NBER WORKING PAPER SERIES

PERCEPTIONS OF EQUITY AND THE DISTRIBUTION OF INCOME

Julio J. Rotemberg

Working Paper 5624

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1996

I wish to thank Mauro Guillen, Frank Levy, Mancur Olson, and Michael Woodford and the participants at seminars at Columbia, Harvard, Johns Hopkins, the NBER and the University of Maryland for helpful suggestions and the NSF for research support. This paper is part of NBER's research program in Economic Fluctuations and Growth. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper 5624
June 1996

PERCEPTIONS OF EQUITY AND THE DISTRIBUTION OF INCOME

ABSTRACT

This paper builds a simple model where there is a connection between employees' perception of the "fairness" of employers and the actual distribution of income. Wages are based in part on employers' assessments of the productivity of individual employees. I show that the equilibrium distribution of income depends on the beliefs of employees concerning the accuracy of these evaluations. I give conditions under which the distribution of income across employees of the same vintage is more equal if employees believe that these evaluations are generally inaccurate (so that they are skeptical of capitalists in general) than when they believe that these evaluations are accurate. The model is consistent with the fact that, in a sample of seven countries, the distribution of income is more unequal in countries where people feel that income inequality is not too large.

Julio J. Rotemberg
Sloan School of Management
E52-432
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02139
and NBER

People have opinions about the degree to which the differences in pay reflect differences in productivity. In particular, they have views on the extent to which differences in income reflect, instead, personal connections, are the result of favoritism or result from capricious decisions by people with influence. In this paper, I show that opinions of this sort can affect the actual distribution of income even when they do not affect individuals' abilities. In particular, a change in the degree to which people believe that differences in pay reflect differences in productivity can change the actual income of people even if their marginal product remains unchanged. [1]

My focus on the effects of these beliefs raises the question of where these beliefs come from. In particular, there is the issue of whether people's perceptions about the degree to which individual income reflects individual productivity shouldn't be identical to the actual connection between income and pay. In this paper, I allow the actual connection between these variables to differ from the perceived connection and I do this for two reasons.

First, it is very difficult to know what the actual connection between productivity and pay actually is. If one leaves aside the possibility of experimenting by varying one's own productivity, it is very hard to learn about this connection unless one actually observes the output of different people. [2] But, in modern industrial firms where workers are extremely interdependent and where the aggregate value produced by the firm is subject to large random disturbances, measuring the output of individuals is very hard. There are, of course, indirect measurements of the relationship between productivity and pay. Under the assumptions of the stripped down model I present, for instance, this relationship can be learned by measuring the changes in the income of people who change jobs. However, the extent to which statistics of this sort shed light on the issue depends on auxiliary assumptions, the validity of which is itself difficult to establish empirically. That is not to say that different

[1]The paper is closely related to Piketty (1995) where changes in the perceived connection between effort and income change the equilibrium level of effort and thus change actual income. The difference is that in Piketty (1995) these changed beliefs affect actual income only by first affecting the potential productivity of individuals.

[2]This too is closely related to Piketty's (1995) argument that it is difficult to learn the connection between effort and pay.

individuals do not contribute differently to firm output or that people do not have estimates of what these differences are. It is just that these estimates are based on pieces of information that are hard to communicate between people so that it is difficult to obtain "objective" estimates of individual productivity.

The second reason to consider models where the perceived connection between pay and productivity differs from the actual one is that people do not appear to agree on this connection. A particularly telling set of examples of this disagreement can be found in Hochschild (1981), who reports on twenty eight in-depth interviews about the distribution of resources. Summarizing she says (p. 140), "The poor ... often argue that if productivity were truly rewarded, this would create *more* equal incomes (italics in original)." And goes on to say (p. 141) "On the other hand, wealthy respondents ... often argue that if productivity were truly rewarded, this would create *less* equal incomes." [3] This sort of disagreement is consistent with the earlier claim that people find it difficult to communicate effectively their information about the way the economy works.

Because firms perceive the productivity of different people to be different, they will pay different people different amounts. Indeed, in the case of a frictionless competitive labor market, the wage of an individual is equal to that individual's marginal product. Opinions about whether people are or are not paid their marginal product do not alter this so that the have no effect on the actual distribution of income.

In the model I present, there are two departures from the assumptions of the standard competitive labor market which ensure that these opinions matter. First, I suppose that a worker's departure from a firm depends on his outside opportunities and that the firm has imperfect information about these opportunities. The result is that employers have some monopsony power and that wages depend both on the firm's perception of their employees' marginal products, and on the firm's assessment of the likelihood that these employees will

[3] These differences are consistent with those found in larger surveys. McClosky and Zaller (1984) report that 84% of respondents whose family income is above $35000 view the capitalist system as fair and efficient while only 51% of respondents with lower income do so. This type of statement may simply reflect different tastes for redistribution but it also could reflect differences in opinions about the determinants of individual incomes.

leave. Second, there are search costs so that workers have to form an opinion about what they will be paid on the outside before they leave a firm.

This combination of assumptions yields a simple mechanism by which opinions affect the distribution of income. If workers regard the evaluations of their employers as inaccurate, they are tempted to quit if they receive a low evaluation while they are inclined to stay if they receive a favorable one. The reason is that, if capitalists are inaccurate, it is good to work for a firm that has an inaccurately high estimate of one's own ability. Employees who receive high ratings from their current employers conclude that they have found a good match while those who receive a low rating expect alternate employers to pay higher wages.

Thus, the perception that capitalists have become more accurate makes workers that have received a low rating from their firm more likely to stay (because they now expect a low rating from other firms as well). A monopsonistic firm then tends to cut their wages. Similarly, this change in perception leads workers that have gotten a high rating to be more willing because they are no longer afraid of getting low wages from alternate employers. This, in turn, tends to raise their wages. Since the wages of highly rated employees rise while those of employees with low ratings fall, inequality increases.

The paper proceeds as follows. Section 1 presents the model and studies how the wage distribution changes in response to changed perceptions by workers of the accuracy of firm's evaluations. Section 2 considers social efficiency, and, in particular, studies whether overall output in the economy rises or falls when workers deem firms to be more accurate. The measure of overall output I consider is, in some sense, a welfare measure because it includes the nonpecuniary utility that people get in their jobs. One difficulty with evaluating aggregate outcomes is that one needs an "objective" measure of firms' ability to evaluate workers and, for the reasons I mentioned it is hard to know how one would go about obtaining one. For the purpose at hand, I pretend that firms have an accurate assessment of their own ability to judge individuals' productivity. This leads me to use the firms' estimates of their accuracy in computing the model's predictions regarding aggregate output.

The main conclusion of this section is a negative one. For reasons I explain, there appears to be no robust connection between overall output and the beliefs of workers. Even when workers' beliefs become more accurate in the sense of getting closer to the capitalists beliefs that are used to evaluate social output, overall output can fall. Thus, my model is one where there not be a trade-off between equity and efficiency: for certain parameter values it is possible to increase both at the same time. I also show that simpler measures of output like the amount of output that is produced to be sold in the market (which, like GNP, is readily observable) do not give an accurate picture of whether overall output (which includes the enjoyment people derive from their jobs) increases. Thus it is possible that an increase in workers' perception of distributional equity leads to an increase in inequality together with a fall in overall output, even though a measure of market output such as GNP rises.

Sections 1 and 2 focus on inequality across individuals whose publicly observable characteristics are identical. The individuals differ only in the evaluation made of them by their employers and this information is held privately by individuals and their employers. In Section 3, I show that one can reinterpret the model so that it applies also to the inequality of wages between people who differ in observable characteristics like education. Section 4 discusses the empirical applicability of the model. In particular, it discusses measurements of people's opinions about income inequality and studies the connection of these opinions with the actual distribution of income. It shows that, across a sample of seven countries, actual inequality is negatively correlated with the perception that income inequality is excessive. This is in broad accordance with the predictions of my model. Section 5 concludes.

1 The Model

1.1 The Basic Structure

Employees are of two types. I suppose that those for whom a parameter T is equal to one have higher productivity, at least in certain occupations, than those for whom the parameter T is equal to zero. I will suppose that T is not observable, though certain firms do observe signals that are correlated with the true value of T. The focus of my study is the connection between employee beliefs about the accuracy of firms signals about T and the wages that firms pay in equilibrium.

There are two periods. In the first period, each employee starts out attached to a firm. This attachment implies only that the firm has observed a signal S that is correlated with T. After seeing this signal, the employee's initial firm can make him an offer. This offer specifies a wage that the employee earns in the first period assuming that he works at the original firm (which I label o) during this period. If the employee stays in firm o, his contribution to the firm's first period output depends on his type T. In particular, the value of the output he produces is given by $q + vT$ where q and v are parameters. Firm o, however, does not observe these individual outputs even after the first period is over. To rationalize this, one can imagine that the firm's actual output is contaminated by so much noise that it contains negligible information about the employee's productivity. Alternatively, one can imagine that the value of the employee's output accrues over time and that the effect on firm output of employees of different productivities is not detected until much after the employees have left the firm. The result, in any event, is that the firm's only indication of the employee's productivity at the start of the second period is the realization of S.

