Income Inequality and the Overworked American

Michael Carr

November 9, 2008

Abstract

Over the last 30 years there has been a substantial increase in both wage inequality and the share of individuals working long weekly hours. Recent experimental and theoretical research suggests that the widening income distribution could play an important causal role in determining labor supply patterns. Using a sensitive version of the Panel Study on Income Dynamics linked to the Current Population Survey, this claim is tested. Little support is found for the claim that either inequality in the work or inequality in the community increases weekly work hours. The results are robust to different measures of inequality, different specifications of the reference group, and different panel data estimation techniques including a new method for multi-way clustering of errors. The conclusion is that the positive correlation between work hours and inequality at the occupational level in cross-section analyses found by previous researchers is capturing variations in the underlying institutional structure of occupations and self-selection into occupations.

Introduction

Schor's (1993) book *The Overworked American* claimed that the average American worker experienced an almost unprecedented increase in average weekly work hours over the 1970's and 1980's. Some years later, it has become clear that Schor was half right[1]. Mean weekly hours for women have increased somewhat, from 33 hours per week in 1981 to about 37 hours per week in 2005. Mean weekly work hours for men, however, remained essentially unchanged at approximately 42 hours per week. But, as others have documented, the big change in labor supply during the last two and a half decades is in the dispersion of work hours among individuals and households. The share of the work force that works at least 49 or less than 35 hours per week, respectively, has increased significantly, especially for men. This development mirrors the increase in income inequality in the United States that also began around 1980 (Mishel, Bernstein and Allegretto 2007).

On the face of it, the presence of these parallel trends seems self explanatory. The income distribution is a result of the labor supply distribution. But, most of the increase in inequality came from high incomes pulling away from the rest of the distribution. High income individuals are almost exclusively salaried, thus their income does not mechanically increase with their work hours, at least in the short-run. This suggests that the income distribution may be both a result of and a cause of the distribution of work hours.[2] The assertion that inequality can have a causal effect on work effort and work hours is widely supported by an ever growing body of theoretical and empirical work suggesting that individuals have preferences over both the distribution of income and their position in the distribution (Akerlof 1982, Akerlof and Yellen 1990, Fehr and Schmidt 1999, Fehr, Kirchler, Weichbold and Gachter 1998, Lazear and Rosen 1981). The results of these studies have lead many people to attempt to find support for the mechanisms described, in particular, in the theoretical work. Very few of the studies looking for evidence of a causal link from wage inequality to effort or work hours use real world empirical data. The empirical work thus far has been almost exclusively experimental, and has focused on work effort not hours.

This paper examines the potential causal links running from wage inequality to labor supply. In so doing, it covers a number of limitations to the existing work. First, the different hypotheses for why a causal link might exist hinge on two different dimensions of inequality. I consider work versus community inequality. Empirical work based on Lazear and Rosen's (1981) tournament model of the firm argues that wage inequality determines work hours through the incentive that inequality provides to work hard and get a promotion, thus it requires a measure of inequality defined by the work. Empirical studies loosely deriving

[1] See (Jacobs and Gerson 2005) for a very detailed description of the patterns of individual and family labor supply from the 1970's to the present.

[2] This fact has been suggested by a number of authors up to this point (Bell and Freeman 2001, Bowles and Park 2005, Frank 2005, Frank and Levine 2007, Landers, Rebitzer and Taylor 1996, Schor 1998). Some of these arguments will be discussed in more detail later.

from Veblen's (1899) and Duesenberry's (1949) arguments that a given individual's consumption is partially determined by the individual's consumption relative to the consumption of others requires a measure of income inequality defined over an individual's social reference group. One of the potential groups is the community the individual resides in. This is the first analysis, as far as the author knows, that distinguishes between these two dimensions of inequality. Second, unlike other analyses, this paper uses a true micro panel data set, the Panel Study of Income Dynamics (PSID). The PSID is combined with the Current Population Survey (CPS) by Metropolitan Statistical Area for community and (experience, education) cells or occupation for work, using a sensitive[3] set of geographic identifiers made available upon request by the administers of the PSID. This allows the reliable calculation of group statistics using the CPS data, while retaining the panel data propoerties of the PSID.

It is found that in the cross section there is a positive correlation between work inequality and both work hours and the share of men working long hours, respectively. And, in the cross section, there is a positive correlation between community inequality and the share of men working long hours, though not mean work hours. This correlation disappears with the inclusion of basic controls such as education and age. In the fixed effects regressions there is no evidence of a positive correlation between either community or work inequality. This is in stark contrast with the results of other research which has consistently estimated a positive coefficient on wage inequality in a labor supply regression.

The conclusion is that studies that aggregate up above the individual return a positive coefficient because of underlying institutional and "cultural" differences across subgroups that determine both work hours and the wage distribution as well as self-selection into the reference group. There is likely to be considerably more standardization of work arrangements within occupations than across. In addition, different types of occupations have differring degrees of standardization, leading to a positive correlation between the wage distribution and the hours distribution across occupations. Further, individuals self-select into a given reference group. They choose an occupation often knowing what the work arrangements are like ahead of time. This means that although lawyers may have chosen to be a lawyer because long relative hours are rewarded with high incomes, they are not working long hours because of the income distribution conditional on having chosen to be a lawyer. In Manski's (1993) terms, the cross-section regressions suffer from the reflection problem. Individuals choose occupations that reflect their underlying preferences, their behavior is not caused by the behavior of others in the occupation.

The paper is organized as follows. Section I reviews the existing literature, including the theoretical underpinnings of the directly relevant empirical work. Section II addresses some of the empirical difficulties

[3]The Panel Study of Income Dynamics provides, upon request, a set of geographic identifiers which are not publicly available. This allows the identification of the location of a given household down to the census tract. Because of the nature of the identifiers, this PSID does not allow researchers using them to share any of the data resulting from their use.

inherent in estimating these regressions, especially the Reflection Problem (Manski 1993), and describes the main variables of interest. Section III describes the data and presents some background statistics, some of which have been addressed by other authors as well. Section IV presents and discusses the regression results in the context of the existing hypotheses for why a causal relationship between inequality and labor supply might exist. Section V investigates the short-comings of the existing literature, and explains why the results found here differ. Section VI concludes.

I. Existing Literature

Schor (1993) claimed that, in 1993, the U.S. workforce was in the middle of a historically unprecedented increase in the amount of weekly and annual hours worked. However, Jacobs and Gerson (2005, p.19) demonstrate that, with the benefit of hindsight, the upward trend that Schor identifies is an artifact of beginning the analysis in the mid-1970's. The 1970's appears to be an anomoly. It is the only period in the post World War II era in which mean weekly work hours are significantly below 40 hours per week. But, what Schor did accomplish was to draw attention to the fact that there were meaningful changes in labor supply patterns happening during this time. Specifically, as Jacobs and Gerson (2005) document but do not explain, what happened was not an increase in mean work hours, but instead an increase in the dispersion of work hours. Beginning in the 1980's, there was a sharp increase in the share of individuals working at least 49 hours per week (long hours) or less than 35 hours per week (short hours), respectively. These simultaneous trends are most pronounced for well-educated male household heads from middle and high income households (Jacobs and Gerson 2005, Kuhn and Lozano 2008). Also beginning in the late 1970's was the now well documented increase in wage inequality in the U.S.. In 1980, the U.S. family income gini coefficient was approximately 0.36, since then it has increased to approximately 0.44, a 22% increase (Mishel et al. 2007).

Most likely because of the increase in income inequality, research on the behavioral implications of inequality gained a new life. The purpose of this paper is to further investigate the causal role that income inequality plays in determining labor supply, holding the wage level fixed. In other words, this paper is an attempt to isolate the elasticity of labor supply with respect to the relative wage, holding the wage level fixed. The modern incarnation of the questions surrounding income inequality, effort and work hours choices, and labor demand begins with three papers: Akerlof (1982), Akerlof and Yellen (1990), and Lazear and Rosen (1981). To the author's knowledge, there are no papers that attempt to apply the Akerlof (1982) partial gift exchange model or the Akerlof and Yellen (1990) "Fair-Wage Effort Hypothesis" to real world empirical applications and the question of labor supply. Because of this, these two papers will be set aside

at least for now. The Lazear and Rosen paper, on the other hand, provides the main theoretical foundation for one of the two strands of research that attempts to explicitly link income inequality to labor supply. The second strand of research, and its theoretical foundations, will be addressed next.

Lazear and Rosen (1981) argue that firms, and by extension the labor market, resemble a tournament. Or, more precisely, if firms and the labor market resemble a tournament, then the problem of perfectly monitoring a worker's effort is solved. The thrust of the model is as follows. A firm cannot perfectly monitor a worker's investment in a job, where investment can be understood to be the aggregation of a number of behaviors, one of which is effort expended on the job. The firm can only observe output. Output is a function of two unobservable terms: investment and a random term. The firm maximizes profit, which amounts to inducing maximum investment in all of its workers, so that any variation in output is random. If the firm pays a piece-rate according to the level of output, then a worker can provide low levels of effort when the random term is high, and earn the same wage as when investment is high but the random term is low. But, if the firm pays a fixed wage according to a worker's relative output, this problem is eliminated. This structure is known as a tournament. The wage structure is predetermined, and workers are paid the wage associated with their respective ranks in the tournament. The worker's rank is determined by relative output, and the workers know this.

For a given worker, the best response to the investment decisions of all other workers is increasing in the dispersion of the payoffs. The more unequal the distribution of payoffs in the tournament, the higher the marginal return to moving from rank n to $n+1$, assuming a given probability of moving between ranks and that the number of ranks in the tournament that has a positive payoffs also remains unchanged. An increase in the dispersion of payoffs causes an increase in worker effort, all else equal, because the marginal return to effort increases. In the real world, it is often the case that neither output nor effort are (perfectly) observable, in this case the firm will impute a worker's effort or output based on observable signals. One of the likely candidates is labor supply.

This is precisely what Bell and Freeman (2001) and Kuhn and Lozano (2008) argue. Bell and Freeman show that, in Germany, the probability of promotion is positively correlated with work hours across a wide variety of occupations. The interpretation being, again, that work hours are a signal of effort. In a very different context, and for very different reasons, Landers, Rebitzer, and Taylor (1996) also draw the same conclusion. In this case, they find that lawyers who work longer hours are the most likely to be promoted. The interpretation they give is not that lawyers work long hours because of the spot-market pay structure within law firms, but rather that because firms earn money based on the number of billable hours worked, partners have an incentive to promote the lawyers who work the longest hours. None-the-less, if promotions for lawyers do not also come with a pay increase, it is quite unlikely that they would still put in such long

hours. Therefore, it is still the case that the distribution of wages in the firm causes long work hours.

