Incentive Effects of Inequality and Economic Development*

Hyejin Ku†
Florida State University

Timothy C. Salmon‡
Florida State University

July 2009

Abstract

The effect of inequality on economic development is often debated in the development literature with many finding that inequality among different groups in a society has a negative impact on economic growth while others find the opposite result. Our study investigates a behavioral phenomenon which may not only be associated with the existence of such inequality but may also significantly exacerbate any adverse consequences of inequality for growth. In particular we investigate whether or not individuals exhibit a discouragement effect in the face of inequality that leads to lower work effort. If such an effect exists it provides a mechanism for converting even idiosyncratic inequality into sustained inequality with adverse consequences for the individuals being affected by the inequality and the economy as a whole. We investigate this phenomenon using an economic experiment to allow us to cleanly vary the nature of inequality and to allow us to directly observe several characteristics of the workers. We find robust support for the existence of a discouragement effect lending credibility to the claims that such an effect would exist in external situations among workers confronted with disadvantageous inequality.

JEL Codes: C90, D61, D63, J24, O15, O40
Key Words: inequality, efficiency, productivity, experiment

1 Introduction

Economists have established many ways in which initial inequality may influence the paths of economic development. In a series of historical studies, Engerman and Sokoloff (1997, 2002 and 2005) argue that the initial differences in factor endowments between the North and South America contributed to the emergence of different institutions between the two regions, which in turn led to the divergence in the rates of growth. While the typical finding in the growth literature may be a negative correlation between inequality and growth in

*The authors would like to thank the National Science Foundation for research support in funding these experiments.

†Florida State University, Department of Economics, 113 Collegiate Loop, Room 263, PO Box 3062180, Tallahassee, FL 32306-2180, hku@fsu.edu. Phone: 850-644-7208 Fax: 850-644-4535.

‡Florida State University, Department of Economics, 113 Collegiate Loop, Room 263, PO Box 3062180, Tallahassee, FL 32306-2180, tsalmon@fsu.edu. Phone: 850-644-7207 Fax: 850-644-4535.

a cross-section of countries, there appear to be no conclusive findings on the relationship between growth and inequality from recent studies based on panel data analysis.[1] Due to the importance of this relationship between inequality and economic growth in the design of a broad range of economic development policies there is a need to provide a better understanding of the mechanics behind some of the linkages between growth and inequality. There is a substantial theoretical literature which proposes a number of different linkages between growth and inequality including Loury (1981), Galor and Zeira (1993), Banerjee and Newman (1993), Durlauf (1996), Benabou (2000) and Mookherjee and Ray (2003). This literature is aimed at doing two things. Some of these papers attempt to explain a link between inequality and growth while others focus mainly on attempting to explain how initial inequality may lead to persistent inequality that remains in the long run. By and large the linkage between inequality and either growth or persistent inequality is modeled as being based on differential investment opportunities of agents typically in the form of human capital and occupational choice.

We propose an additional linkage between inequality and growth that can also be useful in explaining the persistence of inequality which is based on a behavioral response individuals may exhibit when confronted with inequality. If there are members of a group who experience lesser opportunities for earnings and advancement than members of other more advantaged groups then one might expect that the disadvantaged individuals will exhibit some degree of discouragement from their situation and exhibit a variety of responses ranging from decreased on the job work effort, decreased human capital acquisition, decreased contributions to public goods and in general engaging in a lower level over a range of other activities that are important for individual as well as societal advancement. In the context of a wage inequality, Akerlof and Yellen (1990) propose a similar phenomenon as a theoretical construct which they refer to as the "fair-wage effort hypothesis" and it is described by stating that "workers have a conception of a fair wage; insofar as the actual wage is less than the fair wage, workers supply a corresponding fraction of normal effort."

If such a behavioral response to inequality exists then there are a number of consequences associated with it. First, this effect *ceteris paribus* would generate a negative relationship between inequality and economic growth. In general there are a number of forces at play in an aggregate economy so this is not to claim that the overall effect must be negative if this effect exists but the partial effect of inequality on growth due to this behavioral effect would be negative. Second, separately or in addition to the theoretical models cited above this

[1] Benabou (1996) and Banerjee and Duflo (2003) provide a comprehensive review of the literature on the issue of inequality and growth.

effect could generate persistent inequality from initial idiosyncratic inequality. The reason is that even if the initial inequality were idiosyncratic, if a worker responds by exerting lesser effort or investing less in human capital, then this response has the potential to legitimize and then perpetuate the initially low opportunities. Thus initial idiosyncratic inequality can be converted into longer term inequality. Any potential discouragement effect is, of course, likely to be worsened if the inequality in opportunity is related to a more generalized social phenomenon of race, ethnic, gender or caste based discrimination in which individuals from these disadvantaged groups expect to be confronted with unequal treatment in many aspects of their lives.

There is of course also the possibility of a counter effect which we might refer to as an "encouragement effect" in which those who benefit from the inequality may work harder in response to their advantaged position. The ultimate question of empirical interest is if either of these two effects exist and if both do what the net effect is on overall productivity. The goal of the present study is to examine exactly this issue in order to determine the likely effect on overall productivity due to the existence of inequality. We present the results of a laboratory based experimental study designed to allow clean inference on whether or not individuals become discouraged or encouraged in the face of inequality. The environment we construct in our experiment should be a very strong test of the discouragement effect in particular because many of the aspects of discrimination and long run experience with the inequality will not be present in the experiment. Consequently, if we find a discouragement effect in our simplified laboratory environment then it will be robust evidence that such an effect could exist in a broader context in which the inequality is based on overt discrimination and individuals would have long term experience with feeling the effects of the inequality.

There is prior evidence on the issue but both the sign and robustness of any effect is unclear. One of the main reasons for this lack of clarity is due to the difficulty of identifying the discouragement effect in naturally occurring data. Much of the work in this area has been focused on identifying an effect due to wage inequality in workplaces. While there may be unfairness in naturally occurring wage schedules it is difficult to separate cleanly between a wage that is unfairly low and one that is deservedly low due to skill differences or to traits which would be unobserved by a researcher. The summary of the attempt to do this for identifying race based discrimination in wages against blacks in the US provided in Heckman, Lyons, and Todd (2000) explains many of the inherent problems in such attempts. Proper inference on the effects of inequality on effort though requires cleanly severing the link between wages and characteristics of a worker to find situations in which it is clear that "unfair inequality" is present. Were we able to find situations in which

wages are clearly unequal based on something other than traits of workers related to their productivity then there is also an empirical difficulty in observing work effort. Effort itself is rarely observable and the measurement of most proxies for work effort in common use (rate of promotion, turnover etc. . .) are potentially related to any discrimination that generated the wage inequality which adds structural bias to the data and limits inference. Even were a viable proxy for effort to be available, there is still a problem of separating out effort differences due to pure wage effect and effort differences due to the existence of inequality. Separating these effects requires observing workers exerting effort when faced with the same wage in similar environments that differ only in regard to whether or not wage inequality is present. Due to these difficulties in inference from naturally occurring data, we present evidence from a laboratory experiment in which we will be able to observe effort directly while implementing exogenously set wage schedules with and without inequality to overcome these limits on inference.

In our experiment workers will be exogenously assigned wages as a proxy for general opportunity and they will then engage in a real effort task (encoding random sequences of letters into numerical code) for piece rate earnings. The subjects will also have an alternative task which they can engage in for very low earnings to represent an outside option which might be very loosely interpreted as a "shirking" option. This eliminates a potential confound present in many laboratory real effort tasks as with nothing else to do subjects could well engage in the main task out of boredom. Our experiment addresses another important confound which is necessary to distinguish between a wage effect and an inequality effect. Simply observing that lower paid workers worked less would not be conclusive in demonstrating an inequality effect as it could simply be due to the incentive effect of lower wages. Consequently our design will include control sessions in which all subjects receive equal wages so that we can identify the inequality effect by comparing the performance of workers receiving a particular wage in sessions with no inequality to that of workers receiving the same wage in the presence of inequality.

Our experimental design will also address a less commonly investigated issue regarding how the response to inequality might vary with the relative sizes of the advantaged and disadvantaged groups. The relative group size may be important because the existence of an incentive effect hinges on workers perceiving a low wage as being unfair. That perception of unfairness may be diminished if the majority of workers are receiving the same low wage and may be heightened if only a small number do. To examine this possibility our experiments will vary the size of the high and low wage groups to determine if there is a systematic response of individuals to the status of their groups as being a minority or majority.

In examining the data from the experiments we will be interested in the overall effects on productivity as this will inform us on the efficiency aspects of inequality but we will also be looking at distributional issues. As argued in Ravallion (2001) it is important to look "beyond the averages" to understand the effects on subpopulations. Consequently, in addition to the overall average response we are also interested in determining if there are any specific groups who may respond more strongly to the existence of inequality than others. Towards that end we will be obtaining some basic demographic information on our subjects as well as using a few simple ability tests to measure various forms of cognitive ability.

In the end we do find that a discouragement effect exists and it is of non-trivial size but we do not find evidence in favor of an encouragement effect. This finding is quite important because it provides a demonstration of one channel through which inequality can lead to a decrease in economic efficiency. While our experiment can certainly not provided a calibrated measure of how large such an effect might be external to the laboratory setting, the fact that we find any measurable discouragement effect in this laboratory setting that might be considered *a priori* as unlikely to generate much of a response, should lend substantial credibility to the claim that a similar effect is likely to exist in external situations in which discrimination is present and persistent.

In addition to this main finding we also show evidence in regard to the types of individuals who might be more likely to respond to inequality. While we find no differential response to inequality by gender we do find that individuals with differing cognitive capacities do respond in measurably different ways. We also find results suggesting that the source of the inequality can be an important factor, at least for some individuals. Previous research in Bolton, Brandts, and Ockenfels (2005) shows that some individuals will judge the fairness of an allocation based on the procedure used to generate it rather than just the allocation itself. We use what is sometimes considered a procedurally fair procedure for generating the inequality in our experiment (i.e. pure randomization) and find that subjects who we can measure separately as being more likely to view the allocation as unfair generate a much stronger response.

In section 2 we will provide an overview of our experiment design. Section 3 will present a series of hypotheses regarding what one might expect to observe in the experiments based on prior literature. In section 4 we will present our results and we will provide a concluding discussion in section 5.

2 Experimental Design

The base task in this experiment involves subjects taking random strings of 4-letter "words" and using a code key to translate those letters into a numerical code. The subjects were shown the encoding key and the string of letters on a screen and would enter the encoded version below the word. The same code key was used for the entire session. They had a button which would allow them to submit a word and after doing so they would be immediately given a new word. For every word they encoded correctly they were paid a piece rate wage which was exogenously set by the experimenters at the beginning of the session. This production task is very similar to the one used in Erkal, Gangadharan, and Nikiforakis (2008).