Before the employee either accepts or rejects the first period offer from firm o, he also receives an outside offer from an alternative firm that I will label a_1. I assume that this outside firm offers him a wage equal to $\bar{q}$. [4] By being in this alternative firm, the employee also gains nonpecuniary utility which is equivalent to an additional compensation of n. Abusing the language somewhat, I thus treat him as receiving a total compensation of $\bar{q} + n$, where I am including the nonpecuniary utility in his compensation. The value of n is drawn from a distribution whose cumulative density function is $F(n)$. For much of the analysis, I will focus on the special case where F represents the uniform distribution, with density

[4] The simplest rationalization for this is there are actually two alternative firms in which the individual has a productivity $\bar{q}$ and that the employee gets the same nonpecuniary utility in both. If each of these firms gets to make one wage offer to the employee then the logic of Bertrand competition leads them to offer him $\bar{q}$. An alternative interpretation would include in $\bar{q}$ not only the value of the employee's first period output but also the value of the information about the employee's type that the alternative firm can expect to acquire. I neglect this because it would complicate the analysis.

equal to h for n between 0 and $1/h$. [5]

If the employee works for a_1, a new firm which I label b observes a new signal R which is also correlated with the employee's T. If the employee works for firm b in the second period, his output is, once again, given by $q + vT$. Because this firm cares about the employee's type, it will make him an offer that depends on its best estimate of the probability that the employee has a T equal to one. I assume that the nonpecuniary utility that the employee derives at this new firm is equal to zero.

In the second period, the employee also gets an offer from an alternative firm that I denote by a_2. I assume that he gets this offer whether he worked for o or a_1 in the first period. Firm a_2, just like a_1 offers a wage $\bar{q}$. [6] I also use n to denote the value (in terms of wages) of the employee's nonpecuniary utility at a_2 and this too is drawn from $F(n)$ though this realization is independent of the one that gives nonpecuniary utility at a_1.

Thus, in period 1, the employee can work at either o or a_1. If he works at o, he can work at either o or a_2 in the second period. If he works at a_1, he can work at either b or a_2 in the second period. Note that I have assumed that the employee enjoys nonpecuniary utility only at those jobs (a_1 and a_2) where his productivity does not depend on his type. This, as well as well as the assumption that nonpecuniary utility is uncorrelated over time in jobs where the employee type is unimportant, simplifies the analysis considerably without, hopefully, having a material effect on the results.

My analysis of the model proceeds as follows. In the next subsection, I consider equilibrium wages in the second period at firms o and b. Then, I study expected compensation in the second period. This differs from expected wages at o and b because the individuals have the option of working in alternate firms where they also get nonpecuniary compensation. Finally, the fourth subsection describes the equilibrium wages at firm o in the first period.

[5] Using zero as the lower bound for the realizations of n does not reduce the generality of the analysis since one can think of the lowest value for n as being incorporated in $\bar{q}$.

[6] Again this can be rationalized by the assumption that there are two such firms and that their circumstances are identical.

1.2 Second Period Wages

Employers o and b know that employees will leave and go to a_2 if they are offered a wage w that is lower than the employee's realization of $\bar{q} + n$. On the other hand, the expected benefit from keeping an employee to which the firm pays w is $(q + vET - w)$, where ET represents the firm's expectation of the employee's T. Thus, these employers set their wage offers w to maximize

$$F(w - \bar{q})[q + vET - w] \tag{1}$$

where E takes expectations conditional on the information available to the firm.

Assuming an interior solution, the optimal w, w^*, satisfies the first order condition

$$f(w^* - \bar{q})[q + vET - w^*] - F(w^* - \bar{q}) = 0 \tag{2}$$

The second order condition requires that

$$f'[q + vET - w^*] - 2f < 0 \tag{3}$$

Differentiating (2) with respect to ET, $\bar{q}$ and w^*, one obtains

$$dw^* = \frac{f}{2f - f'[q + vET - w^*]} vdET + \frac{f - f'[q + vET - w^*]}{2f - f'[q + vET - w^*]} d\bar{q} \tag{4}$$

The second order condition (3) implies that the denominator in these expressions is positive so that an increase in the employee's expected productivity T raises the wage offer. To ensure that increases in the employee's outside wage, $\bar{q}$, also the firm's wage offer, it is necessary and sufficient that

$$f'[q + vET - w^*] - f < 0 \tag{5}$$

which is consistent with (3) but imposes a more stringent requirement on f'.

As can be seen from (4), there is a special case where the wage offer is linear in expected productivity. This is the case where F is uniform so that f' is zero and f is a constant equal to h (at least in the range where F is between 0 and 1). In this case, (2) reduces to

$$w^* = \frac{q + \bar{q} + vET}{2} \tag{6}$$

Because f' is zero and both (3) and (5) are trivially satisfied, (6) implies that the wage is increasing in both ET and $\bar{q}$. For future reference it is useful to compute the resulting profits using (1). These are

$$h\Big(\frac{q-\bar{q}+vET}{2}\Big)^2 \tag{7}$$

I assume throughout the analysis that the worker's productivity is no smaller in firm o (or b) than in firm a_2 so that q is no smaller than $\bar{q}$. This implies that, regardless of the distribution F, the optimal wage always exceeds $\bar{q}$ and that F is positive so that there is a strictly positive probability that the employee stays at firm o. It then follows that the equilibrium wage is always below the employee's expected productivity $E(q+vT)$. Note also that the solution is interior as long as the highest possible n, which I label $\bar{n}$ and which satisfies $F(\bar{n})=1$, also satisfies

$$f(\bar{q}+\bar{n})[q-\bar{q}-\bar{n}+vET]<1 \tag{8}$$

If condition (8) is satisfied, individuals leave firm o with positive probability. If it is violated, the optimal wage is $\bar{q}+\bar{n}$ and the employee stays with probability one.

To compute the optimal wage, one must know the value of ET. I now compute these values concentrating first on the value for firm o and then on firm b. Firm o computes this expectation knowing only the value of S. I suppose that, just like T, the signal S can take on two values, zero and one. I let the unconditional probability that S is equal to one be equal to the unconditional probability that T is equal to one and denote both these probabilities by ϕ. Thus ϕ denotes both the fraction of high productivity individuals and the fraction of individuals who get a high S.

Since I wish the signal S to be informative about T, I assume that the probability that S is equal to one conditional on T being equal to one, which I denote by σ, exceeds ϕ. It is important to note that this is a subjective probability concerning the accuracy of a signal. Thus there is little to prevent different agents from using different values of σ in their calculations. In particular, I denote by σ the subjective conditional probability of S held by firms and by σ_w the subjective conditional probability held by workers. To simplify the analysis, I also use σ, the conditional probability believed by firms, to evaluate overall output in this economy. Thus, in effect, I am pretending that σ also represents the true frequency with which a T equal to one is associated with an S equal to one. However, if one is willing not to compute average values of output one would not need to decide whether any given value of σ is true and one could still learn how changes in people's subjective probabilities affect equilibrium wages.

Letting P denote probabilities, the expected value of T conditional on observing that S is equal to one is

$$P(T=1|S=1) = \frac{P(S=1|T=1)P(T=1)}{P(S=1)} = \sigma \tag{9}$$

while the expected value of T conditional on observing that S is equal to zero is

$$P(T=1|S=0) = \frac{P(S=0|T=1)P(T=1)}{P(S=0)} = \frac{(1-\sigma)\phi}{1-\phi} \tag{10}$$

Given that σ exceeds ϕ, this latter probability is smaller and, indeed, it equals zero if the signal S is perfectly discriminating so that σ is equal to one. Even when it is not, the fact that σ exceeds ϕ together with (4) implies that the equilibrium wage paid by o to employees with an S equal to one, w_1^o, exceeds the wage paid employees with an S equal to zero, w_0^o.

I now consider the wage offered by firm b in the second period. I assume that R, the signal observed by b, has characteristics very similar to those of S. In particular, the unconditional probability that R is equal to one is ϕ. Also, firms believe that the probability that R is one conditional on T being equal to one is σ while workers believe this conditional probability to be σ_w. I also assume that everyone believes that the realization of R is independent of the realization of S once one conditions on the true value of T. In other words

$$P(R=1|S=j,T=k) = P(R=1|T=k) \qquad \text{where } k,j=0,1$$

Before computing b's expected value of T for individual workers, I study the probabilities that, for a given worker, a signal S equal to one (or zero) will be followed by a signal R equal to one (or zero). These conditional probabilities are important in determining whether individual workers quit their firm o jobs in the first period. They are also important ingredients in calculating the expected value of T by a b firm that has observed a particular realization of R.