Bell and Freeman (2001) run a set of regressions using occupation level data from Germany and the United States. Their measure of wage inequality, the gini coefficient, is defined within a given occupation. They regress mean work hours in an occupation on the gini coefficient of wages for that occupation, and find that they are positively correlated, holding the wage level fixed. Using the evidence that work hours and the probability of promotion are positively correlated, they conclude that the positive relationship between mean occupation work hours and occupation wage inequality is due to the promotion incentive that inequality provides. From now on, this explanation will be referred to as the "promotion incentive" argument.

Kuhn and Lozano (2008) do something similar to Bell and Freeman, but with three meaningful differences. First, rather than trying to explain *mean* weely work hours in an occupation, they analyze the share of men ages 25 to 64 working greater than 49 hours per week in a given occupation out of all employed men in that occupation. Second, they use data only on the United States, derived from the Current Population Survey. Third, their analysis adds a time-series element by tracing labor supply patterns of population subgroups through time. They find three important results. First, as with Bell and Freeman, they find a positive correlation between the degree of inequality in an occupation and the share of men working long hours. This is not surprising given that there is likely a positive correlation between mean work hours and the share of men working long hours, especially in occupations dominated by men.

More interesting, however, is the finding that the share of men working long hours increased more in occupations where real wages were growing slowly than in occupations where they were growing quickly, but only for those occupations where real wage growth actually occured during this period. And, the share of men working long hours increased more during the 1980's when real wage growth was slow, then in the 1990's when real wage growth picked up somewhat for many men. This is important because it suggests that changes in the level of inequality, or the level of inequality itself, may not universally provide an incentive to work long hours. Think of the degree of wage inequality as representing the possibility for advancement for a given worker in that occupation. When real wages are growing quickly, for a given income distribution, workers do not need to get promoted in order to increase their income. However, when real wages are growing slowly, the only way to increase their income is to get a promotion, especially for salaried workers. If the degree of income inequality reflects the marginal return to promotion, then workers will only work towards a promotion if the possibility of promotion exists. Thus, the level of inequality only provides an incentive to work long hours when real wages are growing slowly. Although this does not undermine the thrust of Kuhn and Lozano's argument, it at least suggests it is more complicated than simply arguing that inequality causes an increase in work hours.

The second strand of literature that specifically addresses the causal link running from income inequal-

ity to labor supply derives from the work of Veblen (1899) and Duesenberry (1949), and more recently Frank (1985a, 2005, 2007, 2007) and Schor (1998). This strand consists primarily of two papers: Bowles and Park (2005) and Park (2005). The empirical components of the two papers are quite different from each other, but the theoretical underpinnings are the same. The basic idea is that utility depends not only on the level of one's own consumption but also on own consumption relative to the consumption of the given individual's reference group. [4] Where the reference group is defined loosely as a group of individuals located above a given individual in the income distribution. An increase in inequality causes an increase in the gap between a given individual's consumption level and that of her reference group. This causes the individual to consume more.

In Bowles and Park (2005) a cross-country panel data set of OECD countries is used. Because of this, all variables must be defined at the level of the country. In the cross-country panel regression, with mean annual work hours as the dependent variable, Bowles and Park find a positive coefficient on the gini coefficient for wages in the respective country. They interpret the positive correlation between the gini coefficient and mean work hours as evidence of consumption emulation. Park (2005), on the other hand, looks at changes in married female labor force participation. Park finds that the probability of a married women being in the labor force is negatively correlated to the husband's relative wage. However, in this case, the reference group is defined along demographic and geographic characteristics. The husband's wage is compared to other married men who are "similar" to the person in question, rather than simply to those wealthier than him. Although this finding is in support of the consumption emulation hypothesis, these results are not directly relevant here because it is difficult to compare results describing changes at the extensive margin with those describing changes at the intensive margin.

The explosion in behavioral and experimental research also provides a number of reasons why there might be a causal link running from the distribution of income to labor supply. Theoretical work such as Akerlof and Yellen's (1990) extended model of gift exchange, called the "Fair-Wage Effort Hypothesis," suggests that effort is partially determined by an individual's wage relative to the wage the individual believes to be fair. If the distribution of wages the individual uses to determine the fair wage changes, the individual's relative wage will also change. The direction of the change is determined by the individual's position in the wage distribution, as well as the position of the reference wage. Experimental research inspired by the pioneering work of Fehr and Schmidt (1999) largely supports the claims made by Akerlof (1982) and Akerlof and Yellen (1990). Specifically, these experiments consistently document a positive correlation between chosen effort and the wage level, and chosen effort and the relative wage.

[4] See Appendix A for a simple version of a model with relative consumption where an increase in inequality increases consumption and labor supply.

II. Empirical Methodology

The only channel through which income inequality can determine labor supply is interpersonal comparison. This is not to say that it is the process of interpersonal comparison itself that determines labor supply, rather that for an individual to know the degree of inequality and her position in the income distribution, the individual must compare her outcome to the outcomes of the others she deems relevant. Bell and Freeman (2001) present evidence that the probability of promotion is positively correlated with labor supply, suggesting that hours worked over and above the minimum required serves as a signal of the amount of effort expended on the job. The only way that long hours can serve as a signal that an individual is working harder than others, and deserves a promotion, is if the individual is working longer hours than everyone else. This feature implies that it is more correct to think of the effect of income inequality as being embedded in a model of social interactions, in which a given individual's labor supply choices are a function of the labor market outcomes of the individual's reference group.

The basic social interactions model posits that an individual's choice is a function of the mean behavior of her reference group. To continue the above example, the direct empirical representation of the Lazear and Rosen model, when effort and output are not directly observable, is not whether work hours tend to increase with wage inequality. Instead, three conditions must hold. First, an individual's labor supply must be an increasing function of mean labor supply of her reference group. If this is not the case, then work hours cannot be a signal of effort. Second, the probability of promotion must be an increasing function of the individual's labor supply relative to the reference group. Only if these two conditions hold does a positive correlation between labor supply and wage inequality suggest that the Lazear and Rosen model holds.

In the Lazear and Rosen model, the utility maximizing level of investment is determined by the marginal utility of investment, which is a function of the income distribution. For a given individual, the process of determining how much effort to provide does not require interpersonal comparison. However, when using work hours as the signal of relative output, the individual must *actually* work longer hours then her reference group to demonstrate that her effort and/or output is higher. This is precisely the structure of a social interactions model. Therefore, if a positive correlation between work wage inequality and labor supply is found, the next step must be to establish a positive correlation between mean work work hours and the individual's work hours. If this condition holds, then we investigate the relationship between promotion and labor supply. However, as mentioned in the introduction, steps two and three are not necessary here because there is not a significant positive correlation between work wage inequality and labor supply.

Similarly, the ideal empirical representation of the consumption emulation hypothesis is whether individual consumption is an increasing function of mean and/or the dispersion of reference group consump-

tion. This condition establishes that consumption emulation exists, and that the degree of emulation varies with the level of inequality. Although it has been established elsewhere that consumption emulation exists, it has not been established that the degree of emulation varies positively with the degree of inequality (Frank 1985a, Frank 2005, Frank and Levine 2007, Hopkins and Kornienko 2004, Luttmer 2005). The next step is to link consumption emulation to labor supply by establishing a positive correlation between reference group consumption and individual labor supply. It does not matter in which order these two conditions are established. This paper examines the second condition. If this condition is established, further research must investigate the first. Estimating social interactions equations like this, however, comes with a whole host of empirical difficulties, the most important of which is the "Reflection Problem" (Manski 1993). Before moving to the reflection problem, reference groups must first be defined.

A. *Reference Groups*

A reference group is simply a group of people that a given individual compares herself to. The individual in question can be a member of the group and wants to appear to be a member (conformism), is not a member of the group but wants to a appear as if she is (emulation), is a member of the group and wants to appear as if she is not (nonconformism), and is not a member and wants to appear as if she is not a member (distancing). Specifying a reference group for consumption emulation that is consistent with Veblen's (1899) arguments falls under the second category: a given individual is assumed not to be a member of the reference group but wants to appear as if she is, she is an emulator. The reference group for the promotion incentive argument falls under the third category: an individual who is a member of a group but wants to appear as if she is not, she is a nonconformist. The latter will be dealt with first, followed by the former.

The primary difficulty one confronts when defining a reference group is that it is context specific. As a general rule of thumb, social psychologists argue that individuals tend to compare themselves to others who are, on average, similar along lines relevant to the context (Burke 2004). Definition of the work reference group is somewhat complicated. The previous studies mentioned above that have looked specifically at the causal link running from income inequality to labor supply have defined the work reference group as the individual's occupation. Although this is reasonable, it is far too narrow a definition, and likely underestimates the amount of inequality a given individual encounters in work life. For example, a janitor will look to what other janitors earn, but will also likely look to other occupations such as groundskeepers, house cleaners, building maintenance workers, and the like. This process is similar to how Levine (1993) describes pay setting behavior by executives in large corporations. Levine shows, through surveys of compensation setting executives, that pay for a given job is set by comparing pay rates both within the firm across occupations and within similar occupations across firms. It is also how Akerlof and Yellen (1990) claim that

individuals form a conception of what is a fair wage for their respective jobs.

By definition, a given individual's reference group will include all individuals residing in jobs that the individual feels are similar to her own. This is precisely the set of wages that should determine effort in a tournament. A wage associated with a job that is completely unattainable for a given individual should have no bearing on the individual's incentive to get a promotion. This set of wages is certainly larger than all wages in the individual's occupation. It should include comparisons across occupations within the individual's firm, across firms within the individual's occupation, and across occupations across firms. The occupation captures one aspect of this set, comparison across firms within occupations, but entirely misses the remaining two. There are a number of reasons why this omission is problematic. Logic suggests that the coefficient on wage inequality should be biased downwards by defining the group too narrowly. However, as will be discussed below, this may not be the case.

Second, and most important, many promotions entail a change of occupation. Table 1 shows the share of wage changes greater than 5%, 10%, and 15%, respectively, from year $t-1$ to year t that are associated with a change of occupation. Assuming that the very large majority, if not all, significant pay increases are the result of a promotion, 52% of all 5% real wage increases, 55% of real wage increases over 10%, and 58% of all real wage increases over 15% entailed a promotion. Because more than half of all promotions entail a change of occupation, the occupation wage distribution is only a small portion of the wage distribution that is relevant for a determinig the return to a promotion.

Table 1. Share of heads who changed occupations when real hourly wage changed

	Percent Real Wage Change		
	0.05	0.1	0.15
Different Occ.	28452	22290	18352
Total	54316	40378	31855
Share	0.52	0.55	0.58

Source: Author's calculations based on the Panel Study of Income Dynamics.
Notes: (i)Universe is all employed individuals ages 25-64. (ii) Occupation is defined by the 3 digit 1970 Census occupation codes.