Each session consisted of 16 subjects and these 16 subjects were divided into two cells of 8 at random. Inside of those cells of 8, subjects were further randomly assigned into what we labeled "blue" and "green" groups. The subjects were informed of the color of their group before anything else happened in the experiment as an attempt to make it clear that group assignment was exogenous. The meaning of the group assignment was only explained later to the subjects and the actual meaning referred to the wage rate they would be assigned. Table 1 provides a summary of the experimental treatments as well as the number of observations of cells for each treatment. There are a total of four treatments contained in this design. There are two control treatments called Common Wage - Low and Common Wage - High in which all subjects received the same common wage regardless of group assignment. For consistency with the other treatments subjects were still divided such that there were 6 members of the blue group and 2 of the green in each treatment but the wages across groups were constant. The high wage was \$0.09 per correctly encoded word while the low wage was \$0.03. The other two treatments will be called the Disadvantaged Minority treatment (Minority) and the Disadvantaged Majority treatment (Majority). These two treatments introduce inequality by having one group possess the high wage and the other the low such that in the Minority treatment, the members of the 2-person group are assigned the low wage while in the Majority treatment the members of the 6-person group are assigned the low wage.

The experiment was programmed using z-Tree, Fischbacher (2007). After subjects saw an initial screen indicating the color of the group to which they were assigned, they were presented with a sample of the main screen for the experiment showing them the encoding task. In the course of explaining this screen to the subjects, they were explicitly informed of the wage rates that would be in effect for both groups. Also, before each round of

Table 1: Summary of experimental design.

	Common Low	Common High	Disadvantaged Minority	Disadvantaged Majority
Blue Group (6 persons)	\$0.03	\$0.09	\$0.09	\$0.03
Green (2 persons)	\$0.03	\$0.09	\$0.03	\$0.09
Number of Cells	4	2	8	8

production, subjects would see a screen which included a table showing them a column for each subject in their 8-subject cell indicating their group color and corresponding wage rate. The idea of stating this information to them repeatedly was to ensure that they clearly understood both the wage rate differential as well as the relative size of both groups. After the first round, this screen also showed them information on their own past earnings. They were not shown the earnings or production levels of other individuals in the experiment at any time. The only information they see about other subjects is their wage rates. As such there are no interactions between members of a group or members of a cell making each subject independent of the others. This choice of feedback was made so that the only things that should be salient to the subjects that might affect their behavior are the treatment variables themselves and any session or group effects should be minimal.

In order to give subjects an outside option should they wish to avoid the production task, we included another task on their screen. This other task was the option of playing Tic-Tac-Toe (TTT) against the computer. The computerized opponent was programmed to be moderately difficult but beat-able. This task was only minimally incentivized in that it paid a subject \$0.01 per win. Due to the difference between this wage rate as well as the time it would take to win a game and the wage rate and time to correctly encode a word for the encoding task, it should have been quite clear to subjects that TTT would never compete in financial terms with the main task. It was designed mainly to be at least mildly more "fun" than the encoding task and allow subjects who did not want to engage in encoding another activity so that they would not just have to stare at the screen in boredom.

There were 12 rounds of production in each session with each round lasting four minutes. Subjects were not instructed on which task to engage in but rather they were told how both worked, the wage rates of both and then told they were free to allocate their time between the two tasks as they wished.[2] At the end of each of the 12 rounds, each subject was presented with the screen summarizing her output in the encoding task, earnings from the

[2] Full instruction scripts are available from the authors upon request.

encoding task, earnings from TTT, and cumulative earnings. On the practice screen they were also allowed to practice TTT as well as the encoding task for a few times before moving on to the first paying round.

At the end of the 12 rounds we had subjects fill out a short demographic questionnaire and complete two short sets of questions intended to measure various aspects of cognitive decision making. The first of these is the Cognitive Reflection Test (CRT) described in Frederick (2005). This three question test is designed to determine the degree to which subjects engage in thoughtful and reflective versus quick and impulsive decision making. For example, one of the questions is:

> "A bat and a ball cost $1.10. The bat costs $1.00 more than the ball. How much does the ball cost?"

The most common quick answer one might come up with is $0.10 but this is clearly wrong as upon further reflection the correct answer is $0.05. We used this measure because those who score low on it may be more inclined to make impulsive decisions and that impulsiveness might well lead to exhibiting greater effects due to the treatments.[3]

We also used a second set of questions involving pattern matching problems in which we asked subjects to fill in the number that fits best in sequences such as: 3 6 9 12 (_). We had 10 sequences of this sort that varied in difficulty. This test should measure a general facility with numerate tasks (thus we will refer to this as measuring "numeracy") and therefore might correlate with performance on tasks external to the lab. The importance of this will be explored in the next section. Subjects received no payment for completing these questions and we gave them 100 seconds per test to complete as many of the questions as they could in the time frame. For the CRT the majority of the subjects finished well before the time limit while for the pattern matching test the time constraint was binding. Since speed was an issue in the main production task the fact that the constraint was binding for the pattern recognition task is not a problem since our goal was to measure performance under time pressure. For the CRT one could imagine that a time constraint could generate even less reflective thinking than normal but since the constraint was so rarely even close to binding this should not be a problem.

[3] Oechssler, Roider, and Schmitz (2008) investigate whether the CRT is good overall measure of cognitive ability and find that it is correlated with other aspects of decision making such as risk and time preferences as well as likelihood of making certain types of reasoning mistakes. While they do find some correlation with these other elements of decision making, our focus is on using the CRT for the purpose it was designed which is to discriminate between impulsive and reflective decision makers.

In addition to these cognitive measures we obtained inside the lab we also wanted to obtain some measure of cognitive performance or intellectual ability external to the lab. We therefore asked the subjects for permission to allow us to access their educational background data from the University registrar's office. The data we obtained included various performance measures including college GPA as well as total SAT/ACT scores.[4] Due to the potential lack of comparability of GPA across students of different majors we decided to use SAT/ACT scores for this purpose. By default we used a subject's SAT total score and for those students who took only the ACT we used the conversion tables provided by the ACT organization for converting the ACT scores into their SAT equivalent.[5]

We have conducted 11 sessions with 176 subjects generating the number of cells per treatment as noted in table 1. All subjects received $10 for showing up to the session and sessions last a little over an hour. Subjects earned on average $28.89 ($39.72 for high wage workers and $19.86 for low wage workers) including their show-up fee.

3 Hypotheses

As a way of providing a framework for evaluating the results of the experiments and explaining the relevant insights from prior literature on these issues we will provide a series of hypotheses regarding what one might expect to observe from these experiments. The most basic hypothesis serving as a motivation for these experiments can be thought of as a specialized version of the Fair Wage-Effort hypothesis as described in Akerlof and Yellen (1990) and we will use this as our first hypothesis:

> **Hypothesis 1 (Discouragement Hypothesis) -** *The productivity of low wage workers should be lower in the two treatment conditions involving inequality than in the control treatment without inequality.*

While Akerlof and Yellen (1990) state and examine this hypothesis mostly on theoretical grounds there are a number of prior empirical results which might be thought of as suggestive that this effect will exist. An immediate thought might be that the extensive literature on trust and reciprocity in the form of a trust game, Berg, Dickhaut, and McCabe (1995), or gift-exchange game, Fehr, Kirchsteiger, and Riedl (1993), would support this hypothesis but on a deeper investigation that is less clear. In the standard gift exchange game the only

[4] An important note about this is that we were unable to obtain disaggregated scores by sections of these tests. We only have the combined total scores.

[5] http://www.act.org/aap/concordance/index.html

players involved in a transaction are the worker and employer with the gift-exchange effect occurring when an employer gives the worker a high fixed wage and the worker responds with high effort as a reciprocal response to that gift. Any labor supply response in this relationship is a pure wage effect while the discouragement hypothesis is based on the idea that there is an inequality effect. For an inequality effect to exist the worker needs to see the wages of others in order to base a judgement of the fairness of the wage. There is at least one paper, Charness and Kuhn (2007), which does allow workers to see the wages of others in a gift exchange game and the authors find no effort difference between making the wages of others public versus private. These results can not really be interpreted as providing inference on this discouragement hypothesis, though, because the wage differentials subjects saw in those experiments could have been reasonably interpreted as fair. The reason is that the wages of the workers were set endogenously by a subject in the role of employer and the workers were heterogeneous with respect to the earnings their effort generated to the employer. So if another subject was receiving a higher wage there was a clear justification for that which is that the other subject may be one who generates more money per unit of effort to the employer. Burchett and Willoughby (2004) present data from a related experiment which also tests whether or not publicly announced and different compensation rates matter in that they had subjects engage in a real effort task and the subjects received randomly assigned wage structures for performing the same task. The researchers found that subjects generated quite different work efforts depending on whether information about the wages of others was public with those earning a low per unit or fixed wage decreasing their output while those with the high wage increased their output when the compensation information was public. While the researchers found clear evidence that public knowledge of the compensation schemes had an impact on effort, it is not certain to what the impacts can be attributed. The confound is due to the fact that some of the subjects received fixed wages which were substantially higher than the potential earnings of either of the piece rate wage earners which causes any judgement of unfairly unequal compensation among the piece rate earners to be unclear.

In addition to the laboratory evidence in favor of this hypothesis there is also field evidence regarding the claim but the results appear mixed. Brown, Gardner, Oswald, and Qian (2008) provides evidence that worker happiness is potentially negatively related to wage inequality. That study showed in part that the rate at which individuals quit their jobs is correlated with the level of positive skewness in the pay distribution which is consistent with the claim that employees at the low end of the wage distribution will quit (an extreme form of effort reduction perhaps) in response to wages that are low in a relative

sense to others in their firm even though they might be less likely to quit given the same wage at a firm with a less skewed distribution. On the other hand, Carr (2008) uses PSID data to examine these issues and finds little support for a claim that unequal earnings in a community impact any of his measures of workplace performance. The inference in both of these studies may be limited though due to many of the issues explained before regarding why we are pursuing a laboratory approach.

There is a separate strand of research involving laboratory data, naturally occurring data and data from field experiments all suggesting that group identity can effect work effort which also allows for the possibility of wage inequality to do so if there is a connection between group identity and wage inequality. Hoff and Pandey (2006) conduct a study in which members of high and low castes perform a production task with a common piece rate wage. When caste is not mentioned prior to the task, members of both castes perform equivalently but when caste membership is made salient members of the low caste generate lower production. The authors argue that the explanation for this result is that low caste members may have believed that their effort would be rewarded less than that of others due to prior experience with discrimination even though the rewards were in fact common.[6] This result is part of a larger literature on how social identity and organizational structure can effect behavior as explored in Akerlof and Kranton (2000) and Akerlof and Kranton (2005). There are multiple ways in which this result is important for the current study. The first important implication is that this result seems to confirm the claim that members of lower castes who feel discriminated against will exhibit a strong discouragement effect and in fact it can be so strong that even when the differential treatment is removed the effects may still persist. The additional implications will be explored in regard to two later hypotheses.