The probability that a signal S equal to one will be followed by an R equal to one is

$$\begin{aligned} P(R=1|S=1) &= P(R=1,T=1|S=1)+P(R=1,T=0|S=1) \\ &= P(R=1|S=1,T=1)P(T=1|S=1)+P(R=1|S=1,T=0)P(T=0|S=1) \\ &= \sigma^2+\frac{(1-\sigma)^2\phi}{1-\phi}=\frac{\sigma^2+\phi-2\sigma\phi}{1-\phi} \end{aligned} \tag{11}$$

while the probability that R will equal zero conditional on S being equal to one is equal to 1 minus the expression in (11). Similarly, the probability that R will equal zero given that S equals zero equals

$$\begin{aligned} P(R=0|S=0) &= P(R=0,T=1|S=0)+P(R=0,T=0|S=0) \\ &= P(R=0|S=0,T=1)P(T=1|S=0)+P(R=0|S=0,T=0)P(T=0|S=0) \\ &= \frac{\phi(1-\sigma)^2}{1-\phi}+\left[1-\frac{\phi(1-\sigma)}{1-\phi}\right]^2=1-\phi+\phi\Big(\frac{\sigma-\phi}{1-\phi}\Big)^2 \end{aligned} \tag{12}$$

while the probability that R will equal one given that S is equal to zero is equal to 1 minus the expression in (12).

Both the expression in (11) and that in (12) are increasing in σ (given that σ exceeds ϕ). The intuitive reason for this is that a higher σ implies that both S and R are more highly correlated with T. As a result, a higher σ ensures that S and R are more highly correlated with each other. Thus, if workers perceive that the value of σ is high, they expect the perceptions of future employers about their ability to be highly correlated with the perception of their current employer. By contrast, if workers perceive σ to be low, they regard this correlation as low.

I now turn to firm b's expectation of T for individuals with R equal to either zero or one. Firm b does not observe the realization of S. However, it can use the fact that the worker whose R it observes has agreed to accept a job with a_1 to make inferences about the probability that the worker has had a signal S equal to one. If, in particular, the equilibrium is one where only workers whose S is equal to zero accept jobs with alternative firms in the first period, firm b should infer that any worker whose R it observes must have had an S equal to zero. Thus, the probability z that b should attach to finding an employee whose S was equal to one depends on the type of equilibrium that prevails in period 1. I study the form of this equilibrium further below. For the moment, I compute the expected values of T under the alternate assumptions that S is known to equal 1 and that S is known to equal zero. The expected value assuming that there is a probability z that S is equal to one is then a simple weighted average of the two numbers that I compute here.

If S is known to equal one, the expected value of T of an employee whose R is equal to one is

$$\begin{aligned} P(T=1|S=1,R=1) &= \frac{P(T=1,R=1|S=1)}{P(R=1|S=1)} = \frac{P(R=1|T=1,S=1)P(T=1|S=1)}{P(R=1|S=1)} \\ &= \frac{\sigma^2(1-\phi)}{\sigma(\sigma-\phi)+\phi(1-\phi)} \end{aligned} \tag{13}$$

while the expected value of T for an employee whose R is equal to zero is

$$\begin{aligned} P(T=1|S=1,R=0) &= \frac{P(T=1,R=0|S=1)}{P(R=0|S=1)} = \frac{P(R=0|T=1,S=1)P(T=1|S=1)}{P(R=0|S=1)} \\ &= \frac{\sigma(1-\sigma)(1-\phi)}{(1-\phi)^2-(\sigma-\phi)^2} \end{aligned} \tag{14}$$

On the other hand, if S is known to equal zero, the expected value of an employee whose R is equal to zero is

$$\begin{aligned} P(T=1|S=0,R=0) &= \frac{P(R=0|T=1,S=0)P(T=1|S=0)}{P(R=0|S=0)} \\ &= \frac{(1-\sigma)^2\phi(1-\phi)}{(1-\phi)^3+\phi(\sigma-\phi)^2} \end{aligned} \tag{15}$$

Since both the conditional and the unconditional distributions of S and R are identical, the expected value of T for an employee whose R is equal to one while his S is equal to zero is the same as that of an employee whose S equals one and his R is equal to zero.

Simple calculations establish that the expression in (13) is no smaller than the expression in (14) as long as σ is no smaller than ϕ and that the inequality is strict if σ is strictly larger than ϕ. Similarly, the expression in (14) is strictly larger than the expression in (15) if σ is strictly larger than ϕ. Put simply, employees who get signals equal to one are deemed more productive than individuals who get signals equal to zero.

Now suppose that firm b attaches a probability z to the event that the individual got an S equal to one. If it sees an R equal to one, its posterior expectation of T is

$$P(T=1|R=1) = z\frac{\sigma^2(1-\phi)}{\sigma(\sigma-\phi)+\phi(1-\phi)} + (1-z)\frac{\sigma(1-\sigma)(1-\phi)}{(1-\phi)^2-(\sigma-\phi)^2} \tag{16}$$

while, if it sees a signal R equal to zero the expected T is

$$P(T=1|R=0) = z\frac{\sigma(1-\sigma)(1-\phi)}{(1-\phi)^2-(\sigma-\phi)^2} + (1-z)\frac{(1-\sigma)^2\phi(1-\phi)}{(1-\phi)^3+\phi(\sigma-\phi)^2} \tag{17}$$

The expression in (16) exceeds that in (17) so that the expected value of T is larger for employees whose R equals one. Thus, (4) together with the second order condition implies that the wage paid by b to an employee with R equal to one, $w_1^b(z)$ exceeds the wage paid to an employee whose R equals zero, $w_0^b(z)$. Moreover, both these wages are increasing in z, the probability that an employee who takes a job at a_1 gets a signal S equal to one.

1.3 Expected Second Period Compensation

Second period compensation differs from the second period wages paid by o and b because the employee has the option of working at a_2. In particular, if an employee gets a wage offer w from either of these firms which is below $\bar{q}+n$, he goes to work for a_2 in the second period. Thus, an employee who knows that his second period wage offer by either o or b is equal to w has an expected compensation (including the value of the nonpecuniary utility the employee gets on the job) of $c(w)$ where the function c is given by

$$c(w) = F(w-\bar{q})w + \int_{w-\bar{q}}^{\infty}(\bar{q}+n)dF(n) \tag{18}$$

The derivative of $c(w)$ with respect to w is $F(w-\bar{q})$ which is strictly positive given that q exceeds $\bar{q}$. Thus, $c(w)$ is strictly increasing in w. This is not surprising, an employee who gets a higher wage offer can expect his compensation to be higher because he can afford to be choosier with the outside offers he accepts. Because $F(w - \bar{q})$ is an increasing function of w, $c(w)$ is actually a convex function of w. This convexity is particularly obvious in the case where F is uniform. In this case, (18) becomes

$$c(w) = h(w - \bar{q})w + \int_{w-\bar{q}}^{1/h} (\bar{q} + n)hdn$$

Carrying out the integration, this is

$$c(w) = \bar{q} + \frac{h}{2}(w - \bar{q})^2 + \frac{1}{2h} \tag{19}$$

which is convex in w because when w is small the individual is relatively unlikely to stay at the firm so that an increase in w has a relatively small effect. The reason this function falls with h when w is equal to $\bar{q}$ is that an increase in h lowers the mean of the value of outside offers.

In the first period, an employee's expected second period compensation depends also on his expectation concerning the wage he will be offered by either o or b. To simplify this analysis, I assume that the worker observes the realization of his signal S. [7] This means that the employee knows the wage that o will offer him in the second period. [8] I label the employee's expectation of this wage $w^o(S)$ where depending on the realization of S, this equals either w^o_0 or w^o_1.

Now consider an employee who works at a_1 in the first period. I denote his expected compensation by $c^a(z, S)$. This compensation depends on the fraction z of employees in firms like a_1 whose S is equal to one because this fraction determines the wages paid by b in the second period. It also depends on the employee's own S because this determines the employee's subjective probability that his signal R will equal one. If his signal S equals one,

[7] It is important, however, that he be unable to demonstrate to firm b that his S is equal to one.

