

Our Bark is Bigger than Our Bite: Stated Preferences over the Distribution of Wages and the Effect on Productivity

Eric Cardella[†]
Texas Tech University

Alex Roomets[‡]
Franklin & Marshall College

May 1, 2018

Abstract:

Wage dispersion and the associated level of wage inequality, within firms and across society more generally, has received substantial attention across both academic research and popular press. The objective of this study is to investigate individuals' preferences over different distributions of wages across groups of workers with heterogeneous ability, and to identify the possible productivity effects associated with different distributions of wages. We develop an experimental design where we elicit participant workers' preferences over a set of different wage distributions, and then have the groups of workers complete a real-effort work task under the different possible wage distributions. Importantly, the set of possible wage distributions vary in the degree of inequality, overall efficiency, and implied level of income inequality. Overall, we find that workers reveal strong preferences over how wages should be distributed, and these preferences are non-stationary and depend on the worker's relative rank within the group. However, the actual distribution of wages has little impact on the *quantity* of output produced. On the other hand, we do observe some impact of the wage distribution on output *quality*. The results from this study can have important implications for firms in making the determination of how to distribute wages among employees to generate the most productivity, as well as the most employee satisfaction. More generally, our results could have important implications for policies aimed at limiting wage dispersion within firms, or policies aimed at providing more wage transparency within firms.

Keywords: Wage Dispersion; Wage Equality, Wage Preferences; Productivity; Work Quality; Redistribution

JEL Codes: C91, D31, D61, D63, D71, J30, M52

[†]Rawls College of Business, Texas Tech University, Lubbock, TX 79409; Telephone: (806) 834-7482; Email: eric.cardella@ttu.edu

[‡] Economics Department, Franklin & Marshall College, Lancaster, PA 17604; Telephone (717) 358-4869; Email: alex.roomets@fandm.edu

1 Introduction

The distribution of wages within firms and the associated degree of dispersion and inequality is a hot button issues that continues to garner substantial attention in both academic research and popular press. In particular, there is a rich literature promoting the idea that the distribution of wages can influence the productivity of workers via wage comparisons and corresponding inferences regarding fairness (e.g., Frank, 1984; Lazear, 1989; 1991; Akerlof & Yellen, 1990; Bewley, 1999 for seminal work on this topic). At the same time, there has been a lot of discussion, especially recently, on policies aimed at compressing wages and reducing inequality (e.g., raising minimum wage or regulating executive compensation).¹ There has also been substantial discussion and even some reform (e.g., Dodd-Frank regulation) on pay secrecy, and the extent to which employee pay information within firms should be made more transparent, which would subsequently reveal information about the distribution of wages. The motivation of this study is twofold. First, we investigate worker preferences regarding how wages are distributed across a group of workers with heterogeneous ability. Second, we examine how different distributions of wages impact worker productivity.

We conduct an experimental study using a real-effort, data entry task. As part of the design, participant workers are arranged into groups of three. Prior to completing the task, we present participant workers with a menu of four different wage distributions—specifying a rank-based, piece-rate wage for each worker in the group—that vary in terms of the degree of: wage inequality, efficiency, and overall level of implied income inequality. After we elicit preferences over the various wage distributions, participant workers learn their rank within the group; we vary whether relative rank within the group is randomly assigned or merit based (i.e., determined by productive ability). We then have the groups of workers complete the real-effort task under one of the realized wage distributions, and measure group-level and worker-level productivity.

Our experimental study enables us to answer several important questions: (i) what type of wage distribution do people prefer be assigned to a group of workers with heterogeneous abilities; (ii) are wage distribution preferences stationary, or do preferences differ based on the individual's relative rank within the affected group, whether they are behind the “veil of ignorance”, or whether

¹ For example, in 2009 the Obama administration mandated a CEO pay cap of \$500,000 for firms receiving federal aid. In 2013 Switzerland put forth a referendum to limit the CEO pay to a maximum of 12 times the pay of the lowest earning employee, although the referendum was not passed into legislation. Similarly, in 2012 France implemented a policy that limited executive compensation at state-owned companies (and possibly later to private companies).

the individual acts the role of a planner and does not belong to the impacted group of workers, (iii) how do different wage distributions, and the corresponding characteristics, impact aggregate productivity of the group, (iv) are there possible heterogeneous productivity effects associated with the different wage distributions based on relative rank within the group, and (v) does the manner by which workers are ranked within the group, either randomly or merit-based, impact wage distribution preferences and the productivity response to the different wage distributions.

In terms of wage distribution preferences, our results suggest there is significant heterogeneity. Moreover, we also find that wage distribution preferences are not stationary and vary significantly based on the perspective from which they are stating their preference. Specifically, when participants are choosing from a planner's perspective, the majority of people prefer a distribution with equal wages (roughly 50%) or very unequal, but collectively efficient, wages (roughly 30%). Similar results are observed when participants state their preference from a "veil of ignorance" perspective (Harsanyi, 1953; Rawls, 1971), where they belong to the group of workers but don't know, ex-ante, their relative position in the group. However, when workers are asked their preferred distribution assuming they know their position within the group, we observe significant changes in preferences based on position. Namely, when choosing from the perspective of being the top ranked worker in the group, participants overwhelmingly prefer wage inequality (roughly 70%), which gives the highest payoff to the top ranked worker. Conversely, from the perspective of being the lowest ranked worker, participants overwhelmingly prefer a wage distribution that gives them the highest wage and provides more final income equality (roughly 50%). Thus, workers appear to have a preference for a distribution of equal wages, so long as they are not negatively impacted by it, or could not benefit from a more redistributive distribution.

Existing experimental literature on "social welfare preferences" indicates that most people generally have a preference to (re)distribute unequal *income* distributions in the direction of greater income equality (e.g., Durante et al., 2014; Lefgren et al., 2016; Gee et al., 2017; Deffains et al., 2016 for a review), even if there is some efficiency loss in doing so (Mitchel et al., 1993; Scott et al., 2001; Michelbach et al., 2003; Beckman et al., 2004; Krawczyk, 2010; Durante et al., 2014; Hong et al., 2015; Beckman et al., 2016). Importantly, these studies consider the distributions of income and often the redistribution choices come ex-post (after the incomes have been earned). In our study, we look specifically at ex-ante preferences for the distribution of wages and not final incomes. That said, there is a clear connection between our study and the prior literature on income

redistribution; in particular, the distribution of wages ultimately plays a large role in shaping the distribution of final incomes, which then potentially necessitates redistribution. We document a similar pattern regarding wage distribution preferences; specifically, many people (when acting as a planner or from behind the veil of ignorance) prefer the wage equality distribution over the wage inequality distribution, even though it is less efficient. Importantly, though, because of substantial heterogeneity in worker ability, the wage equality distribution generates inequality in the distribution of final incomes. Thus, our results suggest that people seem to express a preference for *ex-ante* equality in the distribution of wages. This result accords with previous studies suggesting that people are less inclined to support costly income redistribution when income (and the corresponding inequality) is largely determined by effort rather than luck (e.g., Piketty, 1995; Fong, 2001; Alesina & Angeletos, 2005; Alesina & La Ferrara, 2005; Benabou & Tirole, 2006; Krawczyk, 2010; Cojocaru, 2014; Lefgren et al., 2016; Deffains et al., 2016), which implies *ex-ante* fairness and equal opportunities. By considering upstream preference for the distribution of wages, our study can help us better understand what motivates social welfare preferences for income redistribution, and our results suggest that *ex-ante* wage equality might play an important role in the degree of desired, *ex-post*, redistribution.

In terms of how the wage distribution impacts productivity, we document negligible differences in output levels across the four different wage distributions when relative ranking within the group is merit based. More specifically, we find no significant differences in either group-level or worker-level output (by relative ranking within the group) across wage distributions. When rank within the group is determined randomly, we do observe a little more variation in output across the different wage distributions, with the lowest group-level productivity being observed when wages are equal across workers. Thus, there may be some small productivity effects resulting from differences in the distribution of wages when workers' relative position are randomly determined. However, we do have some small sample sizes in some distributions, as a byproduct of the heterogeneity in the preference data that impacts the realized wage distributions, and the significant effects are muted after controlling for ability. As a result, our findings largely suggest that the wage distribution has, at most, a small effect on output. Moreover, in settings where workers' positions are merit based, which strikes us as the more representative case, the wage distribution has essentially no effect on output.

Recently, there is a growing body of experimental literature exploring how wage comparisons and resulting perceptions of wage fairness impact effort provision; typically these papers use some type of multi-lateral, gift-exchange paradigm with a fixed wage, and the results are somewhat mixed. In terms of chosen-effort studies, Abeler et al. (2010), Clark et al. (2010), Gächter and Thoni (2010), Angelova et al. (2012), Nosenzo (2013), Gross et al. (2015), and Charness et al. (2016) document evidence that horizontal wage comparisons can impact productivity, while Charness and Kuhn (2007), Bartling et al. (2011), and Bolton and Werner (2016) find little evidence of a wage comparison effect on effort provision.² The results from real-effort studies are also mixed with Cohn et al. (2014) finding evidence of relative wage comparisons impacting productivity, whereas Hennig-Schmidt et al. (2010), Greiner et al. (2011), and Butler (2016) finding no significant wage comparison effect on productivity. While the main focus of our study is not to examine how social wage comparisons impact productivity, per se, it is the case that by considering four different wage distributions (which are publicly known to workers), we are able to identify possible inherent wage comparison effects; moreover, because we implement a piece-rate wage, we are able to examine this effect under a performance-pay scheme, which differs from the prior literature (cf. Greiner et al., 2011 and Butler, 2016 who also have conditions involving piece-rates). In line with the findings from Greiner et al. (2011) and Butler (2016), we document very little wage comparison effect on amount of output produced. In fact, even in the perverse case where the highest ranked worker (based on a proxy measure for productivity) is paid the lowest piece rate, we see no impact on the productivity of the highest ranked worker, which is contrary to the finding in Gross et al. (2015) under a fixed wage. Importantly, the result from our study suggest that performance-pay may act as a strong enough incentive to motivate effort and mitigate any possible negative wage comparison effects on productivity.

By using a real-effort work task, not only are we able to examine how the distribution of wages impacts the amount of output produced, but we are also able to examine possible impacts on output *quality*. This is important given that prior literature has documented possible tradeoffs between quality and quantity (e.g., Paarsch & Shearer, 1999; Bellemare et al., 2010; Ederer & Manso, 2013; Rubin et al., 2017). Our results suggest that the wage distribution might impact productivity more subtly through the quality of output produced. Most notably, when high ability workers are paid

² Relatedly, Gill and Prowse (2012) document experimental evidence that effort provision is impacted by comparisons of peer effort provision.

the least, they produce significantly lower quality. In a sense, this is a subtler manifestation of the effect found by Gross et al. (2015) where high ability workers reduce effort if they are paid less than the low ability worker. We also find some evidence cutting the other way; low ability workers produce higher quality when they are given a higher wage. Lastly, we find some evidence implying that if worker ranks/positions within a group are determined randomly, then unequal pay schemes might induce lower quality output by those workers receiving the unfair wages.

Topics surrounding the degree of income inequality continue to remain in focus of important scholarly debate (see Piketty, 2014; Piketty & Saez, 2003; 2014; Piketty & Zucman, 2014). Recent large-scale survey studies by Alesina & La Ferrara (2005), Norton & Ariely (2011), Cruces et al. (2013), and Kuziemko et al. (2015) show that people seem to have a preference for less income inequality, relative to both the existing level, as well as their perception of the current level. Given that wages, possibly in combination with the level of productivity, play a large role in shaping income distributions, better understanding preferences over wage distributions and their possible impact on productivity strikes us as an important step in the process of better understanding views on income inequality. Moreover, in terms of actually achieving less income inequality, there are at least two broad approaches: (i) somewhat indirectly through ex-post redistributive policies like taxes and subsidies, or (ii) more directly through ex-ante policies that alter the distribution of wages and promote greater equality in wages. Kiatpongsan & Norton (2014) provide survey evidence that people want more equal pay within firms (relative to the actual severe inequality that currently exists). Our results from the preference elicitation echo this finding. Although, our results also suggest that wage distribution preferences may be more nuanced, in the sense that a desire for more wage equality within a firm might hinge crucially on an individual's position within the firm (or society) and the degree and direction by which they would be impacted by more wage equality.

We think this study makes several contributions to the existing literature. First, by considering four different wage distributions—that differ in terms of their degree of wage equality, efficiency, and implied income equality—and groups of three workers, we can gain more robust inferences into the effect of how relative wages and social wage comparisons impact productivity beyond just a two-worker, high/low wage setting. Second, we elicit preference data on the type of wage distribution people most prefer, which, when combined with productivity data, can be very informative for wage reform policies aimed at combating income inequality, both at the firm-level and aggregately at the industry-level. This is particularly true of reforms intended to combat

income inequality and/or promote wage transparency. Third, we consider the productivity effects associated with relative wage comparisons under a piece-rate scheme; given the prevalence and increasing trend toward the inclusion of performance pay in the economy (Lemieux et al., 2009), it is important to understand the interaction between relative wage comparisons and performance pay. In particular, our results suggest that performance pay can possibly mitigate some of the negative productivity effects previously shown to arise from relatively unfair fixed-wages. Fourth, we are able to explore possible heterogeneous productivity effects arising from the distribution of wages based on the relative rank/position of the worker, as well as how relative rank/position is determined within the group (random assignment vs. merit based). Overall, the result from our study can help to deepen our understanding of preferences over the distribution of wages across heterogeneous groups of workers, and the corresponding impact of relative pay comparisons on productivity, both in terms of quantity and quality of output produced.

2 Experimental Design

We conducted an experimental study involving a real-effort work task. In the experiment, participants are first asked to state their preferences for how wages be distributed among a group of workers with heterogeneous ability. The set of possible wage distributions vary in their degree of inequality, efficiency, and implied level of final income inequality. Groups of workers then completed a compensated real-effort task under one of the possible wage distributions. All experimental sessions were conducted at the Rawls College of Business at Texas Tech University. All participants were recruited from a college-maintained, subject-pool database. In total, 27 experimental sessions were conducted with 480 total participants (46% were female; $M_{\text{age}} = 21$). We implemented a between-groups design where each participant partook in just one session of one of three experimental conditions. Each session lasted approximately 60 minutes and participants earned an average of \$22.4. The experiment was computerized, and the software was programmed in z-Tree (Fischbacher, 2007). A copy of the experimental instructions for all conditions and sample screen shots are provided in the supplemental Appendix.