Taken as a whole, these various strains of literature build a compelling case that there is at a minimum an expectation that wage inequality could have a significant impact on productivity even if the effect has not been robustly demonstrated. In addition to the *discouragement* effect derived from the Fair Wage-Effort Hypothesis there is also the potential for an opposite *encouragement* effect which might influence the behavior of those earning high wages. Though this effect is certainly discussed much less in these literatures, the gift-exchange literature discussed above provides one justification for this hypothesis in that workers receiving a high wage may perceive it to be a "gift" and respond with high effort.

[6] The production task involved the subjects solving mazes and there were graders who determined if a maze had been correctly solved. So the expectation of differential returns could have been based on expected differential treatment from the grader eventhough caste status was never communicated to the graders.

Again, one must be careful to separate the pure wage effect from the inequality effect but if the existence of inequality enhances the view of the high wage as a gift then this could lead to the observance of an encouragement effect among the high paid workers. We mention this hypothesis mainly because of the economic efficiency implications which are that even if the discouragement effect exists then if this encouragement effect also exists it could counterbalance the discouragement effect leaving any negative efficiency consequences from the discouragement effect ameliorated or possibly even counterbalanced completely. This leads to our second hypothesis:

> **Hypothesis 2 (Encouragement Hypothesis) -** *The productivity of the high wage earners should be higher in the two treatment conditions involving inequality than in the control treatment without inequality.*

In discussing the possibility for this discouragement effect it is important to understand what forces might be generating it. One source that might at first seem like a possibility for generating it can be found in the literature on social preferences in which there are common findings supporting the claim that individuals dislike inequality as demonstrated theoretically in the model of inequality aversion proposed by Fehr and Schmidt (1999) and by the ERC model in Bolton and Ockenfels (2000) among others. To determine if inequality aversion can explain the discouragement effect we need to think about how to extend these models into the production domain because the existing models of inequality aversion deal only with preferences over final allocations and not decisions over wealth production. The most direct way to extend these models is to assume that inequality aversion manifests itself as a reduction in utility when final earnings among a total of I members of a group are unequal. Consider the utility function specified in equation 1 where w_i is the wage received by individual i, n_i is the productivity of individual i, $c(n_i)$ represents the cost to individual i of producing at the level of n_i, $h(\cdot)$ measures the level of inequality in final wealth among all the relevant agents, and β_i is a parameter indicating the degree to which agent i is averse to inequality.

$$u(w_1, n_1, ..., w_I, n_I) = w_i * n_i - c(n_i) - \beta_i h(w_1 * n_1, ..., w_I * n_I) \tag{1}$$

Cost of effort should clearly be increasing in effort, or $c' > 0$, while the effect of one individual's effort on the overall level of inequality is contingent on their earnings relative to the overall average. If we consider $h(\cdot)$ to just be a measure based on the variance in the earnings distribution, which is one reasonable way of measuring the degree of inequality,

then in general any individual whose earnings are below the average will decrease $h(\cdot)$ by increasing their own production,n_i, while anyone with earnings above average earnings will increase $h(\cdot)$ by increasing n_i. If we let if $\overline{WN}$ be the average earnings then the properties of $h(\cdot)$ are summarize in equations 2 and 3.

$$h' < 0 \text{ if } w_i n_i < \overline{WN} \tag{2}$$

$$h' > 0 \text{ if } w_i n_i > \overline{WN} \tag{3}$$

The effect of adding $h(\cdot)$ to the standard utility function is that for low wage earners higher productivity decreases expected inequality and this could be seen as offsetting the effort cost of increased productivity. This should lead to low wage earners producing at a higher level than if they were unconcerned about inequality in final earnings. High wage earners who are averse to inequality would face an additional cost from high production above a certain level in addition to any effort cost because increased productivity would worsen inequality in final earnings. Consequently they should be predicted to decrease work effort when faced with inequality in order to make the final wealth distribution more equitable. Predicting exactly how efforts would be chosen according to this model would require solving for a Nash equilibrium among all of the workers but since we have no control over the functional form of $c(\cdot)$ or $h(\cdot)$ nor do we know the distribution of β_i in this experiment our focus will be only on the comparative static effects of moving from a condition with no wage inequality to one with wage inequality.

For any approach to modeling inequality aversion similar[7] to equation 1, we will find behavioral predictions opposite to those in hypotheses 1 and 2 which demonstrates that this way of extending inequality aversion to production decisions is actually not consistent with the discouragement effect. The key point that this demonstrates is that where the concern for inequality is focused is important for the direction of the predicted effect of inequality. In the motivation behind the discouragement effect the concern is placed at the intermediate wage level while here the focus is on final wealth allocations. The fact that these constructions yield conflicting directional predictions is quite interesting and may help later in interpreting any behavioral response we may see due to inequality. Towards that end we can state a third hypothesis based on this model of inequality aversion.

[7] We note that ours is a primitive specification of inequality aversion which does not allow for such things as individuals caring more about inequality when their earnings are below average than when they are above. Augmenting the model to include aspects like this one or others would not change the fundamental comparative statics so long as utility is decreasing in inequality and inequality is measured based on some measure similar to the variance in final earnings.

Hypothesis 3 (Inequity Aversion Hypothesis) - *The productivity of high (low) wage earners should be lower (higher) in the two treatment conditions involving inequality than in the control treatment without inequality.*

The previous hypotheses deal only with the effects of inequality itself and not with the group size treatments. Our intention with the group size treatments was to provide different contexts through which subjects might view the inequality. From the point of view of the low wage subjects, one might think that if they were almost singled out to be in the minority of the subjects in the experiment to receive a low wage that this might be viewed as less fair than were they just one among many who received the low wage. Consequently, if the group size differential matters we are led to make our fourth hypothesis:

Hypothesis 4 (Group Status Hypothesis) - *The productivity will be more responsive to the existence of inequality in the Disadvantaged Minority Treatment than the Disadvantaged Majority Treatment.*

The main motivation behind this hypothesis is to help understand the nature of any discouragement effect we might find and the specific prediction if the discouragement effect is operative is that if low wage workers view their position as less fair when they are in the minority then there will be a larger discouragement effect in that treatment. Due to the differential effects of a high/low wage workers effort on the overall level of earnings inequality in the different group configurations one can also construct an argument in which the inequality aversion effects are similarly magnified.

In addition to these broad treatment effects, we also envision that different groups might be more likely to exhibit a different response to the presence of inequality than others. In particular, groups who might be more familiar with expecting that their returns from effort would be rewarded less than that of others should be more likely to exhibit a response to inequality that those less familiar with that position. This potential effect is drawn from the results in Hoff and Pandey (2006) which showed that low caste members exhibited a discouragement effect in their study potentially related to their prior experience with discrimination. This suggests that in our setting in which we are inducing differential wages that members of groups who have experienced work force or wage discrimination may be more likely to exhibit responses to the existence of those differential wages. Due to the size and make-up of our sample we are unable to test for this effect along ethnic or racial lines which might be more likely to generate an effect but we can test for it along gender lines which leads to our fifth hypothesis:

Hypothesis 5 (Gender Hypothesis) - *Women will exhibit a stronger discouragement effect than men.*

The basis for hypothesis 5 is the presumption that many women may feel as though women are compensated less than men for equal effort as this is reported quite often in the popular news media**.** To the extent that women in our sample may have internalized this past discrimination we may find that they are more likely to be discouraged in the face of inequality similar to the members of the low caste in Hoff and Pandey (2006).

Our next hypothesis is related to the performance on the CRT measure of the degree to which our subjects engage in reflective versus impulsive decision making. The tendency of a person to think deeply about matters related to the experiment is important to the potential judgement of whether or not the wage assignment is unfair. The wage rate assignments are unequal which could lead to an initial response by impulsive decision makers of judging that inequality to be "unfair." More reflective decision makers though might think more deeply about the procedure used to assign those wage rates, pure randomization, and view the procedure itself as fair and then judge the outcome of that procedure to also be fair. The notion that individuals might judge the fairness of a final outcome on the basis of the procedure which generated it is discussed in detail in Bolton, Brandts, and Ockenfels (2005) in which the authors provide evidence that some individuals will view the outcome of a procedurally fair process as fair even if the outcome of the process is inherently unfair while others will judge the fairness solely on the basis of the outcome itself. Since looking past the allocation itself to the process requires a more reflective thought process, it is a reasonable claim that subjects exhibiting more reflective decision making tendencies as measured by the CRT will be less likely to view the unequal wage allocations as unfair while those measured to be more impulsive will be more likely to see the allocation as unfair. Given that a behavioral response to the wage inequality is predicated on the judgement that the inequality is to some degree unfair, this reasoning leads to our sixth hypothesis.

Hypothesis 6 (CRT Hypothesis) - *Subjects who score lower on the CRT and are therefore measured to be more impulsive should exhibit a stronger discouragement effect than the subjects scoring higher on the CRT.*

Our final two hypotheses are strongly related and are based on the results from the pattern matching questions that were designed to measure the facility of our subjects with numerate tasks and the subjects' SAT score. To the extent to which these tasks can be claimed to measure something about the cognitive skills of our subjects, the scores on

them can be used as a general measure of cognitive ability and those with differing levels of cognitive ability may be expected to respond differently to inequality for a variety of reasons. One way of generating a testable hypothesis on this issue is to presume that subjects who are generally less able may also be used to facing situations in which they expect less reward for a similar level of effort as others and this may predispose them to exhibiting a discouragement effect. We use these two different approaches to measuring cognitive ability because they may measure different aspects of cognitive capacity. We find that this is possible as the two measures turn out to have a relatively low correlation of 0.125.[8] Our final two hypotheses can be stated as follows:

Hypothesis 7 (Numeracy Hypothesis) - *Subjects who score lower on the numeracy measure should exhibit a stronger discouragement effect than those subjects scoring higher on the measure.*

Hypothesis 8 (SAT Hypothesis) - *Subjects who score lower on the SAT should exhibit a stronger discouragement effect than those subjects scoring higher on the SAT.*

4 Results

4.1 Data Overview

As an initial look at the data from the experiment, table 2 displays some raw summary statistics regarding average per round output and average number of times per round workers chose to play TTT by treatment with low wage and high wage workers broken out separately. At this level of aggregation it appears that the general directional effects predicted by hypotheses 1 and 2 are present in that for low (high) wage workers the average productivity is lower (higher) in the two sessions with inequality. The other interesting point about aggregate productivity is that if there is a wage effect it appears to be negative because high wage workers produce less than low wage workers in the common wage treatments. We can also note that while neither low nor high wage workers play TTT with great frequency, there is some tendency for low wage workers to play more often (though with high variance) and high wage workers to play less when confronted with inequality.