[8] Alternatively, one can assume that he has to infer his S, and thus his second period wage at o, from his first period wage. This gives the same equilibrium allocation as in the case where the employee observes S directly as long as employees with different values for S get different first period wages. What complicates the analysis of this case is that the firm can now choose to pay all employees the same first period wage so as to make it impossible for them to learn their realization of S. In the numerical simulations I have done, I have not found any parameter combinations for which the firm would profit from this concealment of information. Thus, for the parameters I have studied the equilibrium is the same whether employees observe S directly or not.

his expected compensation in the second period after working at a_1, $c^a(z,1)$ is

$$\begin{aligned} c^a(z,1) &= P_w(R=1|S=1)c(w_1^b(z)) + (1-P_w(R=1|S=1))c(w_0^b(z)) \\ &= c(w_1^b(z))\frac{\sigma_w^2+\phi-2\sigma_w\phi}{1-\phi} + c(w_0^b(z))\frac{1-2\phi-\sigma_w^2+2\sigma_w\phi}{1-\phi} \end{aligned} \tag{20}$$

where P^w denotes probabilities computed using worker's perceptions of σ.

If, instead, his S is equal to zero, his expected compensation, $c^a(z,0)$ is

$$\begin{aligned} c^a(z,0) &= (1-P_w(R=0|S=0))c(w_1^b(z)) + P_w(R=0|S=0))c(w_0^b(z)) \\ &= c(w_1^b(z))\left[\phi-\phi\left(\frac{\sigma_w-\phi}{1-\phi}\right)^2\right] + c(w_0^b(z))\left[1-\phi+\phi\left(\frac{\sigma_w-\phi}{1-\phi}\right)^2\right] \end{aligned} \tag{21}$$

Rather obviously, $c^a(z,1)$ exceeds $c^a(z,0)$. The main interest in these formulas, however, lies in the way that $c^a(z,1)$ and $c^a(z,0)$ change with σ_w. Differentiating (20) and (21), we obtain

$$\frac{dc^a(z,1)}{d\sigma_w} = 2\frac{\sigma_w-\phi}{1-\phi}[c(w_1^b(z)) - c(w_0^b(z))] > 0$$

and

$$\frac{dc^a(z,0)}{d\sigma_w} = -2\phi\frac{\sigma_w-\phi}{1-\phi}[c(w_1^b(z)) - c(w_0^b(z))] < 0$$

Thus, the belief that capitalists are more inaccurate (so that σ_w is low) leads workers who get an S equal to one to expect lower wages on the outside while it leads workers who get an S equal to zero to expect higher wages on the outside. In other words, a low value of σ_w leads workers to regard both the current realization of S and the future realization of R as likely to be mistaken. As σ_w falls, workers deem it more likely that a high realization of S will be followed by a low realization of R. Thus, an employee who gets a high S value becomes less optimistic about his compensation on the outside. Similarly, an employee with a low S becomes less pessimistic about his compensation by b.

Before concluding this section, it is worth comparing the compensation an employee can expect to get in the second period when he stays at o with what he can expect to get if he leaves and joins a_1. This difference depends on what outsiders expect S to be for employees who leave as well as on the employee's actual value of S. To start with the simplest case, suppose the employee gets an S equal to zero and outsiders expect all employees who leave to have an S equal to zero. In this case there is no difference between perception of S inside and outside firm o. I will show that, when F is uniform and the employee's perception of σ, σ_w, is the same as the firms', the employee's expected second period wage is the same at o than outside. However, the convexity of the $c(w)$ function implies that expected compensation is larger on the outside.

If F is uniform and R is equal to one, (6) and (14) imply that the wage offer by b is equal to

$$\frac{q+\bar{q}}{2} + (v/2)\frac{\sigma(1-\sigma)(1-\phi)}{(1-\phi)^2-(\sigma-\phi)^2} \tag{22}$$

On the other hand, if R is equal to zero, (6) and (15) imply that this wage offer equals

$$\frac{q+\bar{q}}{2} + (v/2)\frac{(1-\sigma)^2\phi(1-\phi)}{(1-\phi)^3+\phi(\sigma-\phi)^2} \tag{23}$$

Using (12) the expected wage offer is

$$\frac{q+\bar{q}}{2} + (v/2)\frac{(1-\sigma)\phi}{1-\phi} \tag{24}$$

which is the same as the wage offer that o will make in the second period. The reason the expected wage offers are the same is that both o and b start out with the same information. Thus, by the law of iterated expectations the expected value of b's expectation of T is the same as o's expectation of T. Since, in addition, wages are linear in firm's expectations of T, the two expected wage offers are the same. While the expected wage offers are the same, the employee gets a more variable wage offer by going to a_1 than by staying at o. Since this variability raises (19), the employee has a higher expected compensation on the outside.

An essentially identical analysis applies to the case where the employee draws an S equal to one and in which outside employers expect the employee's S to be equal to one. On the other hand, when the employee's S is equal to one and outside firms expect it to equal one only with probability z where z is less than one, the expected second period wage is lower at b than at o. Similarly, when the employee's S is equal to zero and outside firms expect it to be equal one with positive probability, the employee's expected wage is higher on the outside.

1.4 First Period Wages

I now turn to the way wages offered by o in the first period depend on the realization of S. Any given realization of S leads, as we saw earlier, to an estimate of T that is not changed from the first to the second period. Thus, if paying a wage w leads an employee to stay, and supposing the firm does not discount the future, the firm gains

$$q + vET(S) + \{F(w^o(S) - \bar{q})[q + vET(S) - w^o(S)]\} - w. \qquad (25)$$

where $w^o(S)$ represents, once again, the optimal second period wage as a function of S that I computed in Section 1. In this expression, the term in curly brackets is equal to the maximized value of (1), it represents the expected second period profits from being able to offer an employee with a particular S a wage $w^o(S)$ in the second period.

Assuming employees do not discount the future either, an employee who stays at firm o expects a present value of total compensation over two periods equal to

$$w + c(w^o(S))$$

By leaving, the employee can expect to earn a present value of total compensation equal to

$$\bar{q} + n + c^a(z, S)$$

The employee thus leaves whenever his n is above the cutoff value $n^*(z, S, w)$ which is given by

$$n^*(z, S, w) = w + c(w^o(S)) - \bar{q} - c^a(z, S). \qquad (26)$$

Thus, firm o sets first period wages in order to maximize

$$F(w + c(w^o(S)) - \bar{q} - c^a(z, S))\Big\{q + vET(S) + F(w^o(S) - \bar{q})[q + vET(S) - w^o(S)] - w\Big\} \quad (27)$$

In general, the wage that maximizes (27) depends both on z and on S so that I denote this optimal wage by $w^*(z, S)$. Assuming an interior solution, the first order condition that characterizes $w^*(z, S)$, is

$$f(n(z, S, w^*(z, S)))\Big\{q + vET(S) + F(w^o(S) - \bar{q})[q + vET(S) - w^o(S)] - w^*(z, S)\Big\} - F(n(z, S, w^*(z, S))) = 0 \qquad (28)$$

I am now in a position to describe the equilibrium for this model. Suppose we start with a candidate value of z, which I call $\tilde{z}$. Given this candidate z, one can solve (28) for $w^*(\tilde{z}, S)$ and use (26) to compute the optimal cutoffs for the two types. This allows one to compute $\hat{z}$ as

$$\hat{z}(\tilde{z}) = \frac{\phi F(n^*(\tilde{z}, 1, w^*(\tilde{z}, 1)))}{\phi F(n^*(\tilde{z}, 1, w^*(\tilde{z}, 1))) + (1 - \phi) F(n^*(\tilde{z}, 0, w^*(\tilde{z}, 0)))} \tag{29}$$

The denominator of this expression gives the fraction of all employees that leaves firm o for given $\tilde{z}$ and wages $w^*(\tilde{z}, S)$. The numerator represents the fraction of all employees who both leave firm o and have an S equal to one. The variable $\hat{z}$ therefore equals the fraction z that results from $\tilde{z}$ and the wage policy $w^*(\tilde{z}, S)$. In equilibrium, $\hat{z}(\tilde{z})$ must thus equal $\tilde{z}$ itself. Because $\tilde{z}$ and $\hat{z}$ take values between zero and one and the $\hat{z}$ function is continuous, the Brower fixed point theorem implies that there exists an equilibrium $\tilde{z}$ which equals the resulting $\hat{z}$. In general, there is no reason to suppose that there exists only one equilibrium though, generically, the equilibria will be locally distinct. [9]

To study how the equilibrium wages vary with σ_w, one thus generally has to take into account how the equilibrium z varies with σ_w as well. However, it is of interest to see how wages would vary if z were fixed both to develop intuition and because there is a special case where this applies directly. If z were fixed, one can see how wages vary simply by differentiating (28) with respect to the wage and to c^a. This gives

$$\frac{dw(z,S)}{dc^a} = \frac{f - f'\{q + vET(S) + F(w^o(S) - \bar{q})[q + vET(S) - w^o(S)] - w\}}{2f - f'\{q + vET(S) + F(w^o(S) - \bar{q})[q + vET(S) - w^o(S)] - w\}} \tag{30}$$

The second order condition implies that the denominator of this expression is positive and, as long as a condition that is analogous to (5) is satisfied, the numerator is positive as well. Then, an increase in expected outside compensation in the second period leads firm o to raise its first period wage. Thus, in this case, a reduction in σ_w, which lowers c_1^a and raises c_0^a lowers the wage associated with S equal to one, $w^*(z, 1)$, and raises the wage associated

[9]Indeed, Acemoglu and Pischke (1996) display multiplicities in a very similar model. The function $\hat{z}$ is increasing because a high value of $\tilde{z}$ leads to high second period wages, which in turn leads to a high value of $\hat{z}$. Multiplicities arise if the slope of this function is greater then one for certain values of $\hat{z}$ while it is smaller than one for others. In the numerical simulations reported below, the equilibrium turns out to be globally unique.

with S equal to zero, $w^*(z,0)$. Thus, as long as $w^*(0,1)$ starts out being higher than $w^*(0,0)$ an increase in σ_w raises the gap between the highest and lowest wages paid inside firm o.

