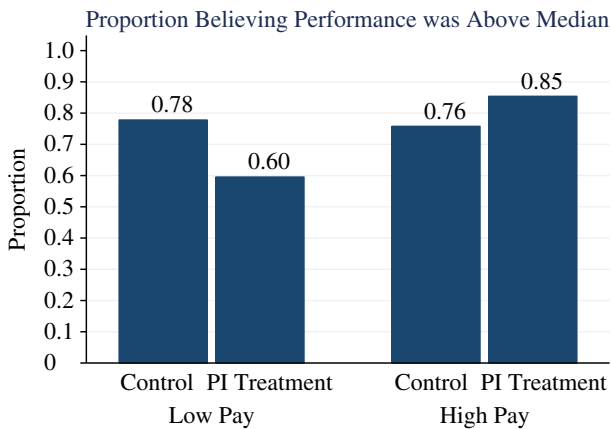


Fig. 2. Relative Performance Beliefs, Experiment 1

Notes. The Figure reports the proportion of participants believing their performance ranked above the median in terms of accuracy, by pay group and treatment.

versus 77%; $p = 0.027$). At the same time, the proportion of high-pay PI participants believing their performance ranked in the top half was higher in the PI treatment than in the pooled control treatments (85% versus 77%), although not significantly so in the raw data.

Table 3
Relative Performance Beliefs, Experiment 1

Dependent variable = Participant believed own performance was in top 50%						
	(1)	(2)	(3)	(4)	(5)	(6)
	PI treatment only		C-HP + C-LP		PI + C-HP + C-LP	
High pay	0.26*** (0.09)	0.30*** (0.06)	-0.02 (0.12)	0.01 (0.07)	-0.02 (0.12)	0.01 (0.06)
PI treatment	-	-	-	-	-0.16*** (0.05)	-0.18*** (0.05)
PI treatment × high pay	-	-	-	-	0.23** (0.06)	0.23***
Proportion correct		0.93 (0.57)		1.19** (0.47)		1.14*** (0.40)
Proportion seen correct		-0.33 (0.21)		-0.48 (0.38)		-0.39** (0.19)
Above median performance		0.16 (0.14)		0.11 (0.16)		0.09 (0.12)
Average guessing order		-0.02 (0.06)		0.09 (0.09)		0.04 (0.06)
Male		-0.02 (0.13)		-0.19*** (0.03)		-0.10 (0.06)
Age		0.02 (0.01)		0.00 (0.02)		0.01 (0.01)
Observations	83	79	69	68	152	147

Notes. Each column presents marginal effects estimates from a separate probit model. Columns 1 and 2 use only data from PI treatment sessions, columns 3 and 4 use only data from the pooled control sessions (C-HP and C-LP), while columns 5 and 6 use all data from all sessions. The dependent variable in each column is the same: a dummy variable taking the value one whenever that the participant stated believing their performance was in the top 50% compared to other participants in the urn-guessing task, and taking the value zero otherwise. Robust standard errors clustered by session appear in parentheses. *** Significant at 1%, ** significant at 5%, * significant at 10%. The independent variables are: 'high pay' is a dummy taking the value of one if the participants' pay structure was \$4 for a correct urn-guess and \$2 for an incorrect guess, and 0 otherwise; 'PI treatment' is an indicator variable taking the value one if the participant was assigned to the PI treatment, and 0 if the participant was assigned to either the C-HP or the C-LP treatment; 'proportion correct' is the proportion of the participant's 10 urn guesses that were correct; 'proportion seen correct' is the proportion of correct guesses made by others that the participant observed, averaged over all 10 rounds of the urn-guessing task; 'above median performance' is a dummy variable indicating whether the participants' performance in the urn-guessing task was actually (weakly) above the median – this variable was obviously unknown to the participant; 'average guessing order' is the participant's average guessing sequence order over all 10 rounds of the urn-guessing task.

Moving beyond raw means, Table 3 reports marginal effects estimates from six separate probit models of the decision to report an above median performance belief. As a baseline for comparison, columns 1, 3 and 5 report reduced form estimates which include no additional controls beyond the bare minimum: a dummy for pay level and, where applicable, a dummy for the PI treatment as well as an interaction with pay level. Columns 2, 4 and 6, on the other hand, include controls for several potentially important confounds: participants' actual performance (proportion correct); a measure of the accuracy of the other participants' guesses that an individual witnessed during the course of the experiment (proportion seen correct); and a dummy variable indicating whether the participant actually performed above the median, as a simple

specification check (above median performance).¹⁸ Since performance may be affected by the sequence order in which an individual submits his or her guess, I also include a control for each individual's average guessing sequence order (average guessing order). Finally, I control for the limited set of demographics I have at my disposal: gender and age.¹⁹

Columns 1 and 2 restrict attention to observations from the PI treatment. Without controls (column 1), the impact of being assigned to the high-pay group is exactly that observed in the raw means: a 26 percentage point increase in the likelihood of reporting an above median performance belief. Inserting controls (column 2) increases the estimated magnitude of this effect slightly, to 30 percentage points. Interestingly, none of the other controls – not even actual performance – are by themselves significant.

Considering C-HP and C-LP by themselves (columns 3 and 4) and repeating the same exercise, the data suggest that pay level by itself has no effect on performance beliefs. The estimated marginal effect of high pay is both close to zero and non-significant. Instead, actual performance is a significant determinant of beliefs, as is gender: those who perform better in an absolute sense, as well as females, are more likely to report above-median performance beliefs.

Finally, in columns 5 and 6 I pool all observations and formally compare treatment to control to estimate the impact of inequality on relative performance beliefs. Both columns tell a similar story. The negative and significant coefficient on the dummy for PI treatment demonstrates that being exposed to disadvantageous inequality reduces relative performance beliefs: low-pay participants are from 16 (column 5) to 18 (column 6) percentage points less likely to report above-median performance beliefs in the PI treatment than in the relevant control sessions (C-LP). At the same time, the positive and significant interaction between high pay and the PI treatment dummy, coupled with the small and non-significant coefficient on the dummy for high pay, show that high pay bolsters relative performance beliefs. The implied effect of high pay in the context of inequality is to increase the probability of reporting an above-median performance belief by about 20 percentage points.

RESULT 2. Salient inequality impacts relative performance beliefs, undermining the beliefs of the disadvantaged and bolstering the beliefs of the advantaged.

¹⁸ Recall that, after the urn-guessing phase of the experiment was over, each subject was asked the following question: 'compared to others, how accurate were your urn predictions?' Their answers could be either 'top 50%' or 'bottom 50%'. One minor technical issue that must be noted is that the median number of correct guesses varies depending on how widely one interprets the reference group 'others'. The narrowest definition I considered was defined with respect to subjects within each session. I also considered two slightly broader definitions that turned out to be less problematic as they coincided at a value of 60% of guesses correct. These broader medians incorporated either all guesses within each version, or all guesses across all versions. Within the control sessions of the experiment (C-LP and C-HP) these distinctions were unnecessary as the various definitions of the median coincided. Results are robust to using different definitions of the median.

¹⁹ Demographics were collected as part of a non-incentivised survey conducted at the very end of each session. The survey was conducted after all rounds of the urn-guessing game were completed and after beliefs were elicited in order to avoid priming effects.

2.3. Discussion

Overall, the results from Experiment 1 show that salient inequality significantly colours beliefs about relative performance. This raises three main questions. First of all, it is not clear from the results whether inequality primarily affects beliefs about ability or effort. This is an important distinction. If pay inequality operates primarily on effort beliefs, then to the extent that effort provision here reflects transient situational/behavioural factors rather than more permanent dispositional/characterological factors the effects of inequality may be short-lived and peculiar to retrospective performance beliefs: if one performed poorly in the past simply because of low effort, past performance may be uninformative about future performance and consequently have limited impact on decisions about competing in the future. On the other hand, if inequality primarily affects beliefs about a stable trait such as ability, then the effect of inequality on beliefs may be persistent enough to affect decisions involving future ability-based competition. A second open question is whether opportunities for learning about one's own and others' performances are necessary to produce the belief patterns observed. Since I only elicited beliefs after many rounds of the urn-guessing task, the data cannot tell me whether learning exacerbated or dampened inequality's impact on beliefs. Finally, a third question raised by Experiment 1 is whether past unequal treatment does in fact affect future competitiveness. While the belief patterns observed would suggest as much, the data provide no direct evidence.

To address all three of these questions directly I conduct a second experiment, detailed below. Before moving on, though, let me take a moment to address one aspect of the effort *versus* ability question indirectly. One may be concerned that the patterns in beliefs are driven not by a personal reaction to salient inequality but instead by beliefs about the effects of inequality on performance at the group level. Essentially, participants may have believed that pay inequality would demoralise the disadvantaged group and that demoralisation (or bad mood, more generally) lowered effort provision and hence performance. While there is no evidence that disadvantage actually lowered effort provision – recall, participants disadvantaged by inequality performed no worse and, if anything, better than the advantaged group – participants may still have expected mood or morale to affect performance which could have affected their reported beliefs.

If this were the case, one would expect to find a strong consensus among all participants that the high-pay group, as a whole, performed better. What would be needed to shed light on this implication would be an incentive compatible question about which group performed better. As it turns out, I have just such a question: after reporting their own relative performance beliefs, participants were asked to state their beliefs about which group, as a whole, performed better. This group performance belief was elicited using the same incentive compatible mechanism used to elicit own relative performance beliefs.²⁰ Figure 3 reports histograms of the responses. Overall, only 57% of PI treatment participants reported believing the high-pay group

²⁰ Participants were asked which group answered a larger proportion of their urn guesses correctly. They could choose between one of two options, presented in random order: the high-pay group or the low-pay group. Participants received one additional dollar for a correct answer and no additional dollars for an incorrect answer.