2.1 Real-Effort Work Task

For the real-effort work component, participants completed a data entry task where they entered information from printed job applications into a digital database. Participants were provided with

a packet of printed job application forms that were populated with fictitious applicant information.³ Participants were required to type select information from each application into a computerized, digital template. The required fields included: the application number, date, the applicant's name, birthdate, phone number, email, zip code, education level, and some questions about materials that had been submitted with the original application (e.g., cover letter, resume, references, and recommendation letters). For the remainder of the paper, the output level of a participant will be in reference to the number of applications they entered into the digital database. Applications were not checked for accuracy, so total output includes entries with errors. Participants were paid for erroneous entries, which enables us to measure output quality as we discuss in more detail later.

We think this data entry task is well-suited for examining how the distribution of wages among workers can impact productivity for several reasons. First, there is really no analytical component to completing the task (e.g., anagram, word unscrambling, or multiplication tasks), which implies that productivity is an increasing function of effort. This is important because it will enable us to identify how the wage distribution impacts productivity specifically through the provision of effort. Second, because of the open-endedness of inputting hand-typed information and the format with which the information is entered (e.g. dates and phone numbers), there is both a quantity and quality dimension associated with productivity. More detail about how we measure quality is provided in Section 5.3. Being able to measure both quantity and quality provides a more robust analysis of how productivity is impacted by the wage distribution. Such a tradeoff between quantity and quality of output under performance-pay schemes is discussed theoretically by Stiglitz (1975), Lazear (1986) and Holmstrom & Milgrom (1991), and documented empirically by Paarsch & Shearer (1999), Bellemare et al. (2010), Ederer & Manso (2013), Johnson et al. (2015), and Rubin et al. (2017). Third, there is likely to be substantial variation in the skill or ability of participants. Importantly, this *natural* variation allows us to explore possible heterogeneous treatment effects regarding both wage distribution preferences and their effect on productivity.

³ In the instructions, participants were not told that the job applications were fictitious. As the same time, to ensure that no deception was used, participants were also not told that the applications were real. Rather, participants were truthfully informed that their task was to enter information from printed job application forms into a digital database. For stronger causal identification of potential treatment effects, we wanted all participants to have the same packet of printed job application forms, which is why we opted to use fictitious applications. However, the opacity in the instructions created an environment where it seems plausible that participants perceived the data entry task as regular, economically valuable work (Falk & Ichino, 2006), intended to add to the credibility of our findings.

2.2 Procedure and Sequencing

The experiment consisted of three parts, which are separately discussed below. The first part was a practice work period, where each participant had an opportunity to enter “practice” applications, as a way to familiarize themselves with the task. The second part was the voting stage, where participants stated their preferred wage allocation under some different scenarios. The third part was the paid work stage, where participants entered applications and were compensated based on the given wage distribution. At the onset, participants were provided instructions for the first part and informed that there would be multiple parts to the study, but were not provided instructions for later parts until after the first part had concluded. After completing the paid work period, participants filled out a questionnaire to elicit general demographic characteristics, and other attitudes regarding: politics, merit pay, income equality, wage inequality, and redistribution.

2.2.1 Part 1 – Practice Work Period

In Part 1, participants completed an 8-minute, non-incentivized practice work period. Participants were informed that this was a practice period, and they received a separate packet of job applications that were labeled as “Practice Applications”, which were formatted very similarly to the ones used in the paid work period. We informed participants that later in study their total compensation would depend on how many applications they are able to enter into the database in an allotted period of time, and that the practice period provides them with an opportunity to become familiar with the applications and using the digital template. The primary function of the practice work period was to provide us with a relative measure of ability of participants. We did not incentivize production in this practice period to ensure that there was no difference in accrued earning entering the paid work period, which could have confounded our results regarding the effect of different wage distributions on productivity.

2.2.2 Part 2 – Voting Stage over Possible Wage Distributions

In Part 2, we elicited information about each participant’s preferred wage distribution. Participants were first informed that they would be randomly paired with two other participants in the same session to form a 3-person group, and that each person in the group would receive a ranking within the group of either 1, 2, or 3. We implemented the following two conditions for determining assignment of worker ranks within the group:

Earned Rank (EARNED) Condition – Each participant in the 3-person group is ranked based on their level of output in the practice period. The person with the most output was assigned rank 1, the second most assigned rank 2, and the least output assigned rank 3.⁴

Random Rank (RANDOM) Condition – Each participant in the 3-person group is ranked randomly, independent from their level of output in the practice period.

The motivation for implementing these two conditions is to identify whether earning one's rank within the group based on merit affects wage distribution preferences and the resulting productivity response, compared to random assignment. Prior survey evidence suggest that people tend to be more favorable to redistribution when income is less merit based, while many laboratory experiments confirm this notion that people are more redistributive (across different domains) when income is determined more by chance/luck.⁵ Moreover, a recent field experiment by Breza et al. (2018) suggests that effort responses to wage inequality may be less pronounced when the pay inequality is merit based and determined by productivity. In essence, the EARNED condition is more meritocratic in the way rank is determined. As such, these two conditions allow us to tease out how merit possibly interacts with how people believe wages should be distributed, and how the productivity of workers reacts to different wage distribution when rank is merit based.

After being informed about how the ranking within the 3-person group of workers would be determined, participants were instructed that they would view a table that listed different possible wage allocations. All participants faced the same set of 4 wage distributions. Table 1 depicts the 4 different wage distributions, the corresponding piece-rate wage to each worker based on their rank, the average wage, and the label indicating the “type” of wage distribution each option was intended to model. In the study, however, the four different wage distributions were instead generically labeled A, B, C, D and the columns of the table were randomly oriented for each participant, to avoid any order effects, demand effects, anchoring effect, or focal point effects.

⁴ In the event of a tie, the rank was randomly determined among the set of participants that tied. Participants were informed of this tie breaking procedure in the instructions.

⁵For examples of survey evidence see: Alesina et al. (2001), Fong (2001), Corneo & Grüner (2002), Alesina & Angeletos (2005), Cojocaru (2014). For examples of experimental evidence see: Mitchell et al. (1993), Hoffman et al. (1994), Clark (1998), Ruffle (1998), Cherry et al. (2002) Michelbach et al. (2003), Oxoby & Spraggon (2008), Krawczyk (2010), Balafoutas et al. (2013), Durante et al. (2014), Lefgren et al. (2016). Relatedly, Gill and Prowse (2014) find that effort provision in a tournament competition is impacted by prior outcome when the outcome is determined by luck/chance. Although, there are some prior experimental papers that find less support for the claim that merit plays a role in redistributive preferences (e.g., Esarey et al., 2012; Ku & Salmon, 2013; Gee et al., 2017).

[Table 1 – Possible Wage Distributions]

When rank is based on productivity (EARNED condition), the four distinct wage distributions are meant to represent the following types of wage allocations: (i) a *Wage Equality* environment where the wage is the same for all participants, (ii) a *Wage Inequality* environment where wages are positively correlated with rank, (iii) an *Income Equality* environment where the wages are negatively correlated with rank such that the highest rank worker received the lowest wage while the lowest rank worker received the highest wage, thus potentially compressing the final incomes among the workers⁶, and (iv) a *Minimum Wage* environment where the wages are still positively correlated with rank, but the wage to the lowest rank worker is raised to a higher level than in the *Wage Inequality* distribution and the level of inequality is lower.

In addition to differences in the allocation of wages across the different ranked workers, the wage distributions also vary, as expected, in terms of their efficiency; namely, the average wage per worker differs across wage distributions. As depicted in Table 1, the *Wage Inequality* distribution is the most efficient with the average wage being \$1.2 per application. Importantly, in moving to the *Wage Equality* or *Income Equality* distributions, there is a loss in efficiency with the average wage dropping to \$1 per application. Finally, in the *Minimum Wage* distribution, there is an average wage of \$1.1, which falls between the efficiency of the *Income Equality* and *Wage Inequality* distributions. Of primary interest is better understanding preferences over different types of wage distributions, and the consequential effect of different types of wage distributions on productivity. As such we will take as given that there are efficiency tradeoffs associated with the wage distribution becoming more equal, and even going as far as to invert the wage scale to make incomes more equal. This idea is consistent with Okun's (1975) "leaky bucket" idea capturing the tradeoff between equality versus efficiency.⁷ It is beyond the scope of this paper to

⁶ Since final income depends on productivity, which varies substantially across workers, it is not the case that this *Income Equality* distribution will generate the exact same level of final income across all three group members for every group. So while we denote this distribution the *Income Equality* distribution, we do not mean this in the literal sense that every worker in the group will receive, ex-post, the same amount of final income. Rather, this is a distribution that, ex-ante, is intended to produce more implied final income equality via the negative correlation between relative ranks within the group and the wage. In the RANDOM condition when rank within the group is determined randomly, then the interpretation of *Income Equality* distribution is different; namely, this distribution can now be viewed as a more compressed wage distribution, relative to the *Wage Inequality* distribution.

⁷ It is common among experimental studies on redistribution to incorporate a tradeoff between equality and efficiency (e.g., Mitchell et al., 1993; Scott et al., 2001; Andreoni & Miller, 2002; Michelbach et al., 2003; Beckman et al., 2004; Traub et al., 2009; Krawczyk, 2010; Hong et al., 2015; Beckman et al., 2016). Indeed, when making redistributive choices, people seem to have a preference for both equality and efficiency, amongst other possible

delve into why these inefficiencies might arise; however, our modeling is based off the premise that, in practice, altering the wage profile would likely require some form of regulatory intervention (e.g., wage controls, subsidies, transfer payments, etc.), that would reduce total welfare and, thus, shrink the overall size of the pie that can be divided among workers.

For the elicitation of each participant's preferred wage distribution, we used the strategy method to elicit preferences under several different scenarios. In particular, participants don't know their actual rank within the group when they state their preferred wage distribution. Rather, we asked participants to choose their preferred wage distribution under the following five scenarios:

Rank 1– preferred distribution for *own* group assuming you are Rank 1

Rank 2– preferred distribution for *own* group assuming you are Rank 2

Rank 3– preferred distribution for *own* group assuming you are Rank 3

Veil – preferred distribution for *own* group assuming your rank is unknown

Planner– preferred distribution for *another* group within the session

We made it clear to participants that there was a chance that their selected wage distribution in each of the five scenarios might actually be implemented in the work stage (depending on some combination of random chance and their realized rank). As such, participants were encouraged to consider their decisions in each scenario carefully and report truthfully.

The motivation for asking each worker's preferred wage distributions under these five scenarios is to better characterize wage distribution preferences and to examine if, and to what extent, such preferences are stable. The Veil scenario is meant to represent the situation of choosing a preferred wage distribution behind the “veil of ignorance” (Harsanyi, 1953; Rawls, 1971), where participants know they will be a part of the group for which the wage distribution will be applied, but they don't know, ex-ante, where in the group they will be positioned. The Planner scenario is meant to capture each participant's preferences when playing the role of a manager or social planner. Prior studies have shown that income redistribution choices can differ based on whether people act as planners, from behind the veil, or known position/rank (e.g., Traub et al., 2009; Schilberg-Horish, 2010; Durante et al., 2014; Beckman et al., 2016). Hence, we anticipate wage distribution preference to exhibit a similar type of non-stationarity and differ based on the voting scenario.

factors (e.g., Mitchell et al., 1993; Andreoni & Miller, 2002; Konow, 2003; Michelbach et al., 2003; Engelmann & Strobel, 2004; Bolton & Ockenfels, 2006; Fehr et al., 2006; Traub et al., 2009).

To summarize, our voting stage enables us to identify each participant's preferred wage distribution (from the set of 4 different wage profiles) conditional on their rank within the group; thus, we allow for the possibility that each participant's preferences might actually be a function of their relative position within the group. We also identify each participant's unconditional preferred wage distribution, when their position within the group is, ex-ante, unknown to the participant. Lastly, we identify their wage distribution preference when they are deciding for another group of workers (that they are not a part of), i.e., when they are acting the role of a benevolent social planner. By considering a condition where rank is based on productivity and another where rank is determined randomly, we are also able to identify how wage preferences possibly interact with merit based ranking within the group of workers.

2.2.3 Part 3 – Paid Work Period

In Part 3 of the experiment, participants performed a paid work task for a 22-minute period. First, a wage distribution is selected for each group (in part by random chance and in part based on the various wage distribution choices made during Part 2).⁸ Prior to beginning the 22-minute paid work period, participants received information about their realized rank, the selected wage distribution for the group, and their corresponding piece-rate wage (determined based on their rank and the wage distribution). However, to ensure the cleanest test of how the wage distribution impacts productivity, participants were not informed about whether the selection of the wage distribution had been based on random chance or the result of the implementation of a participant's actual choice from Part 2.⁹ To ensure that the rank and wage information was salient, it was continually displayed on the same screen as the digital database template. During the work period, participants were free to work at their own pace and complete as many applications in their stack as they were able to or chose to. Moreover, participants were seated at individual computer carrels

⁸ The reason why we included the random component of the wage distribution selection was to ensure that we generated sufficient productivity data for each of the four possible wage distributions. In the instructions, we informed participants of how exactly how this selection process would be implemented. In particular, we informed them that a computer program would first randomly choose a number from 1 to 10. If the number is 1 through 4 (40% chance), then the wage allocation for the group will be randomly selected from the possible options presented in Part 2, with each possible option being equally likely to be chosen. Each participant in your group will then receive their designated wage, based on their actual rank in the group, from the chosen wage allocation. If, alternatively, the computer program chooses a number 5 through 10 (60% chance), then the wage allocation in Part 3 for your group will be determined according to one of the actual wage allocations chosen in Part 2 by one of the participants, selected at random.

⁹ For example, prior work by Charness (2004) and Charness & Levine (2007) show that attribution impacts how agents respond to wages; namely the effort level of agents is impacted, via reciprocal motivations, by whether their wage was chosen randomly or by an actual human principle. Therefore, to ensure that possible treatment effects are not contaminated by how the wage distribution was chosen, we do not reveal this information to participants.

and were unable perfectly observe the progress and output of other workers; as such, there is little scope for peer effects to influence productivity (e.g., Falk & Ichino, 2006; Mas & Moretti, 2009; and Herbst & Mas, 2015 for a review).