This gap between the highest and lowest firm paid inside a firm corresponds, broadly, to the variability of wages inside firms. The analysis above also implies that an even simpler measure of inequality increases with σ_w. This simpler measure is the gap between the highest and lowest first period wage paid in the economy which, given that the economy contains only three wages also corresponds broadly to measures such as the interdecile range. In the model, the lowest wage paid is $\bar{q}$ so this measure of inequality rises whenever $w^*(z,1)$ rises.

Because the denominator of (30) is positive, it follows that the optimal wage cannot rise by as much as c^a itself. Given (26), this implies that a rise in c^a unambiguously lowers the cutoff n at which employees leave and, as a result, it raises the number of quits. Therefore, a reduction in σ_w increases the turnover of employees whose S is equal to zero while it raises the turnover of employees whose S is equal to one.

I now turn to the question of whether, as required above for inequality within firms to rise, the first period wage of employees whose S is equal to one exceeds that of employees whose S is equal to zero. A higher S implies a higher ET and this unambiguously raises the terms in curly brackets in both (27) and in (28). As in the analysis of (4), this raises the optimal wage. On the other hand, employees whose S is equal to one can expect a higher second period wage in firm o if they remain and, because their probability of getting an R equal to one is larger, they also anticipate a higher value for $c^a(z,S)$. The former raises the cutoff n in (26) while the latter lowers it. Thus, with (5) satisfied, the former tends to reduce $w^*(z,1)$ relative to $w^*(z,0)$ while the latter tends to increase it. Thus, the fact that $c^o(w^o(S))$ goes up for employees who infer that their S is equal to one introduces some ambiguity in the relationship between $w^*(z,1)$ and $w^*(z,0)$.

I now show that, in the case of F uniform, the effect of the increase in $c^o(w^o(S))$ is quantitatively less important than the other two effects so that, $w^*(z,1)$ is indeed larger than $w^*(z,0)$. In the uniform case, substituting the wage in (6) into (19) and using (7), (28) becomes

$$h\Big[q + vET + h\Big(\frac{q - \bar{q} + vET}{2}\Big)^2 - w\Big] - h\Big[w + \frac{h}{2}\Big(\frac{q - \bar{q} + vET}{2}\Big)^2 + \frac{1}{2h} - c^a(z,S)\Big] = 0$$

so that the optimal wage is

$$2w = q + vET - \frac{1}{2h} + c^a(z,S) + \frac{h}{2}\Big(\frac{q - \bar{q} + vET}{2}\Big)^2 \tag{31}$$

For a given $c^a(z,S)$, this is clearly increasing in ET so that $w^*(z,1)$ would exceed $w^*(z,0)$ if employees with different wages expected their outside compensation in the second period to be the same. In fact, employees with S equal to one expect a higher second period compensation on the outside and this too raises $w^*(z,1)$ relative to $w^*(z,0)$.

I now provide some intuition for the result that, for constant z and under the conditions I have laid out, increases in σ_w raise the relative turnover of high S individuals together with their relative wages. Figure 1 helps one interpret these results. This figure shows the labor supply curve given, implicitly, by (26). This is a labor supply curve because a higher wage leads workers to have a higher cutoff n^c so that the fraction of this type of workers staying at the firm, $F(n^c)$, rises. The firm then has a conventional monopsony problem in which there is a marginal cost of labor that lies above the labor supply curve and a marginal revenue from an additional unit of labor which, in this case, is horizontal. The equilibrium cutoff is the one that ensures that marginal cost is equal to marginal revenue and the equilibrium wage ensures that this cutoff is on the labor supply curve.

An increase in c^a for this worker (which would result if the employee had an S equal to zero and σ_w fell) would shift the labor supply curve to the left because it raises what the worker can expect to earn outside and thus it makes departures more attractive. This leads to a reduction in the cutoff, so that more workers of this type leave in equilibrium and, at least for linear labor supply, also leads to an increase in the equilibrium wage. An increase in σ_w lowers c^a for workers with S equal to zero while raising it for those with S equal to one, so that it moves the labor supply curve of employees with S equal zero to the right while it moves that of employees with S equal to one to the left. It thus increases the relative turnover and wages of high S workers.

The question is then whether taking into account the endogeneity of z changes these results. It is easy to see that, even with endogenous z, either the labor supply curve of workers with S equal to one must move to the left or that of employees with S equal to zero must move to the right. The reason is that changes in z move both curves in the same direction. For instance, increases in z move both labor supply curves to the left since they raise the value of c^a for both sets of employees.

There is actually one interesting case where the assumption that z does not change when σ_w changes is appropriate. This is the case where employees with S equal to one have such high productivity that firm o is led to pay them a wage that keeps them at firm o with probability one. In other words, in this case, $w^*(0,1)$ is not an interior solution but rather, is the wage that ensures that the cutoff (26) is given by $\bar{n}$. Thus,

$$w^*(0,1) = \bar{n} + \bar{q} + c^a(0,1) - c(w^o(1))$$

As long as the cutoff for employees with S equal to zero is lower, z is zero and remains zero for arbitrarily small changes in σ_w. The earlier analysis implies that a small reduction in σ_w raises $c^a(0,0)$ and lowers $c^a(0,1)$. The above equation now implies immediately that the wage for employees with S equal to one falls while, given the conditions discussed earlier, the wage $w^*(0,0)$ must rise so that inequality within the firm falls. While turnover for high S employees does not change in this case, those with S equal to zero leave firm o with increased probability.

2 Social Efficiency

In this section I ask whether the allocation of labor is more efficient when workers have a higher value of σ_w. For the purposes of this analysis, I ignore distributional issues altogether. In other words, I ask simply whether equilibria with different values of σ_w produce a larger amount of total social output. The natural definition of social output in the model, which I will label Ω includes both the output produced in firms for sale in the market, which I will label Q and the nonpecuniary benefits earned by workers, particularly since the two are comparable in the worker's eyes.

An increase in Ω does not, of course, guarantee that the allocation is Pareto superior and, indeed, it generally will not be since some workers will see their wages rise while others will see their wages decline. An increase in this output measure suggests that a Pareto superior allocation would be obtainable if one had access to lump sum transfers but, if these really were possible, the distributional concerns which are the heart of this paper would not be very important. Nonetheless, it is useful to know whether increasing trust in capitalism by increasing σ_w raises, in some sense, the total value of the pie that is produced. Because these are more readily observable measures, it is also of interest to study the effect of changes in σ_w on the value of market output Q and on the value of the average first period wage, which equals

$$W = \bar{q} + \phi\Big[hn^*(z,1)(w^*(z,1) - \bar{q})\Big] + (1-\phi)\Big[hn^*(z,0)(w^*(z,0) - \bar{q})\Big] \tag{32}$$

In particular, it is worth knowing whether the change in market output Q or in the average first period wage is a good indicator of the change in Ω.

When a worker with a given ET works at either o or b, market output is $q + vET$. I suppose that when he works at either a_1 or a_2, his market output is $\bar{q}$ so that his social output inclusive of nonpecuniary benefits is $\bar{q} + n$. [10] Because workers voluntarily decide whether to stay at o or b, their decisions involve cutoff values for n; workers whose n is below the cutoff stay while the others leave. From a social point of view it is also optimal to have cutoffs for n. In particular, if it is optimal for workers with given values for their signals and a given n to leave their employers, it follows that all workers with the same signals and higher values of n should leave as well. Similarly, if workers with given signals and a given n should stay, so should all workers with lower values of n.