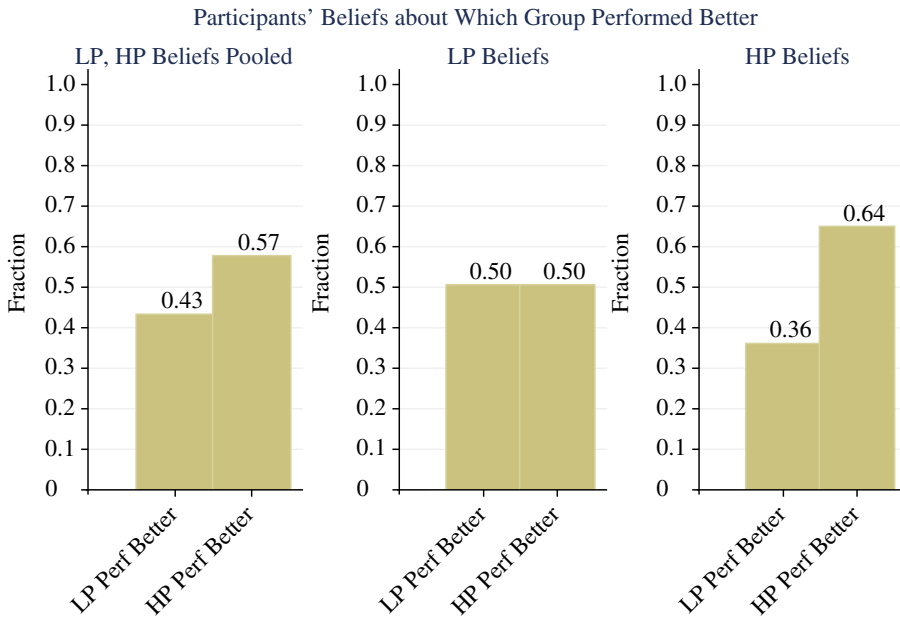


Fig. 3. *Beliefs About Relative Group Performance, Experiment 1*

Notes. The Figure presents histograms of the proportion of participants reporting believing the low-pay group (LP) or high-pay group (HP) performed better, as a whole, in terms of the accuracy of their urn guesses.

performed better than the low-pay group – hardly a landslide and only marginally significantly larger than 50% ($p = 0.08$). Low-pay group participants were, in fact, exactly evenly split on this question. Consequently, the data suggest that a story relying primarily on beliefs about the effect of inequality on group-level effort cannot account for all of the patterns observed and, in particular, is an unlikely explanation for the negative effect of inequality on the beliefs of participants disadvantaged by inequality.

3. Experiment 2

To address more directly whether inequality colours beliefs about effort or ability, I conduct a second experiment. In this second experiment, I vary participants' pay level in the same way as before, i.e. two pay levels with different total compensation but identical marginal performance incentives. In addition, I vary the nature of the task involved. Participants are randomly assigned one of two possible tasks: either an ability-intensive task (ten questions from an IQ test) or an effort-intensive task (10 letter-counting exercises). Intuitively, if the effect of inequality on relative performance beliefs primarily operates through beliefs about relative effort provision, then one would expect a strong and qualitatively similar response of beliefs to inequality on both tasks. On the other hand, if inequality primarily colours beliefs about relative ability, then one would predict a strong effect of inequality on performance beliefs only on the ability-intensive task. Furthermore, in order for either channel to be a plausible

explanation for the belief patterns documented in the previous experiment, one would also expect the relationship between inequality and performance beliefs to be qualitatively similar to the relationship documented there.

More formally, in a theory appendix available (online Appendix B) I construct a simple model providing a theoretical basis for these intuitive conjectures. The model accounts for the fact that performance on any task will typically require effort so that it may never be possible to obtain a pure measure of (cognitive) ability. However, because effort is a shared determinant of performance in the ability-intensive and effort-intensive tasks I use, comparing the effect of inequality on beliefs across these two tasks can shed some light on how salient inequality affects beliefs about ability and effort provision separately.

Before moving on, two more differences between this and the previous experiment deserve mention. First of all, tasks here are completed individually and without feedback. There are no opportunities to learn about one's own or about others' performance. Comparing Experiment 1 to Experiment 2 may, therefore, provide evidence about the role that explicit relative performance feedback plays, if any, in determining the effect of inequality on beliefs. Second, to investigate whether past inequality affects subsequent competition decisions, participants here must choose which of two possible pay schemes – piece-rate pay or tournament incentives – will determine their compensation on a second task.

3.1. *Design and Procedures*

Experiment 2 was conducted on-line using participants recruited from Mechanical Turk, an online labour market for small tasks provided by Amazon.com. The experiment paid a fixed participation fee of \$1 which is high but not extreme, by Mechanical Turk standards. The recruitment advertisement specified that in addition to this fixed fee, each participant had a 10% chance of being paid their otherwise-hypothetical earnings from the experiment, which could be quite a bit higher than the participation fee. Random selection for payment was made credible by basing it on an official state lottery: the immediately subsequent drawing of California's mid-day 'Daily 3'. (for details, see online Appendix A). Both of these factors were intended to garner a high number of participants quickly in order to minimise word-of-mouth learning about the design of the experiment. Participation was restricted to the US in order to keep incentives comparable across individuals. In total, six separate sessions were conducted and 518 individuals completed the experiment.²¹

Experiment 2 is comprised of three phases. Before the first phase begins, participants are instructed that only one phase will be randomly chosen to count towards their potential experimental earnings. They are told nothing in advance about each of the three phases except that each new phase will be clearly labelled as such. The first phase features a $2 \times 2 \times 2$ design that varies: pay level (high pay, low pay);

²¹ When conducting experiments on Mechanical Turk, one may worry about attrition introducing selection effects. In the current context, the main concern would be that low pay induced participants to exit more readily than high pay did. Reassuringly, attrition was low overall and, more importantly, did not vary by pay level. Overall attrition was 6.8%. Among those assigned high pay (low pay), attrition was 7.0% (6.5%).

inequality salience (informed/not informed about other pay level); and task type (effort-intensive, ability-intensive). Participants assigned high pay earn \$2 for each correct answer and \$1 for each incorrect answer, while participants assigned low pay earn \$1 for each correct answer and nothing for incorrect answers. The treatment sessions can be thought of as those featuring salient inequality. In these sessions, before pay levels are assigned, participants are told about both possible pay levels and instructed that they have an equal chance of being assigned either one. In the control sessions – not featuring salient inequality – participants are only told about the pay level they are assigned. No participant is ever informed about the existence of the task he or she is not assigned.

The ability-intensive task consists of ten questions from a Raven's progressive matrices booklet. This test is widely recognised as a culture-free measure of general intelligence and consists of guessing which of eight pieces best completes a picture. The effort-intensive task also features ten separate, multiple choice, questions. Each effort-intensive task question presents participants with a different 200-character string of upper and lower case letters and asks how many times a particular letter appears in the string. This latter task requires essentially no ability; instead performance depends solely on effort. To make the two tasks as similar as possible, each question on both tasks is presented as a picture of roughly the same size and each task provides participants with eight possible answers.²²

The second phase of Experiment 2 consists of belief elicitation. After completing all ten questions of their assigned phase 1 task, participants are asked to state how they believe they performed on the task relative to other participants in their session. Valid responses consist of three mutually exclusive categories: strictly above the median, strictly below the median or exactly at the median. This mechanism provides proper incentives to report beliefs truthfully by the same reasoning used to justify the mechanism used in Experiment 1. Relative to Experiment 1 the incentives for reporting beliefs truthfully are ramped up in two ways. First of all, here a correct belief pays \$20 instead of \$1, while an incorrect belief still pays nothing. Second, because only one phase of the experiment is chosen to determine each participant's potential earnings, the beliefs participants state here are more consequential: they may determine an individual's entire earnings.

In the final phase – phase three – participants learn that they must now complete another task which will be similar to their phase 1 task. Before beginning the task, however, they must choose their pay structure. The choice is between piece-rate pay or tournament incentives.²³ The piece-rate option, which pays \$1.50(\$0.50) for each correct (incorrect) answer, provides the same marginal performance incentives as the two pay levels used in phase 1: each correct answer pays one dollar more than each incorrect answer. Tournament incentives are implemented as follows. After completing the phase three task, participant i 's score (s_i) is compared with the score of one other, randomly selected, participant (s_j). If $s_i \geq s_j$ ($s_i < s_j$), participant i earns \$30(\$0).

²² For the effort-intensive task, presenting the string of characters as a picture rather than simple text has an added benefit. It makes it cumbersome, if not impossible, to cut-and-paste the string into a text editor which would otherwise be one easy way to cheat on the task.

²³ Neutral wording is used to describe these two options.

The parameters of the incentive structures were chosen to ensure a substantial fraction of participants chose each option: a risk neutral individual who believes s/he will answer eight questions correctly and that this performance will be placed below the median with a 60% chance is indifferent between the two options. To minimise the possibility that extraneous considerations such as other-regarding preferences affect the choice of pay structure, it is explained to participants that their choice of pay structure in no way affects any other participant's earnings.

Finally, after participants completed the experiment they were asked to provide some demographic details about themselves. I use these self-reported demographics as controls where appropriate in the formal analysis of the results, below. The outcomes of interest in Experiment 2 are two:

- (i) retrospective relative performance beliefs related to the initial, phase 1, task; and
- (ii) the choice of pay structure for the phase 3 task. Let us now turn to the results.

3.2. Results

I begin by considering performance on the phase 1 task. Table 4 reports the number of questions participants answered correctly, on average, by condition, task and pay level. As in the previous experiment, actual performance is never affected by pay level. For example, participants assigned high pay and the ability-intensive task in the context of salient inequality (top left quadrant) answered 8.20 of the ten questions correctly, while those assigned low pay performed slightly better at 8.31 – although this difference is obviously not significant. On the effort-intensive task (top right quadrant), the picture is similar: participants assigned low pay performed slightly better than those assigned high pay, on average answering 8.78 *versus* 8.55 questions correctly respectively. The only difference worth noting is that the introduction of salient inequality on the ability-intensive task generally increased performance compared to the control condition (8.20 *versus* 7.48 for high pay, $p = 0.02$; 8.31 *versus* 7.28 for low pay $p = 0.03$). The effect was roughly uniform across pay levels, however,

Table 4
Actual Task 1 Performance, Experiment 2

	Ability-intensive task		Effort-intensive task	
	High pay	Low pay	High pay	Low pay
Treatment				
Number of correct answers	8.20	8.31	8.55	8.78
SE	(0.23)	(0.25)	(0.19)	(0.14)
Observations	64	68	75	79
Control				
Number of correct answers	7.48	7.28	8.46	8.69
SE	(0.35)	(0.34)	(0.18)	(0.29)
Observations	68	67	61	36

Notes. The Table reports the average number of correct answers out of the 10 questions asked on the first task of Experiment 2, by pay level and treatment. Standard errors appear in parentheses.

making it an unlikely explanation for the type of belief patterns observed in the previous experiment. Nevertheless, in the analysis of results below I control for task performance where possible.

Turning from means to distributions, Figure 4 presents histograms of the distribution of performance on the ability-intensive task by pay level and condition, while Figure 5 presents these histograms for effort-intensive task performance. These histograms reveal substantial individual heterogeneity in performance on both tasks but no glaring differences in performance across pay levels in either the treatment or control sessions. The most consistent difference observed in all histograms is a shift in mass towards perfect performance (10 correct answers) among the low-pay group as compared to the high-pay group. In both treatment and control sessions, on both the ability-intensive task (Figure 4) and the effort-intensive task (Figure 5), the proportion of participants with perfect scores is higher among the low-pay group than among the high-pay group. The performance histograms therefore support the notion that, if anything, low pay tends to increase performance rather than undermine it.