[8] One might be surprised at the low correlation but the likely reason is that while a subject's score on the math section of the SAT should be expected to be highly correlated with the numeracy measure their score on the verbal section should be less strongly correlated. Since we have only the combined SAT score this likely explains the diminution of the correlation.

Table 2: Summary Statistics

		Common Wage	Disad Minority	Disad Majority
Low Wage Workers	Output	28.05 (8.65)	27.54 (8.37)	26.59 (8.36)
	TTT	1.58 (4.62)	1.35 (3.95)	1.31 (4.06)
High Wage Workers	Output	26.39 (9.25)	27.58 (8.16)	28.27 (6.56)
	TTT	1.49 (2.66)	0.99 (2.52)	0.43 (1.55)

Standard errors in parentheses.

In addition to the basic summary statistics we can examine figures 1-5 for a look at any differences over time and whether different subsets of the population respond differently to the existence of inequality. Figure 1 shows the average output per round separated out by treatment and by high wage versus low wage earners for the entire sample. The first characteristic to note is that there is definitely a learning trend by the subjects as output increases over time due to subjects gaining greater facility with the task and perhaps memorizing portions of the code key. Further, we also see that for low wage earners, the increase in productivity appears to be less in the treatments with inequality than in the control treatment without inequality while the reverse is true for the high wage workers.

Figures 2-5 display similar looks at the data but with the subjects separated into groups by gender, CRT score, Numeracy score or SAT score. The gender separation should be self-explanatory and for the cognitive measures we have separated the subjects into high and low scoring groups on each measure to investigate whether these different groups respond differently to the treatments. A subjects is placed in the low CRT category if they gave 0 correct answers on the CRT instrument. If they had any correct answers they are placed in the high CRT category. Out of our 176 subjects 57 got at least 1 answer correct on the CRT and the other 119 had 0 correct answers.[9] For the Numeracy and SAT measures, we have put those scoring below the median in the low category on that measure and those above the median are in the high category. These breakdowns of the figures are a way of visualizing any treatment effects that may differ across groups as suggested by several of our hypotheses. In these figures we see many of the same characteristics of the overall data such as the overall upward trend but we also see several cases in which there appears to

[9] We note that this is lower performance than found in Frederick (2005), but in our case this test was administered after a long and tedious experiment and so some performance drop is to be expected. Since all we really wish to do is separate out the most deliberate thinkers from the others, this should not be much of a problem.

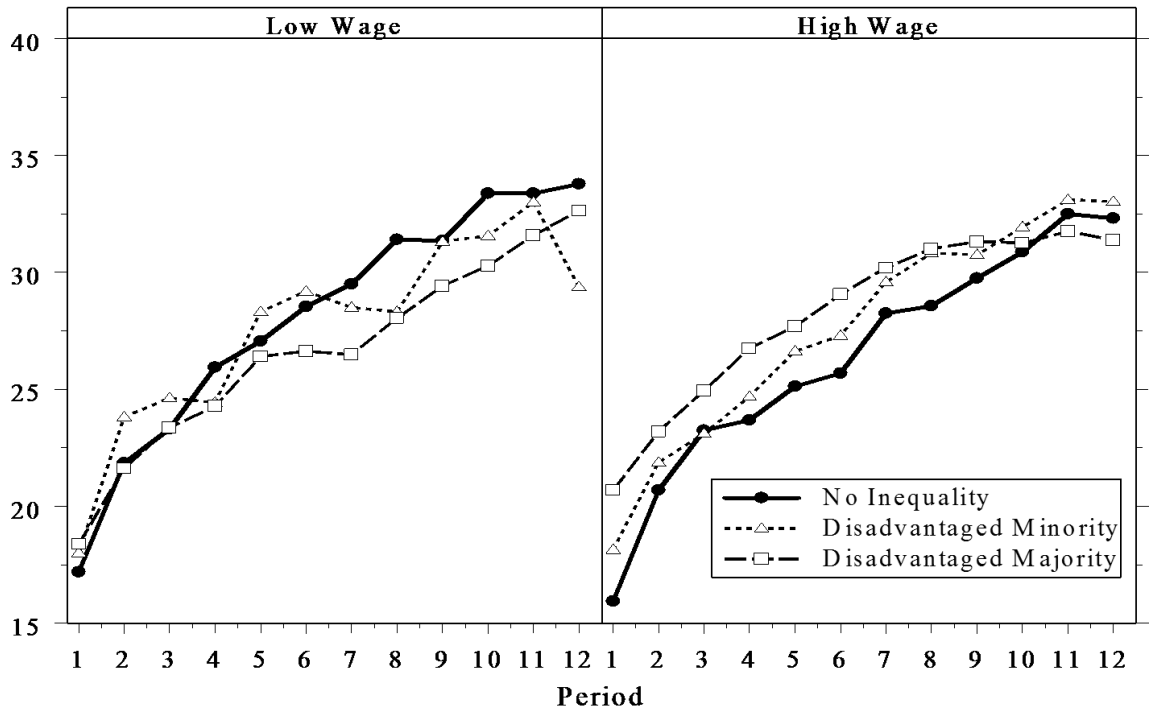

Figure 1: Average production per round.

be potential discouragement and encouragement effects. We also see some indications that the existence and or strength of these effects may differ across groupings. We of course can not draw hard conclusions from these visual inspections though so we will refrain from further discussion of any particular effects until we provide the statistical analysis in the next sections.

Prior to engaging in the more careful tests of the hypotheses though it is also useful to examine some baseline results on the determinants of productivity that will be important for framing the later analysis. Table 3 contains random effects panel regressions with robust standard errors clustered on the subject of output per period on a dummy variable for the wage rate, a dummy variable for whether or not the subject is in the small group and then all of the demographic and cognitive control measures we will be using in later regressions.[10] It also includes a similar set of regressions with the dependent variable equal to 1 if the subject chose to play TTT in a round. The only data considered in these regressions is derived from

[10] A linear trend is included in the regressions to account for the change in performance over time. We note that while the time trend in the data as seen in figures 1-5 is not exactly linear, we have conducted this and all later regressions we will show using dummy variables for each period instead and the interpretations remain the same. We have chosen to present results from the more parsimonious specification.

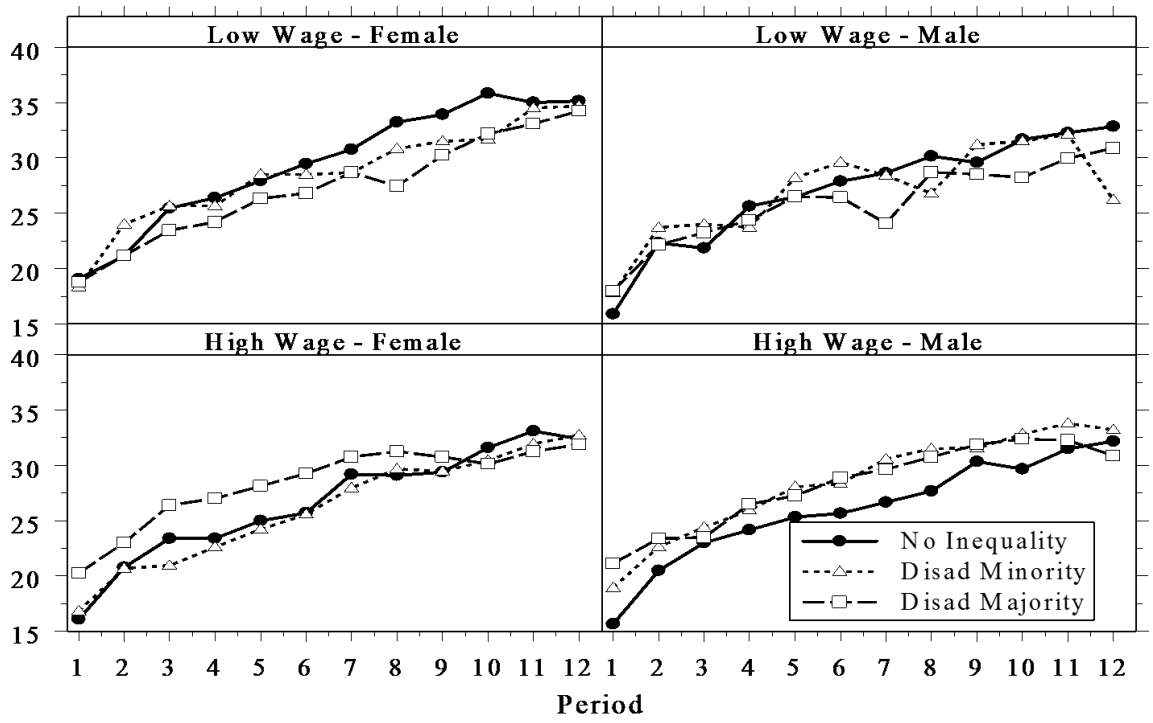

Figure 2: Average production over time by gender.

subjects in the common wage control treatments as our goal here is to obtain some baseline measures on the determinants of overall performance. The question these regressions are intended to address is whether any of the demographic and cognitive measures we are using have any substantial impact on productivity or propensity to play TTT. For the cognitive measures we will be using binary measures to indicate whether or not a subject has been categorized as a member of the low group on each measure. Thus L-CRT is equal to 1 if a subject has been sorted into the low CRT group and the same for L-Numeracy and L-SAT.