In equilibrium, there are six cutoff values of n. First, there are two cutoffs at b in the second period associated with R being equal to either zero or one. I label these n_0^b and n_1^b respectively. Second, there are two cutoffs at o in the second period associated with S being

[10] As noted in footnote 5, $\bar{q}$ could in principle also include the rents from learning the value of the signal R. If it did include these rents, however, $\bar{q}$ would change with z so the earlier analysis would have to be modified.

equal to either zero or one. I label these n_0^o and n_1^o respectively. Finally, there are the two cutoffs at o in period one associated with S being either zero or one and these I label n_0 and n_1.

Before computing Ω and Q as a function of these six cutoffs, it is worth computing social and market output of a worker at o (or b) in the second period assuming that the worker's expected T is given by ET and that the cutoff is given by n^c. Market output is simply

$$\bar{q} + F(n^c)(q + vET - \bar{q})$$

while social output is

$$\bar{q} + F(n^c)(q + vET) - \bar{q}) + \int_{n^c} ndF(n)$$

In the uniform case this equals

$$\omega(n^c, ET) = \bar{q} + \frac{1}{2h} + hn^c(q + vET - \bar{q}) - h(n^c)^2/2 \tag{33}$$

so that the derivative of social output with respect to n^c is

$$h(q + vET - \bar{q} - n^c) \tag{34}$$

Total output thus increases with the cutoff n^c as long as output inside the firm, $q + vET$, is larger than $\bar{q} + n^c$, the output outside the firm for the cutoff value of n^c. Because social output is a concave function of n^c, the optimal cutoff is the one where these two values of output are equal.

Now consider the social output of an employee with S equal to one. If he stays in o, his output in the first period is $q + v\sigma$ while his expected output in the second period is $\omega(n_1^o, \sigma)$. If he goes to a_1, his output in the first period is $\bar{q} + n$. With probability $P(R = 1|S = 1)$ his expected output in the second period is $\omega(n_1^b, E(T|R = 1, S = 1))$ while with probability $P(R = 0|S = 1)$ it is $\omega(n_0^b, E(T|R = 0, S = 1))$. Thus, in the uniform case, overall expected output for such an employee is

$$\begin{aligned}\Omega_1 &= hn_1\Big(q + v\sigma + \omega(n_1^o, \sigma)\Big) + \int_{n_1}^{1/h} h\Big(\bar{q} + n + P(R = 1|S = 1)\omega(n_1^b, E(T|R = 1, S = 1)) \\ &\quad + P(R = 0|S = 1)\omega(n_0^b, E(T|R = 0, S = 1))\Big)dn \end{aligned} \tag{35}$$

while his market output equals

$$
\begin{aligned}
Q_1 &= 2\bar{q} + hn_1(q + v\sigma - \bar{q} + hn_1^o(q + v\sigma - \bar{q})) \\
&\quad +(1 - hn_1)[P(R = 1|S = 1)hn_1^b(q + vE(T|R = 1, S = 1) - \bar{q}) \\
&\quad +P(R = 0|S = 1)hn_0^b(q + vE(T|R = 0, S = 1) - \bar{q})]
\end{aligned}
$$

The derivative of Ω_1 with respect to n_1 is

$$
\begin{aligned}
h(q + v\sigma - \bar{q} - n_1) &+ \{h[\omega(n_1^o, \sigma) - P(R = 1|S = 1)\omega(n_1^b, E(T|R = 1, S = 1)) - \\
&\quad P(R = 0|S = 1)\omega(n_0^b, E(T|R = 0, S = 1))]\} \qquad (36)
\end{aligned}
$$

The first term in this expression is the efficiency gain that occurs in period 1 when the value of output inside the firm, $q + v\sigma$, exceeds the output outside at the cutoff value of n, $\bar{q} + n_1$. The term in curly brackets represents the benefits in terms of second period output that result from having an additional employee at o instead of at a_1. An analogous argument to the one used to derive (35) establishes that overall expected output from an employee with an S equal to zero is

$$
\begin{aligned}
\Omega_0 &= hn_0\Big(q + v\frac{(1-\sigma)\phi}{1-\phi} + \omega(n_0^o, \frac{(1-\sigma)\phi}{1-\phi})\Big) \\
&\quad + \int_{n_0}^{1/h} h\Big(\bar{q} + n + P(R = 1|S = 0)\omega(n_1^b, E(T|R = 1, S = 0)) \\
&\quad +P(R = 0|S = 0)\omega(n_0^b, E(T|R = 0, S = 0))\Big) \qquad (37)
\end{aligned}
$$

which has a derivative with respect to n_0 equal to

$$
\begin{aligned}
h(q + v\frac{(1-\sigma)\phi}{1-\phi} - \bar{q} - n_1) &+ \{h[\omega(n_1^o, \frac{(1-\sigma)\phi}{1-\phi}) - P(R = 1|S = 0)\omega(n_1^b, E(T|R = 1, S = 0)) - \\
&\quad P(R = 0|S = 0)\omega(n_0^b, E(T|R = 0, S = 0))]\} \qquad (38)
\end{aligned}
$$

Total expected output is then

$$\Omega = \phi\Omega_1 + (1-\phi)\Omega_0$$

while total market output is

$$Q = \phi Q_1 + (1-\phi)Q_0$$

where Q_0 is computed analogously to Q_1.

We saw earlier that changes in σ_w affect n_1 and n_2 and this obviously results in changes in overall output. Below, I compute these effect on output but it is important to realize that, in general, there are additional output effects from changes in σ_w. In particular, whenever changes in σ_w affect z, firms of type b change their wage offers in the second period. This, in effect, changes the cutoffs values n_1^b and n_0^b and this has additional welfare effects. Thus, my numerical analysis looks at three different measures to assess the effect of σ_w on efficiency. First, I look at $h(q + vET - \bar{q} - n_1)$ and $h(q + vET - \bar{q} - n_0)$ to see the period one output effects of changes in σ_w. Then, I look at the derivatives of the form of (36). I also compute the derivative of Ω with respect to σ_w taking into account the changes in z. [11] Using the same method, I also compute the derivative of Q with respect to σ_w.

The use of simulations also allows me to compute changes in different measures of inequality. So far, I have concentrated on the gap between the highest and lowest wage paid inside firm o and the gap between the highest and lowest first period wage paid in the economy. Here, I consider two additional measures. The first is the ratio of the average wage earned by employees whose S equals one, which I denote by $w_{S=1}$, to $w_{S=0}$, the average wage earned by employees whose S equals zero. If S were observable, this would be a natural estimate of the inequality between people whose S differs. The second measure I consider is the variance of the logarithm of first period wages. Both of these measures can be computed using variants of the method used to compute average wages in (32).

Table 1 presents the results for four combinations of parameters. All these numerical exercises are carried out assuming that q and $\bar{q}$ are equal to one, while σ is equal to .9 and h is equal to 1.5. As can be seen from the table, the first three experiments involve a ϕ equal to .7 while the last involves a smaller ϕ. The parameter v equals 1 in the first column, 2 in the second and 3 in the last two. In all four simulations, the first period wage is higher for employees with S equals to one and rises with σ_w, even when taking into account the

[11] I compute this derivative by first computing the equilibrium z and the resulting Ω when σ is equal to σ_w and then recomputing these values for a slightly smaller value of σ_w.

endogenous response of z. Thus, the ratio of the highest to lowest first period wage paid in the economy rises. In addition, increases in σ_w raise the ratio of the highest to the lowest wage paid inside firm o and also raise the average wage earned by people whose S equals one relative to the average wage earned by people whose S equals zero. The variance of log wages, on the other hand, does not always increase though it increases in three of the four combinations considered.

The main qualitative difference between these simulations is that increases in σ_w lead to different output effects for different parameter values. As can be seen from the information in the first two rows of the table, the first column is unique in representing a situation where v is sufficiently low that even high S employees leave firm o with positive probability. I report only one such simulation because all of those I have computed have qualitatively similar properties to that reported in column one. In particular, they involve declines in both overall output and in market output in response to increases in σ_w.

The reason for this turns out to be that an increase in σ_w reduces the turnover of employees with S equal to zero while raising the turnover of employees with S equal to one. Unfortunately, the turnover of these latter employees is socially more costly. This can be seen by analyzing the first period benefits from increasing the cutoff n for both sets of employees. The general formula for these first period benefits is given in (34) though it reduces to the simpler formula in the eighth and ninth rows in the case where q is equal to $\bar{q}$. The entries in these rows show that social output rises in the first period when the cutoff for either type of employee is increased, though it rises more with the cutoff for employees whose S equals one. The reason for this discrepancy is that (31) implies that, leaving everything else equal, firm o only increases the wages of employees with S equal to one by one half of their extra productivity. This means that the gap between the actual productivity of an employee and the wage he is paid (which is what determines whether he leaves) is larger for more productive employees. The result is that there is a bigger social gain when the retention rate for these employees goes up and this can be accomplished by lowering σ_w.