A better reference group for this particular issue is all individuals residing in jobs that are attainable for the given person. This group should consist of all individuals who share a similar set of characteristics that are relevant to the labor market. This set of characteristics is quite large. It should include the standard human capital variables: education, experience, and ability. But, it likely also includes other behavioral traits less commonly discussed and more difficult to measure (Bowles, Gintis and Osborne 2001). Delineating the work reference group for a tournament in this way is consistent with social psychological research. An individual would never choose to compete with a group of individuals she has no chance of beating. For example, individuals that hold a job consistent with having a high school diploma would not consider the

pay rates of jobs that require a Ph.D. when determining work effort. Similarly, an individual with extensive experience in a given type of job is more likely to compare pay rates with other experienced individuals, not inexperience individuals.

Consumption emulation, on the other hand, requires a quite different specification of the reference group. Veblen's (1899) argument is that everyone emulates the consumption of the wealthiest members of society. Psychological research suggests, however, that this assumption is too extreme. For most people the consumption level of the wealthy is completely unattainable. In this case, according to psychologists, individuals will adjust their referent to avoid comparisons with something they can never attain (Austin 1977, Burke 2004). In line with this research, most have taken a less extreme form of referent choice.[5] As mentioned, Bowles and Park (2005) simply use the gini coefficient for a country, and Park (2005) uses a measure of income relative to a reference group defined by demographic and geographic characteristics. The approach taken here is similar to the latter.

There are any number of ways one can delineate a reference group for consumption. Certainly, one's coworkers are a relevant reference group, suggesting the potential for a considerable amount of overlap between the work reference group and the consumption reference group. But, one's neighbors and wider community is at least as important a reference group as one's coworkers. Many daily activities involve interacting with people from different neighborhoods in a given city, with varied consumption patterns and income. Going to the mall, the movies, or the grocery store all involve coming into contact with people who might be judging others based on their observable consumption patterns. Even if one does not know a single person in the grocery store, people seem to not like going to a place where they look like they do not belong. This is inline with Frank's (2000) and Schor's (1998, 2004) work on consumption emulation, all of which suggests that casual interaction between individuals who do not know each other and will likely never see each other again (or never literally came into contact with each other) can impact a given individual's consumption decisions. Although one's coworkers likely fall within this group, the breadth of the group is much different. A work reference group may cross wide geographic boundaries but will likely stay within a small band of the income distribution, relative to the distribution as a whole. The community reference group is strictly defined by the local geography, but covers a much wider segment of the income distribution.

A good representation of a community based consumption reference group is all individuals residing in the same region of a city. Defining the reference group as a unit as small as a neighborhood ignores the fact that the neighborhood itself may have been chosen partly because of consumption emulation motives. But, defining a region of a city as the reference group has its own pitfalls. First, without knowing the city very

[5]See (Frank 1985a, Frank 2005, Frank and Levine 2007, Hopkins and Kornienko 2004, Luttmer 2005) for examples of less extreme assumptions about referent choice in consumption choices that are not directly relevant to the labor supply question.

well, what appears to be a reasonable definition of a region by observable characteristics may turn out not to be. In other words, dividing a city into regions would most likely result in completely arbitrary definitions of a region. Second, in many cities there are central locations where individuals from across the entire income spectrum interact. Because of these interactions, even individuals from spatially distant regions of a city will know, and potentially be affected by, the visible consumption decisions of the entire city.

Therefore, three reference groups will be used. The consumption emulation reference group is defined by the city the individual resides in. The work reference group has two definitions. For comparison, one set of regressions will include occupation as the work reference group. The other set of regressions will include the definition of work reference group, referred to as the "workgroup" from now on, will consist of the wages of all individuals residing in jobs that can loosely be considered as attainable for a given individual. The set of attainable jobs will be defined by the characteristics of the individual holding the job, such as education, experience, and ability. How the three reference groups are measured will be discussed in section III.

B. *The Reflection Problem*

Besides the definitions of the reference groups, the other main methodological issue is dealing with the Reflection Problem. The Reflection Problem is a term coined by Manski (1993), and is simply a problem of correlation versus causation, or endogeneity bias, depending on how the problem is framed. For example, when we see a group of teenagers, all of whom smoke, we are faced with the following conundrum: are they friends because they smoke or does a given group member smoke because the individual's friends–the reference group–smoke? That is, it is difficult to disentangle whether a group of individuals share an observable behavior or characteristic because they intentionally joined a group of people that exhibit the same behavior or characteristic, or if the given individual exhibits a behavior or characteristic because other group members do, but the individual is a member of the group of for other reasons. If the former is true, then a regression estimating the effect of reference group behavior on individual behavior is capturing correlation, not causation. If the latter is true, then it is causation, not correlation.

The formal specification of the reflection problem laid out in Manski (1993) technically only applies to a regression that regresses an individual's outcome on the mean of that outcome in the individual's reference group. In this case, it technically only applies if the regression equation specified below regressed individual labor supply on mean reference group labor supply. Manski demonstrates that the coefficient on mean group behavior is only identified under a certain set of conditions when the regression is specified in this manner, often referred to as the linear-in-means model. In a nutshell, if mean group behavior is a linear function of unmeasurable group characteristics and/or mean group behavior is uncorrelated with the same unobservable group characteristics, then the coefficient on mean group behavior is not identified. If the reference group

is correctly specified for the behavior in question, then any observable characteristics of the reference group should be correlated with the unobservable characteristics in the error term, otherwise there is no reason to claim that the individual is a member of that group. The real concern is whether mean group behavior is a linear function of unobservable characteristics. Of course, the chance of anything being a perfect linear function of another variable is very, very small. But, if it holds in approximation then identification is weak, and the closer the correlation is to a linear relationship the weaker identification becomes.

This problem is attenuated somewhat when the variable measuring reference group behavior is either a higher order moment of reference group behavior, such as the standard deviation instead of the mean, or if the variable measuring reference group behavior is not the same as the individual behavior being measured. In this case, the reference group variables are measures of income inequality, and the individual variable is labor supply, so the strict form of the the reflection problem is eliminated. However, the essential point of the argument still stands. In the absence of the ability to control for self-selection into groups, there is always the possibility that the coefficient on the reference group variable(s) will suffer from the reflection problem. And unfortunately, most of the time the coefficient will be biased upwards, over estimating the true causal impact of reference group behavior on individual behavior.

There are some proposed solutions to this problem, depending on the structure of the data set.[6] There are two commonly used methods for panel data. The first, proposed by Durlauf (2004), utilizes the fact that in some applications the individual will choose to change reference groups (e.g. changing neighborhoods). This information can be used to determine an individual's preferences over certain group characteristics. The second method is fixed effects, which will be used here (2004). There are a lot of well known draw backs to fixed effects. In this application, one of these drawbacks is potentially problematic. For fixed effects to adequately control for reference group selection, it must be the case that the process by which individuals choose their group is stable over time. If it is not, then we must be able to control for the set of processes that might not be stable. However, given that a simple mean of group behavior is not being used as a measure of group behavior, this particular issue will most likely not arise. Additionally, some variables that represent an individual's location in her life-cycle will be included, such as age.

C. *Endogeneity Issues*

The second of the potential endogeneity issues could also bias the coefficient on the measures of inequality. The purpose of this paper is to examine the possibility of a causal link running from wage inequality to labor supply. But, the distribution of income is also a result of labor supply decisions. This is particularly true for wage workers if the distribution of weekly or yearly income is used. To avoid this problem, the distribution

[6] See Brock and Durlauf (2001a), Brock and Durlauf (2001b), and Durlauf (2004) for a comprehensive review of the issue.

of hourly wages will be used. There is no immediate effect on the wage distribution when hours worked changes.

There is a second potential endogeneity issue, however. If it is the case that working longer hours increases the probability of a promotion, then hours worked partially determines the wage distribution in the future. From an econometric standpoint, this does not pose a problem unless workers take into account their future impact on the wage distribution. This is implausible, and thus is not a concern. But, even if it were plausible we would not expect it to be a problem. When an individual gets a promotion, for the most part, they are moving from one job with a well defined and stable relative pay to another. The wage distribution does not change when individuals change jobs, the only thing that changes is the recipient of the wage income. The wage distribution is largely determined independently of who specifically occupies the job. If this were not true, then Lazear and Rosen's tournament model would not be applicable at all, and much of the literature on the effect of inequality would be unfounded.

III. Data and Descriptive Statistics

The analysis presented below uses two data sets: the outgoing rotation groups from the Current Population Survey (CPS), and the Panel Study of Income Dynamics (PSID). The individuals in the PSID are the unit of observation for the regression analysis. The PSID is a panel data set covering a group of households chosen in 1968, and expands each year to include all households that evolved from the original set. The survey was administered once a year until 1997, and every other year since. The PSID is used for a number of reasons. First, it is a longitudinal data set. This is an advantage because all of the studies mentioned above that use micro data and have a time-series component use quasi-panels constructed from the CPS. Because of this, individual fixed effects cannot be used, making it impossible to control for individual specific but unmeasurable characteristics, most importantly those that describe how an individual chooses a reference group. Second, there are very few missing or otherwise unusable observations relative to the CPS (as a share of the total number of observations). Finally, it has very detailed income and hours data. The version of the PSID used here has additional geographic identifiers not included in the publicly available version. These identifiers allow the PSID to be linked to the CPS by Metropolitan Statistical Area.

The data for calculating descriptive statistics of overall labor supply patterns and for calculating reference group statistics comes from the CPS. The CPS only asks about usual weekly work hours and earnings in the fourth and eighth month that a household participates in the survey. The outgoing rotation groups pull out all households in their fourth or eighth month in a given year, and compiles them into one data set. There are three main reasons for doing this. First, this method provides more detailed information on work hours

and earnings. Second, the data set is much larger than the March Supplement, the alternative source within the CPS for this type of data. Third, although Census data is much larger and carries the same information, there is not a publicly available version of the Census that has individual data and includes the MSA as a geographic identifier. Fourth, the CPS is collected yearly, unlike the Census.

As mentioned, the first of the two work reference groups is occupation. In line with Kuhn and Lozano (2008), three digit census occupation categories are used. [7] Occupation inequality is measured by the standard deviation of wages within an occupation. In order to maximize the number of usable occupations, the standard deviation was calculated over a three year moving window. Occupation inequality in a given year is measured as the standard deviation of real wages from years $t-1$, t, and $t+1$. Any occupation whose three year moving window does not average at least 75 observations over the entire period is dropped. This leaves 380 occupations. Although this process excludes a large number of occupations in the CPS, it only drops 515 otherwise usable observations from the PSID.