2.3 Baseline Condition to Measure of Productivity

In order to examine how the distribution of wages impacts productivity, it is important to have an accurate estimate of the distribution of *baseline* productivity of our sample of participant workers; namely, the distribution of output *given* a wage level and *absent* any potential impacts from the distribution of wages across members in the group or the manner by which ranks are assigned. Because of the different possible wages and the natural variation in ability among participants, we need the distribution of productivity for each possible wage rate. Thus, to estimate baseline productivity, we conducted a BASELINE condition where workers simply completed the work task under a fixed wage, which removed the group dynamics and the possible wage distribution effects on productivity. Specifically, in the BASELINE condition, workers first completed the 8-minute practice period, as in the EARNED and RANDOM conditions. Then, after the practice period, participants in the BASELINE condition completed a 22-minute work period under a similar piece-rate wage. Moreover, they were informed that their wage would be randomly determined and could vary across participants in the session. The set of possible wages that participants could receive was {\$.5, \$.75, \$1, \$1.25, \$1.8, \$2.1}, which corresponded to the six unique values that span the four different wage distributions depicted in Table 1. This approach of using a fixed piece-rate scheme to estimate the baseline productivity potential of workers has been similarly implemented by Abeler et al. (2011), Kube et al. (2013), and Gneezy et al. (2017).

Importantly, participant workers in the BASELINE condition were not working within a group, they only knew their own piece-rate wage, and they were not informed of the other possible wages that other workers were receiving. As such, we maintain that the observed distribution of output in the BASELINE condition provides a consistent estimate of productivity of our sample of participant workers for each of the six unique wage levels. Having an estimate of baseline productivity for each wage level is instrumental in our ability to causally identify any potential productivity effects associated with the different types of wage distributions we consider. Moreover, observing the entire distribution of productivity enables us to investigate possible heterogeneous productivity effects based on worker ability. Our BASELINE condition consisted of 126 participants – roughly 20 observations per wage level.

3 Preliminary Findings

Prior to discussing our main results regarding stated preferences over wage distributions and the subsequent effect on productivity, we discuss some preliminary patterns in the data that are necessary for identifying our main treatment effects.

3.1 Correlation between Output in Practice Period and Paid Work Period

Table 2 presents the summary statistics for productivity in the 8-minute practice period and the 22-minute paid work period for each of the three conditions—BASELINE, EARNED, RANDOM—aggregated over all possible wage levels.

[Table 2 – Aggregate Summary Statistics for Productivity by Condition]

Looking first at productivity in the 8-minute (un-paid) practice period in Panel A, we see that, across all three conditions, participant workers produced non-zero amounts of output. Moreover, the productivity levels in the practice period are very similar across three conditions, as we would expect given that the three conditions are identical until after the practice period. The average practice output was 5.5 applications in the BASELINE condition, 5.6 in EARNED, and 5.21 in RANDOM; a series of pairwise Mann-Whitney U-tests yield no statistically significant differences in practice period output between the three conditions. The similarity in practice period output across the three conditions suggests we have adequate randomization to condition; thus, we can identify how different wage distributions, along with the manner in which they are assigned, impact productivity by focusing on productivity differences in the paid work period.

In terms of productivity in the paid work period, Panel B of Table 2 displays the summary statistics for total output across the three conditions, aggregated over all wages. From Table 2 we see that the average output was 19.4 in the BASELINE, 20.1 in EARNED, and 20.1 in RANDOM; testing all pairwise comparisons using a Mann-Whitney U-test, there are no significant differences between the three conditions. Importantly, there is a very strong positive correlation between productivity in the practice work period and productivity in the paid work period ($r = .616$; $p < .001$).¹⁰ Recall, in the EARNED condition, we assign worker rankings based on productivity in the practice period. Given that the intent of this feature is to assign ranking within the group based on worker ability, it is vital that a worker's output in the practice period be strongly positively

¹⁰Looking at each condition separately, the correlations between practice period output and work period output are $r = .845$ ($p < .001$) in the BASELINE, $r = .819$ ($p < .001$) in EARNED, and $r = .418$ ($p < .001$) in RANDOM.

correlated with their output in the paid work period; this implies that, on average, workers who were higher ability and produced more in the paid work period were the same workers who produced more in the practice period and, correspondingly, assigned a higher rank. Furthermore, the strong positive correlation between output in the practice and paid work period enables us to use the observed output in the practice period as a reasonable proxy for an ordinal measure of ability for each worker—with higher practice period output corresponding to higher ability.

3.2 Pure Wage Effect on Productivity

We also examine output from the BASELINE condition to get an estimate of the *baseline* distribution of productivity, and identify any potential pure wage effects. If needed, we can control for pure wage effects in our analysis, and separately identify any effects results from the specific type of wage distribution. Figure 1 displays the average output in the paid work period in the BASELINE condition for each of the six possible wages.

[Figure 1 – Average Output by Wage Level: BASELINE Condition]

From Figure 1 it is evident that there is very little pure wage effect on productivity. Specifically, for wages of \$.5, \$.75, \$1, \$1.25, \$1.8, and \$2.1, the average output levels were 19.9, 19.4, 19.0, 20.7, 18.7, and 19.2, respectively. A joint test of the effect of wage on output reveals no significant wage effect (ANOVA: $p = .552$). Similarly, there are no significant pairwise differences in output across the six different wage levels. Our results suggest that in the work environment we consider, there appears to be very little pure wage effect over the range of wages we consider.¹¹ This enables us to cleanly identify the effects of our selected wage distributions on productivity. In particular, if we observe differences in worker productivity in the paid work period across the different wage

¹¹ While beyond the scope of this paper to explore pure wage effects on effort and output, we speculate that there are several factors that could have contributed to the observed null wage effect. The first is the process by which wages are determined. In our setting, wages are essentially determined from a random external process and not by a real human principal; as such, this could have mitigated reciprocal motivations of the worker and, consequently, muted the effort response to different wages (e.g., Charness, 2004; Charness & Levine, 2007; Gächter & Thoni, 2010). Second, in the BASELINE, while workers are aware that other workers might be receiving a different wage, they are unaware of the set of possible wages. Hence, there is no reference point for determining wage fairness and, thus, no scope for judgements about wage fairness to factor into the worker's effort choice (Akerlof & Yellen, 1990). We further note that this null wage effect is consistent with some other recent experimental studies. For example, Hennig-Schmidt et al. (2010), Carpenter (2016), Carpenter & Gong (2016), and Cardella & Depew (2017) document negligible pure wage effects in real-effort, mailer assembly tasks, while Greiner et al. (2011) find no significant wage effects in a numerical data entry task and Sliwka & Werner (2017) find very little wage effect in an abstract real-effort task.

distributions, then this can be attributed solely to differences in the distribution of wages across workers and how wages compare across workers.

4 Experimental Results: Stated Wage Distribution Preferences

To examine preferences for the distribution of wages across workers, we first present the data from the voting stage for the EARNED (186 participants) and RANDOM (168 participants) conditions. In the next section we present the productivity data. Recall that when voting on their preferred wage distribution, workers were unaware of their actual rank within the group. We used a strategy method approach, asking workers to make their selection as if they were in each of the five different possible scenarios: (i) Rank 1, (ii) Rank 2, (iii) Rank 3, (iv) Veil, and (v) Planner.

Overall, the voting data reveals significant variation in the preferred wage distribution across the five different voting scenarios; only 10% of participants (37 out of 354) voted for the same preferred distribution in all five scenarios. The remaining participants displayed very little consistency in wage distribution preferences across the voting scenarios. Specifically, Table 3 presents the percentage of participants that chose the same preferred distribution across all the different voting scenario pairs for both the EARNED and RANDOM conditions. Table 3 reveals substantial inconsistency in preferences between different scenarios (the highest level of consistency across any two scenarios was 61% between the Veil and Planner in the EARNED condition). As would be suspected, preferences between voting scenarios that presented very different incentives, like Rank 1 versus Rank 3, were not very consistent. In voting scenarios that presented more similar incentives, like Planner versus Veil, there was more consistency. On the whole, the voting data suggests that workers explicitly considered the different scenarios they were under; in other words, the different voting scenarios seemed to be a salient and important determinant in participant's wage distribution preferences.

[Table 3 – Preference Consistency across Voting Scenario by Condition]

4.1 Wage Distribution Preferences for EARNED Condition

We first present the wage distribution preference data for workers in the EARNED condition. Figure 2 displays the relative frequency that each of the four different wage distributions was chosen in each of the five different voting scenarios.

[Figure 2 – Breakdown of Wage Distribution Preferences in EARNED Condition]

From Figure 2, we see that 55% of workers in the Planner scenario preferred the *Wage Equality* distribution, 18% the *Wage Inequality*, 8% the *Income Equality*, and 19% the *Minimum Wage*. The voting data from the Planner scenario reveals that there is substantial heterogeneity in terms of peoples' preferred wage distribution; moreover, the majority of people preferred the *Wage Equality* distribution, forgoing some overall efficiency (compared to the *Wage Inequality* and *Minimum Wage* distributions).

Next, we consider preferences in the Veil scenarios, where the worker is part of the affected group but is unsure, ex-ante, of their rank within the group. In the Veil scenario, we see that 54% of workers preferred the *Wage Equality* distribution, 27% *Wage Inequality*, 6% *Income Equality*, and 13% *Minimum Wage*. Comparing to the Planner scenario, a similar proportion of worker prefer the *Wage Equality* distribution, and it remains the most preferred distribution. However, in the Veil scenario significantly more workers preferred the *Wage Inequality* distribution compared to in the Planner scenario (Chi-Squared test: $p = .047$). Thus, when workers are personally affected by the wage distribution, there is a shift away from the *Minimum Wage* distribution toward the *Wage Inequality* distribution that is, overall, more efficient but has greater wage dispersion.

Moving to the three scenarios where the workers assume they know their rank within the group, Figure 2 reveals significant shifts in the distribution of preferences based on the worker's rank, and several interesting results emerge. In the Rank 1 condition, we see that 11% preferred the *Wage Equality* distribution, 70% the *Wage Inequality*, 1% the *Income Equality*, and 18% the *Minimum Wage*. Thus, the overwhelming majority of workers (70%) prefer the *Wage Inequality* distribution, and this proportion is significantly higher compared with either the Planner scenario (Chi-Squared test: $p < .001$) or the Veil scenarios (Chi-Squared test: $p < .001$). This shift is primarily driven by choices in the Rank 1 scenario moving away from the *Wage Equality* distribution, where only 11% of workers in the Rank 1 scenario preferred the *Wage Equality* distribution, which is significantly less than in the Planner scenario (Chi-Squared test: $p < .001$) and the Veil scenario (Chi-Squared test: $p < .001$). Additionally, a non-trivial fraction of workers (18%) prefer the *Minimum Wage* distribution, which is significantly different from zero (Chi-Squared test: $p < .001$); this implies that people do care about wage dispersion since the *Minimum Wage* distribution features a lower payment to the rank 1 worker and is less efficient compared to the *Wage Inequality* distribution, but also features less dispersion in wages across the workers.

In the Rank 2 scenario, 43% preferred the *Wage Equality* distribution, 36% the *Wage Inequality*, 13% the *Income Equality*, and 8% the *Minimum Wage*. Not surprisingly, very few people chose the *Minimum Wage* distribution, where the wage for the rank 2 worker is the lowest. Yet, there is substantial heterogeneity in the preferred wage distribution among the other three distributions – *Wage Equality*, *Wage Inequality*, *Income Equality* – all of which have the same \$1 wage for the rank 2 worker. Interestingly, *Wage Equality* is still the most preferred at 43%, although this percentage is lower than either the Planner (Chi-Squared test: $p = .022$) or Veil (Chi-Squared test: $p = .029$) scenarios. While some in the Rank 2 scenario shift to the *Wage Inequality* distribution, which has the same payout for the rank 2 worker and is more efficient, many workers choose to forgo the efficiency gain and corresponding large wage dispersion for equal wages.

Lastly, in the Rank 3 scenario, 35% preferred the *Wage Equality* distribution, 6% the *Wage Inequality*, 49% the *Income Equality*, and 10% the *Minimum Wage*. Again, not surprisingly, the most preferred distribution is the *Income Equality* distribution that has the highest wage for the rank 3 worker. In fact, workers in the Rank 3 scenario chose the *Income Equality* condition significantly more compared to all other voting scenarios (Chi-Squared test: $p < .001$ for all 4 other scenarios). This is a result of there being significantly less choices for both the *Wage Equality* and *Wage Inequality* distributions in the Rank 3 scenario compared to the Planner (Chi-Squared test: $p < .001$) or Veil (Chi-Squared test: $p < .001$) scenarios. At the same time, a substantial fraction of workers (35%) continue to choose the *Wage Equality* distribution, even with this distribution having the same level of efficiency but a lower payment to the rank 3 worker compared to the *Income Equality* distribution.

4.2 Wage Distribution Preferences for RANDOM Condition

Next, we consider the wage distribution preference data for workers in the RANDOM condition. It is useful to point out that because rank is not based on productivity, the interpretation and final income implications of the wage distributions are different compared to the EARNED condition. Notably, the *Income Equality* distribution essentially becomes a moderate wage inequality treatment. Additionally, neither the *Wage Inequality* nor *Minimum Wage* distribution would generate the same degree of final income inequality in the RANDOM condition compared to the EARNED condition. Figure 3 displays the relative frequency that each of the four different wage distributions was chosen as being preferred for each of the five different voting scenarios. Overall a similar pattern emerges in the RANDOM condition as the EARNED condition.

Regarding the Planner scenario, from Figure 3, we see that 45% of workers preferred the *Wage Equality* distribution, 20% the *Wage Inequality*, 16% the *Income Equality*, and 19% the *Minimum Wage*. Again, the most preferred distribution is the *Wage Equality* distribution. Interestingly, roughly 16% of the people in the RANDOM condition prefer the *Income Equality* distribution, where there is some wage inequality, over the *Wage Equality* distribution even though both have the same degree of efficiency; this suggests that even if worker rank is determined randomly, a non-negligible fraction of the worker had a preference for some level of wage inequality. For the Veil scenario, 50% of workers preferred the *Wage Equality* distribution, 29% the *Wage Inequality*, 11% the *Income Equality*, and 10% the *Minimum Wage*. Similar to as in the EARNED condition, the main difference in the RANDOM condition between the Planner and Veil scenarios is that more people prefer the more efficient, but more unequal, *Wage Inequality* distribution (Chi-Squared test: $p = .042$) when they belong to the group to which the wage distribution applies.