RESULT 3. *Salient inequality does not decrease performance of the low-pay group on either the ability-intensive or the effort-intensive task.*

Next, I examine whether inequality affects relative performance beliefs. Since the belief elicitation mechanism used here admits three categories rather than two – above, below or exactly median – I begin by examining histograms of participants’ beliefs. Considering observations from the treatment sessions first, in Figure 6 I report beliefs from participants assigned the ability-intensive task in the top row and from




Fig. 4. *Actual Performance, Ability-intensive Task, Experiment 2*

Notes. The Figure reports histograms of participants’ total number of correct answers (max possible = 10) in Experiment 2, by pay group and treatment for the ability-intensive task. The x-axis of each histogram refers to the number of correct guesses.

Performance On Effort Intensive Task

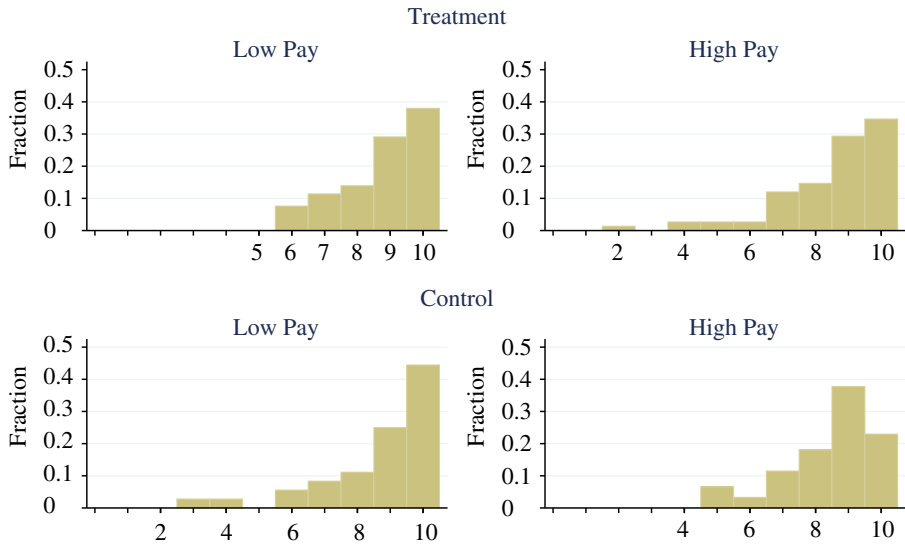


Fig. 5. Actual Performance, Effort-intensive Task, Experiment 2

Notes. The Figure reports histograms of participants' total number of correct answers (max possible = 10) in Experiment 2, by pay group and treatment for the effort-intensive task. The x-axis of each histogram refers to the number of correct guesses.

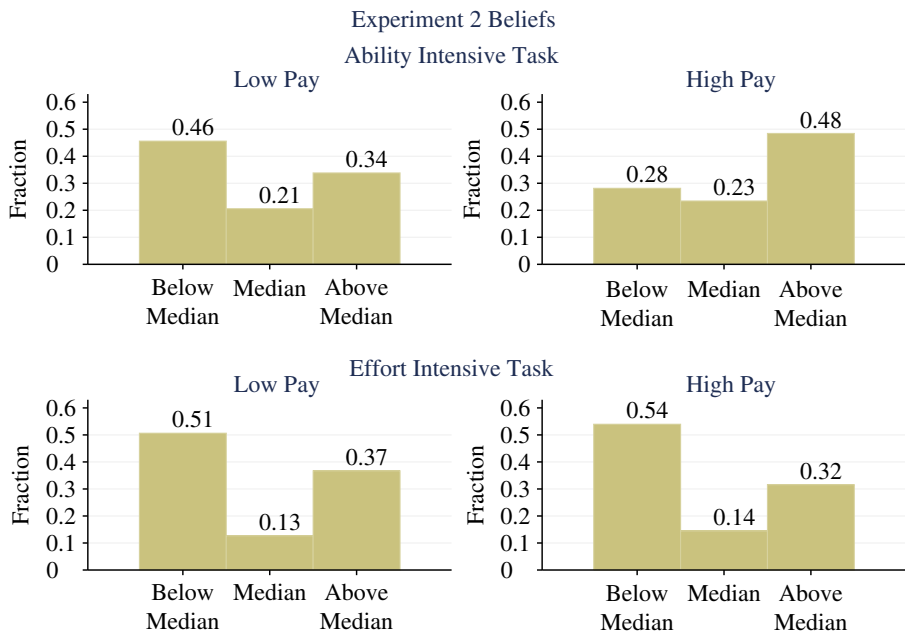


Fig. 6. Relative Performance Beliefs Across Tasks in the Treatment Sessions, Experiment 2

participants assigned the effort-intensive task in the bottom row. The overall message from these histograms is that in the presence of salient inequality beliefs react more strongly to pay when the task is ability-intensive. Beliefs appear qualitatively similar across pay levels when the task is effort-intensive. However, when the task is ability-intensive there is a qualitative and quantitative shift towards believing one performed below the median (46% versus 28%, $p = 0.02$) and away from believing one performed above the median (34% versus 48%, $p = 0.04$) among disadvantaged participants. This lends credence to the notion that salient inequality affects beliefs about ability rather than effort in a way that justifies and may reinforce inequality.

Another way to test this intuition is to compare beliefs in the treatment sessions to beliefs in the control sessions. Since the only difference between treatment and control is the presence of salient inequality, this comparison sheds light on how the introduction of salient inequality itself affects beliefs and whether this effect differs by the nature of the task involved. In Figure 7, I present histograms of participants' beliefs in the treatment session (solid bars) overlaid with histograms of beliefs in the control sessions (transparent, dashed-line, bars). Considering first the ability-intensive task (left panel), there is no qualitative difference in beliefs by pay level in the control sessions. For both pay levels, roughly equal mass is placed on above and below-median beliefs. The proportion of participants reporting an above-median belief is slightly higher among the high-pay group than among the low-pay group, but not significantly so ($p = 0.27$). The same can be said about the proportion reporting a below-median

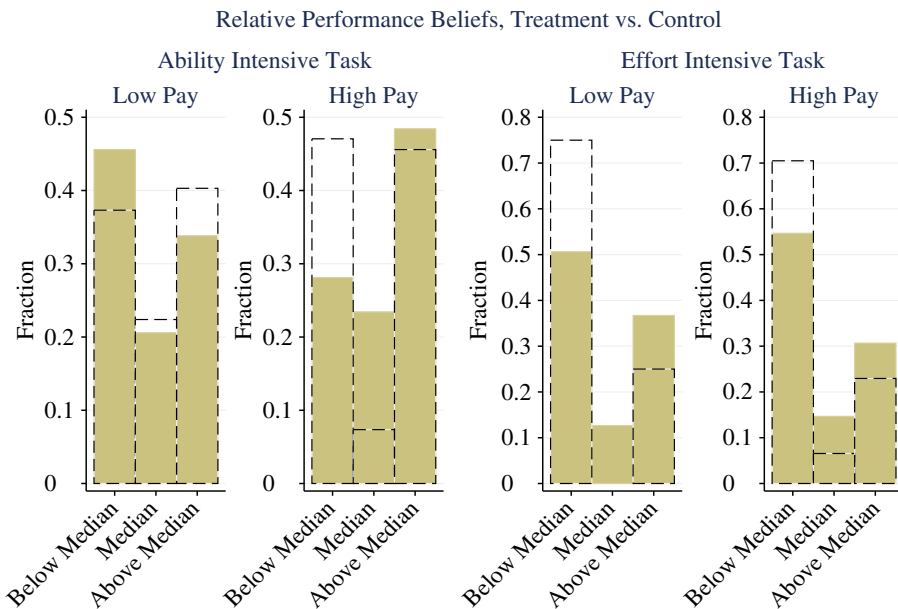


Fig. 7. Relative Performance Beliefs in Treatment and Control, Experiment 2

Notes. The Figure reports the proportion of participants reporting believing their performance ranked above, below or at the median in terms of accuracy in Experiment 2, by pay group and treatment on the ability-intensive task. The solid bars refer to the treatment session data; the dashed-line bars refer to the control sessions.

performance belief ($p = 0.13$). This comparison suggests that pay level itself does not affect relative performance beliefs qualitatively and consequently is an unlikely explanation for the qualitative differences in beliefs by pay level observed in the context of salient inequality.

Considering next the effort-intensive task (Figure 7, right panel), I find no qualitative nor quantitative difference in the control session beliefs (transparent, dashed-line, bars) across pay levels. For instance, the proportion of participants reporting an above-median or below-median belief does not differ significantly across pay levels at anywhere near conventional significance levels. The same can be said for the treatment sessions (solid bars) considered separately so that, apparently, pay level does not affect performance beliefs related to effort-intensive tasks with or without salient inequality. However, the histograms do reveal a possible bright side of inequality. Irrespective of pay level, participants are significantly less likely to believe they performed below the median in the treatment sessions than in the control sessions ($p = 0.01$ low pay; $p = 0.03$ high pay). Because this effect is uniformly positive, however, inequality's effect on relative effort beliefs remains an unlikely explanation for the qualitative inequality-justifying shift in beliefs observed in the previous experiment.

Digging a bit deeper, I estimate formal models of participants' beliefs which allow me to control for potentially confounding factors. In Table 5 I restrict attention to participants assigned the ability-intensive task and report six separate model estimates. The basic set of controls present in all six columns includes a dummy for high pay (high pay), a dummy for whether the session involved salient inequality (treatment) and an interaction between these two indicators. Because actual performance may determine beliefs and average performance differed across treatment and control sessions, I also insert as an explanatory variable the number of questions a participant answered correctly. As a proxy for familiarity with the California state lottery used to select participants for payment which may affect perceived incentives, I also insert a full set of US time zone dummies.²⁴ All even columns feature, in addition, controls for available demographics: income, age, gender and education.

Columns 1 and 2 of Table 5 present estimates of participants' full three-category relative performance beliefs measure. Because this measure is ordinal and has more than two categories, the appropriate model is an ordered probit. The small and non-significant coefficient on the dummy for high pay in both columns demonstrates that high pay *sans* salient inequality has no effect on performance beliefs. On the other hand, the positive and significant coefficient on the interaction between high pay and treatment shows that introducing salient inequality bolsters the performance beliefs of the highly paid. Finally, after controlling for demographics (column 2) the negative and significant coefficient associated with treatment documents that low pay (the excluded group) undermines the beliefs of the disadvantaged in the context of salient inequality relative to low pay in the control sessions lacking inequality.

While the ordered probit estimates establish the significance of the effect of inequality on beliefs, they are notoriously difficult to interpret. Consequently, I

²⁴ I infer the participant's actual location from the location associated with his or her ip address, which I have in my data. These dummies also absorb any region-specific fixed effects which may vary across time zones, examples of which may include the level of education of the population or differences in culture.