The second definition of work reference group is the set of workers who share similar characteristics. This set is assumed to be delineated by human capital and other behavioral traits that are rewarded in the labor market. The CPS has limited measures of human capital and no measures of behavioral traits. In order to provide a sufficient number of observations in any given cell, only two delineators of the reference group will be used: age and education. Education is the only human capital variable that can be directly measured in the CPS. Age is a proxy for experience. The most common method of measuring experience, age minus education minus sixteen, actually measures potential experience. Because, after education is complete all variation in experience comes from age, age is a suitable proxy for experience. Education is divided into four groups: high school degree or less, some college including associates degree, bachelor's degree, and some advanced degree. Age is also divided into four categories: 25-34, 35-44, 45-54, and 54-64. Thus, a given individual in the PSID that has valid age and education observations is assumed to compare to one of 16 reference groups delineated by age and education in a given year. The analysis only includes ages 25 to 64 because too many individuals younger than 25 are still in school and too many individuals older than 64 are starting to retire or work part-time. This is also the age range used by Kuhn and Lozano (2008).

This is a very broad definition of reference group. The primary issue with a definition this broad is that many advanced degrees do not endow the individual with the same earnings potential. A 35 year old person with a master's degree in social work does not face the same wage schedule as 35 year old with an M.B.A, yet they are considered part of the same group. There is no way to avoid this dilemma with the CPS. The advantage of this measure is that a quick look at any job advertisement reveals that, besides skill

[7] The CPS does not use a consistent definition of occupation across the entire time period. David Autor has made available on his website an occupation crosswalk valid for 1983-2003. The crosswalk was extend backwards to 1981 and forwards to 2005 by the author. David Autor's crosswalk is available at http://econ-www.mit.edu/ dautor/alm/occ8090.zip.

requirements, education and experience are the most commonly listed requirements of a job. An individual can also use this information to approximate the wage distribution of the jobs she is qualified for.

The reference group for consumption emulation is the city the individual resides in. The sensitive set of geographic identifiers used here allows identification of the Metropolitan Statistical Area (MSA) a given individual resides in. An MSA is any urban area that contains a core with a population of at least 50,000 people. Besides being the most commonly used definition of city, the data on MSAs is dense enough in a given year to allow reliable calculation of MSA statistics without dropping a significant number of MSAs that exist in both the PSID and the CPS. There are two main problems with using this measure. First, it likely over-estimates the degree of inequality a given individual actually experiences. This is especially true for the wealthier individuals in an MSA. While lower income residents will most likely work in areas where the wealthier residents reside, the wealthy residents have much less reason to frequent the areas that lower income residents do, particularly in large cities. Second, the Current Population Survey only identifies the 57 largest MSAs in the United States prior to 1986. Therefore, prior to 1986 this definition of community excludes some observations that later become observable because of more complete coding of MSAs.

Because the two existing hypotheses, promotion incentive and consumption emulation, have very different reference group specifications, using two measures of inequality that are not highly correlated allows the regression analysis to differentiate between the opposing hypotheses. It also provides the ability to more accurately describe the differences between subgroups within the sample. Specifically, how the effect of the two measures of income inequality varies between salaried and hourly workers,[8] and how the magnitude of the estimated coefficients on inequality varies across the income distribution.

The regression analysis covers the years 1981 to 2005. The regression sample is limited to male household heads ages 25-64. This is done both for comparability with Kuhn and Lozano (2008) who also limit the sample to employed men ages 25-64, and for reasons discussed below. The final data set has a total of 41,958 observations when using occupation as the reference group and 34,467 observations when using workgroup as the reference group. The primary source of lost observations comes from linking the PSID with the CPS data by MSA. Prior to 1986, the CPS only identified the largest 57 MSAs in the U.S.. From 1986-2005, all MSAs that have observations are identified, but any household that either does not reside in an MSA or resides in an MSA that was not included in the CPS can not be matched either. In total, this process drops 22920 observations, or approximately 35% of the otherwise usable observations. Also, there are some unusable observations for either age or education due either to coding error or the respondent did not know the answer or refused to answer the question. This accounts for the slightly smaller data set when using work inequality instead of occupation inequality.

[8]This approach is also used by Kuhn and Lozano (2008).

Figure 1 shows mean weekly work hours for all employed men and women ages 25-64, respectively, from 1981-2005 based on the CPS. Mean weekly work hours remained essentially unchanged for men, but increased somewhat for women. From figure 2, we can see that the rise in mean work hours for women is a result of the small increase in the share of women working at least 49 hours per week, and the somewhat larger increase in the share of women working full-time, indicated by the downward sloping line for the share of women working less than 35 hours per week. In other words, for women the distribution of work hours conditional on being employed shifted to the right, but maintains a similar shape. The story is different for men. Mean weekly work hours for men have remained essentially unchanged since 1981. However, the share of men working at least 49 hours has increased considerably. In 1981, the share was 15%. This increased steadily to just under 23% in 1999, then declined somewhat to 2005. At its peak, the share increased by roughly 53%.

Figure 1. Mean weekly work hours by gender: 1981-2005

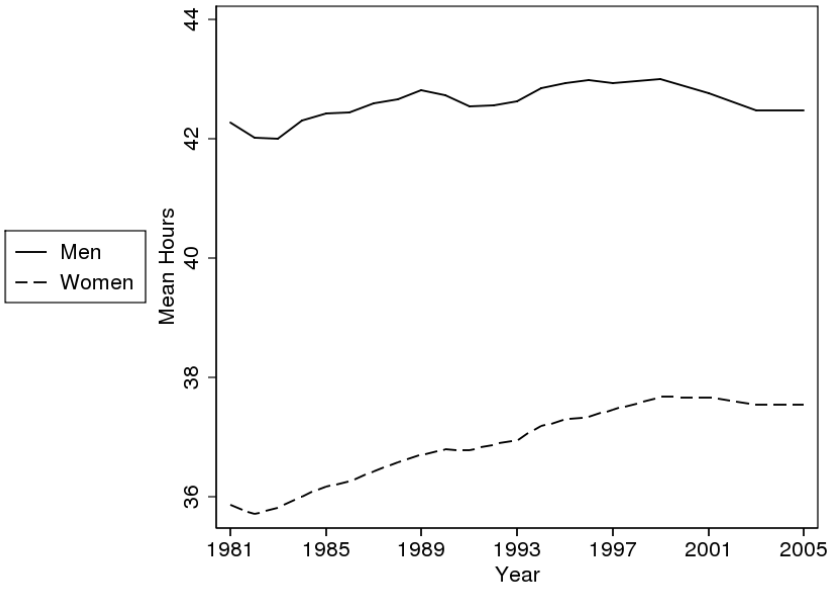

Source: Author's calculations based on the Current Population Survey.
Notes: (i) Data covers years 1981-2005, excluding 1998, 2000, 2002, and 2004. (ii) Universe is all employed individuals ages 25-64.

The essential message of figure 2 is that both men and women experienced significant changes in labor supply patterns over this time, but the patterns of change are quite different. Women simply worked more. Both the share of women working at least 35 hours and the share of women working at least 49 hours increased. Male labor supply, on the other hand, became more unequal. The majority of the increase in the share of men working long hours came in the 1980's and early 1990's. At the same time, the share of men working short hours was also growing, albeit much slower.

It is likely the case that the reasons for the changes in the share of men working long hours and in the

Figure 2. Share of individuals working long or short hours by gender: 1981-2005

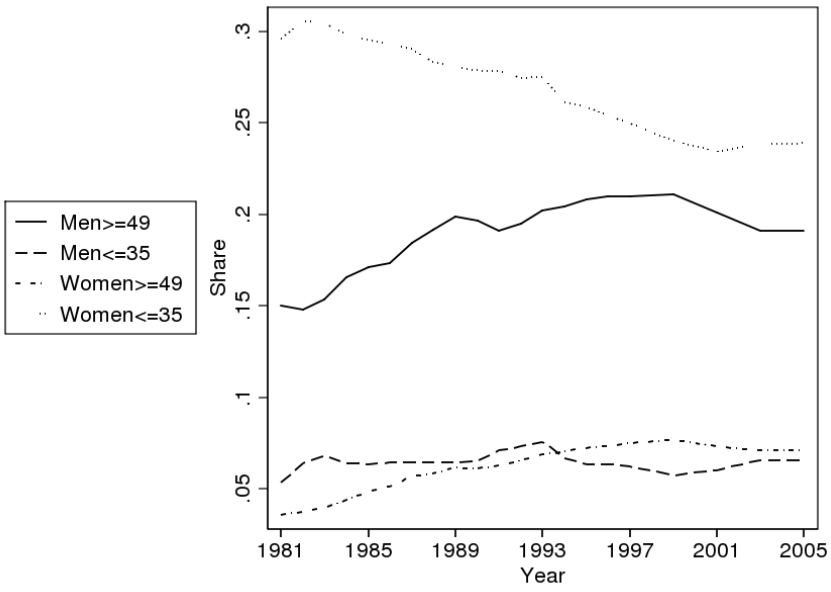

Source: Author's calculations based on the Current Population Survey.
Notes: (i) Data covers years 1981-2005, excluding 1998, 2000, 2002, and 2004. (ii) Universe is all employed individuals ages 25-64.

share of women working full-time are somewhat different, or at least require separate analyses. This does not mean that the phenomena are not linked. There is likely some substitution of female work hours for male work hours within households, though Juhn and Murhpy (1997) and Devereux (2004) suggest that this substitution is limited. It does mean that it is inappropriate to run a single regression with data on both men and women. Further, there are very few instances of two adult households that have a female head in the PSID. Limiting the sample to household heads with another adult present effectively limits the sample to men.

Bell and Freeman (2001) and Kuhn and Lozano (2008) both estimate cross-section regressions. Tables 2, 3, and 4 show the cross-section patterns that hold for MSA and work reference group inequality. Table 2 shows mean weekly work hours by quantiles of MSA inequality, where quantile 1 is the lowest level of inequality. There is a clear upward trend in the share of men working long hours as the degree of inequality within an MSA increases. Also, the share of men working short hours is higher in the higher inequality MSAs, though this relationship is not monotonic. Given that mean work hours is constant across all ten quantiles, the degree of dispersion in labor supply increases as the degree of wage inequality increases within an MSA. It is also the case that higher inequality MSAs have lower unemployment rates.

Table 3 shows similar trends to table 2. Each row represents a work reference group, delineated by age and education. The table is sorted in ascending order according to inequality within workgroups, measured

Table 2. Hours Statistics by Quantiles of MSA Inequality

Degree of Ineq.	Quantile	Work>=49	Work<35	Mean Hours	Unemployment
Least Unequal	1	0.150	0.061	42.105	0.064
	2	0.166	0.060	42.427	0.061
	3	0.173	0.061	42.473	0.058
	4	0.174	0.064	42.405	0.049
	5	0.188	0.068	42.600	0.046
	6	0.199	0.063	42.810	0.044
	7	0.202	0.065	42.816	0.045
	8	0.200	0.067	42.712	0.044
	9	0.201	0.066	42.730	0.042
Most Unequal	10	0.205	0.070	42.690	0.041

Source: Author's calculations based on the Current Population Survey.
Notes: (i) Data covers years 1981-2005, excluding 1998, 2000, 2002, and 2004. (ii) Universe for hours statistics is all employed men ages 25-64. (iii) Universe for unemployment rate is all individuals ages 25-64 residing in the MSA who are in the labor force.

by the standard deviation of real wages in 1981 dollars. The share of men working long hours increases with the level of inequality, though the trend is not nearly as clear as with MSA inequality. Although there is a considerable amount of variation in the share of men working short hours across workgroups, there is no clear trend. And, again, the level of unemployment seems to decrease as the level of inequality increases. The primary difference between the two tables is that there is noticeable variation in mean hours worked across work groups. Men ages 35 to 44 with an advanced degree work the most, with average weekly work hours of 45.08. This is followed closely by men ages 45-54 with an advanced degree, who work an average of 45.06 hours per week. The group with the lowest work hours are men ages 55-65 with a high school diploma or less. They only work an average of 40.5 hours per week.