[Figure 3 – Breakdown of Preferred Wage Distribution: RANDOM Condition]

Regarding rank-specific preferences, in the Rank 1 scenario, the majority of workers (70%) prefer the *Wage Inequality* distribution, where the rank 1 worker receives the highest wage but also has the most inequality, and this proportion is significantly higher compared with either the Planner scenario (Chi-Squared test: $p < .001$) or the Veil scenarios (Chi-Squared test: $p < .001$). In the Rank 2 scenario, worker preferences are more evenly distributed among the *Wage Equality* (39%), *Wage Inequality* (30%), and *Income Equality* (24%) distributions where the pay to the rank 2 worker is the highest; although, the *Wage Equality* distribution remains the most preferred choice despite the lower efficiency. Again, the fact that 24% of rank 2 workers prefer the *Income Equality* distribution over the *Wage Equality* distribution indicates a preference for some moderate level of wage inequality. In the Rank 3 condition, 71% of workers prefer the *Income Equality* distribution, where the rank 3 worker received the highest wage, which is significantly higher than in the Planner (Chi-Squared test: $p < .001$) or Veil scenarios (Chi-Squared test: $p < .001$).

4.3 Selfish Voting Patterns

Another aspect of the wage distribution preference data that is of interest is purely selfish or self-interested voting. Not too surprisingly, of the various motives related to redistribution preferences, self-interest has been shown to be important, if not the dominant motive (e.g., Esarey et al., 2011;

Durante et al., 2014). To investigate further, we observe how many participants voted, in each of the first three scenarios, to maximize their own wage.¹²

Overall, we classify 43% of participants (80 of 186) in the EARNED condition and 58% (98 of 168) in the RANDOM condition as selfish voters. This difference is statistically significant (Chi-Squared test: $p = .004$). Hence, selfish voting is much more likely in the RANDOM condition. This result is in concert with results about voting consistency, both in the sense that selfish voting requires a lack of consistency, and that without countervailing social norms related to earning a wage, self-interest seems to dominate as a source of distributional preferences. Interestingly, we also find that participants who voted for *Wage Equality* in the Planner scenario were significantly less likely to be classified as a selfish voter (41% compared to 60%). This suggests that norms expressed while acting as a planner (absent direct monetary incentives) correlate with preferences in the Rank 1 and Rank 3 conditions, when monetary self-interest plays a bigger role.¹³

Based on the selfish voting classification, we can re-examine the wage distribution preference data specifically for those participants who *were not* classified as selfish; namely, what pattern emerges for those participants who exhibited some other redistributive motivation besides just maximizing their own wage. Table 4 presents the breakdown of preferred wage distribution for both the EARNED and RANDOM conditions, and several interesting findings emerge. First, from Table 4, we see that votes for *Income Equality* fall substantially. This suggests that most of the votes for *Income Equality*, in both treatments, are driven by narrow self-interest. Our motivation for including the *Income Equality* distribution was to provide, for the EARNED condition, some way for participants to express a preference (through wages) for more final income equality. By lowering the wage of the most-productive worker and raising the wage of the least-productive worker, a plausibly more-equal income distribution could be obtained. This rationale seems to have been largely de-motivating for participants given the similar drop in *Income Equality* votes in both conditions when removing selfish voters.

¹² Maximize one's own wage requires voting for *Wage Inequality* in the Rank 1 scenario, *Income Equality* in the Rank 3 scenario, and anything but *Minimum Wage* in the Rank 2 scenario. It could be argued that it requires voting for *Wage Inequality* in the Veil scenario, but this would technically depend on a participant's belief about where they are likely to fall in the distribution, and so we don't include this requirement.

¹³ We also examined the role of gender and ability (as measured by the participant's output in the practice period), and neither seemed to contribute significantly to selfish voting. We wouldn't necessarily expect to find gender differences so a non-result there is no surprise. Ability, on the other hand, might be expected to be relevant. This is because high-ability individuals have more to gain or lose from a given wage difference. This suggests participants may have been unaware of their relative ability in this domain of data entry.

[Table 4 - Breakdown of Preferred Wage Distribution for Non-Selfish Voters]

While it's tempting to declare that participants' norms and morals are more sensitive to wage differences than income differences, we must caution that our version of income equality is not exacting in the way that our wage equality is. The *Wage Equality* distribution successfully fixes all wages equal to each other, while the *Income Equality* distribution only promotes more equal incomes under certain reasonable conditions. Still, our analysis suggests that sensitivity to wage differences is at least of a similar magnitude as sensitivity to income differences in certain contexts. Table 4 also reveals that preferences for *Wage Inequality* in the Rank 1 condition are largely driven by narrow self-interest (especially in the RANDOM condition). While this is less surprising than results about *Income Equality*, it is interesting that *Wage Inequality* remains reasonably popular in the EARNED treatment after removing selfish voters. This could be a result of norms related to "earning" a higher wage or preferences for efficiency, and speaks to the contexts where people might prefer more unequal wages, and, in turn, more unequal incomes.

4.4 Discussion of Wage Distribution Preferences between EARNED and RANDOM

In summary, we document substantial heterogeneity in wage distribution preferences across our sample of participant workers, which is consistent with the prior literature on preference for redistribution (e.g., Andreoni & Miller, 2002). Furthermore, we find that wage distribution preferences are not stable; rather, they are contingent on the perspective from which the participant is voting. When acting from the perspective of an uninvolved planner (and to a lesser extent from behind the veil), the majority of participants prefer the distribution with equal wages, despite it being the least efficient. That said, many participants also prefer the most efficient but most unequal distribution of wages. Thus, there appears to be a very salient tradeoff between equality and efficiency in the distribution of wages, which has also been documented in preferences for how income be distributed (see references in Footnote 8). However, when acting from the perspective of a known rank/position within the group of workers, we see that preferences tend to become much more selfish. In particular, when acting in the advantaged position of being rank 1, participants overwhelmingly prefer the distribution with the highest payoff for the rank 1 worker, despite the large degree of wage inequality. Conversely, when acting in the disadvantaged position of being rank 3, participants overwhelmingly prefer the somewhat perverse distribution with the highest payoff for the rank 3 worker, despite this being the most inefficient.

In general, a similar pattern in wage distribution preferences emerges between the EARNED and RANDOM conditions; however, there are some key differences. Notably, in the Planner scenario the percentage who vote for *Wage Equality* is marginally lower than in RANDOM, 45%, than in the EARNED, 55% (Chi-Squared test: $p = .071$), with more people instead choosing the *Income Equality* distribution that has more moderate level of wage dispersion (but the same level of efficiency). This further supports the notion that some of the workers prefer a level of wage inequality over equal wages. Another difference emerges under the Rank 3 scenario where the fraction of workers in the RANDOM condition that choose the *Income Equality* distribution is over 20 percentage points higher compared to the EARNED condition, 71% vs 50% (Chi-Squared test: $p < .001$), while the fraction that choose the *Wage Equality* distribution is 12 percentage points lower than in EARNED condition, 23% vs. 35% (Chi-Squared test: $p = .013$). This indicates that workers are much more willing to forgo their own wage for wage equality when their low rank in the group is a result of their lower ability (compared to other workers in the group) than when their low rank is a result of random chance. Alternatively stated, in the RANDOM condition, workers who are randomly assigned rank 3 feel more entitled to receive the highest wage (from the *Income Equality* distribution) because their disadvantageous rank was by chance and not because of low ability. Overall, our results suggest that merit can play some role in how people think wages ought to be distributed among workers.

As part of the post-work survey, we elicited some self-reported information on a variety of different characteristics including: (i) attitudes toward income inequality, (ii) attitudes toward wage inequality, (iii) deservingness of pay, (iv) attitudes toward welfare, and (v) political affiliation. We conducted some post-hoc analysis to examine any possible moderating effects of these characteristics on wage distribution preferences. For brevity we omit reporting the full set of results regarding these moderating effects. However, we note that in general, there was very little difference in the voting patterns that emerged based on differences in these characteristics, and any effects that emerged were small and in the obvious and expected direction. For example, one pattern that seemed to consistently emerge is that participants who reported more strongly that income inequality was a serious problem were more likely to vote for the *Wage Equality* distribution. Similarly, participants who reported that current wage inequity in large firms was too high were more likely to vote for the *Wage Equality* distribution. Alternatively, participants who reported more strongly that merit is an important component in determining pay were more likely

to vote for the *Wage Inequality* distribution. Lastly, more politically conservative participants were more likely to be classified as selfish voters; this result could be due to the emphasis that conservative policies place on individual self-determination, as well as the emphasis that liberal policies place on equality. These results are generally consistent with findings document in the prior literature related to preferences for income redistribution.

5 Experimental Results: Effect of Wage Distribution on Productivity

In the prior section we found that there is substantial heterogeneity in the type of wage distribution that workers prefer. We now turn to the second main objective of this study; to examine how different types of wage distributions impact the productivity of the affected workers. In this section, we present the data on worker productivity from the 22-minute paid work period that followed the wage-voting stage. When appropriate, we decompose the data by rank-assignment condition (EARNED vs RANDOM), wage distribution, and worker rank. We begin by focusing on the quantity of output produced by each worker, which we measure as the number of application entered. Later, we turn our attention to quality and further investigate how the distribution of wages possibly impacts production quality. Table 5 presents the average output of workers in the paid work period, broken down by the wage profile for the group, the workers rank within the group. Panel A displays the results for the EARNED condition and Panel B for the RANDOM condition.

[Table 5 – Average Output in Paid Work Period by Condition and Worker Ranking]

5.1 Group-Level Output across Wage Distributions

Looking first at output from the EARNED condition, Panel A of Table 5 reveals that there is very little difference across the four different wage distributions. In particular, the average output level (across all workers under the given wage distribution) was 19.96 under *Wage Equality*, 20.29 under *Wage Inequality*, 19.47 under *Income Equality*, and 20.33 under *Minimum Wage*. Testing for an overall treatment effect of the wage distribution on output level, we find no significant effect (ANOVA: $p = .718$).¹⁴ The distribution of wages (based on the four possibilities we consider), had little effect on the aggregate amount of output produced within the group of workers when worker rank was earned based on productive ability.

¹⁴Similarly, if we run a regression of the productivity in the paid work period on the wage distributions, while controlling for output in the practice period, none of the coefficients on the wage distribution dummies are significantly different from zero, further indicating a null effect of the wage distribution on productivity in EARNED.

Looking next at aggregate output in the RANDOM condition, Panel B of Table 5 reveals that there are some differences across the four wage distributions. The average output (across all workers under the given wage distribution) was 18.66 under *Wage Equality*, 19.98 under *Wage Inequality*, 19.78 under *Income Equality*, and 22.78 under *Minimum Wage*. Testing for an overall treatment effect of the wage distribution on the aggregate level of production within the group, we do find a significant treatment effect (ANOVA: $p < .001$). Specifically, the *Wage Equality* distribution results in lower average output than under *Wage Inequality* (Mann-Whitney U-test: $p = .088$) and the *Minimum Wage* distribution (Mann-Whitney U-test: $p < .001$). Additionally, average output is significantly higher under the *Minimum Wage* condition compared to the *Wage Inequality* condition (Mann-Whitney U-test: $p = .007$) and the *Income Equality* distribution (Mann-Whitney U-test: $p = .008$).¹⁵

5.2 Worker-Level Output across Wage Distributions

Our design not only allows us to identify how the different wage distributions impact aggregate output of the group, but also possible heterogeneous impacts on the different-ranked workers in the group. Even in the EARNED condition where we found essentially no significant group-level output differences across wage distributions, it's still possible that there might be differential effects for certain types of workers. Some recent papers have documented heterogeneous effects of pay comparisons on effort. For example, Gross et al. (2015) find that high ability worker reduce effort when getting paid less than or equal to low ability workers, while the same is not true when low ability workers get paid less than, or equal to, high ability workers. Similarly, Abeler et al. (2010) find that high ability workers choose lower effort when getting paid the same as low ability workers. There may also be asymmetric effects in how workers respond to wage comparisons.¹⁶ Cohn et al. (2014) find that workers lower effort when getting paid less than another person, but the workers do not increase effort when getting paid more than another person. Mas (2006) finds that reductions in effort among workers are larger the lower the wage relative to a reference point, suggesting that effort responses to less fair wages might be larger.

¹⁵ Similar to in the EARNED condition, if we regress output on wage distribution, while controlling for output in the practice period, the wage distribution effect is muted. In particular, only one marginally significant effect (at the 10% level) emerges and that is between the *Wage Equality* distribution and the *Minimum Wage* distribution.

¹⁶ Relatedly, Card et al. (2012) find an asymmetric effect of wage comparisons on job satisfaction, where employees who earn less than the median report lower job satisfaction, while workers who earn above the median do not report higher satisfaction.

In the EARNED condition, Panel A of Table 5 shows that no sizable heterogeneous effects on output based on worker ranking. For the rank 1 workers in the group, their average output was 22.31 under *Wage Equality*, 22.08 under *Wage Inequality*, 22.00 under *Income Equality*, and 21.30 under *Minimum Wage*, and there is no significant treatment effect of the wage distribution (ANOVA: $p = .881$). For the rank 2 workers, their average output was 19.44 under *Wage Equality*, 20.31 under *Wage Inequality*, 19.20 under *Income Equality*, and 20.60 under *Minimum Wage*, and there is no significant treatment effect of the wage distribution (ANOVA: $p = .719$). Lastly, for the rank 3 workers, their average output was 18.13 under *Wage Equality*, 18.50 under *Wage Inequality*, 17.20 under *Income Equality*, and 19.10 under *Minimum Wage*, and there is also no significant treatment effect of the wage distribution (ANOVA: $p = .582$). Overall, the data reveals that the implemented wage distribution had virtually no impact on worker output, regardless of their rank within the group, when rank was earned.¹⁷

In the RANDOM condition, Panel B of Table 5 reveals more variation in worker-level output across wage distributions. In particular, for the rank 1 workers, their average output was 20.58 under *Wage Equality*, 19.55 under *Wage Inequality*, 19.67 under *Income Equality*, and 22.11 under *Minimum Wage*, although the overall treatment effect of the wage distribution is not significant (ANOVA: $p = .266$). For rank 2 workers, their average output was 17.17 under *Wage Equality*, 20.05 under *Wage Inequality*, 20.13 under *Income Equality*, and 23.00 under *Minimum Wage*, and there is significant treatment effect of the wage distribution (ANOVA: $p = .028$). For rank 3 workers, their average output was 18.25 under *Wage Equality*, 20.35 under *Wage Inequality*, 19.53 under *Income Equality*, and 22.78 under *Minimum Wage*, and the treatment effect of the wage distribution is significant (ANOVA: $p = .046$). While it is the case that the wage distribution seems to impact worker-level output more in the RANDOM condition than in the EARNED condition, it is prudent to exercise some caution regarding this finding. Namely, the significant treatment effect of the wage distribution is largely being driven by the higher output levels under the

¹⁷ While workers in the BASELINE condition did not work in groups of three, we programmed the software such that groups of three were randomly created and workers were assigned counter-factual worker ranks based on their productivity in the practice work period, just as in the EARNED condition. As such, we are able to identify a “baseline” measure of output for each of the worker ranks absent any group dynamics and social comparisons. When we do this, the average output for rank 1 workers is 22.09, for rank 2 workers is 19.00, and for rank 3 workers is 16.86. In comparing the EARNED condition to the BASELINE, there are no significant differences for either rank 1 or rank 3 workers, and only a marginal difference for rank 2 workers ($p = .097$). Thus, not only does the type of wage distribution not have a large impact on output produced, we also find little evidence that working in a group and having information about the distribution of wages across the group impacts productivity.