Table 5
Relative Task 1 Performance Beliefs, Ability-intensive Task, Experiment 2

	Relative ability beliefs, ability-intensive task					
	(1)	(2)	(3)	(4)	(5)	(6)
	3-category belief measure		Below median belief dummy		Above median belief dummy	
High pay	-0.07 (0.19)	-0.09 (0.18)	0.10 (0.10)	0.11 (0.09)	0.05 (0.05)	0.05 (0.05)
Treatment	-0.24 (0.16)	-0.25** (0.12)	0.10 (0.08)	0.11* (0.06)	-0.09 (0.08)	-0.08 (0.08)
High pay × treatment	0.49** (0.23)	0.56** (0.22)	-0.26** (0.12)	-0.29*** (0.11)	0.10* (0.06)	0.12* (0.07)
Number of correct answers	0.04*** (0.02)	0.05*** (0.02)	-0.01* (0.01)	-0.01** (0.01)	0.02** (0.01)	0.02** (0.01)
Demographic controls	N	Y	N	Y	N	Y
Observations	267	265	267	265	267	265

Notes. Columns 1 and 2 present ordered probit estimates using as the dependent variable the 3-category relative performance belief measure described in the text, taking values 1 = 'believed performance was below median' to 3 = 'believed performance was above median'. Columns 3–6 report marginal effects estimates from probit models using as the dependent variable the dummy listed in the column heading. Robust standard errors clustered by session appear in parentheses. *** Significant at 1%, ** significant at 5%, * significant at 10%. The explanatory variables are: 'high pay' = a dummy taking the value of one if the participants' pay structure was \$2 for each correct guess and \$1 for each incorrect guess; 'treatment' = a dummy indicating whether participants took part in a session where they were informed about both potential pay scales; 'Number of correct answers' = the participant's actual performance on task 1. This variable is unknown to participants, of course; Demographic controls are all self-reported and include: income category dummies based on participant's last-year's total income; age and age squared, which are constructed from the participant's answer to the question 'in what year were you born?'; an indicator for whether the participant has a college degree; and gender. Demographic controls were typically non-significant, so they have been omitted for readability. In addition, each estimate includes a full set of dummies for the US time zone in which the participant's ip address was located – 'Mountain Time Zone', 'Central Time Zone' and 'Eastern Time Zone', the excluded category being 'Pacific Time Zone' – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives.

estimate the effect of inequality on the propensity to report a below-median belief (columns 3 and 4) and the propensity to report an above-median belief (columns 5 and 6), separately. Because the dependent variable in each of these estimates is binary, I can use a probit model and report the associated marginal effects estimates which are easier to interpret: they correspond to the increase in the probability of reporting a below-median or above median belief from increasing the explanatory variable by one unit. While these estimates are easier to interpret, they also throw away information about the magnitude of changes in beliefs – i.e. whether beliefs move by one category or two categories – so the associated significance levels should be interpreted as lower bounds. Considering first below-median beliefs (columns 3 and 4), I find that high pay significantly decreases the probability that an individual reports a below median belief, but only when inequality is salient. After controlling for demographics, the estimates also suggest that introducing salient inequality (marginally) significantly increases the probability that a disadvantaged participant reports a below-median belief. Putting this in perspective, the estimated effect of salient inequality on the performance beliefs of the disadvantaged (11 percentage points) is roughly equivalent to the effect of

decreasing one's actual performance from a perfect task score (10) to an abysmal task score (0). Finally, considering the propensity to report an above-median score (columns 5 and 6) I find a consistent picture: salient inequality bolsters the performance beliefs of the advantaged, increasing their propensity of reporting an above-median belief; the estimated effect of inequality on the beliefs of the disadvantaged is again deleterious, reducing the propensity to report an above-median belief. The magnitude of this latter effect is again sizable, being roughly equivalent to four extra questions correct out of ten, albeit non-significant.

Next, I repeat the same exercise restricting attention to the effort-intensive task (Table 6). The estimates again paint a picture consistent with the histograms above even after controlling for several potentially confounding factors. The first two columns present ordered probit estimates of participants' full three-category relative performance belief measure. Here, I find no differential effect of pay level when inequality is made salient: the coefficient on the interaction between treatment and high pay is non-significant. The positive and significant coefficient associated with treatment by itself, coupled with all other coefficients being non-significant, suggests that introducing salient inequality results in a uniform increase in performance beliefs. Columns 3 and 4 suggest this increase comes through a significant decrease in the likelihood of believing one performed below the median. This latter effect is uniform

Table 6
Relative Task 1 Performance Beliefs, Effort-intensive Task, Experiment 2

Relative ability beliefs, effort-intensive task						
	(1)	(2)	(3)	(4)	(5)	(6)
	3-category belief measure		Below median belief dummy		Above median belief dummy	
High pay	0.09 (0.11)	0.09 (0.13)	-0.06 (0.04)	-0.06 (0.04)	-0.02 (0.04)	-0.01 (0.05)
Treatment	0.55* (0.29)	0.55** (0.28)	-0.25*** (0.09)	-0.25*** (0.08)	0.11 (0.08)	0.12 (0.08)
High pay × treatment	-0.21 (0.19)	-0.20 (0.23)	0.10 (0.06)	0.10 (0.08)	-0.04 (0.07)	-0.04 (0.08)
Number of correct answers	0.01 (0.04)	0.01 (0.04)	0.00 (0.02)	0.01 (0.02)	0.01 (0.01)	0.02 (0.02)
Demographic controls	N	Y	N	Y	N	Y
Observations	251	247	251	247	251	247

Notes. Columns 1 and 2 present ordered probit estimates using as the dependent variable the 3-category relative performance belief measure described in the text, taking values 1 = 'believed performance was below median' to 3 = 'believed performance was above median'. Columns 3 and 4 report marginal effects estimates from probit models using as the dependent variable a dummy taking the value of one if the participant reported believing they performed above the median (i.e. a value of 3 in the 3-category scheme). Robust standard errors clustered by session appear in parentheses. *** Significant at 1%, ** significant at 5%, * significant at 10%. In addition, each estimate includes a full set of dummies for the US time zone in which the participant's ip address was located – 'Mountain Time Zone', 'Central Time Zone' and 'Eastern Time Zone', the excluded category being 'Pacific Time Zone' – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives. For a description of all control variables, both reported and omitted, see Table 5, above.

across pay levels, as in the histograms above. Meanwhile, above-median performance beliefs (columns 5 and 6) are apparently not affected.

RESULT 4. Salient inequality significantly affects relative performance beliefs in a way consistent with inequality only when task performance depends on ability.

All together, the visual evidence provided by the histograms of participants' performance beliefs and the more formal evidence obtained by estimating models of participants' beliefs tell a consistent story. The data suggest that the experience of salient inequality colours beliefs about relative ability in a way that justifies the inequality. At the same time, when performance depends solely on effort provision the data suggest that experienced inequality has, if anything, a uniformly positive effect on relative performance beliefs. Importantly, this makes beliefs about relative effort provision an unlikely explanation for the belief patterns observed in Experiment 1. Because ability is a stable trait, it is plausible to expect some persistence in the differential effects of inequality on relative ability beliefs and to suspect that prior inequality may continue to shape the subsequent decisions which depend upon these beliefs in a way that reinforces initial inequality.

Before considering competitiveness, I must present one additional piece of evidence. Recall, one reason for conducting the second experiment was to shed light on the effort *versus* ability confound, i.e. whether the effect of salient inequality on performance beliefs is primarily driven by beliefs, albeit mistaken, about the group-level effect of inequality on effort provision rather than individuals' beliefs about their own traits. However, it could be that the ability-intensive and effort-intensive tasks are simply too different to provide convincing evidence on this question. Even though, as before, actual performance never differs by pay level it is still possible that participants expected disadvantage to demoralise the low-pay group as a whole. To test one implication of this confound, in the treatment sessions I elicited beliefs about which group, as a whole, performed better.²⁵ If beliefs were driven by expectations about group-level demoralisation, one would expect responses to this group-level performance beliefs question to be highly skewed towards believing the pay-disadvantaged group perform strictly worse. In Figure 8 I report histograms of responses to this question, restricting attention to the most relevant case: treatment session participants assigned the ability-intensive task. I find little evidence that beliefs reflect a strong consensus about the low-pay group performing worse, either overall or within each pay group. The proportion believing the high-pay group performed better is never significantly different from the proportion believing the low-pay group performed better and, moreover, the mode of each distribution is that both groups performed equally well.

²⁵ Participants were asked which group, on average, will have the highest percentage of correct answers on the phase 1 task. Three mutually exclusive responses were provided: the high-pay group performed better; the low-pay group performed better; or both groups performed equally well. Participants could earn \$20 for a correct response and nothing for an incorrect response. Participants were instructed that if phase 2 was chosen to count towards their earnings, then one of the two questions comprising phase 2 – this question or the individual relative ability belief question – would be chosen randomly to determine their entire earnings. Thus, this question is both incentive compatible and consequential.

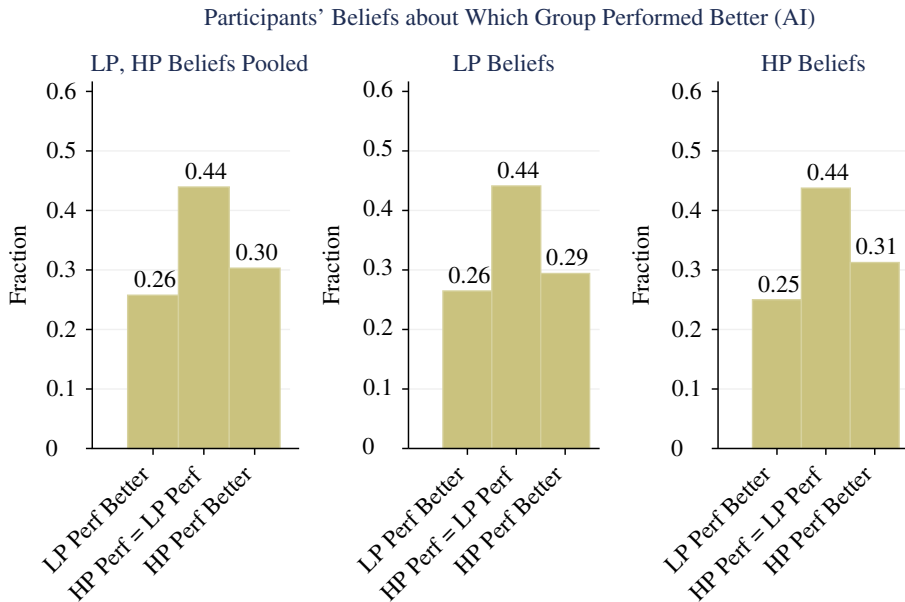


Fig. 8. *Beliefs About Relative Group Performance, Ability-intensive Task, Experiment 2*

Notes. The Figure presents histograms of the proportion of participants reporting believing the low-pay group (LP) or high-pay group (HP) performed better, as a whole, in terms of number of questions answered correctly. Observations are restricted to ability-intensive task participants.