The inverse relationship between unemployment on the one hand, and inequality and the share of men working long hours on the other is potentially important for three reasons. First, efficiency wage models of unemployment such as Shapiro and Stiglitz (1984) and Bowles (1985) argue that unemployment has a disciplining effect on workers. The higher the level of unemployment, the higher the probability that, if a worker becomes unemployed, they will not be able to find a new job. This induces workers to provide higher levels of effort and/or work hours. The inverse relationship between inequality and unemployment rules out the possibility that inequality is picking up the effect of unemployment in regressions that fail to control for it, such as the work of Kuhn and Lozano. To directly address this problem, reference group unemployment will be included in the regressions as well as reference group wage inequality.

Second, it may also suggest that the demand constraints workers face when choosing their work hours are less likely to be binding in groups where inequality is high. Some researchers have pointed out that demand constraints, or employer preferences over their workers' work hours, can bias estimated wage elasticities towards zero (Altonji and Paxson 1988, William T. Dickens 1993, Kahn and Lang 1991). The same problem

Table 3. Hours statistics by work place reference group inequality

Degree of Ineq.	Workgroup Number	Education	Age	Work>=49	Work<35	Weekly Hours	Unemployment	Work Ineq.
Least Uneq.	1	High Sch.	25-34	0.143	0.070	41.955	0.089	4.464
	2	High Sch.	35-44	0.146	0.052	42.226	0.064	5.140
	3	Some Coll.	25-34	0.172	0.080	42.168	0.054	5.335
	4	High Sch.	45-54	0.131	0.051	41.972	0.052	5.436
	5	High Sch.	55-65	0.104	0.099	40.489	0.046	5.585
	6	Some Coll.	35-44	0.189	0.043	42.981	0.041	6.560
	7	Bachelor's Deg.	25-34	0.225	0.070	42.963	0.032	6.944
	8	Some Coll.	45-54	0.177	0.044	42.756	0.037	7.284
	9	Advanced Deg.	25-34	0.289	0.100	43.863	0.028	8.008
	10	Some Coll.	55-65	0.146	0.096	41.212	0.035	8.313
	11	Advanced Deg.	35-44	0.331	0.054	45.086	0.020	8.466
	12	Bachelor's Deg.	35-44	0.266	0.042	44.062	0.027	8.938
	13	Advanced Deg.	45-54	0.332	0.056	45.061	0.018	10.074
	14	Bachelor's Deg.	45-54	0.254	0.044	43.781	0.026	10.430
	15	Bachelor's Deg.	55-65	0.198	0.088	42.084	0.029	10.998
Most Uneq.	16	Advanced Deg.	55-65	0.295	0.106	43.519	0.020	12.113

Source: Author's calculations based on the Current Population Survey.
Notes: (i) Data covers years 1981-2005, excluding 1998, 2000, 2002, and 2004. (ii) Universe for hours statistics is all employed men ages 25-64. (iii) Universe for unemployment rate is all individuals ages 25-64 residing in the work group who are in the labor force. (iv) Each row of the table represents one workgroup. Workgroup ID numbers are assigned according to the level of inequality in the workgroup.

can apply to the estimated coefficients on measures of inequality. Further, Stewart and Swaffield (1997) report that 47.2% of employed British men ages 21-64 desire to work fewer hours than they do. And also that "60% of male full-time employees identified fear of losing their job as an important reason for working longer hours than they used to and 47% said that 'their employer takes advantage of the fact that people are afraid of losing their jobs' (quotation marks in original) (Stewart and Swaffield 1997, p524-5)." The second point, that workers work longer hours to avoid job loss is an important insight. First, it suggests that workers also believe that work hours are a signal of effort expended at a job. It also suggests that the fact that 47.2% of male workers want to work fewer hours does not necessarily reflect a hard constraint imposed on the worker by the employer, but could reflect constraints imposed on them by their environment. One of these constraints could be a desire to work less but the employer is unwilling. It could also reflect a desire to work less but the worker cannot because she needs a promotion, or she needs the money to maintain consumption, or she is afraid of losing her job. Because all that is needed for the estimation is that workers have some freedom of choice at the margin, even if the full 47.2% of workers who want to work fewer hours face a hard constraint imposed by their boss, that still leaves 52.8% of the population who is satisfied with their work hours choices. Thus, identification may be weak, but it is still possible.

However, the inverse relationship between unemployment and the incidence of long work hours does not indicate that fear of losing one's job is not at play. An individual should only fear losing her job if it is costly to do so. Besides the psychological costs associated with job loss, the two most important determinants of job loss are the gap between the current wage and the unemployment wage (or the next best alternative) and the length of time the individual expects to stay at the lower wage. This quantity can be quite high for individuals in high wage jobs, even if unemployment rates are low because both the gap between the given wage and the next best alternative is large and the time spent at the lower wage level could be quite high. This suggests that the cost of job loss could be an important determinant of labor supply. The correlation between an individual's wage level and the cost of job loss for that individual will likely be large, implying that the inclusion of the wage level should pick up much of the variation due to changes in the cost of job loss. Because we are not interested in the magnitude of the coefficient on the wage, this does not interfere with interpretation of the estimation results.

Third, for work reference groups, the inverse relationship between work hours and unemployment could simply reflect the breadth of how work reference groups are defined here. If using occupation as the reference group is at one end of the breadth spectrum, using (age,education) cells is at the other. The "correct" definition certainly lies somewhere in between. However, identifying the single correct measure is a futile effort. Thus, the broad measure of (age,education) cells and the narrow measure of occupation cells will be used to bracket the results. The problem with the broad measure is that the dispersion in types of jobs an

individual can hold increases with the level of education and experience. For example, the group with the highest level of inequality is men ages 55-65 with an advanced degree. This group could potentially run the entire range of employment relationships, it could include substitute teachers as well as wall street analysts, semi-retired individuals as well as men making a last ditch effort to increase their retirement savings. If this is the case, it is not surprising that the dispersion of work hours increases as inequality increases. Kuhn and Lozano discuss this issue when moving from two digit occupation groups to three digit occupations. They rule out the possibility that the results they find are driven entirely by how broadly they define an occupation. This suggests that a similar result may hold here.

For the sake of comparability with existing work, table 4 repeats the exercise in table 2 by quantiles of occupation inequality measured by the standard deviation of wages within an occupation. Occupations are sorted into 10 quantiles based on the degree of wage inequality within the occupation, with quantile 1 representing the lowest level of inequality. There is a clear trend towards higher shares of men working long hours and longer mean hours as occupational wage inequality increases. There is also a downward trend in mean occupational unemployment as inequality increases. Therefore, by defining the work reference group more narrowly, the trends that are present for both MSAs and (age,education) groups are magnified. This suggests that comparing the estimated coefficients on occupation inequality and (age,education) inequality does in fact resemble an upper bound and a lower bound on the effect on inequality.

Table 4. Hours statistics by quantiles of occupational inequality

Degree of Ineq.	Quantile	Work>=49	Work<35	Weekly Hours	Unemployment
Lowest Ineq.	1	0.108	0.127	40.550	0.087
	2	0.103	0.077	41.029	0.078
	3	0.127	0.076	41.497	0.067
	4	0.145	0.065	41.996	0.063
	5	0.145	0.050	42.225	0.054
	6	0.178	0.056	42.620	0.048
	7	0.224	0.053	43.300	0.037
	8	0.247	0.047	43.820	0.028
	9	0.304	0.054	44.426	0.032
Highest Ineq.	10	0.239	0.071	43.307	0.038

Source: Author's calculations based on the Current Population Survey.
Notes: (i) Data covers years 1981-2005, excluding 1998, 2000, 2002, and 2004. (ii) Universe for hours statistics is all employed men ages 25-64. (iii) Universe for unemployment rate is all individuals ages 25-64 residing in the occupation who are in the labor force.

IV. Regressions

Before proceeding, it would be useful to briefly review the predictions of the two hypotheses that currently exist in the literature: the promotion incentive and consumption emulation arguments. The promotion incentive argument is built on inequality in the workplace and more generally an individual's work reference

group, defined here by either (age,education) cells or occupations, with the former referred to as the 'work-group'. The degree of inequality, measured by the standard deviation of wages in the reference group, should be positively correlated work hours. No explicit claims have been made about whether the strength of this effect should be related to the income level. Thus, if the promotion incentive argument holds, there should be a positive coefficient on workgroup inequality, and either a positive or negative coefficient on the interaction between workgroup inequality and an individual's wage. Further, this effect should be considerably larger than the effect of MSA inequality.

The consumption emulation argument hypothesizes that inequality within an individual's community, possibly including work peers, but also including many more people, should be positively correlated with work hours. The theoretical models that underlie this hypothesis make no claim about how the effect of this relationship should vary with the wage. An example of this class of models is in appendix A. Therefore, if the consumption emulation argument holds, there should be a positive coefficient on MSA inequality.

Table 5 shows the first set of regressions. The dependent variable of all of the regressions is the log of usually weekly work hours. The universe is all employed men ages 25-64 from 1981 to 2005. The first two columns show the results for the pooled regressions, the remaining columns are fixed effects regressions using the dummy variable technique. This facilitates the estimation of the clustered standard errors. As Moulton (1990) points out, when a regression contains multiple observations from within the same group and group level variables are included as independent variables in the regression, the standard errors will be biased downwards (upwards) if the standard errors are positively (negatively) correlated. In this case, each individual is a member of three non-nested groups: (i) the individual's set of observations across time, the MSA the individual resides in, and either the individual's workgroup or the occupation depending on the specification.