Minimum Wage distribution, but we only observe a total of nine groups under this wage distribution (because of the infrequency with which this distribution was voted as the preferred distribution). Given the small sample size, it is possible that the presence of a few high ability outliers could be driving this result.¹⁸

Indeed, for robustness we examined the impact of wage distribution on output via regression analysis, while controlling for ability with each participant's total output in the practice period. For brevity we omit reporting the full set of results as they largely consistent with the results presented in Table 5. To summarize, ability (as measured by practice period output) is strongly positively related to output in the work period. Moreover, the regression analysis reveals no significant effects (at the 5% level) of the wage distribution on output in the EARNED condition for the entire sample or any of specific ranks. Similarly, for the RANDOM condition, the regression analysis yields no significant wage distribution effects (at the 5% level) for the entire sample or any of the specific ranks. Thus, after controlling for ability, we confirm that the wage distribution has no significant effects on total output in the EARNED condition and, at most, only a small effect in the RANDOM condition.

5.3 Quality of Output across Wage Distributions.

Even though we don't observe large effects of the different wage distributions on the *quantity* of output, worker productivity can be multidimensional. Specifically, in our data entry task, as well as many real-world work tasks, there is a quality dimension; worker productivity encompasses not only how much output the worker produced but also the quality of that output. Thus, there is the possibility for the difference in wage distributions to impact the *quality* of work produced, even in the absence of significant differences in quantity. For example, Greiner et al. (2011) find that when workers are paid a piece-rate and are paid a lower wage compared to another paired worker in the same role, then they produce significantly lower quality work.

In our task, the possibility arises for workers to produce lower quality output in the form of more error-laden entries. Similar to the approach implemented in Greiner et al. (2011), we construct a quality measure based on the accuracy of the information entered from the application; naturally, more errors correspond to lower quality output. Recall that by design, participant workers did not have to enter applications correctly to receive their compensation. Thus, there was

¹⁸In fact, if we omit the *Minimum Wage* distribution and just look across the other three, the wage distribution effect for the remaining three distributions is no longer statistically significant for either rank 1 or rank 2 workers.

certainly the opportunity for workers to produce lower quality output via inaccurate data entry. For each application entry, we check accuracy based on following fields of entry:

- First Name
- Last Name
- Zip code
- Set of 6 radio buttons indicating education background and inclusion of documents

For each application entered, we check if each of these 9 different fields were entered correctly. If any of the 9 fields were entered incorrectly, that application is designated as being inaccurate. Then, we construct a worker-level measure of quality, which we denote as *inaccuracy*, as the proportion of total forms entered that were inaccurate.¹⁹ In essence, higher rates of *inaccuracy* would represent lower levels of effort, manifested through being more careless and less focused. Table 6 displays the average *inaccuracy* rate, broken down by worker ranking across each of the four wage distributions, separately for the EARNED and RANDOM conditions.

[Table 6 – Output Quality: Comparison of *Inaccuracy* Rates]

We begin by looking at output quality in the EARNED condition.²⁰ From Panel A, we see that the overall *inaccuracy* rate is relatively stable across the four wage distributions; in particular, the average *inaccuracy* rates are 13.3% under *Wage Equality* distribution, 12.9% under *Wage Inequality*, 10.2% under *Income Equality*, and 12.8% under *Minimum Wage*, which are not jointly significantly different (ANOVA: $p = .732$). In terms of comparisons at the worker rank level, there are no discernable differences in *inaccuracy* rates for either rank 1 or rank 2 workers across wage distributions. On the contrary, for rank 3 workers, the *inaccuracy* rate is noticeable lower under the *Income Equality* condition (7.1%) compared to the other three distributions (15.7%, 15.3%, 17.3%, respectively); this difference is significant when comparing to *Wage Equality* (Mann-Whitney U-test: $p = .019$) or *Wage Inequality* (Mann-Whitney U-test: $p = .023$) distributions. Thus, when the low ability rank 3 workers are provided with the highest wage, they seem to take the task more seriously and produce higher quality output.

¹⁹ Our results are qualitatively robust if we instead look at the total number of inaccurate fields entered per application. We prefer the simple binary measure of *inaccuracy* because it ensures that any differences we document are not being inflated by applications that are haphazardly entered with all 9 fields being incorrect. Moreover, the proportional nature of the *inaccuracy* rate measure makes for ease of interpretation.

²⁰ As we did with output quantity, we first tested for differences in output quality (based on the measures we use in our analysis) across the six different piece-rate wage levels in the BASELINE conditions. Importantly, we found no statistically significant differences in output quality across the six different wage levels. This suggests that there is little pure wage effect on output quantity, which mirrors the results we found above in Section 3.2 on quantity.

Panel B of Table 6, presents the *inaccuracy* rates for the RANDOM condition. Similar to with the output level data, we see much more variation in quality in the RANDOM condition. Namely, the overall average *inaccuracy* rate was 7.2% under *Wage Equality* distribution, 12.1% under *Wage Inequality*, 16.0% under *Income Equality*, and 13.8% under *Minimum Wage*, and these differences are jointly significantly different (ANOVA: $p = .008$). Thus, in the aggregate, the *Wage Equality* distribution seems to generate the lowest rate of inaccurate entries (i.e., the highest quality), and this difference is significant when comparing to *Income Equality* distribution (Mann-Whitney U-test: $p < .001$), and marginally so to both the *Minimum Wage* and *Wage Inequality* distributions (Mann-Whitney U-test: $p = .082$ and $p = .061$, respectively). In the aggregate, quality appears to be higher under *Wage Equality*.

Some interesting results also emerge within the worker ranks in the RANDOM condition. In particular, for the rank 1 workers there is a noticeably higher *inaccuracy* rate under the *Income Equality* condition (16.7%), especially compared to the *Wage Inequality* condition (Mann-Whitney U-test: $p = .079$), and marginally so to the *Wage Inequality* condition (Mann-Whitney U-test: $p = .106$); thus, rank 1 workers seem to produce the lowest quality output when they are paid the least (and their low rank is a result of random chance). For rank 2 workers, we also see differences in quality across distributions. In particular, they have the lowest *inaccuracy* rate under the *Wage Equality* distribution (5.2%), compared to the other three distributions (15.5%, 11.7%, 12.0%, respectively), which are all significantly different (Mann-Whitney U-test: $p = .036$; $p = .011$; $p = .053$; respectively). Thus, rank 2 workers appear to produce lower quality output when the wage distribution is not equal, there is another worker who is unjustifiably making more than they are. Lastly, for the rank 3 workers, we see that similar to rank 2 workers, the *Wage Equality* distribution produces the lowest *inaccuracy* rate (7.0%). Also, we see that the *Income Equality* condition, where rank 3 workers are paid the most, induced the highest *inaccuracy* rate (19.7%); this rate is significantly higher than under *Wage Equality* ($p = .009$) and the *Wage Inequality* ($p = .049$). Thus, the higher wage, when it is randomly assigned, possibly entices the rank 3 workers to work faster and make more errors.

While inaccurate data entry is a clear signal of producing lower quality output, our specific task allows for an alternative means of reducing effort and producing lower quality. In particular, as part of the digital template there was some *suggested* formatting embedded in certain data fields. Namely, for the phone number we suggested that it be entered of the form (####) ###-####, as

opposed to a string of 10 numbers; similarly, the application # took the form ###-####-#####. We assert that another possible manifestation of lower effort, which produces lower quality output, is not following the suggested or implied formatting of the data entry fields. This is a subtler form of shirking where the worker is essentially “cutting corners” to complete each entry in a more expedient manner. Given the piece-rate compensation scheme, this seems like a plausible channel with which to see lower quality. We evaluate proper formatting using four different checks: (i) phone # entered using the suggest format, (ii) application # entered as formatted with the dashes, (iii) first name capitalized, and (iv) last name capitalized. If at least one of these fields was not properly formatted, we denote that application has having improper formatting. Similar to the *inaccuracy* measure, we construct a worker-level measure, denote as *formatting error*, as the proportion of improperly formatted applications.²¹ Table 7 presents the average *formatting error* rates across conditions, wage distributions, and worker ranks.

[Table 7 – Output Quality: Comparison of *Formatting Error* Rates]

From Panel A of Table 7, we see that in the EARNED condition there are no significant differences in the overall *formatting error* rates across the four distributions, which range from 33.9% to 45.2%. However, we do see some sizable differences emerge in *formatting error* rates specifically for rank 1 workers. In particular, *formatting error* rates are noticeable higher in both the *Wage Inequality* (61.5%) and *Income Equality* (70.3%) distributions, compared to *Wage Equality* (18.1%), which are both significantly different (Mann-Whitney U-test: $p = .020$ and $p = .010$, respectively). Thus, rank 1 workers appear to enter the data comparably accurately, but they take shortcuts and cut corners with formatting to get the job done faster when they are either paid a high wage (under *Wage Inequality*) or an unfairly low wage (under *Income Equality*) compared to the *Wage Equality* distribution. For rank 3 workers, we see a similar pattern as with *inaccuracy*; namely, the *formatting error* rate is lowest in the *Income Equality* distribution (20.5%), and this difference is marginally significant compared to both the *Wage Inequality* and *Minimum Wage* distributions (Mann-Whitney U-test: $p = .058$ and $p = .089$, respectively). Thus, rank 3 workers

²¹ As we would expect, the *inaccuracy* and *formatting error* measures are positively and statistically significantly correlated with each other ($r = .168$; $p < .001$). Importantly, however, the magnitude of the correlation coefficient indicates that each of these measures is capturing district information about the quality of output. Thus, considering both measures provides a more robust picture of how the wage distributions impact output quality.

appear to be more careful, in terms of both accuracy and formatting, when they are paid a relatively higher (and undeserved) wage than their group members.

In the RANDOM condition, we see from Panel B of Table 7 that overall *formatting error* rate is higher in the *Income Equality* distribution (62%) compared with the other distributions, especially the *Wage Equality* and *Wage Inequality* distributions (Mann-Whitney U-test: $p = .008$ and $p = .014$, respectively). We also see that rank 1 workers have a much high *formatting error* rate in the *Income Equality* distribution (60.8%), especially compared to the *Wage Inequality* distribution (26.2%), which is significantly different (Mann-Whitney U-test: $p = .043$); For rank 2 workers, we see a parallel effect emerge as with *inaccuracy*; namely; they have the lowest *formatting error* rate under the *Wage Equality* distribution (17.8%), compared to the other three distributions (40.7%, 57.0%, 44.4%, respectively). Thus, rank 2 workers appear to be less diligent and cut corners when there is another worker who is undeservingly making more than they are. Overall, the results suggest that merit based ranking can play an important moderating role in terms of the way output quality of workers of given ranks respond to the distribution of wages; both in terms of the accuracy with which data is entered and the care with which the task is completed.

5.4 Realization of Preferred Wage Distribution

We document substantial heterogeneity in wage distribution preferences. A natural question is whether workers are differentially impacted based on whether the realized wage distribution (under which they are working) is consistent with their stated preferences. To shed light on this question, we first stratify workers in the paid working period based on whether the implemented wage distribution matches their stated preference, given their actual rank within the group.²² Next, we examine if this impacts productivity, both in terms of level of output and quality of output. To do so, we regress both the level of output and the quality of output on worker rank and the interaction with a dummy variable, $PWD = 1$, when the implemented wage distribution matched the workers stated preferred wage distribution (PWD). Table 8 displays the regression results.²³

²² Note, within each group it is possible for zero up to all three workers to have the realized wage distribution match their stated preference. As a practical example of how this was determined, suppose the realized wage distribution was the *Wage Equality* distribution. If, in the group, the worker ranked #1 voted for *Wage Inequality* in the Rank 1 scenario, while the workers ranked #2 and #3 voted for *Wage Equality* in the Rank 2 and Rank 3 scenarios, respectively, then we would indicate that both the rank 2 and rank 3 workers got their preferred distribution.

²³ In order to improve statistical power for these regressions, we look at the total number of observed accuracy errors (or formatting deviations) on the application level. We use robust standard errors, clustered by participant, to account for the interdependence across individuals. For quantity results we use participant-level analysis.

[Table 8 - Productivity under a Preferred Wage Distribution]

From column 1 of Table 8, we see that working under one's stated preferred wage distribution has little impact on the output level in the work period; this is true for both the EARNED and RANDOM conditions, as well as for each of the worker ranks. Namely, the interaction term of *PWD* and worker rank is not statistically significant for any of the three worker ranks in either condition. However, looking at columns 2 and 3, we do see some evidence that working under one's preferred wage distribution does impact output quality. Specifically, in terms of accuracy errors, from column 2 we see that the interaction terms of *PWD* and worker rank are generally negative, and statistically significant for the rank 2 and rank 3 workers in the EARNED condition and for the rank 1 and rank 2 workers in the RANDOM condition; these negative interaction terms indicate the corresponding workers were more accurate in the data entry when working under their stated wage distribution preference. In terms of formatting errors, when looking at column 3, we see a much weaker effect emerge. In particular, only in the EARNED condition and for rank 3 workers does working under one's preferred wage distribution reduce formatting errors.