3.3. *Willingness to Compete*

In this subsection, I ask whether inequality moves beliefs sufficiently strongly to affect subsequent decisions that plausibly depend on individuals' assessments of their own relative ability. While there are many decisions that may fit this description, here I consider only the most obvious and straightforward: the decision to compete on a subsequent ability-intensive task. Ability-based competition may be a particularly important determinant of inequality persistence in the real world as quintessential examples include many paths which are often singled out as crucial for upward mobility. For example, applying to and succeeding at highly selective colleges or pursuing many well remunerated careers all require ability-based competition at some point. Moreover, since these paths require long sequences of costly pre-commitments of time, attention and money as well as repeated critical junctures at which individuals must actively decide to stay the course, the effect of inequality on beliefs need not be extraordinarily persistent. The interaction between being more likely to experience inequality coupled with repeated chances to undermine one's own progress may prove sufficient.

With this in mind, recall that the final phase of Experiment 2 had participants choose whether piece-rate pay or tournament incentives would apply to a subsequent task. Participants were told only that this subsequent task would be similar to the task they had completed in phase 1. I begin my analysis of participants' choices between compensation schemes by remarking that the choice may clearly depend on preference primitives in addition to ability beliefs. Because many prior studies point

to a substantial gender difference in the type of preferences which seem likely to be involved (e.g. risk preferences), to isolate the impact of beliefs on competition I split the data by gender and analyse male and female decisions separately.²⁶

First, I restrict attention to male participants and analyse the raw data graphically. Figure 9 presents the proportion of male participants choosing tournament incentives by task and experiment condition. Considering the ability-intensive task (left panel) where beliefs were strongly affected by pay only in the presence of salient inequality, consistent with the story so far there is no discernible effect of unequal pay *sans* inequality (42% *versus* 45%: $p = 0.39$). Also consistent with the story, in the context of inequality (treatment) low-pay participants are less competitive than high-pay participants: 62% (48%) of high-pay (low-pay) participants selected tournament incentives. Although this difference is not by itself significant ($p = 0.16$), it is substantial in magnitude representing about one-third of the unconditional sample mean.²⁷ The bulk

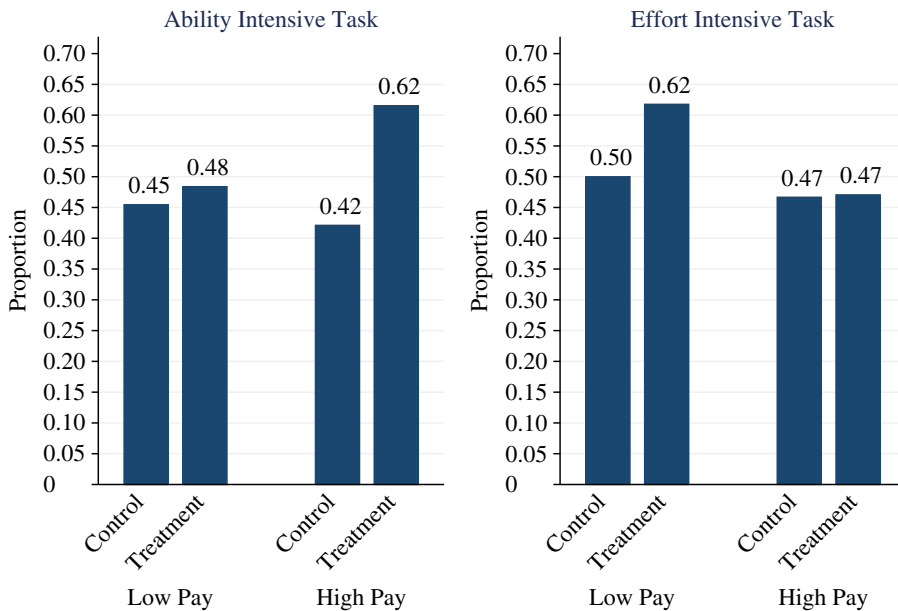


Fig. 9. *Proportion of Male participants Selecting Tournament Incentives, Experiment 2*

²⁶ As an example, consider two individuals who have identical relative ability beliefs before being exposed to inequality but who differ with respect to risk preferences. Assume the piece-rate pay is perceived as being less risky. The more risk averse individual is more likely to prefer the safer (piece-rate pay) option before inequality is introduced and would require a larger boost in ability-beliefs from advantageous inequality to switch to preferring tournament incentives. Consequently, systematic differences in risk preferences make it more difficult to detect an effect of inequality on tournament choice. Among the more risk-averse group, for any given magnitude of the effect of inequality on beliefs it is *ceteris paribus* less likely to observe a switch from preferring the piece-rate option to preferring tournament incentives. Since I collected no separate measure of risk preferences and my beliefs measure is quite coarse, splitting the sample by gender seems reasonable.

²⁷ On the ability-intensive task, the sample mean of tournament is 0.43.

of the difference in competitiveness apparently comes from inequality’s effect on high-pay participants. In the treatment sessions 62% of high-pay participants select tournament incentives, while only 42% of high-pay participants in the control sessions select tournament incentives ($p = 0.07$). Turning next to the effort-intensive task (right panel, Figure 9) where one would expect no substantial effect of inequality on competitiveness since inequality did not affect beliefs, reassuringly I find no significant effect. None of the pairwise differences are significant and, in fact, high-pay participants are exactly equally likely to select tournament incentives in the treatment and control.

In Figure 10 I present histograms of female participants’ choices. Among female participants assigned the ability-intensive task (left panel), the effect of pay level on tournament choice is much milder and never significant for any pairwise comparison.²⁸ The effect is even milder on the effort-intensive task and continues to never be statistically significant.

Turning from graphical evidence to more formal estimates, Table 7 reports marginal effects estimates from eight separate probit models. The first four columns restrict attention to the ability-intensive task, while the last four consider the effort-intensive task only. Each probit model uses participants’ binary tournament incentive choice as its dependent variable.

Consider first the ability-intensive task, the most relevant case. The first two columns restrict attention to male participants, while columns 3 and 4 restrict attention to female participants. Controls are identical to those used when considering performance beliefs




Fig. 10. *Proportion of Female Participants Selecting Tournament Incentives, Experiment 2*

²⁸ The closest to being significant is the difference by pay level within the control sessions, where a one-tailed t-test yields $p = 0.19$. All the other comparisons yield $p > 0.20$.

Table 7
Selecting Tournament Incentives for Task 2, Ability-intensive Task, Experiment 2

Dependent variable = Participant chose tournament incentives								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Ability-intensive task				Effort-intensive task			
	Male		Female		Male		Female	
High pay	-0.04	-0.03	0.09	0.10	-0.02	-0.01	-0.08	-0.07
	-0.09	-0.08	-0.13	-0.13	-0.07	-0.07	-0.20	-0.23
Treatment	-0.03	-0.01	0.04	0.02	0.15*	0.17*	-0.02	-0.02
	-0.07	-0.1	-0.12	-0.12	-0.08	-0.09	-0.13	-0.16
High pay × treatment	0.18**	0.17*	-0.14	-0.11	-0.16	-0.17	0.07	0.04
	-0.09	-0.1	-0.11	-0.12	-0.12	-0.12	-0.22	-0.25
Number of correct answers	0.03***	0.03***	0.03	0.02	0.06	0.05	0.00	0.01
	-0.01	-0.01	-0.02	-0.02	-0.05	-0.04	-0.03	-0.03
Demographic controls	N	Y	N	Y	N	Y	N	Y
Observations	128	128	139	137	120	117	131	130

Notes. Each column presents marginal effects estimates from a probit model using as the dependent variable an indicator taking the value 1 whenever the participant chose tournament incentives to apply to the second task in Experiment 2. Columns 1, 2, 5 and 6 restrict attention to only male participants, while columns 3, 4, 7 and 8 report estimates using only observations from female participants. Robust standard errors clustered by session appear in parentheses. *** Significant at 1%, ** significant at 5%, * significant at 10%. Demographic controls are all self-reported and include: income category dummies based on participant’s last-year’s total income; age and age squared, which are constructed from the participant’s answer to the question ‘in what year were you born?’; an indicator for whether the participant has a college degree; and gender. Demographic controls were typically non-significant, so they have been omitted for readability. ‘Treatment’ is an indicator variable taking the value of 1 if the observation came from a session where participants were informed there were two different pay groups they could be assigned to. In addition, each estimate includes a full set of dummies for the US time zone in which the participant’s ip address was located – ‘Mountain Time Zone’, ‘Central Time Zone’ and ‘Eastern Time Zone’, the excluded category being ‘Pacific Time Zone’ – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives.

above. In the least elaborate specification (column 1), I find that salient inequality increases the propensity to select into a competitive environment among the advantaged. Controlling for actual performance – which participants were never informed of but about which they may have had an intuitive sense – the positive and significant coefficient on the interaction between inequality salience and high pay (high pay × treatment) suggests that being advantaged by prior inequality increases subsequent competitiveness by 14 percentage points – or about 33% of the unrestricted sample mean.²⁹ Moreover, the estimated magnitude of the effect of prior advantage on subsequent competitiveness is not affected by controlling for demographics (column 2). The direction of the effect of disadvantage on competitiveness is consistent with the story so far, but the estimated size is non-significant. Among female participants (columns 3 and 4), as in Figures 9 and 10, I find no discernible effect of either actual performance or past inequality on subsequent competition decisions.

²⁹ The 14 percentage point figure corresponds to the sum of the coefficient on high pay and its interaction with treatment: +0.18 – 0.04 = 0.14.

Repeating the same exercise restricted to participants assigned the effort-intensive task (columns 5–8), I find that high pay never has a differential effect on competitiveness irrespective of gender. However, the estimates do reveal an effect of introducing salient inequality on competitiveness: compared to the control sessions, male participants are marginally significantly more likely to select tournament incentives. This is true whether or not I control for demographics. This finding is reassuring, as it mirrors the patterns in relative ability beliefs documented above – at least for males. It is intriguing because it begins to suggest that salient inequality is not always bad. In environments where performance depends primarily on effort it may be that making inequality salient enhances both relative performance beliefs and subsequent competitiveness for both the advantaged and the disadvantaged.

Summing up, both the raw data and the formal estimates point to a persistent effect of prior inequality on subsequent competitiveness when success depends on ability among male participants. On the ability-intensive task – where inequality coloured ability beliefs – males previously advantaged by inequality receive a significant boost to their willingness to compete on subsequent tasks. Those previously disadvantaged by inequality receive no such boost. Together these patterns imply that prior inequality may cause the advantaged to be more willing to compete than the disadvantaged, which is exactly the type of competitiveness pattern that may turn current inequality into persistent inequality.