These clusters are non-nested, which is not possible to estimate using the conventional clustering techniques. Each individual observation is nested in the cluster of her own observations through time. However, across time at least some of the individuals will change MSAs, workgroups, and/or occupations. If this were not the case, standard errors would be clustered at the highest level of aggregation. But, because the clusters are non-nested, standard errors must be clustered on all groups, simultaneously. This is called multi-way clustering. The technique used here is outlined in Cameron, Gelbach, and Miller (2006). For simplicity, clustering was performed only on the MSA and either the workgroup or occupation, not on the individual. Because a given individual's errors are surely positively correlated, the resulting standard errors are an order of magnitude smaller than if clustering were performed on all three, though some of this is captured by the fixed effect. But, as will become clear, this issue is not of great importance. With an even number of clustering variables, the clustering procedure amounts to first clustering standard errors on one

Table 5. Regressions for All Employed Men Ages 25-64: 1981-2005

Variables	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)
Log Wage	−0.000275	0.009133	−0.005593**	−0.003672	−0.005727**	−0.003905
	0.003072	0.006744	0.002479	0.002882	0.002522	0.002868
Log Age	−0.029780	−0.056549*	0.079680*	0.061711*	0.087321*	0.060975*
	0.029110	0.019338	0.028638	0.019451	0.029180	0.019354
Workgroup Ineq.	0.006735*		−0.000293		−0.000379	
	0.001925		0.001968		0.001903	
Occ. Ineq.		0.003120		0.000616		0.000593
		0.002288		0.000661		0.000566
MSA Ineq.	−0.001919***	−0.000240	−0.000343	−0.000887	−0.000375	−0.000962
	0.001126	0.001461	0.000779	0.001122	0.000778	0.001074
Workgroup Ineq. x Wage	−0.000053		−0.000111***		−0.000111***	
	0.000070		0.000065		0.000066	
Occ. Ineq. x Wage		−0.000055		−0.000049		−0.000050
		0.000055		0.000033		0.000033
MSA Ineq. x Wage	0.000020	0.000007	0.000004	−0.000073***	0.000004	−0.000069***
	0.000069	0.000048	0.000051	0.000039	0.000051	0.000039
Workgroup Unemp.					0.026537	
					0.092861	
Occ. Unemp						−0.213745**
						0.079105
MSA Unemp.					0.145069	0.289152**
					0.115426	0.122390
Constant	3.851476*	3.922842*	3.504753*	3.552682*	3.469681*	3.553780*
	0.103543	0.057551	0.097181	0.067804	0.102627	0.069957
N	34467	41958	34467	41958	34467	41958

Source: Author's calculations based on PSID linked to CPS.
Notes: (i) Significance stars: * 0.1, ** 0.05, and *** 0.1. (ii) Dependent variable in all regression is log of usual weekly work hours. (iii) Universe is all employed men ages 25-64. (iv) Fixed effects are estimated using the dummy variable approach. (v) Inequality within a given reference group is the standard deviation of real hourly wages in that group. (vi) Errors are clustered on MSA and workgroup in columns 1, 3, and 5. Errors are clustered on MSA and occupation for columns 2, 4, and 6.

of the groups, then clustering standard errors on the second group and adding the resulting standard errors together. Then, clustering standard errors on individuals who reside in both groups simultaneously, and subtract this from the previous sum.[9]

In the first column of table 5, workgroup and MSA inequality are included, both measured by the standard deviation of wages within the respective groups. Workgroup has a positive, and statistically significant coefficient, suggesting that the relationship found in table 3 is robust to the inclusion of simple control variables such as age and the log of wages. The coefficient on MSA inequality is negative and significant, immediately undermining the validity of the consumption emulation hypothesis as specified here. The coefficient on occupation inequality in column 2 is positive, but it is insignificant, economically not meaningful, and less than half the size of the coefficient on workgroup inequality. When occupation inequality is included, the coefficient on MSA inequality is larger, however it is negative and insignificant.

In columns 4 through 6 fixed effects are used. The coefficient on workgroup inequality turns negative and insignificant. This switch of signs is important not only because it undermines the support for the promotion incentive argument found in column 1, but also because it suggests that individual specific characteristics that determine the choice of reference group or mean work hours is playing a role here. In particular, individuals with higher levels of education tend to work longer hours, this effect is captured in column 1 by workgroup inequality but by the fixed effect in column 3. The coefficient on MSA inequality remains negative but is no longer significant.

In columns 4 and 6, the coefficient on occupation inequality remains positive, but is approximately 20% as large when fixed effects are included. The coefficient on MSA inequality maintains a negative sign when occupation inequality is included, and is still insignificant. The other results of note are the consistently negative wage elasticity when fixed effects are included, the positive coefficient on log age when fixed effects are included, the large negative coefficient on occupation unemployment, and the large positive coefficient on MSA unemployment when occupation inequality is included. This will be returned to later.

In table 5 little support is found for either the consumption emulation or the promotion incentive hypotheses. Both of these theories, however, have fairly clear implied empirical predictions for the relative magnitudes of the coefficients on inequality between population subgroups. The promotion incentive argument hinges on work hours being a non-noisy signal of effort. For hourly workers, the signal is masked by the inability of the firm to differentiate between a wage worker working longer hours to earn more income, and a worker trying to signal high levels of effort and commitment to the firm. For salaried workers, this problem

[9]The authors provide a simple program for one commonly used statistics software package that implements this clustering technique for any number of clustering variables following a standard OLS regression. That program was not used here because of constraints on the number of variables the statistical package can handle predluding estimation of individual fixed effects with dummy variables.The program is available at http://strategy.sauder.ubc.ca/head/sup/cgmreg.ado.

does not exist. Thus, we would expect hourly workers to be more responsive to inequality than salaried workers if they are responding to the promotion incentive. [10] The relationship for consumption emulation is somewhat different. For consumption emulation, we would expect the coefficient on MSA inequality to be larger for hourly workers than salaried workers precisely because hourly workers can increase their income by working longer hours in the short-run.

Table 6 reports the results of this exercise. Columns 1 and 5 are identical to columns 5 and 6 in table 5. Columns 2 and 4 limit the sample to full-time workers. For a number of reasons, it might be reasonable to consider full-time workers as a distinct group from part-time workers. [11] Most notably, it may be the case that more part-time workers face a binding constraint on maximum work hours for reasons somewhat independent of labor market circumstances. Because of this, full-time workers are likely to be more responsive to labor market incentives than part-time workers. The coefficient on MSA inequality is negative when the sample is limited to full-time workers, and significant when occupation inequality is included in column 6. The coefficient on workgroup inequality is negative and insignificant with essentially the same magnitude as for all male workers. The coefficient on occupation inequality decreases somewhat, remaining positive and insignificant. Thus, the responsiveness of full-time workers to inequality is not qualitatively different than for all workers. Interestingly, full-time workers are less than half as responsive to MSA inequality when occupation inequality is included than all workers.

Turning now to the difference between salaried and hourly workers, again no support is found for either argument. The coefficient on workgroup inequality is larger for salaried workers than hourly workers, but the coefficients are both statistically insignificant, and not statistically distinguishable from each other. The coefficient on MSA inequality is larger for salaried workers than hourly workers, which is the opposite that the consumption emulation hypothesis implies, but is insignificant in both cases and again not distiniguishable from each other. A similar pattern of results holds for occupation inequality. That is, the results are statistically insignificant, and the coefficient is on MSA inequality is slightly larger for salaried workers than hourly workers. In short, except for the coefficients on the unemployment variables, the results are quite robust to changes in specification and sample.

V. Interpreting the Results

Little support is found for either the consumption emulation or promotion incentive hypotheses as they relate to labor supply decisions. After including fixed effects via the dummy variable method, and clustering

[10] Kuhn and Lozano (2008) make a similar point.

[11] This is not to say the full-time workers do not facing binding constraints on maximum work hours from non-labor market sources. For example, there are likely full-time workers who would work even longer hours if they did not have to let the dogs out, pick up the kids from school, or take care of a parent or spouse. All that is being claimed is that the incidence of this *might* be higher among part-time workers.

Table 6. Regressions for All Employed Men, Salaried, and Hourly Men Ages 25-64: 1981-2005

	Workgroup				Occupation			
	All	Full-time	Salaried	Hourly	All	Full-time	Salaried	Hourly
Variables	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)	Coefficient (Std. Err.)
Log Wage	−0.005727**	−0.007525*	−0.003275	0.083444*	−0.003905	−0.008391*	0.001682	0.080777*
	0.002522	0.002017	0.002705	0.017898	0.002868	0.001773	0.003660	0.019226
Log Age	0.087321*	0.111822*	0.125438**	−0.022205	0.060975*	0.095364*	0.108549*	−0.075304**
	0.029180	0.018070	0.050115	0.041731	0.019354	0.013567	0.025489	0.036034
Workgroup Ineq.	−0.000379	−0.000376	0.000464	−0.002896				
	0.001903	0.001104	0.001319	0.003685				
Occ. Ineq.					0.000593	0.000382	0.001352	0.000260
					0.000566	0.000444	0.001170	0.000492
MSA Ineq.	−0.000375	−0.000972	0.000765	−0.000918	−0.000962	−0.001008**	0.000494	0.000119
	0.000778	0.000603	0.001400	0.001263	0.001074	0.000476	0.001809	0.000902
Workgroup Ineq. x Wage	−0.000111***	−0.000042	−0.000074	−0.000250				
	0.000066	0.000040	0.000063	0.000168				
Occ. Ineq. x Wage					−0.000050	−0.000015	−0.000036	−0.000091**
					0.000033	0.000017	0.000074	0.000043
MSA Ineq. x Wage	0.000004	−0.000007	−0.000051	0.000013	−0.000069***	−0.000030***	−0.000143	−0.000095**
	0.000051	0.000030	0.000065	0.000120	0.000039	0.000018	0.000105	0.000043
Workgroup Unemp.	0.026537	0.011730	0.157610***	−0.097821				
	0.092861	0.036687	0.090429	0.089399				
Occ. Unemp					−0.213745**	−0.141878*	−0.112638	−0.147528
					0.079105	0.049800	0.087309	0.093444
MSA Unemp.	0.145069	0.125497	0.082142	−0.038240	0.289152**	0.139077***	0.234630	−0.038246
	0.115426	0.086544	0.204465	0.110015	0.122390	0.081062	0.174839	0.120004
Constant	3.469681*	3.418078*	3.358774*	−0.073720	3.553780*	−0.286007*	−0.387366*	0.094788
	0.102627	0.064644	0.180105	0.110388	0.069957	0.048116	0.090513	0.099236
N	34467	33389	13757	20710	41958	40165	16563	25395

Source: Author's calculations based on PSID linked to CPS.
Notes: (i) Significance stars: * 0.1, ** 0.05, and *** 0.1. (ii) Dependent variable in all regression is log of usual weekly work hours. (iii) Universe is all employed men ages 25-64. (iv) Fixed effects are estimated using the dummy variable approach. (v) Inequality within a given reference group is the standard deviation of real hourly wages in that group. (vi) Errors are clustered on MSA and workgroup in columns 1, 3, and 5. Errors are clustered on MSA and occupation for columns 2, 4, and 6.

the errors to take into account the non-nested clusters the coefficients on all measures of inequality are statistically equal to zero and often have the wrong signs. The coefficient on workgroup wage inequality is consistently negative and insignificant, with the exception of salaried workers where it is positive and insignificant, but not statistically distinguishable from the coefficient for hourly workers. The coefficients on occupation inequality are consistently positive, but are very small and insignificant. The coefficients on MSA inequaltiy range from positive, small and insignificant, to negative and statistically insignificant. Further, when the sample is divided between hourly and salaried workers, the relative magnitudes of the coefficients on the various measures of inequality do not follow the implied patterns of either consumption emulation of promotion incentive effects.