The formatting error results present somewhat of a puzzle. Intuitively, and evidenced by the accuracy results, we would expect participants to put more thought and care into a process they feel is fair, but we see little evidence of this from formatting choices. While it is possible that our formatting-error measurements are noisier than our accuracy-error ones, we believe that "choice" may play a more important role. A formatting shortcut is a conscious choice with predictable results. Ignoring formatting suggestions should allow entries to be made in less time and allow for more total entries and higher final income. This explicit income incentive may have dominated the effect of getting a preferred wage distribution. Overall, even with the weaker formatting results, we believe that participants do respond to getting a preferred wage distribution with higher effort. This effort seems to manifest mostly through the accuracy channel in our design. This result adds support to the idea that a happy worker is a more productive worker.

6 Discussion and Concluding Remarks

The main motivation of this paper is twofold. First, we examine people's preferences over how wages are distributed among groups of workers with heterogeneous ability when faced with a menu of possible wage distributions that differ with regard to their degree of equality, efficiency, and implied final income equality. Second, we examine how the different wage distributions

impact the productivity of the affected group of workers, both in terms of quantity and quality of output. We conduct a stylized, real-effort, data-entry experiment, and we implemented a performance based, piece-rate pay scheme per application they entered. As part of the experimental design, we considered two different conditions for how ranking within the group was determined: (i) earned, which was merit-based, determined by productivity, (ii) random, which was independent of productivity. We presented participants with four different wage distributions intended to represent the following scenarios: *Wage Equality*, *Wage Inequality*, *Income Equality*, and *Minimum Wage*. We first elicited, in an incentivized manner, participants' preferred wage distribution, and then they completed the work task under the realized wage distribution.

Summarizing our main results regarding observed choices in the wage distribution preference task, we find significant heterogeneity in preferences. When participant workers are acting in the role of *social planner*, the majority of people (roughly 55%) prefer the distribution with equal wages, despite it being the least efficient; at the same time, a sizable percentage (roughly 18%) prefer the distribution with the most wage inequality, but also the highest efficiency. Interestingly, we also see a very similar pattern emerge when workers choose from *behind the veil* (when their rank within the group is unknown), where roughly 55% prefer wage equality and roughly 27% prefer wage inequality. However, when workers are asked to choose their preferred distribution assuming they know their rank within the group, we observe significant non-stationarity in preferences that depend on the rank of the worker, and become much more selfish. Specifically, when assuming they are rank 1, workers overwhelmingly prefer the wage inequality distribution (roughly 70%), which gives the highest payoff to the rank 1 worker. Conversely, when assuming they are rank 3, workers overwhelmingly prefer the income equality distribution (roughly 50%), which gives the highest wage to the rank 3 worker. Thus, workers appear to like the idea of wage equality, so long as they are not negatively impacted by it. However, if they are highly ranked, then they instead prefer a more unequal wage distribution that rewards the high-rank worker with a higher wage. In a similar vein, if workers are ranked low, they actually prefer a somewhat perverse wage distribution that provides the low ranked worker with the highest wage. Lastly, we see that merit does impact wage distribution preferences. Most notably, the lowest rank workers are much less likely to vote for the distribution that gives them the higher wage (in favor of wage equality), when rank is earned, compared to randomly assigned; this indicates that these lower ranked workers exhibit less selfish voting when their low rank is deserved.

Substantial research attention has been geared toward understanding peoples' preferences for income redistribution. Given that wages play a large role in shaping the distribution of incomes, it strikes us as an important complement to this literature to consider *upstream* preferences regarding how wages are distributed. Much of the redistribution literature supports the notion that people tend to prefer less income inequality and, hence, are supportive of ex-post redistribution (even if it is costly). Our results regarding wage distribution preferences are largely consistent with the general patterns observed in these prior studies. Namely, when acting as a planner or behind the veil of ignorance, many people are willing to trade-off some overall efficiency in favor of an equal wage distribution, which would consequently generate a more equal income distribution. Thus, our results suggest that in an effort to reduce high levels of income inequality, policies aimed at reducing wage dispersion (e.g., increasing minimum wage or capping executive compensation) may be a plausible alternative to redistributive policies (e.g., taxation and social welfare programs). Moreover, such wage policies might be better supported given they have the appeal of being perceived, ex-ante, as fairer and potentially less costly. At the same time, there is some potential concern that such wage policies could have detrimental impacts on productivity, although our results suggest these concerns might not be substantiated.

In particular, we find very little difference in quantity of output produced across the four different distributions, and this is especially true when worker ranking within the group is merit based. Namely, there are no significant differences in group-level output across the four conditions. Furthermore, we also do not find any sizable rank-level productivity differences across the four distributions. Thus, our findings largely suggest that the wage distribution has, at most, a small effect on output.²⁴

Many of the prior studies that investigate topics related to the effects of relative pay and pay comparisons across workers on productivity have done so under a framework of a fixed wage (cf. Greiner et al., 2011), and several of the studies have documented that relative pay differences

²⁴ Our interest was in determining how different wage distributions impact productivity of the treated workers. As such, we required participant workers to be present for the entire 22-minute work period. Thus, you could interpret this finding as not finding a significant effort response by workers, conditional on working. However, there may be more subtle ways in which the distribution of wages might impact overall productivity, which are beyond the scope of the current study. For example, if participant workers are allowed to choose how much to work and stop working when they want, we might see the distribution of wages impacting how long workers work (e.g., Abeler et al., 2011). Relatedly, if workers are allowed to decide if they want to participate in the labor pool, we might see the distribution of wages impacting labor supply (e.g., Bracha et al., 2015). Perhaps relative pay issues and the distribution of wages impact labor participation decisions more than the effort decision, conditional on working.

across workers can impact productivity.²⁵ Our paper differs from these prior studies in that we implement performance pay in the form of a piece-rate wage. This distinction has very different implications in terms of how changes in effort/productivity impact the payoff of the worker. For example, if we take the standard chosen-effort, gift-exchange game (e.g., Fehr et al., 1993), reducing effort increases the total payoff to the worker, while increasing effort reduces the total payoff to the worker (conditional on the chosen fixed wage by the principal). With piece-rate pay, the opposite relation can emerge; so long as the cost of effort to produce an additional unit is less than the piece-rate, reductions in effort reduce total pay by reducing output, while increases in effort increase total pay by increasing output. As such, with piece-rate pay the motivation to reduce effort in response to an unfair wage, relative to others, might be crowded out by the corresponding reduction in earnings from producing lower output. So, while relative pay and horizontal pay comparison might impact effort under fixed wages, our results suggest that this effect might be moderated under performance pay. We think this is particularly important in light of the fact that performance pay jobs play a large role in the economy (Lemieux et al., 2009).²⁶ The main implication being that for firms whose compensation schemes are (at least part) performance pay, the distribution of wage across employees might play less of a role in shaping productivity.

In many instances, worker productivity can be a multi-dimensional measure of output. Because on the nature of the real-effort work task we use, we are able to also examine how the distribution of wages across workers impacts the quality of output produced. Our results do suggest that relative pay comparisons can impact productivity more subtly through changes in the quality of output produced (measured by data entry errors), which is similar to the finding of Greiner et al. (2011). Most notably, when high ability workers are paid the least (under the income equality condition), they produce significantly lower quality output. This result is similar to the findings in Gross et al. (2015) where high ability workers reduce effort if they are paid less than the low ability worker. Thus, under piece-rate, our results suggest that high ability workers may still shirk in response to low pay, but they will do it in a subtler way through low quality work. We also find some evidence cutting the other way. Specifically, we find that low ability workers produce higher quality output

²⁵ See Abeler et al. (2010), Clark et al. (2010), Gächter & Thoni (2010), Angelova et al. (2012), Nosenzo (2013), Cohn et al. (2014), Gross et al. (2015), and Charness et al. (2016).

²⁶ Lemieux et al. (2009) show that the overall proportion of performance-pay jobs in the U.S. has increased from about 3 percent in the late 1970s to approximately 45 percent in the 1990s. More recently, Kuhn & Lozano (2008) present evidence that incentive pay across fortune 1000 firms has increased over the latter part of the 20th century.

when they are given the highest relative wage. Furthermore, we find some evidence implying that if worker ranks/positions within a group are determined randomly, then unequal pay schemes might induce lower quality output by those workers receiving the unfair wages. Lastly, we document evidence that workers tend to produce significantly higher quality output when the implemented wage distribution is consistent with their stated preferences. Overall, our results regarding the effect of wage distributions on quality can help to deepen our understanding of how wage distributions and relative pay comparisons impact overall productivity, beyond just the output quantity dimension.

A unique component of our study, which we view as an important contribution to the prior literature, is the elicitation of wage distribution preferences. This data allows us to gain insights into workers' attitudes regarding the wage distribution at their current employer, and also how workers might view proposed changes that would impact the wage distribution. In general, we find that people seem to like the idea of more equal wage distributions (Kiatpongsan & Norton, 2014), especially from a planners' perspective or behind the veil of ignorance. At the same time, workers who might be negatively affected by more equal wage distributions are less apt to prefer more wage equality. Better understanding worker preferences over different wage distributions, in combination with productivity effects, can have important implications regarding recent discussions on policies aimed at reducing the pay gap within firms (e.g., limiting executive compensation and/or raising minimum wage), as well as limiting pay secrecy and making wages within firms more transparent. Our results suggest that such wage policies might not result in differences in total output; similarly, exposing a large degree of wage dispersion within a firm might not immediately alter the productivity of worker (e.g., Hibbs Jr. & Locking, 2000), at least in settings where wages are merit based and there is some degree of performance pay. However, there may be longer term, indirect productivity effects that manifest through changes in workers attitudes. Specifically, decreasing wage inequality within a firm could increase cohesiveness (Levine, 1991), satisfaction (Clark & Oswald; 1996; Card et al., 2012; Godechot & Senik, 2015), happiness (Alesina et al., 2004; and Ferrer-i-Carbonell & Ramos, 2014 for a review), and/or the supply of labor (Bracha et al., 2015), which could increase long-run productivity; whereas exposing a large degree of pay inequality through more transparent wage policies might have the opposite effects. Furthermore, output quality might also be negatively impacted.

Lastly, we acknowledge that our design is stylized and the lack of observed differences in output across the different wage distributions may be, in part, an artifact of the design. Namely, Cohn et al. (2015) find that effort responses to wages depend on the workers initial perception of wage, where an effort response to a wage increase is only observed for those who felt underpaid, whereas those who initially felt adequately paid or over-paid don't show an effort response to increases in wage. This might be relevant in our study, especially when wages were merit based, if participant workers felt that even the lowest piece-rates we consider in each distribution constituted adequate pay. In line with the findings from Cohn et al., such perceptions could have muted the possible negative productivity effects to a lower relative wage and also the possible positive productivity effects to a higher relative wage. That said, Charness and Kuhn (2007) also find little effort response to changes in wages and they aptly note that there is a difference between caring about relative wages and actually acting on them. Our wage distribution preference data lends support to this argument put forth by Charness and Kuhn that participant workers in our study do care about relative wages; yet, they don't actually act on differences in relative pay across the different wage distributions in the form of changes in their output. Although, our results suggest that workers may act on these differences through the alternative dimension of quality. Moreover, this dissonance between preferences over wage distribution and output produced is likely to be more salient under performance pay, where changes in output impact income directly.

References

- [1] Abeler, J., Altmann, S., Kube, S., & Wibrall, M. (2010). Gift Exchange and Workers' Fairness Concerns: When Equality is Unfair. *Journal of the European Economic Association*, 8(6), 1299-1324.
- [2] Abeler, J., Falk, A., Goette, L., & Huffman, D. (2011). Reference points and effort provision. *The American Economic Review*, 101(2), 470-492.
- [3] Akerlof, G. A., & Yellen, J. L. (1990). The fair wage-effort hypothesis and unemployment. *The Quarterly Journal of Economics*, 105(2), 255-283.
- [4] Alesina, A., & Angeletos, G. M. (2005). Fairness and redistribution. *The American Economic Review*, 95(4), 960-980.
- [5] Alesina, A., Di Tella, R., & MacCulloch, R. (2004). Inequality and happiness: are Europeans and Americans different?. *Journal of Public Economics*, 88(9), 2009-2042.
- [6] Alesina, A., E. Glaeser, and B. Sacerdote (2001), Why Doesn't the US Have a European-style Welfare State?. *Brookings Papers on Economic Activity*, 2, 187-254
- [7] Alesina, A., & La Ferrara, E. (2005). Preferences for redistribution in the land of opportunities. *Journal of Public Economics*, 89(5), 897-931.
- [8] Andreoni, J., & Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2), 737-753.
- [9] Angelova, V., Güth, W., & Kocher, M. G. (2012). Co-employment of permanently and temporarily employed agents. *Labour Economics*, 19(1), 48-58.
- [10] Balafoutas, L., Kocher, M. G., Putterman, L., & Sutter, M. (2013). Equality, equity and incentives: An experiment. *European Economic Review*, 60, 32-51.
- [11] Bartling, B., & Von Siemens, F. A. (2011). Wage inequality and team production: An experimental analysis. *Journal of Economic Psychology*, 32(1), 1-16.
- [12] Beckman, S. R., DeAngelo, G., Smith, W. J., & Wang, N. (2016). Is social choice gender-neutral? Reference dependence and sexual selection in decisions toward risk and inequality. *Journal of Risk and Uncertainty*, 52(3), 191-211.
- [13] Beckman, S. R., Formby, J. P., & Smith, W. J. (2004). Efficiency, equity and democracy: Experimental evidence on Okun's leaky bucket. *Research on Economic Inequality*, 11, 17-42.
- [14] Bellemare, C., Lepage, P., & Shearer, B. (2010). Peer pressure, incentives, and gender: An experimental analysis of motivation in the workplace. *Labour Economics*, 17(1), 276-283.
- [15] Benabou, R., & Tirole, J. (2006). Belief in a just world and redistributive politics. *The Quarterly Journal of Economics*, 121(2), 699-746.
- [16] Bewley, T. F. (1999). Why wages don't fall during a recession. Cambridge, MA: Harvard University Press.