There are two important notes about the cognitive measures. While it has already been pointed out that L-Numeracy and L-SAT have a low correlation, it is also the case that neither of these variables is strongly correlated with L-CRT. The correlation between L-CRT and L-Numeracy is .027 while the correlation between L-CRT and L-SAT is 0.128. This lack of correlation is important because it demonstrates that each of these cognitive measures is measuring some different aspect of cognitive ability. We must also point out that while we have L-CRT and L-Numeracy scores for all of our subjects, there were 10 subjects in our overall sample for whom we were unable to obtain SAT or SAT-equivalent scores. Consequently when we include L-SAT in our regressions our sample size will decrease. We

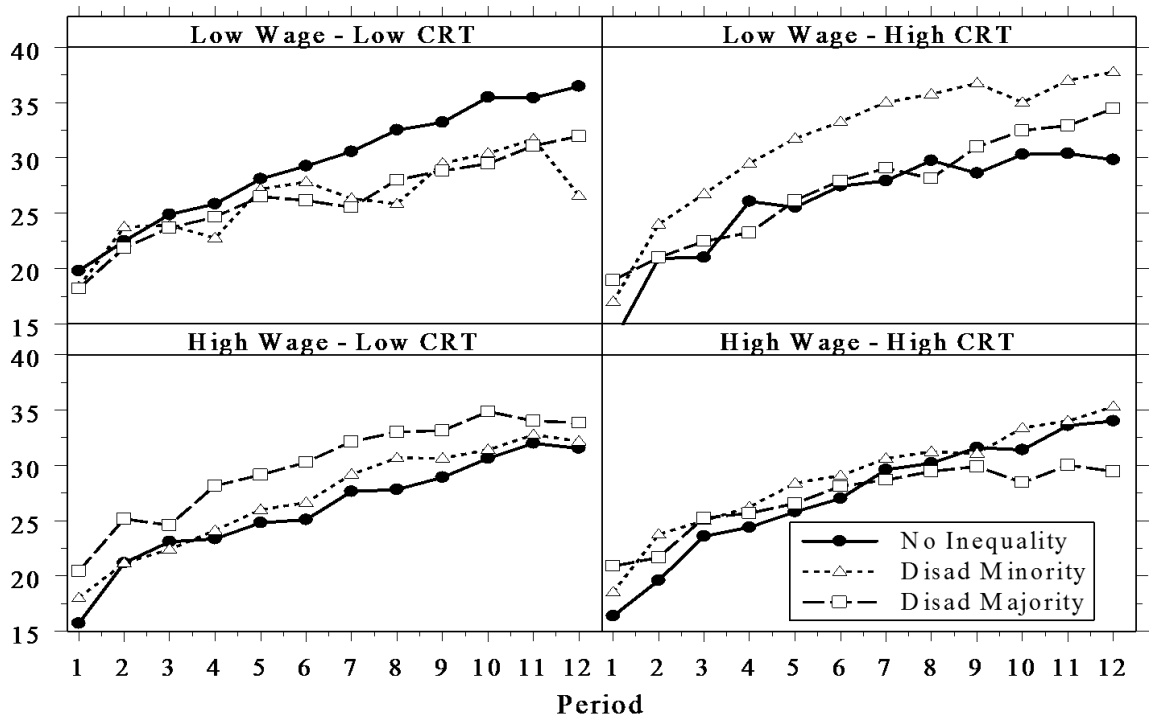

Figure 3: Average production over time by CRT.

have conducted the other regressions not using the L-SAT variable with the lower sample size as well as a robustness check and find the difference to minimally affect any of our interpretations so we have chosen to present the versions with the larger sample size.

The results from the regressions in table 3 demonstrate that for the most part, the demographic and cognitive measures have little impact on the performance of our subjects in the production task. There is evidence that NonWhite subjects may have been less likely to play TTT than White subjects but this difference had no significant impact on productivity. In particular we note that the four variables that we base hypotheses on, Female, L-CRT, L-Numeracy and L-SAT, are all uncorrelated with productivity. These results, of course, have no direct bearing on our stated hypotheses because the hypotheses concern how individuals of different types respond to inequality and the data from the sessions with inequality is not included in these regressions. What these regressions demonstrate is that none of these characteristics can be seen as proxies for ability on this specific task. Were any of our measures such as numeracy or gender correlated with performance on this task then should we see a greater discouragement effect for low numeracy or female subjects and these groups also had lower productivity then it would not be clear if the discouragement effect was

Table 3: Determinants of productivity and probability of playing TTT in common wage treatments.

Dep Var:	Output per Subject Per Period				Binary for if subject Played TTT			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High Wage	-1.669	-2.270	-2.196	-2.270	0.086	0.068	0.100	0.069
	(2.027)	(2.016)	(2.091)	(2.135)	(0.085)	(0.081)	(0.090)	(0.090)
Small Group	3.826**	3.520**	3.084	2.725	-0.125	-0.092	-0.129	-0.106
	(1.718)	(1.779)	(1.919)	(1.953)	(0.078)	(0.080)	(0.092)	(0.091)
L-CRT			1.179	0.512			-0.003	-0.008
			(1.836)	(2.075)			(0.080)	(0.081)
L-Numeracy			-1.320	-1.741			-0.026	-0.015
			(1.936)	(2.165)			(0.085)	(0.099)
L-SAT			-0.426	-2.516			-0.023	-0.042
			(1.855)	(2.224)			(0.082)	(0.102)
Female		1.534		1.828		0.073		0.099
		(1.933)		(1.955)		(0.073)		(0.080)
Age		3.430		25.931		-0.105		0.875
		(3.000)		(19.924)		(0.109)		(0.933)
Age^2		-0.079		-0.635		0.002		-0.022
		(0.058)		(0.501)		(0.002)		(0.023)
NonWhite		1.234		1.492		-0.173**		-0.182**
		(1.668)		(1.690)		(0.069)		(0.085)
Upperclass		0.491		-0.743		-0.118		-0.146*
		(1.797)		(2.171)		(0.075)		(0.088)
Period	1.374***	1.374***	1.362***	1.362***	-0.047***	-0.047***	-0.046***	-0.046***
	(0.111)	(0.112)	(0.114)	(0.114)	(0.006)	(0.006)	(0.006)	(0.006)
Constant	18.170***	-19.780	18.931***	-243.568	0.612***	1.862	0.627***	-7.900
	(1.057)	(36.303)	(2.027)	(197.641)	(0.063)	(1.324)	(0.097)	(9.255)
Observations	576	576	552	552	576	576	552	552
Clusters	48	48	46	46	48	48	46	46

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

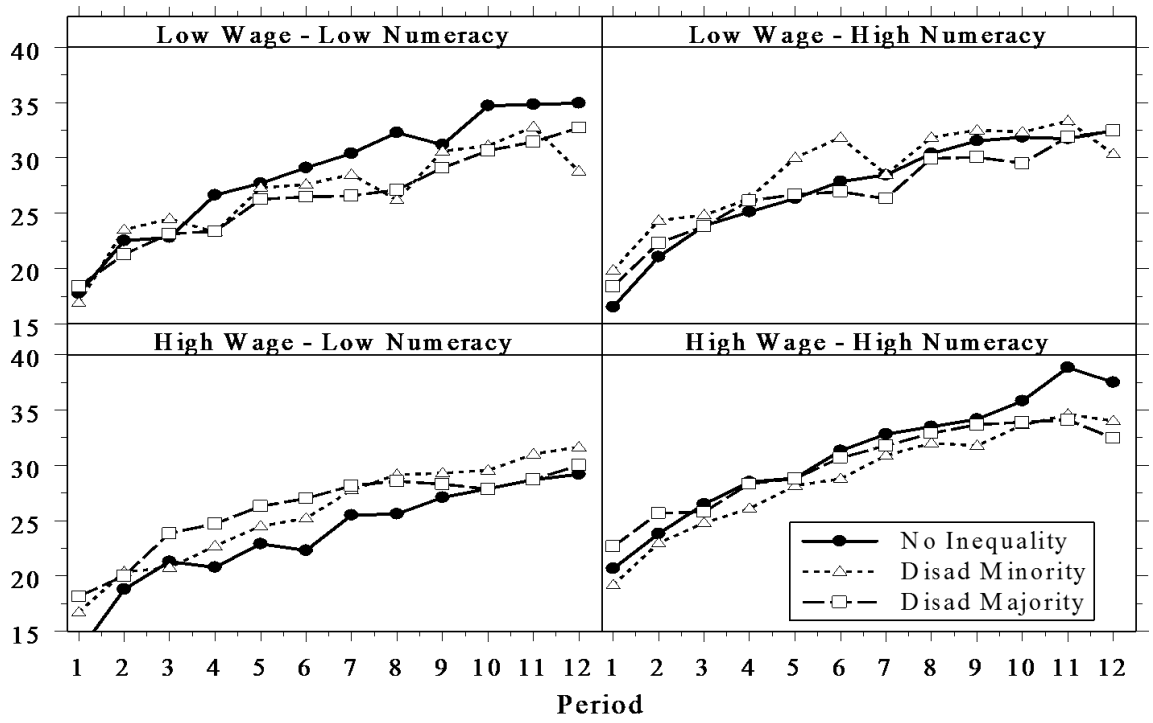

Figure 4: Average production over time by numeracy.

related to a lack of skill or facility with this task or to the other traits we are attempting to measure. Since productivity is in general not correlated with any of our cognitive or demographic characteristics this means that this task is well suited to our purposes.

4.2 Overall Effects of Inequality

The first way we will examine our data to determine what effects may exist due to the existence of inequality will involve looking at the response to inequality over the entire range of the experiment. To pursue this approach we will present a series of regressions to determine whether or not low (high) wage workers in the treatments with inequality produce at a different rate over the entire experiment than low (high) wage workers in the baseline sessions. To investigate our hypotheses suggesting that there may be differential responses to the treatments in different sub-samples, we include additional regressions with the treatment dummy variables interacted with the characteristics we use as the basis for these hypotheses. In each regression we will present here we use data from low and high wage earners separately and will use a standard set of our controls in each regression. All of the regressions conducted use the output of each subject in each period as the dependent variable

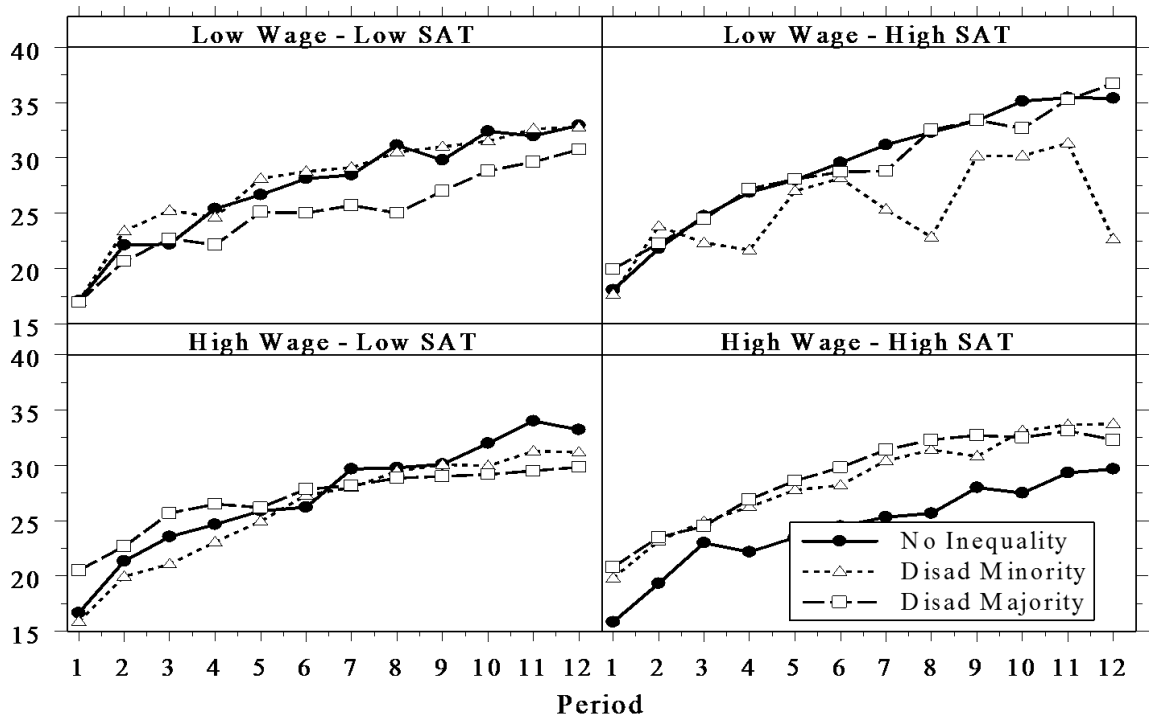

Figure 5: Average production over time by SAT.

and are conducted as panel regressions using a subject level random effects specification with the standard errors clustered by subject. All regressions include a linear trend (i.e. the current period) to account for the general trend observed in figures 1-5 and a constant term. While the observed trend is not exactly linear, all of our results are robust to using separate dummy variables for each period to account for the trend and so we present the more parsimonious specification. Due to the fact that the coefficients on the constant, the linear trend and on the some of the demographic measures are of no intrinsic interest, we will suppress the actual values of the coefficients on those controls to conserve space. This allows us to concentrate on the key variables of interest for inference.