The results found here are quite different from those found by other authors. Bell and Freeman (2001) find that higher levels of occupation inequality are associated with higher mean occupation work hours in a cross-country setting using cross-sectional data, though the data from the two countries is not pooled. Bowles and Park (2005) find a positive relationship between country level inequality and mean annual work hours in the country using a cross-country panel data set. And Kuhn and Lozano (2008), a paper closely related to and written simultaneously with this one, also find a positive correlation between occupation inequality and mean occupation work hours in the United States using a cohort analysis from years 1980 to 2002.

Setting the Bowles and Park paper aside because of the difficulty in comparing cross-country analyses with within country analyses, the regressions that are closest to those of Bell and Freeman and Kuhn and Lozano are the the pooled results in columns 1 and 2 of table 5. But, even in this case, occupation inequality is not statistically significant, although workgroup inequality is. After the fixed effects are included, the positive impact of workgroup inequality goes away. This suggests that the reflection problem does bias the coefficients on the measures of inequality. Because the other studies could not control for self-selection into occupations, the estimated coefficients reflect the combined effect of inequality on labor supply conditional on being a member of the group, and the correlation among members of the group who likely share many unmeasured characteristics. This does not appear to be the case for MSA inequality. Except for column 1 of table 5, the coefficient on MSA inequality is consistently insignificant, regardless of the presence of fixed effects.

This pattern points towards a number of possible explanations, though testing the various possibilities is beyond the scope of the paper. The first possible explanation is that the reference group for work inequality has been misspecified. It may be the case that workers simply compare their wage to every wage in the local area that they work in. If the return to promotion in the local area is generally high, then they work long hours in order to gain a promotion. It could also be the case that salaried workers do work long hours to get

a promotion, but because of consumption emulation motives and not promotion incentive motives, thus the positive coefficient on MSA inequality for salaried workers. This is the preferred explanation if consumption emulation motives are reflected in labor supply decisions and reference groups are specified correctly.

A third possibility is that neither consumption emulation nor promotion incentive motives show up in labor supply decisions, conditional on being in one of the reference groups. If this is the case, then the results that Bell and Freeman (2001) and Kuhn and Lozano (2008) find are simply reflecting differences in wage and hours setting institutions, norms, and "cultures" across occupations. A lawyer does not become lawyer, find herself confronted with a particular pay distribution, and then choose to work long hours. The lawyer chooses to become a lawyer knowing that she will likely have to work long hours. Lawyers also confront a highly unequal distribution of income because the earnings rates vary greatly across the different fields of law, the different types of law firms, and different geographic locations. The same may be true for consumption emulation. An individual may be predisposed, in some sense, to engage in consumption emulation. Regardless of the degree of inequality or their income, they will still engage in consumption emulation. Thus working long hours in order to fulfill the need to have high relative consumption.

Neither of these behaviors would be reflected in the estimated coefficients presented here. They would be captured by a cross-section or cohort analysis. But, this is a fundamentally different question that what either Bell and Freeman (2001) or Kuhn and Lozano (2008) claim to be answering. The question of whether individuals choose work environments with unequal distributions of income because the return to promotion is higher, and then work long hours to get a promotion is different then if individuals work long hours because they are in a work environment that presents them with a large return to promotion. The same goes for consumption emulation. Evidence of self-selection into tournaments has beend documented in experiments.

It is important to note what is not being claimed. These results do not imply that the tournament model is the wrong model of the labor market. The results also do not imply that consumption emulation does not exist. The most commonly referenced model of consumption externalities, Frank (1985a), does not imply that higher levels of inequality should increase consumption on positional goods. And although experimental work on tournaments shows that mean chosen effort in a group is increasing in the degree of inequality, they also show that the number of people who do not actively compete is also increasing in the degree of inequality (2003). How this plays out in the real world is anybody's guess at this point. All that is being claimed is that consumption emulation and promotion incentive effects do not determine labor supply in the way that others have suggested up to this point.

VI. Conclusion

The idea that Americans are overworked is not a myth. Although mean work hours have held steady for men over the last 30 years, the share of men working at least 49 hours or no more than 35 hours per week is rising. This lead to an increase in the dispersion of work hours for men ages 25-64 over the last 30 years. Simultaneously, wage inequality has also increased considerably. Recent theoretical and experimental work suggests that there may be a causal link running from changes in the income distribution to the observed changes in labor supply. This paper further examines the possibility of such a causal link.

Thus far, there are not but a small hand full of papers that attempt to explicitly link income inequality to labor supply. Broadly, they fall under two distinct categories. One line of research is based on Lazear and Rosen's (1981) tournament model of the firm. In this model, the marginal utility of effort is determined by the income distribution. The wider the distribution, the greater the return to a promotion, the more effort the worker expends. In the real world where effort and output might be unobservable, workers devise a signal of effort. One likely candidate is labor supply. Thus, an increase in income inequality raises the marginal return to being promoted, so workers work longer hours. The second strand of research is based on Veblen's (1899) consumption emulation hypothesis. Veblen argues that the consumption standards of a given individual are partially determined by the consumption standard of the individual's reference group. Depending on how the reference group is defined, when inequality increases, individual consumption relative to the reference group decreases. To make up for the loss in utility, individuals consume more.

In order to test the competing hypotheses, it is crucial that the researcher be able to disentangle work inequality for the promotion incentive hypothesis from social group inequality for the consumption emulation hypothesis. To the author's knowledge, this is the first paper that is able to accomplish this. This is also the first paper to use a true micro panel data set to estimate the effect of inequality. Specifically, the Current Population Survey is used to calculate statistics on work and social (community) reference groups. This information is linked to the Panel Study of Income Dynamics using a set of geographic identifiers not publicly available. For the promotion incentive, the PSID is linked to the CPS by occupation as a lower bound and (age,education) cells as an upper bound of the degree of inequality a given worker experiences. For community inequality, the data set is linked by Metropolitan Statistical Area.

Little support is found for either the promotion incentive or the consumption emulation hypothesis. For all employed men ages 25-64, the effect of both types of inequality ranges from statistically insignificant and positive to statistically significant but negative. This remains the case when the sample is limited to only full-time men. When the sample is divided between salaried and hourly workers, some support is found for both hypotheses, though it is weak at best. Namely, the coefficient on work inequality is smaller for hourly

workers than salaried workers, but it is negative in both cases. The coefficient on community inequality is larger for salaried workers than hourly workers, the opposite of the implied prediction of the consumption emulation hypothesis.

The conclusion is that differences in the level of inequality do not determine differences in hours worked in a manner consistent with how consumption emulation or promotion incentives have been specified up to this point. It is most likely the case that there are underlying institutional differences between, in particular, work groups that jointly determine both the distribution of work hours and the distribution of wages in that group. Or, that individuals choose reference groups, particularly for work, because of the distribution of income and the possibility of promotion. They work long hours to gain a promotion, but they are not responsding to the degree of inequality after having chosen the occupation of workgroup. But, this is a different question then if individuals respond to the degree of inequality, haven chosen a reference group.

References

Akerlof, George, "Labor Contracts as Partial Gift Exchange," Quarterly Journal of Economics, 1982, 97 (4), 543–569.

Akerlof, George A., "Social Distance and Social Decisions," Econometrica, 1997, 65 (5), 1005–10027.

___ and Janet L. Yellen, "The Fair Wage-Effort Hypothesis and Unemployment," The Quarterly Journal of Economics, 1990, 105 (2), 255–283.

___ and Rachel E. Kranton, "Identity and the Economics of Organizations," Journal of Economic Perspectives, Winter 2005, 19 (1), 9–32.

Altonji, Joseph G. and Christina H. Paxson, "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs," Journal of Labor Economics, 1988, 6 (2), 254–276.

Austin, William, "Equity Theory and Social Comparison Processes," in Jerry M. Suls and Richard L. Miller, eds., Social Comparison Processes, Washington: Hemisphere Publishing, 1977, chapter 12, pp. 279–306.

Bell, Linda and Richard Freeman, "The Incentive for Working Hard: Explaining Hours Worked Differences in the US and Germany," Labour Economics, 2001, 8 (2), 181–202.

Bowles, Samuel, "The Production Process in a competitive economy: walrasian, neo-hobbesian, and marxian," The American Economic Review, winter 1985, 75 (1), 83–102.

___ and Yongjin Park, "Emulation, Inequality, and Work Hours: Was Thorsten Veblen Right?," Economic Journal, 2005, 115 (507), F397–412.

___ , Herbert Gintis, and Melissa Osborne, "The Determinants of Earnings: A Behavioral Approach," Journal of Economic Literature, December 2001, 39 (4), 1137–76.

Brock, W. and S. Durlauf, "Discrete Choice with Social Interactions," Review of Economic Studies, April 2001, 68 (2), 235–260.

Brock, William A. and Steven N. Durlauf, "Interactions Based Models," in James J. Heckman and Edward Leamer, eds., Handbook of Econometrics, Vol. 5, Amsterdam: North-Holland, 2001.

Burke, Peter J., "Identities and Social Structure: The 2003 Cooley-Mead Award Address," Social Psychology Quarterly, 2004, 67 (1), 5–15.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference with Multi-way Clustering," NBER Technical Working Paper No. 327, September 2006.

carl shapiro and joseph e. stiglitz, "equilibrium unemployment as a worker discipline device," The American Economic Review, june 1984, 74 (3), 433–444.

Devereux, Paul J., "Changes in Relative Wages and Family Labor Supply," Journal of Human Resources,

Summer 2004, 39 (3), 696–722.

Dickens, Shelly J. Lundberg William T., "Hours Restrictions and Labor Supply," International Economic Review, 1993, 34 (1), 163–191.

Duesenberry, James Stemble, Income, Saving, and the Theory of Consumer Behavior, Cambridge: Harvard University Press, 1949.

Durlauf, S., "Neighborhood Effects," in J. V. Henderson and J.F. Thisse, eds., Handbook of Regional and Urban Economics, Vol. 4, Amsterdam: North Holland, 2004.

Fehr, Ernst and Klaus M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," Quarterly Journal of Economics, 1999, 114 (3), 817–868.

__ , Erich Kirchler, Andreas Weichbold, and Simon Gachter, "When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets," Journal of Labor Economics, 1998, 16 (2), 324–351.

Frank, Robert H., Choosing the Right Pond: Human Behavior and the Quest for Status, New York:Oxford University Press, 1985.