- [17] Bolton, G. E., & Ockenfels, A. (2006). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: comment. *The American Economic Review*, 96(5), 1906-1911.
- [18] Bolton, G., & Werner, P. (2016). The influence of potential on wages and effort. *Experimental Economics*, 19(3), 535-561.
- [19] Bracha, A., Gneezy, U., & Loewenstein, G. (2015). Relative pay and labor supply. *Journal of Labor Economics*, 33(2), 297-315.
- [20] Breza, E., Kaur, S., & Shamdasani, Y. (2017). The morale effects of pay inequality. *The Quarterly Journal of Economics*, 133(2), 611-663.
- [21] Butler, J. V. (2016). Inequality and relative ability beliefs. *The Economic Journal*, 126(593), 907-948.
- [22] Card, D., Mas, A., Moretti, E., & Saez, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *The American Economic Review*, 102(6), 2981-3003.
- [23] Cardella, E., & Depew, B. (2017). Output Restriction and the Ratchet Effect: Evidence from a Real-Effort Work Task. *Games and Economic Behavior*, <https://doi.org/10.1016/j.geb.2017.11.005>.
- [24] Carpenter, J. (2016). The labor supply of fixed-wage workers: Estimates from a real effort experiment. *European Economic Review*, 89, 85-95.
- [25] Carpenter, J., & Gong, E. (2016). Motivating Agents: How Much Does the Mission Matter? *Journal of Labor Economics*, 34(1), 211-236.
- [26] Charness, G. (2004). Attribution and reciprocity in an experimental labor market. *Journal of Labor Economics*, 22(3), 665-688.
- [27] Charness, G., Cobo-Reyes, R., Lacomba, J. A., Lagos, F., & Pérez, J. M. (2016). Social comparisons in wage delegation: Experimental evidence. *Experimental Economics*, 19(2), 433-459.
- [28] Charness, G., & Kuhn, P. (2007). Does pay inequality affect worker effort? Experimental evidence. *Journal of Labor Economics*, 25(4), 693-723.
- [29] Charness, G., & Levine, D. I. (2007). Intention and stochastic outcomes: An experimental study. *The Economic Journal*, 117(522), 1051-1072.
- [30] Cherry, T. L., Frykblom, P., & Shogren, J. F. (2002). Hardnose the dictator. *The American Economic Review*, 92(4), 1218-1221.
- [31] Clark, A. E., Masclet, D., & Villeval, M. C. (2010). Effort and comparison income: Experimental and survey evidence. *ILR Review*, 63(3), 407-426.
- [32] Clark, A. E., & Oswald, A. J. (1996). Satisfaction and comparison income. *Journal of Public Economics*, 61(3), 359-381.
- [33] Clark, J. (1998). Fairness in public good provision: an investigation of preferences for equality and proportionality. *Canadian Journal of Economics*, 708-729.

- [34] Cohn, A., Fehr, E., & Goette, L. (2015). Fair wages and effort provision: Combining evidence from a choice experiment and a field experiment. *Management Science*, 61(8), 1777-1794.
- [35] Cohn, A., Fehr, E., Herrmann, B., & Schneider, F. (2014). Social comparison and effort provision: Evidence from a field experiment. *Journal of the European Economic Association*, 12(4), 877-898.
- [36] Cojocaru, A. (2014). Prospects of upward mobility and preferences for redistribution: Evidence from the Life in Transition Survey. *European Journal of Political Economy*, 34, 300-314.
- [37] Corneo, G., & Grüner, H. P. (2002). Individual preferences for political redistribution. *Journal of Public Economics*, 83(1), 83-107.
- [38] Cruces, G., Perez-Truglia, R., & Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, 98, 100-112.
- [39] Deffains, B., Espinosa, R., & Thöni, C. (2016). Political self-serving bias and redistribution. *Journal of Public Economics*, 134, 67-74.
- [40] Durante, R., Putterman, L., & Van der Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4), 1059-1086.
- [41] Ederer, F., & Manso, G. (2013). Is pay for performance detrimental to innovation? *Management Science*, 59(7), 1496-1513.
- [42] Engelmann, D., & Strobel, M. (2004). Inequality Aversion, Efficiency and Maximin Preferences in Simple Distribution experiments. *The American Economic Review*, 94(4), 857-869.
- [43] Esarey, J., Salmon, T. C., & Barrilleaux, C. (2012). What Motivates Political Preferences? Self-Interest, Ideology, and Fairness in a Laboratory Democracy. *Economic Inquiry*, 50(3), 604-624.
- [44] Falk, A., & Ichino, A. (2006). Clean Evidence of Peer Effects. *Journal of Labor Economics*, 24(1), 39-57.
- [45] Fehr, E., Kirchsteiger, G., & Riedl, A. (1993). Does fairness prevent market clearing? An experimental investigation. *The Quarterly Journal of Economics*, 108(2), 437-459.
- [46] Fehr, E., Naef, M., & Schmidt, K. M. (2006). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *The American Economic Review*, 96(5), 1912-1917.
- [47] Ferrer-i-Carbonell, A., & Ramos, X. (2014). Inequality and happiness. *Journal of Economic Surveys*, 28(5), 1016-1027.
- [48] Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- [49] Fong, C. (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics*, 82(2), 225-246.

- [50] Frank, R. H. (1984). Are workers paid their marginal products?. *The American Economic Review*, 74(4), 549-571.
- [51] Gächter, S., & Thöni, C. (2010). Social comparison and performance: Experimental evidence on the fair wage–effort hypothesis. *Journal of Economic Behavior & Organization*, 76(3), 531-543.
- [52] Gee, L. K., Migueis, M., & Parsa, S. (2017). Redistributive choices and increasing income inequality: experimental evidence for income as a signal of deservingness. *Experimental Economics*, 20, 894-923.
- [53] Gill, D., & Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *The American Economic Review*, 102(1), 469-503.
- [54] Gill, D., & Prowse, V. (2014). Gender differences and dynamics in competition: The role of luck. *Quantitative Economics*, 5(2), 351-376.
- [55] Gneezy, U., Goette, L., Sprenger, C., & Zimmermann, F. (2017). The limits of expectations-based reference dependence. *Journal of the European Economic Association* 15(4), 861-876.
- [56] Godechot, O., & Senik, C. (2015). Wage comparisons in and out of the firm. Evidence from a matched employer–employee French database. *Journal of Economic Behavior & Organization*, 117, 395-410.
- [57] Greiner, B., Ockenfels, A., & Werner, P. (2011). Wage transparency and performance: A real-effort experiment. *Economics Letters*, 111(3), 236-238.
- [58] Gross, T., Guo, C., & Charness, G. (2015). Merit pay and wage compression with productivity differences and uncertainty. *Journal of Economic Behavior & Organization*, 117, 233-247.
- [59] Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy*, 61(5), 434-435.
- [60] Herbst, D., & Mas, A. (2015). Peer effects on worker output in the laboratory generalize to the field. *Science*, 350(6260), 545-549.
- [61] Hennig-Schmidt, H., Sadrieh, A., & Rockenbach, B. (2010). In search of workers' real effort reciprocity—a field and a laboratory experiment. *Journal of the European Economic Association*, 8(4), 817-837.
- [62] Hibbs Jr, D. A., & Locking, H. (2000). Wage dispersion and productive efficiency: Evidence for Sweden. *Journal of Labor Economics*, 18(4), 755-782.
- [63] Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7(3), 346-380.
- [64] Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 7, 24-52.
- [65] Hong, H., Ding, J., & Yao, Y. (2015). Individual social welfare preferences: An experimental study. *Journal of Behavioral and Experimental Economics*, 57, 89-97.

- [66] Johnson, R. M., Reiley, D. H., & Muñoz, J. C. (2015). “The war for the fare”: How driver compensation affects bus system performance. *Economic Inquiry*, 53(3), 1401-1419.
- [67] Kiatpongsan, S., & Norton, M. I. (2014). How much (more) should CEOs make? A universal desire for more equal pay. *Perspectives on Psychological Science*, 9(6), 587-593.
- [68] Konow, J. (2001). Fair and square: the four sides of distributive justice. *Journal of Economic Behavior & Organization*, 46(2), 137-164.
- [69] Konow, J. (2003). Which is the fairest one of all? A positive analysis of justice theories. *Journal of Economic Literature*, 41(4), 1188-1239.
- [70] Krawczyk, M. (2010). A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution. *Journal of Public Economics*, 94(1), 131-141.
- [71] Ku, H., & Salmon, T. C. (2013). Procedural fairness and the tolerance for income inequality. *European Economic Review*, 64, 111-128.
- [72] Kube, S., Maréchal, M., & Puppe, C. (2013). Do wage cuts damage work morale? Evidence from a natural field experiment. *Journal of the European Economic Association*, 11(4), 853-870.
- [73] Kuhn, P., & Lozano, F. (2008). The Expanding Workweek? Understanding Trends in Long Work Hours among US Men, 1979–2006. *Journal of Labor Economics*, 26(2), 311-343.
- [74] Kuziemko, I., Norton, M. I., Saez, E., & Stantcheva, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *The American Economic Review*, 105(4), 1478-1508.
- [75] Lazear, E. P. (1986). Salaries and piece rates. *Journal of Business*, 59(3), 405-431.
- [76] Lazear, E. P. (1989). Pay equality and industrial politics. *Journal of Political Economy*, 97(3), 561-580.
- [77] Lazear, E. P. (1991). Labor economics and the psychology of organizations. *The Journal of Economic Perspectives*, 5(2), 89-110.
- [78] Lefgren, L. J., Sims, D. P., & Stoddard, O. B. (2016). Effort, luck, and voting for redistribution. *Journal of Public Economics*, 143, 89-97.
- [79] Lemieux, T., MacLeod, B., & Parent, D. (2009). Performance Pay and Wage Inequality. *The Quarterly Journal of Economics*, 124(1), 1-49.
- [80] Levine, D. I. (1991). Cohesiveness, productivity, and wage dispersion. *Journal of Economic Behavior & Organization*, 15(2), 237-255.
- [81] Mas, A. (2006). Pay, reference points, and police performance. *The Quarterly Journal of Economics*, 121(3), 783-821.
- [82] Mas, A., & Moretti, E. (2009). Peers at work. *The American Economic Review*, 99(1), 112-145.
- [83] Michelbach, P. A., Scott, J. T., Matland, R. E., & Bornstein, B. H. (2003). Doing Rawls justice: An experimental study of income distribution norms. *American Journal of Political Science*, 47(3), 523-539.

- [84] Mitchell, G., Tetlock, P. E., Mellers, B. A., & Ordóñez, L. D. (1993). Judgments of social justice: Compromises between equality and efficiency. *Journal of Personality and Social Psychology*, 65(4), 629.
- [85] Norton, M. I., & Ariely, D. (2011). Building a better America – One wealth quintile at a time. *Perspectives on Psychological Science*, 6(1), 9-12.
- [86] Nosenzo, D. (2013). Pay secrecy and effort provision. *Economic Inquiry*, 51(3), 1779-1794.
- [87] Okun, A. M. (1975). *Equality and Efficiency: The big Tradeoff*. Washington DC: The Brookings Institution.
- [88] Oxoby, R. J., & Spraggon, J. (2008). Mine and yours: Property rights in dictator games. *Journal of Economic Behavior & Organization*, 65(3), 703-713.
- [89] Paarsch, H., & Shearer, B. (1999). The Response of Worker Effort to Piece Rates: Evidence from the British Columbia Tree-Planting Industry. *Journal of Human Resources*, 34(4), 643-667.
- [90] Piketty, T. (1995). Social mobility and redistributive politics. *The Quarterly Journal of Economics*, 110(3), 551-584.
- [91] Piketty, T. (2014). *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.
- [92] Piketty, T., & Saez, E. (2003). Income inequality in the United States, 1913–1998. *The Quarterly Journal of Economics*, 118(1), 1-41.
- [93] Piketty, T., & Saez, E. (2014). Inequality in the long run. *Science*, 344(6186), 838-843.
- [94] Piketty, T., & Zucman, G. (2014). Capital is back: Wealth-income ratios in rich countries 1700–2010. *The Quarterly Journal of Economics*, 129(3), 1255-1310.
- [95] Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA.
- [96] Rubin, J., Samek, A. & Sheremeta, R.M. (2017). Loss aversion and the quantity-quality tradeoff. *Experimental Economics*. <https://doi.org/10.1007/s10683-017-9544-1>
- [97] Ruffle, B. J. (1998). More is better, but fair is fair: Tipping in dictator and ultimatum games. *Games and Economic Behavior*, 23(2), 247-265.
- [98] Schildberg-Hörisch, H. (2010). Is the veil of ignorance only a concept about risk? An experiment. *Journal of Public Economics*, 94(11), 1062-1066.
- [99] Scott, J. T., Matland, R. E., Michelbach, P. A., & Bornstein, B. H. (2001). Just deserts: an experimental study of distributive justice norms. *American Journal of Political Science*, 749-767.
- [100] Sliwka, D., & Werner, P. (2017). Wage Increases and the Dynamics of Reciprocity. *Journal of Labor Economics*, 35(2), 299-344.
- [101] Stiglitz, J. E. (1975). Incentives, risk, and information: notes towards a theory of hierarchy. *The Bell Journal of Economics*, 6(2), 552-579.
- [102] Traub, S., Seidl, C., & Schmidt, U. (2009). An experimental study on individual choice, social welfare, and social preferences. *European Economic Review*, 53(4), 385-400.