The second column shows a situation where the higher value of v ensures that employees with S equal to one are retained with probability one by firm o. Then, social output increases with σ_w because such increases reduce the turnover of employees with S equal to zero (while having no local effect on the turnover of employees whose S equals one). Qualitatively identical results arises when, as in the third column, v is increased further so that it equals 3.

The main purpose of the third column is to permit comparison with the last column, which keeps v equal to 3 while lowering ϕ. This changes the results dramatically because overall output now declines when σ_w goes up. This is true even though first period overall output still rises when turnover by employees with S equal zero is reduced and even though the increase in σ_w does reduce turnover. The reason for the decline in Ω is that the derivative given in (38) is negative because the social value of having employees with S equal to zero outside firm o in the second period is lower than having them inside firm o. The reason for this is that, as I discussed earlier, social output is a concave function of n^b and having the employee be outside the firm induces randomness in his cutoff depending on the realization of R.

Another interesting feature of the last column is that the decline in overall output is accompanied by an increase in market output, as well as by an increase in average first period wages. Thus these simple measures are not good indicators of overall output. In this case, the principal source of loss is that, because employees are more attached to their original employer, they are less likely to get high nonpecuniary benefits from alternate employers.

In conclusion, there are robust examples where an increase in σ_w raises inequality while, at the same time lowering overall output. Thus, the increase in inequality that results from the changes in beliefs that I am considering would be detrimental even if lump sum transfers were available. In addition, there are cases where overall output falls even though measured GNP rises so that simple measures cannot be used to gauge the welfare effects of the changes in beliefs.

3 Observable Characteristics

So far I have treated the signals S and R as observed privately by firms and their employees. Thus the increase in first period wage inequality I derived must be thought of as an increase in the inequality *within* groups with identical observable characteristics. I now show that the analysis developed above can, in principle, apply also to changes in inequality between groups that differ in observable ways. A general treatment of this issue is beyond the scope of this paper. In this section I show only that a special case exists where the analysis in the earlier sections applies unchanged so that an increase in σ_w leads to increased inequality between groups.

I suppose that there is a signal e which, again, takes either a value of zero or one. This signal, which can be thought of as an indicator of education, is correlated with T in the sense that

$$P(T = 1|e = 1) > P(T = 1|e = 0). \tag{39}$$

I also assume that the indicator e is correlated with S so that

$$P(S = 1|e = 1) > P(S = 1|e = 0) \tag{40}$$

Thus, education is both positively correlated with actual productivity and with the signal of productivity S.

I now give an example where, in spite of (39), wages depend only on S and R as before. One feature of this example is that, consistent with Figure 2,

$$P(T = 1|S = 1, e = 1) = P(T = 1|S = 1) \tag{41}$$

In Figure 2, the set of events where S is equal to one is given by the union of the sets whose areas are A, B, D and E. The set of events where e is equal to one is given by the union of the areas A, B and C, while the set of events where T is equal to one is given by the union of A, D and E. Thus, the probability that T is one, conditional on both C and e being equal to one is $A/(A + B)$ while the probability that T is one conditional on just S

being equal to one is $(A+D)/(A+B+D+E)$. Thus, (41) is satisfied as long as

$$\frac{A}{B} = \frac{D}{E}$$

Note that the Figure is consistent with (39) because the probability that T is one conditional on e being one, $A/(A+B+C)$, is larger than the unconditional probability that T is one which is given by $(A+D+E)/(A+D+E+G)$. On the other hand, the probability that T is one conditional on S being one is, given the above equation, $A/(A+B)$. This means that S is a strictly better indicator of T than is e (since the conditional probability that T is one is strictly larger when S is one than when only e is known to equal one. Because the events where e indicates that T is one are also events where S indicates the same thing, e is superfluous so that firm o need look only at S to determine an employee's productivity.

As long as firm b observes S, the wages it pays then depend only on S and R so that e is superfluous for b as well. While I have assumed that S is not directly observable by b, S becomes known in equilibrium to b as long as the parameters are such that only employees with S equal to zero ever leave firm o. Then, firm b knows that any employee it hires has an S equal to zero and can safely neglect information contained in e. Thus, the equilibria studied in earlier sections for these combinations of parameters continue to be equilibria even when the signal e is available.

An econometrician who does not observe S will tend to compute wages conditional on e. The average wage conditional on e being equal to one is

$$w_{e=1} = \frac{A+B}{A+B+C} w_{S=1} + \frac{C}{A+B+C} w_{S=0}$$

while that conditional on e being equal to zero is

$$w_{e=0} = \frac{D}{D+E+F+G} w_{S=1} + \frac{E+F+G}{D+E+F+G} w_{S=0}$$

where G represents the area of points where T, S and e equal one and where, given the way the figure is drawn (40) is satisfied because

$$\frac{A+B}{A+B+C} > \frac{D}{D+E+F+G}$$

This condition implies that the ratio $w_{e=1}/w_{e=0}$ rises with the ratio $w_{S=1}/w_{S=0}$. Since this latter ratio rises with σ_w in all the simulations reported in Table 1, the "education differential" $w_{e=1}/w_{e=0}$ rises as well. [12]

4 Empirical Relevance

One way of demonstrating the empirical validity of the model in this paper would be to show that holding constant the productivity of individuals, the distribution of income is more unequal when people feel that pay is less related to productivity. Unfortunately, neither cross-section nor time series data on peoples' opinions about the connection of pay to productivity appears to exist. What does exist is data on peoples' opinions about the desirability of income dispersion in general and about the extent to which they want redistribution to increase. In addition there is, of course, data on political outcomes in different places at different points in time.

In this section, I discuss data of this sort and relate it to the actual distribution of income. I consider both data across countries at a point in time and, to a much lesser extent data over time for the United States. Figure 3 focuses on seven countries at a point in time. It shows the relationship between the attitudes towards inequality from a 1987 international survey as well as inequality measures from around the same period. [13]

The horizontal axis of this figure gives the fraction of people who agree with the statement that "Differences in income in [respondent's country] are too large". These figures are drawn from Evans (1993) which reports the results of a 1987 survey which asked the same set of questions in the seven countries reported in the figure. [14] On the vertical axis, the figure gives

[12]Note that this analysis is agnostic as to whether education helps in raising S or T. In other words, the analysis takes the joint distribution of T, S and e as given without inquiring as to the determinants of any of these variables. One can thus imagine that the correlation between S, T and e is due in part to the productivity enhancing effect of education, as long as one accepts that after the education is completed, firms have access to a signal of productivity that is better than the individual's educational diploma. Or, one can imagine that e is typically high for able individuals simply because these enjoy education more.

[13]Measures of inequality are not available for every year although this problem may not be extremely severe because inequality changes relatively slowly.

[14]The only other country included in the survey was Hungary which I excluded because I felt that the data on actual inequality was probably not very meaningful in the during its post-communist transformation.

the Gini coefficient for household pre-tax income. These figures are drawn from Deininger and Squire (1995) who report inequality measures from a number of sources. To help make these figures comparable, I used figures from the LIS data base (Atkinson, Rainwater and Smeeding, 1995) for all the countries for which these figures were available, namely for all countries except Austria and Australia. [15] The correlation between these figures is -.72. Such a negative correlation fits with the model presented earlier if one regards the opinion that inequality is too large as closely related to the opinion that pay does not correspond to performance.

There are, of course, alternative interpretations. One possibility is that the perception of large inequities leads to policies which cause actual inequality. This too would make the widespread perception of inequity lead to a relatively even distribution of income. But, if this theory were valid, the perceptions would affect incomes only through their effect on policies.

One way these policies could be leading to the negative correlation in figure 3 is by affecting the measured income distribution itself. The measure of pre tax household income on which my Gini coefficients are based includes some government transfers and cash transfers to poor people might be larger in countries where people dislike inequality more. I have thus also looked at OECD measures of the distribution of labor earnings. The coverage of these data is not uniform for the seven countries in my sample. For Austria, the U.K. and the U.S., the figures appear to be comparable and the ratio D9/D5 is largest in the U.S., intermediate in the U.K. and smallest in Austria. The D5/D1 ratio is highest in the U.S. and lower for the U.K. and Austria, though the latter two numbers are essentially identical. These figures are broadly consistent with the Gini coefficients for pre-tax income.

Another source of potential endogeneity, is that high income people might work less in countries where people dislike inequality more because they are subject to more taxation.

[15] I picked, for each country the observation that was closest to 1987. The data for Switzerland are from 1982, those for Germany are from 1983, those for Norway, the U.K. and the U.S. are from 1986 and those for Ireland and the Netherlands are from 1987. The Australian data come from the 1989 Statistical Yearbook for Australia while those for Austria, whose data are probably the least comparable to the others, comes from 1987.