RESULT 5. Prior inequality significantly affects males' willingness to compete on subsequent ability-intensive tasks in a way that can perpetuate inequality.

3.4. *Do Performance Beliefs Matter for Competitiveness?*

So far, the data suggest that salient inequality affects performance beliefs and that inequality affects competitiveness. The question naturally arises: do performance beliefs directly affect competitiveness? It is possible that salient inequality affects beliefs and, at the same time, affects some other factor – say, a preference for competition – so that the link from beliefs to competitiveness is neither direct nor causal.³⁰ To address this question, I restrict attention to male participants assigned the ability-intensive task and high pay. This is the subset of my data where inequality had a significant effect on the selection of tournament incentives and, so, is the subset of the data where I would like to document a relationship between beliefs and competitiveness.

For this subset of the data, I estimate probit models of the decision to choose tournament incentives for task 2 as a function of beliefs about relative performance on task 1. I estimate three sets of models and report marginal effects in Table 8. The first set of models is the most basic (columns 1 and 2) and simply documents that a positive relationship between salient inequality and tournament choice exists and is significant in this subset of the data. For the second set of models, I use a fact gleaned from the estimates in Table 5 and the graphical evidence in Figure 7: for the advantaged, the primary effect of salient inequality is to lower an individual's propensity to report a

³⁰ I thank an anonymous referee for suggesting this exercise.

Table 8
Beliefs and Tournament Choice, Experiment 2

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.14*	0.20***	0.12	0.17***	0.03	0.08
	(0.08)	(0.07)	(0.08)	(0.07)	(0.07)	(0.05)
Below median belief dummy (BMB)			-0.29*	-0.39**	-0.41*	-0.49**
			(0.17)	(0.16)	(0.21)	(0.20)
Treatment × BMB					0.33	0.34
					(0.26)	(0.26)
Own task 1 performance	0.04	0.02	0.03	0.01	0.03	0.01
	(0.03)	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)
Demographic controls	N	Y	N	Y	N	Y
Observations	62	62	62	62	62	62
[Pseudo] R-squared	0.08	0.13	0.13	0.21	0.14	0.22

Notes. Observations restricted to male participants assigned the ability-intensive task and high pay. Each column presents marginal effects estimates from a probit model using as the dependent variable an indicator taking the value 1 whenever the participant chose tournament incentives to apply to the second task in Experiment 2. Robust standard errors clustered by session appear in parentheses; *** significant at 1%, ** significant at 5%, * significant at 10%. The primary explanatory variables are: ‘treatment’ = a dummy indicating whether participants took part in a session where they were informed about both potential pay scales; ‘Below median belief (BMB)’ = a dummy indicating whether a participant reported believing their performance on task 1 ranked below the median; ‘own task 1 performance’ = the participant’s actual performance on task 1 (max. possible = 10). This variable is unknown to participants, of course. Demographic controls are all self-reported and include: income category dummies based on participant’s last-year’s total income; age and age squared, which are constructed from the participant’s answer to the question ‘in what year were you born?’; an indicator for whether the participant has a college degree; and gender. Demographic controls have been omitted for readability. In addition, each estimate includes a full set of dummies for the US time zone in which the participant’s ip address was located – ‘Mountain Time Zone’, ‘Central Time Zone’ and ‘Eastern Time Zone’, the excluded category being ‘Pacific Time Zone’ – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives.

below-median performance belief. Consequently, in columns 3 and 4 I insert a dummy for reporting a below-median belief (*BMB*) – ignoring for now the interaction between beliefs and inequality. The results are reassuring. Controlling for demographics, being advantaged by salient inequality increases competitiveness, while there is a large and significant negative main effect of *BMB*. For the final set of models, I insert an interaction between *BMB* and the treatment dummy (columns 5 and 6). In this final set of models, the direct effect of salient inequality vanishes, being both non-significant and small in magnitude. Instead, controlling for demographics (column 6) the primary driver of the choice of tournament incentives is now *BMB* whose marginal effect is both large in magnitude and highly statistically significant.

3.5. Which Performance Beliefs Matter – Relative or Absolute?

The data suggest performance beliefs are an important factor in the choice of tournament incentives but still one may wonder whether it is truly relative performance beliefs that matter. An alternative story is that salient inequality bolsters absolute performance beliefs primarily and that the effect on absolute performance beliefs is sufficiently strong to increase the propensity to choose tournament incentives by itself. This could happen if, for example, high pay coupled with salient inequality causes

participants to be certain they will get all ten questions correct on the second task in which case ‘competition’ would pay \$30 for sure, which is better than the \$20 such participants would earn by choosing the piece-rate option.

Unfortunately, I cannot address this concern directly with the data from Experiment 2 as I elicit neither absolute nor relative performance beliefs for task 2. Therefore, I conduct an additional experiment – Experiment 3. I conduct Experiment 3 in exactly the same manner as Experiment 2: on-line, using participants recruited from Mechanical Turk with 10% chosen to be paid according to their choices in the experiment. However, all Experiment 3 participants are assigned the ability-intensive task, since this is the task for which inequality affected competitiveness. Experiment 3 is otherwise identical to Experiment 2 with one more exception. The choice of incentives for task 2 is not between piece-rate pay and tournament incentives but rather between piece-rate pay and lump-sum pay.³¹ Participants who choose the piece-rate scheme are paid according to performance – \$1.50 (\$0.50) for each correct (incorrect) answer. Participants choosing the lump-sum scheme must still complete task 2 but their pay does not depend on their performance. I elicit the minimum lump-sum payment each participant would accept in order to forgo the piece-rate pay option using an incentive compatible Becker–DeGroot–Marschak mechanism (Becker *et al.*, 1964). I interpret this minimum lump-sum as a participant’s valuation for his or her task 2 performance which, in turn, can be viewed as a measure of a participant’s absolute task 2 performance beliefs. Because Experiment 3 is otherwise so similar to Experiment 2, full details are omitted for brevity but are available upon request.

In total 100 individuals took part in Experiment 3. Using these data, I investigate the relationship between inequality salience and absolute performance beliefs. The results are presented in Table 9. There are three points to notice. First of all, the relationship between salient inequality and relative performance beliefs is qualitatively similar to the relationship documented in the previous two experiments, with or without demographic controls (columns 1 and 2). Second, as before, neither pay level nor salient inequality affects actual performance on task 1 (columns 3 and 4). Third and most important, there is no evidence that salient inequality affects absolute performance beliefs in a way that would generate the competitiveness patterns documented in Experiment 2. Either pooling all observations (columns 5 and 6) or restricting attention to male participants only (columns 7 and 8) the estimated interaction between salient inequality and high pay has the wrong sign and is never significant. These data suggest that the relationship between inequality and competitiveness documented in Experiment 2 is unlikely to be attributable to absolute performance beliefs.

3.6. *Earnings Consequences of Competitiveness Patterns*

Having seen that inequality bolsters the competitiveness of male participants on ability-intensive tasks, another question naturally arises: what are the earnings consequences? It could very well be, for instance, that male participants started out being more inclined to compete than is optimal, perhaps being overconfident, so that making

³¹ I thank an anonymous referee for suggesting this modification.

Table 9
Beliefs, Performance and Task 2 Valuations from Experiment 3

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	3-category belief measure		Own task 1 performance		Valuation for own task 2 performance		Valuation for task 2, males only	
High pay	-0.32 (0.38)	-0.58 (0.39)	1.07 (0.81)	0.89 (0.78)	1.86 (1.26)	1.64 (1.28)	2.76* (1.40)	2.31 (1.55)
Treatment	-1.02** (0.47)	-1.19*** (0.46)	0.97 (0.82)	1.00 (0.79)	-0.08 (1.34)	-0.23 (1.49)	1.20 (1.48)	0.33 (1.89)
High pay × treatment	1.08* (0.57)	1.28** (0.57)	-1.00 (1.03)	-1.10 (0.97)	-0.37 (1.76)	-0.28 (1.85)	-2.12 (2.02)	-1.95 (2.28)
Own task 1 performance	0.14** (0.06)	0.17** (0.07)			0.59*** (0.19)	0.64*** (0.22)	0.69*** (0.24)	0.32 (0.35)
Constant			6.22*** (0.77)	-5.58 (4.23)	5.71*** (1.73)	8.24 (6.21)	3.86* (2.02)	-8.05 (6.14)
Demographic controls	N	Y	N	Y	N	Y	N	Y
Observations	100	99	100	99	100	99	53	53
[Pseudo] R-squared	0.12	0.17	0.05	0.24	0.22	0.22	0.29	0.36

Notes. Columns 1 and 2 present ordered probit estimates using as the dependent variable the 3-category relative performance belief measure described in the text, taking values 1 = ‘believed performance was below median’ to 3 = ‘believed performance was above median’. The dependent variable in columns 3 and 4 is how many correct answers the participant provided on Task 1 (max. possible = 10); The dependent variable in columns 5–8 is the minimum certain payment the participant would accept in order to avoid being paid based on their performance on Task 2. Robust standard errors appear in parentheses; *** significant at 1%, ** significant at 5%, * significant at 10%. The primary explanatory variables are: ‘high pay’ = a dummy taking the value of one if the participants’ pay structure was \$2 for each correct guess and \$1 for each incorrect guess; ‘treatment’ = a dummy indicating whether participants took part in a session where they were informed about both potential pay scales; ‘own task 1 performance’ = the participant’s actual performance on Task 1 (max. possible = 10). This variable is unknown to participants, of course. Demographic controls are all self-reported and include: income category dummies based on participant’s last-year’s total income; age and age squared, which are constructed from the participant’s answer to the question ‘in what year were you born?’; an indicator for whether the participant has a college degree; and gender. Demographic controls have been omitted for readability. In addition, each estimate includes a full set of dummies for the US time zone in which the participant’s ip address was located – ‘Mountain Time Zone’, ‘Central Time Zone’ and ‘Eastern Time Zone’, the excluded category being ‘Pacific Time Zone’ – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives.

them even more competitive actually lowers their earnings.³² This is an important question since it speaks to whether salient inequality tends to directly exacerbate initial inequality or, rather, by lowering the earnings of the initially advantaged actually directly dampens inequality.

To address this question, I return to the Experiment 2 data and restrict attention to the 128 observations stemming from male participants engaged in the ability-intensive task. I restrict to this subset because it is here that inequality was shown to affect competitiveness. For each of these participants, I first compute performance on the second task. For individual i , let n_i ($10 - n_i$) denote i ’s number of correct (incorrect) answers. For each participant who chose piece-rate pay, I construct earnings to be: $Earnings_i = 1.50 \times n_i + 0.50 \times (10 - n_i)$.

³² I thank an anonymous referee for suggesting this exercise.