The full set of these regressions is shown in table 4. Columns 1 and 6 of this table examine whether any differences that might have been observed in figure 1 are statistically significant for either low or high wage earners. While the signs of the coefficients are consistent with the discouragement and encouragement hypotheses respectively, the effects in the overall sample were not strong enough over the entire time line of the experiment to be significant. The additional columns which interact these terms with the relevant characteristics for hypotheses 5-8 also indicate relatively weak overall responses to the in-

Table 4: Response of Low and High wage workers to inequality.

	Dependent Variable: Output									
	Low Wage Workers					High Wage Workers				
Char:	None	Female	L-CRT	L-Num	L-SAT	None	Female	L-CRT	L-Num	L-SAT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Disad Min	-0.518	-0.539	5.679*	1.764	-4.141	0.177	0.637	0.954	-3.535	1.045
	(1.716)	(2.308)	(3.150)	(3.300)	(3.851)	(1.918)	(2.469)	(2.538)	(2.500)	(2.970)
Disad Maj	-1.771	-1.745	0.582	-0.905	-1.304	0.791	0.799	0.251	-2.362	2.024
	(1.302)	(1.780)	(2.419)	(1.894)	(1.773)	(1.910)	(2.581)	(2.520)	(2.550)	(2.970)
Disad Min * Char		0.058	-8.727**	-3.767	4.969		-0.888	-1.091	6.136*	-1.312
		(3.359)	(3.689)	(3.735)	(4.594)		(3.591)	(3.477)	(3.410)	(4.177)
Disad Maj * Char		-0.053	-3.737	-1.560	-0.478		0.138	1.491	5.031	-2.293
		(3.069)	(2.836)	(2.569)	(2.908)		(3.855)	(3.867)	(3.946)	(3.913)
Char		2.048	2.319	0.430	-3.470		0.008	-0.837	-8.278***	0.485
		(2.459)	(2.238)	(2.066)	(2.296)		(3.143)	(3.077)	(2.959)	(3.743)
Observations	1152	1152	1152	1152	1032	960	960	960	960	888
Clusters	96	96	96	96	86	80	80	80	80	74

Robust standard errors clustered by individual are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

All regressions include a linear trend and age, age squared, Female, NonWhite, Upperclass, L-CRT, and LNUM. Columns (5) and (10) also include L-SAT.

equality treatments. We do find strong and statistically significant effects due to inequality in the Disadvantaged Minority treatment for low wage workers with a Low CRT score and for high wage workers with a low numeracy score. In these cases we find support for the discouragement and encouragement hypotheses respectively. As we explain next, though there are reasons to believe that while this is a reasonable specification used to examine the fundamental questions, it is perhaps not the best for uncovering any effects due to the treatments. We present these regression results therefore to provide some insight into the nature of the responses but we present our formal tests of our hypotheses in the next section based on a different regression approach.

4.3 Change in Inequality Effects Over Time

Due to the nature of the effects hypothesized it would seem reasonable to assume that any discouragement effect in particular would take time to develop rather than existing from the first period of the experiment. The discouragement effect is predicated on the notion that someone who is faced with a situation in which they are treated unequally will eventually be frustrated by this unfair treatment and will respond to that frustration by exerting lesser effort. This hypothesized effect may therefore require time to develop. Our experiments were designed around the possibility that any impacts of inequality would require time to emerge. This is why we had subjects produce for multiple periods and why we reminded the subjects of their wage rate as well as the wage rate of the others on every result screen. The experience of exerting effort round after round yet seeing that others are consistently better rewarded for their effort is exactly the feedback that one would suspect could lead to the development of a discouragement effect. It is less obvious how this experiment structure would foster or inhibit an encouragement effect, though. That is, any encouragement effect might be more of a response to the initial stimulus and then decay over time as one has acquired substantial enough earnings. On the other hand, after seeing the large reward continue round after round this could lead to an increase in the encouragement effect. The key point in either case is that the types of effects that we are investigating here are ones in which time is almost certain to play an important role. A visual inspection of figures 1-5 will confirm that in many cases there appear to be systematic differences in responses to the treatments over time.

There are a variety of ways that one might examine whether or not there are statistically significant treatments effects that build up over time. We have chosen a specification which involves beginning with a regression structure similar to what was used in the previous

section and then augmenting it to allow for a intercept shift in the second half of the experiment and then interacting that intercept shift with the two treatments. Consider an equation for the following form:

$$y_{it} = \alpha_0 + \alpha_1 D_1 + \alpha_2 D_2 + \alpha_3 (D_1 * SH) + \alpha_4 (D_2 * SH) + \alpha_4 SH + \alpha_5 Period + X_i\beta + c_i + \varepsilon_{it} \quad (4)$$

where D_1 and D_2 represent a treatment dummy variables, SH is a dummy variable indicating the second half of the experiment, X_i is a vector of control variables, c_i is an unobserved individual effect and ε_{it} is a standard error term. The value of α_4 will tell us whether or not in the common wage treatments there is any shift in behavior on average between the first half and second half of the experiment not already accounted for by the linear trend variable.[11] The values of α_1 and α_2 will tell us if there is any difference in the first half of the experiment between the common wage sessions and the treatments involving inequality. The values of α_3 and α_4 will tell us if there are any changes in treatment effects in the second half of the experiment. To investigate our hypotheses regarding differences in responses by sub-populations we will also include regressions with the same structure but will add on another level of interactions with the characteristic of interest, $Char$, as seen in equation 5.

$$\begin{aligned} y_{it} = \alpha_0 + \alpha_1 D_1 + \alpha_2 D_2 + \alpha_3 (D_1 * SH) + \alpha_4 (D_2 * SH) + \alpha_5 SH + \\ + \gamma_0 Char + \gamma_1 (D_1 * Char) + \gamma_2 (D_2 * Char) + \\ \gamma_3 (D_1 * SH * Char) + \gamma_4 (D_2 * SH * Char) + \gamma_5 (SH * Char) + \\ \alpha_5 Period + X_i\beta + c_i + \varepsilon_{it} \end{aligned} \quad (5)$$

We will use combinations of the coefficients on these base and interacted treatment dummies as the foundation for a series of results that represent formal tests of our hypotheses. Due to the complicated combinations of parameters necessary for inference we have constructed table 5 to contain the total treatment effects for the various groups in the first half of the experiment and the shift in the total treatment effect for that group that occurs in the second half of the experiment.[12]

The two base hypotheses underlying this study concern how low wage and high wage workers alter their productivity in the presence of inequality. Columns 1 and 6 in table 5 contain the treatment effects for low and high wage earners separately which allow us

[11] Again, these regressions were also conducted using dummy variables for each period to allow for a more general structure for any temporal effects and no differences in inference are found.

[12] The regressions containing the base coefficients behind the constructions in table 5 can be found in the appendix.

Table 5: Treatment effects from regressions examining response to inequality taking into account differences over time.

	Treatment Effects Over Time and by Characteristics									
	Low Wage Workers					High Wage Workers				
	None	Gender	CRT	Num	SAT	None	Gender	CRT	Num	SAT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
		Male, H-CRT, H-Num, H-SAT					Male, H-CRT, H-Num, H-SAT			
Disad Minority, α_1	0.751	0.786	4.634*	2.930	-1.481	0.210	0.512	1.703	-3.032	0.556
	(1.426)	(1.904)	(2.603)	(2.181)	(3.022)	(1.662)	(2.268)	(2.352)	(2.273)	(2.519)
Disad Majority, α_2	-0.843	-0.464	0.889	-0.089	-0.886	1.900	1.344	2.402	-1.019	1.771
	(1.151)	(1.746)	(2.025)	(1.667)	(1.494)	(1.575)	(2.258)	(2.110)	(2.298)	(2.425)
Disad Min*SH, α_3	-2.536*	-2.651	2.090	-2.333	-5.319**	-0.066	0.250	-1.500	-1.006	0.979
	(1.377)	(1.877)	(2.608)	(2.849)	(2.653)	(1.335)	(1.784)	(2.423)	(2.150)	(2.075)
Disad Maj*SH, α_4	-1.856*	-2.561**	0.987	-1.632	-0.836	-2.219	-1.090	-4.304*	-2.685	0.506
	(1.006)	(1.242)	(1.868)	(1.774)	(1.664)	(1.490)	(2.087)	(2.557)	(2.432)	(2.332)
		Female, L-CRT, L-Num, L-SAT					Female, L-CRT, L-Num, L-SAT			
Disad Minority, $\alpha_1 + \gamma_1$		0.627	-0.820	-0.598	1.224		-0.328	-0.410	2.318	-0.043
		(2.089)	(1.655)	(1.748)	(2.039)		(2.328)	(2.112)	(2.156)	(2.320)
Disad Majority, $\alpha_2 + \gamma_2$		-1.022	-1.531	-1.378	-0.880		2.471	1.907	3.771	1.838
		(1.700)	(1.391)	(1.618)	(1.933)		(2.236)	(2.421)	(2.200)	(2.142)
Disad Min*SH, $\alpha_3 + \gamma_3$		-2.216	-4.455***	-2.811**	-0.793		0.156	0.545	0.567	-0.447
		(1.722)	(1.382)	(1.289)	(1.638)		(1.864)	(1.578)	(1.729)	(1.762)
Disad Maj*SH, $\alpha_4 + \gamma_4$		-1.550	-3.248***	-2.174**	-1.804		-3.067	-0.331	-2.205	-4.213***
		(1.686)	(1.069)	(1.088)	(1.435)		(2.010)	(1.770)	(1.759)	(1.540)
Observations	1152	1152	1152	1152	1032	960	960	960	960	888
Clusters	96	96	96	96	86	80	80	80	80	74

Robust standard errors clustered by individual are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

All regressions include age, age squared, Female, NonWhite, Upper class, L-CRT, LNUM, linear trend and constant. Columns (1) and (5) also include L-SAT. Rows for second half show difference in output from first half.

to examine the validity of these two hypotheses for the overall sample. The key effects to examine for testing these hypotheses are the base treatment effects as well as the change in the effect between the first and second half of the experiment. Examining these effects leads to our first two results which correspond to tests of hypotheses 1 and 2.