Figure 1 – Average Output by Wage Level: BASELINE Condition

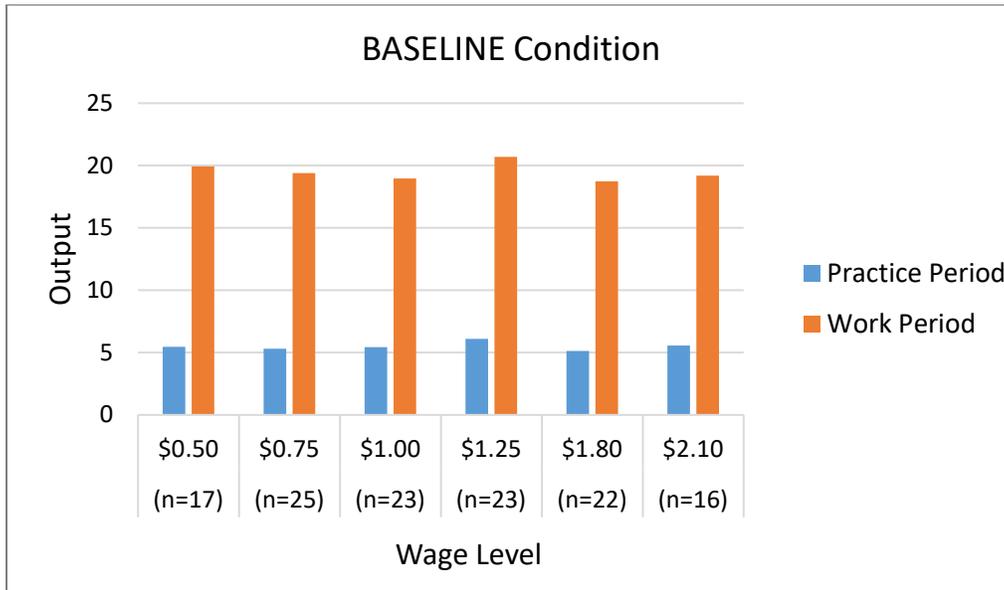


Figure 2 – Breakdown of Wage Distribution Preferences in EARNED Condition

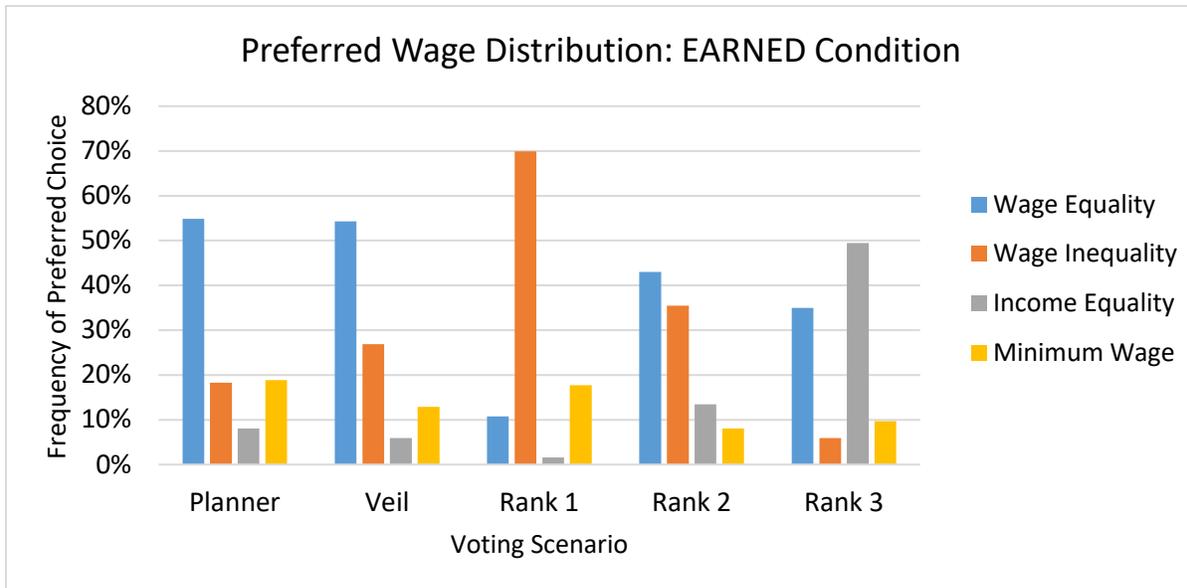


Figure 3 – Breakdown of Preferred Wage Distribution: RANDOM Condition

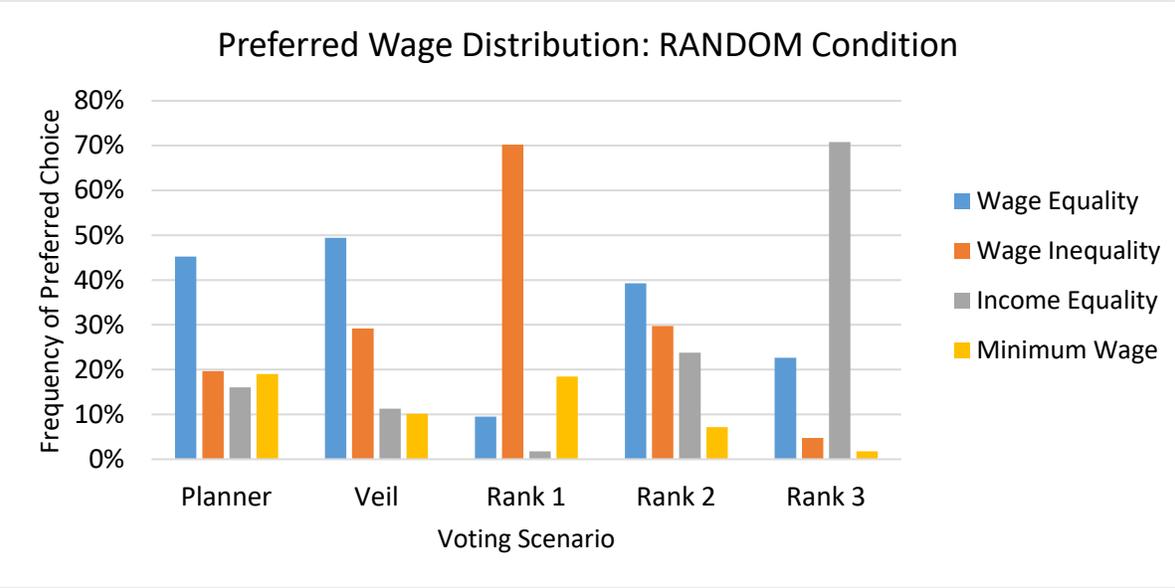


Table 1 – Possible Wage Distributions

Worker Ranking	<u>Wage Distributions</u>			
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum wage</i>
Rank 1	\$1	\$2.1	\$.75	\$1.8
Rank 2	\$1	\$1	\$1	\$.75
Rank 3	\$1	\$.5	\$1.25	\$.75
Average Wage	\$1	\$1.2	\$1	\$1.1

Table 2 – Aggregate Summary Statistics for Productivity by Condition

<i>Panel A: Practice Work Period</i>				
Productivity	<u>Experimental Condition</u>			
	BASLINE (n = 126)	EARNED (n = 186)	RANDOM (n = 168)	
Mean	5.5	5.6	5.2	
Median	5	6	5	
Min	3	3	0	
Max	10	9	9	
<i>Panel B: Paid Work Period</i>				
Productivity	<u>Experimental Condition</u>			
	BASLINE (n = 126)	EARNED (n = 186)	RANDOM (n = 168)	
Mean	19.4	20.1	20.1	
Median	19	20	20	
Min	11	12	10	
Max	30	30	29	
Avg. Earnings	\$23.3	\$22.5	\$21.75	

Table 3 – Preference Consistency across Voting Scenario by Condition

<i>Panel A – EARNED Condition</i>					
	<u>Voting Consistency</u>				
	Rank 1	Rank 2	Rank 3	Veil	Planner
Rank 1	100%	46%	21%	41%	31%
Rank 2	46%	100%	44%	51%	54%
Rank 3	21%	44%	100%	40%	40%
Veil	41%	51%	40%	100%	61%
Planner	31%	54%	40%	61%	100%

<i>Panel B – RANDOM Condition</i>					
	Rank 1	Rank 2	Rank 3	Veil	Planner
Rank 1	100%	35%	10%	36%	29%
Rank 2	35%	100%	37%	52%	40%
Rank 3	10%	37%	100%	30%	27%
Veil	36%	52%	30%	100%	49%
Planner	29%	40%	27%	49%	100%

Table 4 - Breakdown of Preferred Wage Distribution for *Non-Selfish* Sample

<i>Panel A – EARNED Condition</i>						
<u>Sample of Non-Selfish Voters (n = 106)</u>						
Wage Distribution	Rank 1	Rank 2	Rank 3	Veil	Planner	Total
<i>Wage Equality</i>	19%	51%	61%	63%	71%	53%
<i>Wage Inequality</i>	47%	29%	10%	18%	11%	23%
<i>Income Equality</i>	3%	6%	11%	2%	6%	5%
<i>Minimum Wage</i>	31%	14%	17%	17%	12%	18%

<i>Panel B – RANDOM Condition</i>						
<u>Sample of Non-Selfish Voters (n = 70)</u>						
	Rank 1	Rank 2	Rank 3	Veil	Planner	Total
<i>Wage Equality</i>	23%	47%	54%	59%	43%	45%
<i>Wage Inequality</i>	29%	17%	11%	20%	24%	20%
<i>Income Equality</i>	4%	19%	30%	7%	17%	15%
<i>Minimum Wage</i>	44%	17%	4%	14%	16%	19%

Table 5 – Average Output in Paid Work Period by Condition and Worker Ranking

<i>Panel A: EARNED Condition</i>					
	<u>Wage Distribution</u>				
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	<i>p – value</i>
Avg. Rank 1 Output	22.31	22.08	22	21.3	<i>p = .881</i>
Avg. Rank 2 Output	19.44	20.31	19.2	20.6	<i>p = .719</i>
Avg. Rank 3 Output	18.13	18.5	17.2	19.1	<i>p = .582</i>
Avg. Worker Output	19.96	20.29	19.47	20.33	<i>p = .736</i>
Avg. <i>Practice</i> Total	5.63	5.61	5.53	5.67	
# of Workers	48	78	30	30	
# of Groups	16	26	10	10	
<i>Panel B: RANDOM Condition</i>					
	<u>Wage Distribution</u>				
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	<i>p – value</i>
Avg. Rank 1 Output	20.58	19.55	19.67	22.11	<i>p = .266</i>
Avg. Rank 2 Output	17.17	20.05	20.13	23	<i>p = .028</i>
Avg. Rank 3 Output	18.25	20.35	19.53	22.78	<i>p = .046</i>
Avg. Worker Output	18.66	19.98	19.78	22.63	<i>p < .001</i>
Avg. <i>Practice</i> Total	5.25	5.00	5.24	5.60	
# of Workers	36	60	45	27	
# of Groups	12	20	15	9	

Table 6 – Output Quality: Comparison of *Inaccuracy Rates*

Panel A: EARNED Condition

	<u>Wage Distribution</u>				<i>p</i> – value
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	
Rank 1 Avg.	10.7%	10.1%	12.1%	10.7%	<i>p</i> = .967
Rank 2 Avg.	13.7%	13.2%	11.4%	10.3%	<i>p</i> = .897
Rank 3 Avg.	15.7%	15.3%	7.1%	17.3%	<i>p</i> = .400
Overall Group Avg.	13.3%	12.9%	10.2%	12.8%	<i>p</i> = .732
# of Workers	48	78	30	30	
# of Groups	16	26	10	10	

Panel B: RANDOM Condition

	<u>Wage Distribution</u>				<i>p</i> – value
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	
Rank 1 Avg.	9.5%	9.6%	16.7%	12.4%	<i>p</i> = .225
Rank 2 Avg.	5.2%	15.5%	11.7%	12.0%	<i>p</i> = .072
Rank 3 Avg.	7.0%	11.1%	19.7%	16.9%	<i>p</i> = .069
Overall Group Avg.	7.2%	12.1%	16.0%	13.8%	<i>p</i> = .008
# of Workers	36	60	45	27	
# of Groups	12	20	15	9	

Table 7 – Output Quality: Comparison of *Formatting Error Rates*

Panel A: EARNED Condition

	<u>Wage Distribution</u>				<i>p</i> – value
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	
Rank 1 Avg.	18.1%	61.5%	70.3%	36.0%	<i>p</i> = .009
Rank 2 Avg.	36.6%	30.0%	31.3%	40.0%	<i>p</i> = .931
Rank 3 Avg.	47.0%	44.0%	20.5%	51.2%	<i>p</i> = .441
Overall Group Avg.	33.9%	45.2%	40.7%	42.4%	<i>p</i> = .630
# of Workers	48	78	30	30	
# of Groups	16	26	10	10	

Panel B: RANDOM Condition

	<u>Wage Distribution</u>				<i>p</i> – value
	<i>Wage Equality</i>	<i>Wage Inequality</i>	<i>Income Equality</i>	<i>Minimum Wage</i>	
Rank 1 Avg.	36.5%	26.2%	60.8%	35.0%	<i>p</i> = .195
Rank 2 Avg.	17.8%	40.7%	57.0%	44.4%	<i>p</i> = .213
Rank 3 Avg.	51.7%	48.4%	68.1%	55.6%	<i>p</i> = .696
Overall Group Avg.	35.3%	38.5%	62.0%	45.0%	<i>p</i> = .043
# of Workers	36	60	45	27	
# of Groups	12	20	15	9	

Table 8 - Productivity under One's Preferred Wage Distribution

<i>Panel A – EARNED Condition</i>			
	Output Level	Total Accuracy Error	Total Formatting Error
Rank 1	2.065*** (0.793)	-0.344 (0.239)	0.147 (0.242)
Rank 3	-2.064*** (0.725)	0.123 (0.241)	0.382 (0.232)
Rank 1 * PWD	-0.581 (0.835)	0.275 (0.339)	-0.452 (0.503)
Rank 2 * PWD	-0.656 (0.851)	-1.115*** (0.320)	-0.402 (0.412)
Rank 3 * PWD	0.535 (0.953)	-0.800** (0.338)	-2.243*** (0.494)
N	186	3,734	3,734
<i>Panel B – RANDOM Condition</i>			
	Output Level	Total Accuracy Error	Total Formatting Error
Rank 1	0.543 (0.873)	-0.069 (0.222)	0.027 (0.251)
Rank 3	0.210 (0.873)	0.310 (0.333)	0.258 (0.269)
Rank 1 * PWD	-1.237 (1.148)	-1.565** (0.759)	0.343 (0.473)
Rank 2 * PWD	-0.508 (1.250)	-0.871*** (0.318)	0.331 (0.411)
Rank 3 * PWD	-0.609 (1.148)	-0.088 (0.442)	-0.002 (0.446)
N	168	3,372	3,372

Notes: The Output Level regression employs OLS at the subject level. The Error regressions assume a Poisson distribution for the total number of errors (of the specified type) at the application level. Application-level regressions use robust standard errors clustered at the subject level. Ranks are Boolean variables that equal 1 if the subject was the specified rank, and 0 otherwise. PWD (Preferred Wage Distribution) is a subject-level Boolean variable that equals one if the realized wage distribution matches the vote of the subject in the voting scenario corresponding to his or her realized rank, and 0 otherwise. Rank 2 is omitted as a comparison group. Included rank terms can be interpreted as the difference from Rank 2 when PWD=0. Interaction terms can be interpreted as the effect of PWD for the specific rank. (*), (**), and (***) indicate the coefficients are statistically significant at the 10%, 5%, and 1% levels, respectively.