For participants who chose tournament incentives, instead of their actual earnings I compute a measure of their empirical (expected) earnings. For each such individual i , I compute the proportion of all other participants assigned the ability-intensive task, $j \neq i$, for which $n_i \geq n_j$. Call this proportion Win_i , which is the proportion of tournaments individual i would actually ‘win’. Empirical expected earnings for individual i are therefore $Earnings_i = Win \times 30 + (1 - Win) \times 0 = Win \times 30$.

For ease of exposition, I split the sample by pay level and regress earnings on a dummy for whether the participant was exposed to salient inequality. I control for performance on task 1 as a noisy measure of the ‘information’ participants had at their disposal when choosing their incentive scheme, for example, whether they found the ability-intensive task to be easy or difficult. In Table 10 I report four separate regressions. The first two columns restrict attention to participants assigned high pay on task 1. Column 2 includes demographic controls. Mirroring the results on competitiveness, salient inequality increases the earnings of previously advantaged participants substantially. The coefficient on treatment represents an increase of 12.8% of the unrestricted sample mean of earnings (mean = 14.05; SD = 8.04). Controlling for demographics the effect of salient inequality on earnings is, in addition, highly statistically significant. The coefficient on treatment in column 2 implies inequality increases the earnings of the previously advantaged by 15.7% of the unrestricted sample mean of earnings. For the previously disadvantaged, earnings patterns again

Table 10
Earnings from Task 2, Experiment 2

	(1)	(2)	(3)	(4)
	High pay		Low pay	
Treatment	1.80 (0.95)	2.20*** (0.46)	-0.36 (0.35)	0.14 (1.00)
Own task 1 performance	1.51*** (0.23)	1.39*** (0.24)	2.35*** (0.36)	2.38*** (0.45)
Constant	-1.13 (2.26)	-2.46 (6.13)	-2.80 (2.29)	-5.34 (10.28)
Demographic controls	N	Y	N	Y
Observations	64	64	64	64
[Pseudo] R-squared	0.38	0.41	0.54	0.54

Notes. The dependent variable in all columns is Earnings, either actual (piece-rare pay) or imputed (tournament incentives). All estimates are OLS. Robust standard errors clustered by session appear in parentheses; *** significant at 1%, ** significant at 5%, * significant at 10%. The primary explanatory variables are: ‘high pay’ = a dummy taking the value of one if the participants’ pay structure was \$2 for each correct guess and \$1 for each incorrect guess; ‘treatment’ = a dummy indicating whether participants took part in a session where they were informed about both potential pay scales; ‘own task 1 performance’ = the participant’s actual performance on task 1 (max. possible = 10). This variable is unknown to participants, of course. Demographic controls are all self-reported and include: income category dummies based on participant’s last-year’s total income; age and age squared, which are constructed from the participant’s answer to the question ‘in what year were you born?’; an indicator for whether the participant has a college degree; and gender. Demographic controls have been omitted for readability. In addition, each estimate includes a full set of dummies for the US time zone in which the participant’s ip address was located – ‘Mountain Time Zone’, ‘Central Time Zone’ and ‘Eastern Time Zone’, the excluded category being ‘Pacific Time Zone’ – to control for familiarity with the California state lottery used to select participants for payment which could affect performance incentives.

mirror competitiveness patterns: with or without controlling for demographics the estimated effect of salient inequality on earnings is both small in magnitude and statistically non-significant.

On balance, the data again support the story that initial inequality may perpetuate itself by exacerbating initial earnings differences and that, interestingly, even male participants in my experiment were not naturally overly competitive on the ability-intensive task. An important caveat is that selection into competition in Experiment 2 does not determine the pool with whom one competes. While I made this design choice to isolate competitiveness *per se*, it has the drawback of missing an important feature of real-world competition. The effect of having the pool of competitors determined endogenously is an interesting open question which I leave for future research.

3.7. Discussion

The results from Experiments 2 and 3 suggest that unequal treatment affects beliefs about one's relative ability in a way that justifies the inequality: disadvantage undermines the relative performance beliefs of the disadvantaged and bolsters the beliefs of the advantaged on a task where performance transparently depends primarily on (cognitive) ability. Moreover, the effect of inequality on ability beliefs is apparently neither solely dependent on, nor drowned out by, explicit feedback. With (Experiment 1) or without (Experiments 2 and 3) explicit feedback I find a qualitatively similar relationship between inequality and beliefs about relative performance on ability-intensive tasks. On tasks where performance depends primarily on effort, inequality has an unambiguously positive effect on performance beliefs. While intriguing in its own right, this positive effect of inequality on relative effort beliefs is further evidence that the patterns in beliefs found in Experiment 1 and on the ability-intensive task in Experiments 2 and 3 are not primarily about expected demoralisation.

Experiment 2 also sheds light on the consequences of prior inequality on subsequent ability belief dependent choices. Male participants previously advantaged by inequality are subsequently significantly more likely to select into competitive environments – choosing tournament-based pay over piece-rate pay. Those males who select tournament incentives also reap substantially higher expected earnings. The previously disadvantaged receive no such boost. The result is a pattern which may perpetuate inequality: the advantaged become more willing to compete *qua* advantage.

Finally, it should be stressed that I obtain both of these results using a design which pays only one phase. Either the initial task, one stated belief or the subsequent task with the participant's chosen pay structure determines a participant's entire potential experimental earnings. This feature is widely held to ameliorate concerns about wealth or income effects weakening proper incentives across multiple decisions.

4. Concluding Remarks

In this study, I provide the first clean experimental evidence on an under-explored mechanism through which inequality may self-perpetuate. This mechanism is based on

a widely documented psychological phenomenon: just world beliefs (Lerner, 1965; Lerner and Miller, 1978; Bénabou and Tirole, 2006). The mechanism relies on two testable hypotheses: that salient inequality colours beliefs about one's own relative ability in a way that justifies the unequal treatment (Hypothesis 1); and that because of this, prior unequal treatment affects an individual's subsequent willingness to compete (Hypothesis 2). I then go on to present evidence from multiple experiments consistent with these hypotheses.³³ Data from both experiments support the notion that individuals tend to believe they deserve the inequality they experience. While desert/merit may conceptually be a function of both effort and ability – behaviours and traits – my results suggest that it is beliefs about ability in particular which respond differentially to inequality. On two different ability-intensive tasks salient inequality undermined the performance beliefs of the disadvantaged and bolstered the performance beliefs of the advantaged so that, for example, disadvantaged participants were about 20 percentage points more likely to believe they performed below the median than advantaged participants. The data also provide some support for the claim that inequality's effects on ability beliefs are strong enough to colour subsequent decisions. Prior inequality bolstered subsequent competitiveness among male participants by about 14 percentage points, a figure which represents 33% of the sample mean. Considered as a whole, the results from both experiments lend credence to a novel mechanism through which current inequality may become durable inequality.

Documenting this mechanism is of interest for several reasons. First and foremost it may represent a minimal mechanism yielding inequality persistence. In environments like modern meritocratic societies where ability-based competition features prominently at many steps along the path to highly remunerated careers, inequality experienced at crucial junctures may steer the disadvantaged away from these paths and lock in their relative disadvantage. Consequently, even without intentional discrimination, external financial constraints or different underlying preferences or traits, the mechanism documented here may generate durable inequality. A random uninformative shock disproportionately adversely affecting one group more than another may be sufficient to set off a process leading inequality to persist. Second, while I have singled out one type of decision that relies on relative ability beliefs – the decision to enter ability-based competition – there are many other types of decisions relying on these beliefs that may affect economic outcomes. Decisions about when and from whom to seek expert advice, as well as the monetary value placed on such advice or whether or not to heed received advice, plausibly depend on how relatively cognitively skilled an individual perceives himself or herself to be. Heeding bad advice can have sizable consequences as, for example, Bernard Madoff's investors can attest to. Finally, even as economists have recently begun to take seriously the ecology of preference formation – how social interactions and the social environment may shape individuals' preferences – the ecology of economically relevant beliefs formation has not yet garnered the same level of attention. The current inquiry can be seen as

³³ In previous versions of this article, what is referred to there as 'Experiment 2' can be thought of as a pilot for the Experiment 2 reported herein. I omit discussion of the pilot in this version of the article for brevity and clarity of exposition, merely noting that the results in the pilot are consistent with the results of what is called Experiment 2 here.

providing some of the first evidence suggesting that studying how the social environment shapes beliefs may enhance economists' understanding of otherwise-puzzling phenomena.

As a final note, one may wonder whether the patterns documented here in an admittedly artificial context extend to real-world decision making. While estimating the real-life economic consequences of inequality-induced self-discrimination is beyond the scope of this study, suggestive anecdotal evidence may stimulate future research on this topic. With this caveat in mind, consider one of the most economically consequential decisions many individuals ever make: which college to attend. Economist Caroline Hoxby describes the spectacular failure of a well-publicised multi-year campaign conducted by a consortium of elite colleges (e.g. Harvard) to attract talented disadvantaged students:³⁴

After Harvard offered what was, in essence, a free college education to students whose families earned under \$40,000 a year, Hoxby says, 'the number of students whose families had income below that threshold changed by only about 15 students, and the class at Harvard is about 1,650 freshmen'. ... some college administrators (sic) had confided to her that they had reluctantly come to the conclusion that the pool of low-income students with top academic credentials was just limited, and there was not much they could do to change that.

Hoxby and Avery (2012) show, to the contrary, that the pool of high-ability disadvantaged students is large. These students simply do not apply to elite colleges. However, the colleges to which high-ability disadvantaged students apply actually end up costing these students more to attend than elite colleges as they typically offer less financial assistance than elite colleges. Judging solely from their decisions, disadvantaged students apparently place a lower value on attending elite, highly competitive, colleges like Harvard than on attending less prestigious schools. They expend effort to apply to, and spend money to attend, the latter but not the former. This pattern is on its face inconsistent with explanations for persistent inequality relying on labour market discrimination or financial constraints. The study's authors offer an informational story – disadvantaged students may simply not know about the existence of elite colleges or the availability of grants. Information seems unlikely to be the whole story, however, as these highly able disadvantaged students are by definition the cognitive equals of the brightest students in the world with strong incentives to learn about financial aid.³⁵ At the same time, the college choices of the disadvantaged are consistent with the mechanism I document in this study: if elite colleges are perceived as having brighter students and completing college requires ability-based competition – we grade on a scale, after all – then since the pecuniary rewards to attending college depend disproportionately on completing college (sheepskin effect), disadvantage

³⁴ The transcript of the interview can be found here: <http://www.npr.org/2013/01/09/168889785/elite-colleges-struggle-to-recruit-smart-low-income-kids>.

³⁵ Even if lack of information is a proximate cause, this does not resolve the underlying puzzle about why these otherwise brilliant students choose to remain uninformed.

itself may diminish prospective students' assessments of the value of attending elite colleges.