> **Result 1 (Discouragement Hypothesis) -** *Workers facing disadvantageous inequality exhibit a larger drop in output in the second half of the experiment than workers receiving a low wage without the presence of high wage workers.*

> **Result 2 (Encouragement Hypothesis) -** *Workers facing advantageous inequality show no sensitivity in their output to the presence or absence of inequality over the entire duration of the experiment.*

The support for result 1 is based on the fact that the base treatment effects in column 1 of table 5 are not significant indicating that low wage workers evidence no significant initial response to the existence of inequality. However the rows showing the change in output between the first and second half of the experiment for both treatments contain values which are negative and significant. This is an indication that there is a significant discouragement effect but it takes time to develop.

Result 2 is supported by the regression in column 5 in table 5 which shows that neither the base treatment effects nor the changes in treatment effects between the first and second half are significant for those receiving the high wage. The base treatment effects are positive indicating there may be some initial encouragement effect but it is not strong enough to be significant and the negative values for the change in the second half indicate that any initial encouragement effect may dissipate over time in any case though this too is insignificant.

Our next result relates to a test of hypothesis 3, the Inequity Aversion Hypothesis. The support for this result is obvious since hypothesis 3 generated predictions exactly opposite of hypotheses 1 and 2. This leads to our third result which involves a failure to find support for hypothesis 3.

> **Result 3 (Inequity Aversion Hypothesis) -** *The productivity of low wage earners declines in the presence of inequality while that of high wage earners is unchanged in contrast to the predictions of the inequity aversion hypothesis.*

The next hypothesis that we will test is the group status hypothesis which states that the productivity response should be larger in the Disadvantaged Minority treatment than in the Disadvantaged Majority treatment. Since there are no significant productivity effects

for the high wage workers, this hypothesis is not applicable to them but our fourth result is based on testing this hypothesis for the low wage earners.

> **Result 4 (Group Status Hypothesis) -** *While the decrease in production in the second half of the experiment for workers facing disadvantageous inequality is larger (in absolute value) in the Disadvantaged Minority treatment than the Disadvantaged Majority treatment, these differences are not statistically significant.*

Since we find a discouragement effect in the second half of the experiment, we base the test for this hypothesis on whether $|DisadMin * SH| > |DisadMaj * SH|$ for the low wage workers. While we do find that the ordinal ranking holds, the test of whether or not the difference is significant reveals that it is not ($p-$value=0.651) leading to a failure to find evidence in support of this hypothesis. This indicates the possibility that individuals may judge the fairness of relative wages without substantial consideration for the relative sizes of the advantaged and disadvantaged groups or that our sample size was not large enough to allow us to establish the significance of the effect.

The test of the remaining hypotheses, require using the other columns in table 5 which separate out the effects of the treatment by groups. We note that each of the remaining hypotheses deal only with relative magnitudes of any discouragement effect between different groups which means the only relevant results are the ones for subjects receiving a low wage. We have still presented the high wage regressions for completeness but since our hypotheses are not concerned with these results and we find little in the way of systematic reactions by high wage subjects we will refrain from commenting on those results. The first hypothesis that we will test concerns whether or not there is a differential discouragement effect according to the gender of the subjects. The relevant treatment effects are in column 2 of table 5. Hypothesis 5 suggested that female subjects might exhibit more of a discouragement effect than the male subjects but our results suggest otherwise leading to our fifth result.

> **Result 5 (Gender Hypothesis) -** *The productivity of females receiving a low wage displays no statistically significant response to the presence of inequality and there is no statistically significant difference between the responses by females and males.*

In explaining the evidence for this hypothesis we first note that the results reported in table 5 reflect the total and change in the total treatment effects for both groups. Thus the

values reported for females indicate that for females there is no significant initial response to the inequality and there is no significant change over time. To explicitly test the stated version of the hypothesis though, one must compare the magnitude of the female response to the male response. When both effects are negative we then have to determine if the female response is larger in absolute value and then test whether or not that difference is significant. As should be clear by the results in the table, the only possible case for this hypothesis to hold would be the first half of the Disadvantaged Majority treatment as the female response is more negative than the male response in this case but the difference is not significant ($p-$value=0.829). In the other cases, were there any gender based response it would be that men were more discouraged than women but the differences are never significant. This finding suggests that the women in our sample do not see themselves as members of a low caste as did the subjects in the Hoff and Pandey (2006) experiment. An explanation of this might be due to the fact that while there is clear evidence that on average women receive lower wages than men, this gap has been declining over time, Blau and Kahn (2000), and much of the gap can be attributed to occupational choice rather than overt discrimination, Macpherson and Hirsch (1995). Further, more recent data show that more women are now attending college than men with this at least partially attributed to greater expected career prospects for women now than in decades past, Goldin, Katz, and Kuziemko (2006). Consequently, the college age women in our sample may be unlikely to have internalized much of the past discrimination and that would certainly explain why they fail to exhibit any greater sensitivity to inequality than men.

The next hypothesis we will check is the CRT hypothesis and the relevant effects are contained in column 3 of table 5 leading to our sixth result.

> **Result 6 (CRT Hypothesis) -** *Low CRT subjects exhibit a discouragement effect while High CRT subjects do not.*

We find strong support for the CRT hypothesis which was already indicated by the finding in the previous section that in the Disadvantaged Minority treatment the response to the inequality by the low CRT subjects was already strong enough to be detected in the overall regression. In these regressions allowing for differential effects over time, we find a strong drop in productivity in the second half of the experiments involving inequality for the Low CRT subjects but not for the high CRT subjects who have an insignificant increase in productivity. Consequently it is clear that we see stronger discouragement among low CRT subjects than high CRT subjects at least for the second half of the experiment.

If one accepts the claim that the CRT does measure the degree to which our subjects engage in reflective thinking then the differential effect observed between low and high CRT subjects leads to an interesting potential explanation. Since it is the more reflective thinkers who are more likely to consider the fairness of the allocation procedure in judging the fairness of the allocation, as did some of the subjects in Bolton, Brandts, and Ockenfels (2005), it seems the explanation for this differential effect is due to the low CRT subjects being more likely to view the differential wage assignment as unfair. This result is quite useful in helping to understand how our results may generalize outside of this laboratory environment.

The final two hypotheses we will test involve determining whether or not subjects of different levels of demonstrated cognitive ability, as measured by our Numeracy quiz and the subjects' SAT scores, exhibit differential discouragement effects. In particular, the hypotheses propose that the low ability subjects will exhibit a greater discouragement effect than the higher ability subjects. Column 4 of table 5 contains the relevant effects for the Numeracy measure while column 5 contains the effects for the SAT measure leading to our final two results.

> **Result 7 (Numeracy Hypothesis) -** *Low Numeracy subjects exhibit a discouragement effect while high Numeracy subjects do not.*

> **Result 8 (SAT Hypothesis) -** *Low SAT subjects do not evidence a statistically significant discouragement effect in either half of either treatment and their response is no different from that of the high SAT subjects.*

For result 7, we see that there are no significant effects in the high Numeracy subjects but the low Numeracy subjects exhibit a significant decline in productivity over the second half of the experiment. So this provides some support for the Numeracy hypothesis. The specific statement of the hypothesis does however require testing whether the drops observed in both groups in the second half are significantly different from one another and those tests reveal that there is no significant difference ($p-$values of 0.879 and 0.794). For the low SAT subjects we do not find any evidence of a discouragement effect and ordinally the drop in productivity is generally larger for the high SAT subjects which leads to a clear rejection of this hypothesis.

4.4 Inequality and Efficiency

The previous section focussed on testing the hypotheses regarding the behavioral responses of individuals to the presence of inequality. The ultimate reason that those responses are interesting is in how those individual behavioral responses aggregate into overall productivity. Due to the fact that we generally find support for the discouragement effect and fail to find support for the encouragement effect it should be clear that the overall effect we find from inequality on productivity and efficiency is negative.

To help interpret the magnitude of the effect of inequality on efficiency we can provide a few simple calculations of the efficiency drop implied by the coefficients from the regressions. In the overall sample we find that production in the second half of the experiment declines by 1.9-2.6 encoded strings per round for the low wage subjects. Given that in the Common Wage - Low sessions average productivity over the second half of the experiment was 32.13 strings per round per person, the decrease implied by the coefficients represents a 5.9-8.1% drop in productivity. Were we to have found support for the encouragement effect among the high wage workers then this drop might have been offset by increased productivity among high wage earners but we found that inequality had a neutral impact on the output of high wage workers. So the productivity of the high wage subjects does not counteract the productivity drop by low wage subjects.

In the results from the low CRT subjects we have a reasonable argument that these subjects were more likely to see the wage inequality as unfair and since what we really want to measure is the consequences on efficiency from inequality viewed as unfair then it is in the results from this sub-sample which might provide a more reliable benchmark for the efficiency consequences of inequality. The coefficients on the two interaction variables for the low CRT subjects receiving a low wage range from -3.25 to -4.46. In the Common Wage - Low treatment, the subjects in the same low CRT classification produced on average 33.95 encodings per period and so those coefficients suggest a 9.6-13.1% drop in productivity. The low CRT subjects earning a high wage evidence no significant response to the inequality so their response does not offset this productivity decline.

We do want to be clear that we are not proposing that the approximately 10% decline in productivity we find among low wage workers be interpreted as a reliable estimate of the magnitude of the discouragement effect that might exist among, for example, members of a low caste in India. The results from this experiment can certainly not provide an answer that specific or that well calibrated to any specific situation. What we claim is that since we were able to generate a discouragement effect of non-trivial size in this setting, then it

provides strong support that there should exist a discouragement effect, for example, among members of a low caste in India. This is therefore an important issue to consider in the design of policies aimed at enhancing economic development in countries with a substantial amount of discrimination based inequality. Combining this effect with the theoretical work cited in the introduction the indication is that such inequality is likely to become persistent and will lead to negative consequences for economic growth and efficiency.