__ , "The Demand for Unobservable and Other Nonpositional Goods," The American Economic Review, 1985, 75 (1), 101–116.

__ , Luxury Fever, Princeton: Princeton University Press, 2000.

__ , "Positional Externalities Cause Large and Preventable Welfare Losses," The American Economic Review, May 2005, 95 (2), 137–141.

__ , Falling Behind: How Rising Inequality Harms the Middle Class, Berkeley: University of California Press, 2007.

__ and Adam Seth Levine, "Expenditure Cascades," American Economic Review, 2007.

Harbring, Christine and Bernd Irlenbusch, "An Experimental Study on Tournament Design," Labour Economics, 2003, 10, 443–464.

Hopkins, Ed and Tatiana Kornienko, "Running to Keep in the Same Place: Consumer Choice as a Game of Status," The American Economic Review, September 2004, 94 (4), 1085–1107.

Jacobs, Jerry A. and Kathleen Gerson, The Time Divide: Work, Family, and Gender Inequality, Cambridge, MA: Harvard University Press, 2005.

Juhn, Chinhui and Kevin M. Murphy, "Wage Inequality and Family Labor Supply," Journal of Labor Economics, 1997, 15 (1), 72–97.

Kahn, Shumalit and Kevin Lang, "The Effects of Hours Constraints on Labor Supply Estimates," Review of Economics and Statistics, 1991, 73 (4), 605–611.

Kuhn, Peter and Fernando Lozano, "The Expanding Workweek? Understanding Trends in Long Workk Hours among U.S. Men, 1979-2006," Journal of Labor Economics, 2008, 26 (2), 311–343.

Landers, Renee M., James Rebitzer, and Lowell J. Taylor, "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms," The American Economic Review, June 1996, 86 (3), 329–348.

Lazear, Edward P. and Sherwin Rosen, "Rank-Order Tournaments as Optimum Labor Contracts," Journal of Political Economy, 1981, 89 (5), 841–864.

Levine, David I., "Fairness, Markets, and Ability to Pay: Evidence from Compsensation Executives," The American Economic Review, 1993, 83 (5), 1241–59.

Luttmer, Erzo F. P, "Neighbors as Negatives: Relative Earnings and Well-Being," The Quarterly Journal of Economics, August 2005, 120 (3), 963–1002.

Manski, Charles F., "Identification of Endogenous Social Effects: The Reflection Problem," Review of Economic Studies, 1993, 60 (3), 531–542.

Mishel, Lawrence, Jared Bernstein, and Sylvia Allegretto, The State of Working America 2006/2007 An Economic Policy Institute Book, Ithaca, N.Y.: ILR Press, an imprint of Cornell University, 2007.

Moulton, Brent R., "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units," The Review of Economics and Statistics, May 1990, 72 (2).

Park, Yongjin, "The Second Payckeck to Keep up with the Joneses: Relative Income Concerns and Labor Market Decisions of Married Women," University of Massachusetts, Amherst Working Paper, 2005, 1005-10.

Schor, Juliet B., The Overworked American: the Unexpected Decline of Leisure, New York, NY: Basic Books, 1993.

___ , The Overspent American: upscaling, downshifting, and the new consumer, New York: Basic Books, 1998.

___ , Born to Buy: The Commercialized Child and the New Consumer Culture, New York, NY: Scribner, 2004.

Stewart, Mark and Joanna Swaffield, "Constraints on the Desired Hours of Work of British Men," Economic Journal, 1997, 107 (441), 520–535.

Veblen, Thorstein, The Theory of the Leisure Class, New York: Macmillan, 1899.

Weinberg, Bruce, Patricial B. Reagan, and Jeffrey J. Yankow, "Do Neighborhoods Affect Work Behavior? Evidence from the NLSY79," Journal of Labor Economics, 2004, 24 (4), 891–924.

Appendix

A. Model of Veblen Effects

The theoretical underpinnings of the existing work on the causal link running from income inequality to labor supply is based on two basic arguments: the tournament model of Lazear and Rosen (1981) and Veblen's (1899) consumption emulation argument. Because Lazear and Rosen's model is quite well know, no further elaboration is needed. However, there is no agreed upon model representing consumption emulation, or Veblen effects. Therefore, it is useful to lay out a simple model that contains the core arguments put forth by the collection of authors that have written about Veblen effects. This model borrows elements from both Frank (1985b) and Akerlof (1997, 2005).

Assume that the individual solves the choice problem laid out in equation 1, where utility is a function of three arguments: individual consumption (c), leisure (l), and relative consumption ($\hat{c} - c$), $\hat{c}$ is defined in equation 2 and i indexes individuals. Further, maintain the standard assumptions that u is increasing and concave in both c and l, $l \in [0, 1]$, and $u_{(c-\hat{c})} > 0$ but may be either convex or concave. Finally, assume that the individual does not take into account the impact of changes in her consumption on $\hat{c}$, so $\frac{d\hat{c}}{dc} = 0$ even though, given the way $\hat{c}$ is specified, this may not be the case.

$$\max_{c,l} u_i = u(c_i, l_i, c_i - \hat{c_i}) \text{ s.t. } w_i(1 - l_i) = c_i \tag{1}$$

$$\hat{c} = \frac{\sum_1^k B}{k} \text{ where } B = \{c | c_{max} > c > c_i\} \text{ and } k = \# \text{ elements in } B \tag{2}$$

The key term is relative consumption. Reference group consumption, $\hat{c_i}$, is defined as the mean of all consumption greater than c_i. In line with Veblen's arguments and those of other authors, all individuals are assumed to compare upwards regardless of the individual's position in the income distribution. Relative consumption is measured as the difference between individual consumption and reference group consumption, scaled by individual consumption. Note also that defining $\hat{c_i}$ in this manner means that mean reference group consumption is a decreasing function of c_i. This captures the assertion that although individuals do compare upwards, they do not compare themselves to the top of the income distribution.

There are two problematic features of this definition of reference group consumption. First, any increase in consumption inequality due to changes in consumption below c_i will leave $\hat{c_i}$ unchanged. This feature does seem to be inkeeping with Veblen's arguments, but does make the connection between changes in inequality and changes in labor supply less direct. However, for the case of the U.S. since 1980, there is no doubt that $\hat{c_i}$ was increased with the increase in inequality. Second, the consumption gap is not scaled. The definition of

$\hat{c}_i$ implies that the consumption gap is relatively constant across the income distribution because $\hat{c}_i$ increases with c_i. But, it also means that given sized gap causes the same loss in utility regardless of the level of c_i. The most likely choice to normalize the consumption gap is to divide $c_i - \hat{c}_i$ by c_i. But, this adds the stronger assumption that the effect of reference group consumption decreases with c_i because of how $\hat{c}_i$ is defined. Also, scaling in this matter does not change the results in any meaningful ways.

It turns out that the results of the utility maximzation process are not very sensitive to specification of the utility function. Therefore, we will work with the most familiar of the quasi-concave utility functions, the Cobb-Douglas utility function, given in equation 3.

$$\max_{c,l} u_i = c_i^{\alpha} l_i^{\beta} (c_i - \hat{c}_i)^{\gamma} \text{ s.t. } w_i(1 - l_i) - c_i = 0 \tag{3}$$

where everything is defined as above, and $\alpha > 0$, $\beta > 0$, and $\gamma > 0$. The first order condition for this problem is given in equation 4.

$$\frac{\alpha l_i}{\beta c_i} + \frac{\gamma l_i}{\beta(c_i - \hat{c}_i)} = \frac{1}{w} \tag{4}$$

Although the derivative of interest is $\frac{dl}{d\hat{c}}$, it is easier to find the sign of $\frac{dc}{d\hat{c}}$, and use the sign of this derivative to sign the other. If we first solve the first order condition for l in terms of c and $\hat{c}$, we get equation 5.

$$l_i = \frac{1}{w}\left[\frac{\beta c_i}{\alpha} + \frac{\beta(c_i - \hat{c}_i)}{\gamma}\right] = f(c_i, \hat{c}_i) \tag{5}$$

Substituting equation 5 into the budget constraint yields $c_i - w_i(1 - f(c_i, \hat{c}_i)) = F(c_i, \hat{c}_i)$. The implicit function thereom can be applied to this equation, yielding the derivative in equation 6.

$$\frac{dc_i}{d\hat{c}_i} = -\frac{F_{\hat{c}}}{F_c} = -\frac{w f_{\hat{c}}}{1 + w f_c} > 0 \tag{6}$$

where

$$f_c = \left(\frac{\beta}{\alpha} + \frac{\beta}{\gamma}\right)\frac{1}{w} > 0$$
$$f_{\hat{c}} = \frac{-\beta}{\gamma w_i} < 0$$

Because the budget constraint precludes the use of savings or credit to fund consumption, any increase in consumption must be accompanied by an increase in labor supply. Thus, $\frac{dl_i}{d\hat{c}_i} < 0$ is implied by equation 6, meaning that individual i works longer hours in response to an increase in reference group consumption. Any increase in inequality that causes an increase in $\hat{c}_i$ results in increased consumption, and decreased

leisure. But, this only holds when utility in increasing and concave in relative consumption. And, if relative consumption is measured as a simple difference, $\hat{c}_i$ must always be greater than c. If it is a ratio, then it must be specified as $\frac{\hat{c}_i}{c_i}$. For many reasons beyond the scope of either the paper of the appendix, this assumption may not be appropriate.

Finally, the cross-partial derivative $\frac{\partial c_i^2}{\partial \hat{c}_i \partial w_i}$ is indeterminate.

$$\frac{\partial c_i^2}{\partial \hat{c}_i \partial w_i} = -\frac{(f_{\hat{c}_i} + w_i f_{\hat{c}_i w_i})(1 + w_i f_{c_i}) - w_i f_{\hat{c}_i}(f_{c_i} + w_i f_{c_i w_i})}{(1 + w_i f_{c_i})^2} \tag{7}$$

where f_j is the partial derivation with respect to the jth argument, f_{jj} is the second derivative with respect to the jth argument, and f_{jk} for $j \neq k$ is the cross-partial derivative.

The sign of this derivative depends on

$$\frac{\partial c_i^2}{\partial \hat{c}_i \partial w_i} < 0 \text{ iff } -f_{\hat{c}_i}(1 - w^2 f_{c_i w_i}) < w f_{\hat{c}_i w_i}(1 + w_i f_{c_i}) \tag{8}$$

Intuitively, this condition depends on the complicated relationship between the marginal utility of consumption, the marginal disutility of $\hat{c}_i$, and the respective cross-partial derivatives with respect to the wage. If the marginal utility of consumption is sufficiently large, then this sign is always fulfilled. If the marginal utility of consumption is low, or the marginal disutility of $\hat{c}_i$ is sufficiently large then the cross partial is positive, and the effect of relative consumption increases with the wage.