Einaudi Institute for Economics and Finance

Submitted: 17 September 2013

Accepted: 09 May 2014

Additional Supporting Information may be found in the online version of this article:

Appendix A. Instructions.

Appendix B. A Theoretical Predictions about the Relationship between Ability, Effort, Performance and Performance Beliefs.

Data S1.

References

- Afridi, F., Li, S. and Ren, Y. (2012). 'Social identity and inequality: the impact of China's hukou system', Working Paper, Institute for the Study of Labour (IZA).
- Akerlof, G. A. (2002). 'Behavioral macroeconomics and macroeconomic behavior', *American Economic Review*, vol. 92(3), pp. 411–33.
- Akerlof, G. A. and Dickens, W. T. (1982). 'The economic consequences of cognitive dissonance', *American Economic Review*, vol. 72(3), pp. 307–19.
- Allport, G. W. (1954). *The Nature of Prejudice*, Cambridge: Addison-Wesley.
- Anderson, L. and Holt, C. (1997). 'Information cascades in the laboratory', *American Economic Review*, vol. 87(5), pp. 847–62.
- Arrow, K. (1971). *Some Models of Racial Discrimination in the Labor Market*, Santa Monica, CA: Rand Corporation.
- Banerjee, A. V. and Newman, A. F. (1991). 'Risk-bearing and the theory of income distribution', *Review of Economic Studies*, vol. 58(2), pp. 211–35.
- Banerjee, A. V. and Newman, A. F. (1993). 'Occupational choice and the process of development', *Journal of Political Economy*, vol. 101(2), pp. 274–98.
- Becker, G. S. (1971). *The Economics of Discrimination*, 2nd edn., Chicago, IL: University of Chicago Press.
- Becker, G. M., DeGroot, M. H. and Marschak, J. (1964). 'Measuring utility by a single response sequential method', *Behavioral Science*, vol. 9(3), pp. 226–32.
- Bénabou, R. and Tirole, J. (2002). 'Self-confidence and personal motivation', *Quarterly Journal of Economics*, vol. 117(3), pp. 871–915.
- Bénabou, R. and Tirole, J. (2006). 'Belief in a just world and redistributive politics', *Quarterly Journal of Economics*, vol. 121(2), pp. 699–746.
- Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992). 'A theory of fads, fashion, custom, and cultural change as informational cascades', *Journal of Political Economy*, vol. 100(5), pp. 992–1026.
- Bogliacino, F. and Ortoleva, P. (2011). 'Aspirations and growth: a model where the income of others acts as a reference point', Working Paper, California Institute of Technology.
- Bowles, S., Loury, G. C. and Sethi, R. (2014). 'Group inequality', *Journal of the European Economic Association*, vol. 12(1), pp. 129–52.
- Burks, S. V., Carpenter, J. P., Goette, L. and Rustichini, A. (2013). 'Overconfidence and social signalling', *Review of Economic Studies*, vol. 80(3), pp. 949–83.
- Caprariello, P., Cuddy, A. and Fiske, S. (2009). 'Social structure shapes cultural stereotypes and emotions: a causal test of the stereotype content model', *Group Processes and Intergroup Relations*, vol. 12(2), pp. 147–55.
- Cuddy, A. J. C., Fiske, S. T., Kwan, V. S. Y., Glick, P., Demoulin, S., Leyens, J.-P., Bond, M. H., Croizet, J. C., Ellemers, N., Sleeboos, E., Htun, T. T., Kim, H.-J., Maio, G., Perry, J., Petkova, K., Todorov, V., Rodriguez Bailón, R., Morales, E., Moya, M., Palacios, M., Smith, V., Perez, R., Vala, J. and Ziegler, R. (2009). 'Stereotype content model across cultures: towards universal similarities and some differences', *British Journal of Social Psychology*, vol. 48(1), pp. 1–33.
- Di Tella, R., Galiani, S. and Schargrofsky, E. (2007). 'The formation of beliefs: evidence from the allocation of land titles to squatters', *Quarterly Journal of Economics*, vol. 122(1), pp. 209–41.

- Durlauf, S. N. (1996). 'A theory of persistent income inequality', *Journal of Economic Growth*, vol. 1(1), pp. 75–93.
- Eil, D. and Rao, J. M. (2011). 'The good news-bad news effect: asymmetric processing of objective information about yourself', *American Economic Journal: Microeconomics*, vol. 3(2), pp. 114–38.
- Fischbacher, U. (2007). 'z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, vol. 10(2), pp. 171–8.
- Fiske, S. T., Cuddy, A. J. C. and Glick, P. (2007). 'Universal dimensions of social cognition warmth and competence', *Trends in Cognitive Sciences*, vol. 11(2), pp. 77–83.
- Fong, C. (2001). 'Social preferences, self-interest, and the demand for redistribution', *Journal of Public Economics*, vol. 82(2), pp. 225–46.
- Furnham, A. (2003). 'Belief in a just world: research progress over the past decade', *Personality and Individual Differences*, vol. 34(5), pp. 795–817.
- Galor, O. and Zeira, J. (1993). 'Income distribution and macroeconomics', *Review of Economic Studies*, vol. 60(1), pp. 35–52.
- Gaviria, A. (2002). 'Intergenerational mobility, sibling inequality and borrowing constraints', *Economics of Education Review*, vol. 21(4), pp. 331–40.
- Gill, D. and Prowse, V. (2012). 'A structural analysis of disappointment aversion in a real effort competition', *American Economic Review*, vol. 102(1), pp. 469–503.
- Gill, D. and Prowse, V. (2013). 'Gender differences and dynamics in competition: the role of luck', Working Paper, Nuffield Center for Experimental Social Sciences.
- Gill, D., Prowse, V. and Vlassopoulos, M. (2013). 'Cheating in the workplace: an experimental study of the impact of bonuses and productivity', *Journal of Economic Behavior & Organization*, vol. 96(1), pp. 120–34.
- Gneezy, U., Leonard, K. L. and List, J. A. (2009). 'Gender differences in competition: evidence from a matrilineal and a patriarchal society', *Econometrica*, vol. 77(5), pp. 1637–64.
- Gneezy, U., Niederle, M. and Rustichini, A. (2003). 'Performance in competitive environments: gender differences', *Quarterly Journal of Economics*, vol. 118(3), pp. 1049–74.
- Grossman, Z. and Owens, D. (2012). 'An unlucky feeling: overconfidence and noisy feedback', *Journal of Economic Behavior & Organization*, vol. 84(2), pp. 510–24.
- Hafer, C. L. and Bègue, L. (2005). 'Experimental research on just-world theory: problems, developments, and future challenges', *Psychological Bulletin*, vol. 131(1), pp. 128–67.
- Hafer, C. L. and Olson, J. M. (1998). 'Individual differences in the belief in a just world and responses to personal misfortune', in (L. Montada and M. J. Lerner, eds.), *Responses to Victimizations and Belief in a Just World*, pp. 65–86, New York: Plenum Press.
- Hargreaves Heap, S. P., Tan, J. H. W. and Zizzo, D. J. (2013). 'Trust, inequality and the market', *Theory and Decision*, vol. 74(3), pp. 311–33.
- Harrison, G. and Rütstrom, E. E. (2008). 'Risk aversion in the laboratory', in (J. Cox and G. Harrison eds.), *Risk Aversion in Experiments*, vol. 12, pp. 41–196, Bingley: Emerald Group Publishing.
- Hoff, K. and Pandey, P. (2006). 'Discrimination, social identity, and durable inequalities', *American Economic Review*, vol. 96(2), pp. 206–11.
- Hoxby, C. and Avery, C. (2012). 'The missing "one-offs": the hidden supply of high-achieving, low income students', Working Paper, NBER.
- Janoff-Bulman, R. (1979). 'Characterological versus behavioral self-blame: inquiries into depression and rape', *Journal of Personality and Social Psychology*, vol. 37(10), pp. 1798–809.
- Jost, J. T., Pelham, B. W., Sheldon, O. and Sullivan, B. N. (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*, vol. 33(1), pp. 13–36.
- Kahneman, D. and Tversky, A. (1979). 'Prospect theory: an analysis of decision under risk', *Econometrica*, vol. 47(2), pp. 263–91.
- Kelley, H. H. (1967). 'Attribution theory in social psychology', *Nebraska Symposium on Motivation*, vol. 15(1), pp. 192–238.
- Leibbrandt, A., Gneezy, U. and List, J. A. (2013). 'Rise and fall of competitiveness in individualistic and collectivistic societies', *Proceedings of the National Academy of Sciences*, vol. 110(23), pp. 9305–8.
- Lerner, M. J. (1965). 'Evaluation of performance as a function of performer's reward and attractiveness', *Journal of Personality and Social Psychology*, vol. 1(4), pp. 355–60.
- Lerner, M. J. and Miller, D. T. (1978). 'Just world research and the attribution process: looking back and ahead', *Psychological Bulletin*, vol. 85(5), pp. 1030–51.
- Loury, G. (1981). 'Intergenerational transfers and the distribution of earnings', *Econometrica*, vol. 49(4), pp. 843–67.
- Mani, A., Mullainathan, S., Shafir, E. and Zhao, J. (2013). 'Poverty impedes cognitive function', *Science*, vol. 341(6149), pp. 976–80.
- Möbius, M. M., Niederle, M., Niehaus P. and Rosenblat, T. S. (2011). 'Managing self-confidence: theory and experimental evidence', Working Paper, NBER.

- Mookherjee, D., Napel, S. and Ray, D. (2010). 'Aspirations, segregation and occupational choice', *Journal of the European Economic Association*, vol. 8(1), pp. 139–168.
- Moore, D. A. and Healy, P. J. (2008). 'The trouble with overconfidence', *Psychological Review*, vol. 115(2), pp. 502–17.
- Niederle, M. and Vesterlund, L. (2007). 'Do women shy away from competition? Do men compete too much?', *Quarterly Journal of Economics*, vol. 122(3), pp. 1067–101.
- Piketty, T. (2000). 'Theories of persistent inequality and intergenerational mobility', in (A. B. Atkinson and F. Bourguignon eds.), *Handbook of Income Distribution*, vol. 1, edition 1, pp. 429–76, Amsterdam: Elsevier.
- Ray, D. (2006). 'Aspirations, poverty and economic change', in (A. V. Banerjee, R. Bénabou and D. Mookherjee, eds.), *Understanding Poverty*, pp. 409–21, Oxford: Oxford University Press.
- Schlag, K. and van der Weele, J. (2013). 'Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality', *Theoretical Economics Letters*, vol. 3(1), pp. 38–42.
- Steele, C. and Aronson, J. (1995). 'Stereotype threat and the intellectual test performance of the African-Americans', *Journal of Personality and Social Psychology*, vol. 69(5), pp. 797–811.
- Torvik, R. (1993). 'Talent, growth and income distribution', *Scandinavian Journal of Economics*, vol. 95(4), pp. 581–96.