5 Conclusion

The question that motivated this study concerns whether or not individuals who are faced with unequal earnings opportunities will respond to any perceived unfairness by those receiving lesser opportunities decreasing work effort. We find evidence that in our overall sample that while there is no initial response to the inequality, after our subjects have experienced the inequality for several rounds those receiving a lower wage begin exerting less effort than their counterparts in a control group with no inequality. For the subset of our sample who are measured to be impulsive decision makers through their score on the CRT, we find a much stronger discouragement effect which is consistent with the claim that those who are more likely to see the wage allocation as unfair exhibit a stronger discouragement effect.

These results represent clear and compelling evidence that in external situations when there are individuals faced with unequal opportunities which they deem unfair that they will also exhibit lower effort and the consequences of this behavioral response could be substantial. This discouragement effect could show up not only in on the job work effort but also in human capital investment and other activities which help a person advance into higher earnings groups. If the inequality in opportunity is not transient, then this discouragement effect could grow and combined with the forces described in Mookherjee and Ray (2003) and related studies can lead to persistent inequality and poverty. Thus this short term behavioral effect could have long term consequences for the initial generation of workers exhibiting the response as well as their descendents.

At the aggregate level, if there are populations of workers exhibiting lower work effort, investment in human capital and so on then there will certainly be negative impacts on economic growth and development. In our experiment we found output decreases of around 10% and while we will certainly not claim that this number is an accurate estimate of the magnitude of the discouragement effect in any other situation, it does suggest that the effects will be of non-trivial magnitude. The reason is that given the setup in this experi-

ment, there is every reason to suspect that subjects should be immune to a discouragement effect building up over such a short time horizon from purely idiosyncratic wage assignment for performing a trivial task. Given the strength of the response that we find, it seems reasonable that in situations in which opportunity inequality is permanent and tied to an individual's race or caste then a discouragement effect should be even more likely to occur and could well be stronger than what we measure here. Given the long run consequences to the individuals and the aggregate economy, this issue should be of concern to those designing policies to foster economic growth in countries in which such inequality exists.

An immediate follow-up question to the demonstration of this effect, though, is how robust it is to the inclusion of other elements which exist in labor markets such as the opportunity for advancement. The first point in response to that question is that the main environment to which our study is meant to apply are to situations of ethnic, racial, gender or caste based discrimination in which such inequality is (at least perceived to be) permanent and outside the control of any single worker to affect. For these situations, the lack of opportunity for the disadvantaged workers to be promoted to the advantaged group is entirely appropriate. Still there will be other related situations in which there may be some possibility of moving from the disadvantaged to the advantaged group and/or one may consider policies which specifically allow that as a means of overcoming the discouragement effect. This leads to a question of how much of an advancement opportunity is necessary to overcome any discouragement effects? This is not an issue considered in this paper but it is an important question to be investigated in future work.

References

Akerlof, G. A. and R. E. Kranton (2000). Economics and Identity. *The Quarterly Journal of Economics 115*(3), 715–753.

Akerlof, G. A. and R. E. Kranton (2005). Identity and the Economics of Organizations. *Journal of Economic Perspectives 19*(1), 9–32.

Akerlof, G. A. and J. L. Yellen (1990). The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics 105*(2), 255–283.

Banerjee, A. V. and E. Duflo (2003). Inequality and growth: What can the data say? *Journal of Economic Growth 8*, 269–299.

Banerjee, A. V. and A. F. Newman (1993). Occupational choice and the process of development. *Journal of Political Economy 101*(2), 274–298.

Benabou, R. (1996). Inequality and growth. *NBER Macroeconomics Annual 11*, 11–74.

Benabou, R. (2000). Unequal societies: Income distribution and the social contract. *American Economic Review 90*(1), 96–129.

Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, Reciprocity, and Social History. *Games and Economic Behavior 10*, 122–142.

Blau, F. D. and L. M. Kahn (2000, Fall). Gender differences in pay. *Journal of Economic Perspectives 14*(4), 75–99.

Bolton, G. and A. Ockenfels (2000). ERC: A Theory of Equity, Reciprocity and Competition. *American Economic Review 90*, 166–193.

Bolton, G. E., J. Brandts, and A. Ockenfels (2005). Fair procedures: Evidence from games involving lotteries. *Economic Journal 115*, 1054–1076.

Brown, G. D., J. Gardner, A. J. Oswald, and J. Qian (2008). Does Wage Rank Affect Employees' Well-Being? *Industrial Relations 47*(3), 355–389.

Burchett, R. and J. Willoughby (2004). Work productivity when knowledge of different reward systems varies: Report from an economic experiment. *Journal of Economic Psychology 25*(5), 591–600.

Carr, M. (2008). Income Inequality and the Overworked American. Working Paper U Mass-Amherst.

Charness, G. and P. Kuhn (2007, October). Does Pay Inequality Affect Worker Effort? Experimental Evidence. *Journal of Labor Economics 25*(4), 693–723.

Durlauf, S. N. (1996). A Theory of Persistent Income Inequality. *Journal of Economic Growth 1*, 75–93.

Engerman, S. L. and K. L. Sokoloff (1997). Factor endowments, institutions, and differential paths of growth among new world economies: A view from economic historians of the united states. In S. Haber (Ed.), *How Latin America fell behind: Essays on the economic histories of Brazil and Mexico, 1800-1914*. Stanford: Stanford University Press.

Engerman, S. L. and K. L. Sokoloff (2002). Factor endowments, inequality, and paths of development among new world economies. *conomia: Journal of the Latin American and Caribbean Economic Association 3*(3), 41–88.

Engerman, S. L. and K. L. Sokoloff (2005). Colonialism, inequality, and long-run paths of development. NBER Working Paper.

Erkal, N., L. Gangadharan, and N. Nikiforakis (2008). Competition, luck and social preferences: A real-effort experiment. Working Paper, University of Melbourne.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993). Does Fairness Prevent Market Clearing? An Experimental Investigation. *Quarterly Journal of Economics 108*, 437–459.

Fehr, E. and K. Schmidt (1999). A Theory of Fairness, Competition and Cooperation. *Quarterly Journal of Economics 114*(3), 817–868.

Fischbacher, U. (2007). z-Tree: Zurich Toolbox For Readymade Economic Experiments. *Experimental Economics 10*(2), 171–178.

Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives 19*(4), 25–42.

Galor, O. and J. Zeira (1993). Income distribution and macroeconomics. *Review of Ecomic Studies 60*(1), 35–52.

Goldin, C., L. F. Katz, and I. Kuziemko (2006, Fall). The homecoming of american college women: The reversal of the college gender gap. *Journal of Economic Perspectives 20*(4), 133–156.

Heckman, J. J., T. M. Lyons, and P. E. Todd (2000). Understanding black-white wage differentials, 1960-1990. *The American Economic Review Papers and Proceedings 90*(2), 344–349.

Hoff, K. and P. Pandey (2006). Discrimination, Social Identity and Durable Inequalities. *American Economic Review: Papers & Proceedings 96*(2), 206–211.

Loury, G. C. (1981). Intergenerational transfers and the distribution of earnings. *Econometrica 49*(4), 843–867.

Macpherson, D. A. and B. T. Hirsch (1995, July). Wages and gender composition: Why do women's jobs pay less. *Journal of Labor Economics 13*(3), 426–471.

Mookherjee, D. and D. Ray (2003). Persistent inequality. *Review of Economic Studies 70*, 369–393.

Oechssler, J., A. Roider, and P. Schmitz (2008). Cognitive abilities and behavioral biases. Working Paper, University of Heidelberg.

Ravallion, M. (2001). Growth, inequality and poverty: Looking beyond averages. *World Development 29*(11), 1803–1815.

Appendix A

Table 6: Base regressions underlying table 5.

	Dependent Variable: Output									
	Low Wage Workers					High Wage Workers				
Char:	None	Female	L-CRT	L-Num	L-SAT	None	Female	L-CRT	L-Num	L-SAT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Disad Minority	0.751	0.786	4.634*	2.930	-1.481	0.210	0.512	1.703	-3.032	0.556
	(1.426)	(1.904)	(2.603)	(2.181)	(3.022)	(1.662)	(2.268)	(2.352)	(2.273)	(2.519)
Disad Majority	-0.843	-0.464	0.089	-0.089	-0.886	1.900	1.344	2.402	-1.019	1.771
	(1.151)	(1.746)	(2.025)	(1.667)	(1.494)	(1.575)	(2.258)	(2.110)	(2.298)	(2.425)
Second Half	-0.316	-0.949	-1.390	-0.856	0.393	0.753	0.052	1.708	1.608	-0.979
	(0.673)	(0.818)	(1.310)	(1.247)	(0.928)	(1.045)	(1.546)	(2.174)	(1.929)	(1.801)
SH*Disad Min	-2.536*	-2.651	2.090	-2.333	-5.319**	-0.066	0.250	-1.497	-1.006	0.979
	(1.377)	(1.877)	(2.608)	(2.849)	(2.653)	(1.335)	(1.784)	(2.423)	(2.150)	(2.075)
SH*Disad Maj	-1.856*	-2.561**	0.987	-1.632	-0.836	-2.219	-1.090	-4.304*	-2.685	0.506
	(1.006)	(1.242)	(1.868)	(1.774)	(1.664)	(1.490)	(2.087)	(2.557)	(2.432)	(2.332)
Char		1.268	1.415	-0.079	-2.752		-0.553	-0.144	-7.594***	-0.621
		(1.926)	(1.776)	(1.707)	(1.886)		(2.661)	(2.571)	(2.517)	(3.115)
Char*Disad Min		-0.159	-5.454*	-3.528	2.706		-0.841	-2.112	5.350*	-0.599
		(2.826)	(3.118)	(2.798)	(3.776)		(3.132)	(3.036)	(2.965)	(3.513)
Char*Disad Maj		-0.558	-1.620	-1.289	0.006		1.126	-0.496	4.791	0.066
		(2.588)	(2.458)	(2.354)	(2.547)		(3.202)	(3.289)	(3.342)	(3.273)
SH*Char		1.559	1.809	1.016	-1.436		1.122	-1.388	-1.367	2.213
		(1.678)	(1.674)	(1.588)	(1.569)		(2.340)	(2.687)	(2.525)	(2.451)
SH*Char*Disad Min		0.435	-6.545**	-0.477	4.526		-0.094	2.042	1.573	-1.425
		(2.548)	(2.952)	(3.127)	(3.118)		(2.580)	(2.892)	(2.759)	(2.722)
SH*Char*Disad Maj		1.011	-4.235**	-0.542	-0.968		-1.976	3.973	0.480	-4.719*
		(2.093)	(2.153)	(2.081)	(2.198)		(2.898)	(3.109)	(3.002)	(2.795)
Observations	1152	1152	1152	1152	1032	960	960	960	960	888
Clusters	96	96	96	96	86	80	80	80	80	74

Robust standard errors clustered by individual are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

All regressions include age, age squared, Female, NonWhite, Upper class, L-CRT, LNUM, linear trend and constant. Columns (1) and (5) also include L-SAT.