In fact, it is not clear that the countries in my sample in which people dislike income inequality have particularly progressive income taxes. In particular, the correlation of the Gini coefficients based on after tax household income (all of which come from the LIS data base) with the Evans' measure is only -.40.

When people say that they think inequality is excessive, they could have many different things in mind. For instance, they might feel that the differences in income are due to differences in social contributions but that differences in social contributions do not warrant large income disparities. In particular, they might be unhappy if differences in raw abilities (as opposed to in effort) lead to large differences in income. It is also possible that, as in my theory, they dislike income inequality because they feel that differences in income do not reflect differences in productivity. Unfortunately, I am unaware of international surveys which allow one to sort out the relative importance of these perceptions in explaining the differences in attitudes portrayed in Figure 3. In any event, and as I expound on in the conclusion, I regard my theory as only a first attempt at showing that beliefs about inequality matter and hope that related models can explain why other changes in attitudes towards inequality have similar effects.

Because it corresponds most closely to the belief that pay does not reflect performance, I focused attention on the survey which asked whether people thought income disparities were excessive. It is important to note, however, that the answer to this question is strongly positively correlated with the answer to questions that ask whether people favor additional redistribution from the rich to the poor. Evans (1993) reports a measure which gives the mean response to questions asking people whether they would like the unemployed to receive a guaranteed standard of living, whether they support more spending on benefits for the poor, whether they support a guarantee of jobs for everyone and whether they are in favor of a basic income for everyone. With a higher means score representing less enthusiasm for redistribution, the correlation between this mean score and the answer plotted in Figure 3 is -.80. Similarly, the mean score against redistribution has a correlation of .78 with the pre-tax Gini coefficient and of .52 with the after-tax coefficient.

Time series evidence on attitudes towards income inequality seems even harder to obtain. This is particularly surprising given that many political analysts have noted a "shift to the right" in U.S., British and Canadian electoral outcomes. In all three countries, the late 1970's and early 1980's saw the election of leaders like Ronald Reagan, Margaret Thatcher and Brian Mulroney whose rhetoric was substantially more "pro-market" than that of their predecessors. The only U.S. data that confirms that this was accompanied by a change in people's attitude towards redistribution are reported in Kluegel and Smith (1986). They conducted a national survey in 1980 and found that only 18% of their respondents disagreed with the statement that the U.S. was spending too much on welfare. As they say, this is substantially smaller than the 39% of the people who agreed, in a 1969 survey conducted by Feagin (1975), that the U.S. was spending too little on welfare. This comparison suffers from several difficulties not the least of which is that the public perception of government policies might have been different in the two instances.

In spite of the paucity of attitudinal data, it may be reasonable to suppose that the change in electoral outcomes in these three countries is suggestive of an increased confidence in the free market system. If so, it is tempting to think about whether the increased income inequality in these three countries over the last 20 years couple is partially due to this change in attitudes. What makes this explanation attractive is that, as Moss (1995) shows, inequality has worsened disproportionately in these three countries (together with Australia). By contrast countries in continental Europe such as France, Italy and Germany have seen their earnings inequality fall. And, interestingly, leaders in these countries appear not to have shifted nearly as much towards a free market rhetoric. [16]

Thus the increase in United States income inequality surveyed in Levy and Murnane (1992), which involves increased dispersions in wages both within groups of similar educational attainment as well as across groups that differ in their education, could be due to a

[16]This is no place for a thorough survey of the extensive literature on the forces that have recently affected the income distribution in a variety of countries. It is worth mentioning one complementary view here; namely that free-market oriented leaders have worsened the income distribution in their countries by reducing the influence of unions. For a discussion of this possibility see Freeman and Katz (1994).

change in beliefs about the fairness of the free enterprise system. This explanation is consistent with a remarkable change in workers' attitudes towards their employers. As Uchitelle (1994) reports citing a study by Richard Freeman and Joel Rogers, workers no longer view big business as an adversary. They do not blame their employers for the anxiety that they feel about their income and job security. Instead, workers regard their employers as victims of global competition which forces employers to cut costs and lay off workers.

Thus, interestingly, international competition may have increased income inequality via two mechanisms. The first and more traditional one, is that this competition may have reduced the demand for low skilled workers. The second is that the existence of this international competition may have convinced workers that their employers are not being unfair when they cut wages. This, in turn, may have exacerbated inequality through a mechanism akin to the one presented in this paper.

5 Conclusions

This paper presented a first attempt at understanding how changes in workers' conception of the capitalist system affect the distribution of income. The main contribution of the paper is that it shows that changes in beliefs can have a direct effect on the distribution of income even if they do not affect the inherent productivity or the effort of any worker. [17].

This result is based on considering a particular change in beliefs, namely a change in the extent to which people believe employers are accurate in their assessment of individual productivity. This is related to peoples' general attitude towards income inequality because a belief that employers are inaccurate and whimsical in deciding who should get a high wage is presumably associated with a dislike for the income inequality that results from these decisions. The main issue that is left unadressed is whether the results in this paper extend to changes in attitudes towards income inequality that stem from other sources. It would be interesting to know whether a related model would predict that the income distribution

[17]By contrast, most of the modern theoretical literature on income distribution makes the distribution of income vary only as a result of changes in actual the inherent productivity of different workers. For a recent example and some references, see Acemoglu (1995)

depends also on the extent to which people believe that people whose productivity differs ought to earn different levels of income. One possible model along these lines might be constructed by extending Akerlof's (1982) "gift exchange" model. In that model, worker effort depends on the relation between the wage paid and the "reference wage"and firms maximize profits by taking into account this relationship. As social notions of fairness change, the relevant reference wage may change as well. In particular, an increased belief that more productive people should be paid more might lower the reference wage for low skilled workers while raising that for skilled workers. And, in turn, this might widen income disparities.

Another issue that deserves further attention is the way workers modify their beliefs about the appropriateness of the income distribution. In part, this would require modeling the learning of workers along lines similar to Piketty (1995). In addition, it is worth thinking about changes in society that might have relatively rapid effects on these beliefs. Suppose, for instance, that a new technology is introduced and that it is apparent to workers and firms that only some people can use this technology effectively. Or, along the same lines, global competition makes raise the relative value of the output of certain workers. If workers come to feel that employers are able to determine the identity of those workers whose productivity is enhanced by either technical progress or international competition, they are in effect increasing their trust in the connection between pay and productivity.

The other element that has been left out of the model is the effect of ideological changes on the institutions that are determined in the political arena. If workers perceive capitalism as fair, they are more likely to vote for representatives that uphold capitalist institutions, engage in less redistribution and, quite possibly, try to reduce the influence of unions. These institutional changes, particularly the reduction in union power, would presumably affect the income distribution as well. It then becomes important to disentangle the direct effects of changes in ideology from those brought about indirectly through changes in political institutions.

6 References

Acemoglu, Daron,"Matching, Heterogeneity and the Evolution of the Income Distribution," mimeo ,1995.

—— and Jörn-Steffen Pischke, "Why do Firms Train? Theory and Evidence," MIT Department of Economics Working paper 96-7, February 1996.

Akerlof, George A., "Labor Contracts as Partial Gift Exchange," *Quarterly Journal of Economics*, 97, November 1982, 54-69.

Atkinson, Anthony B., Lee Rainwater, and Timothy Smeeding, Income Distribution in European Countries, Working Paper Series, No 121, Luxembourg Income Study, 1995

Deininger, Klaus and L. Squire, "Measuring Income Inequality: A new Data-Base," World Bank mimeo, 1995.

Evans, Geoffrey, "Class conflict and Inequality," in Jowell, Robert, Lindsay Brooks and Lizanne Dowds, coeditors, *International Social Attitudes, the 10th BSA Report*, Brookfield VT: Dartmouth Publishing Co., 1993.

Feagin, J.R. *Subordinating the Poor*, Englewood cliffs, N.J.: Prentice-Hall, 1975.

Freeman, Richard and Lawrence Katz, "Rising Wage Inequality: The United States vs. Other Advanced Countries," in Freedman, R. , ed. *Working under Different Rules*, New York: Russell Sage Foundation, 1994, 29-62.

Hochschild, Jennifer, *What's Fair? American Beliefs about Distributive Justice*, Cambridge, MA: Harvard University Press, 1981.

Kluegel, James R. and Eliot R. Smith, *Beliefs about Inequality*, New York: De Gruyter, 1986.

Levy, Frank and Richard J. Murnane, "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature*, 30, September 1992, 1333-81.

McClosky, Herbert and John Zaller, *The American Ethos*, Cambridge MA: Harvard University Press, 1984.

Moss, David A., *Unemployment in France: "Priority Number One"*, Case 9-795-064, Boston, MA: Harvard Business School Publishing, 1995.

Piketty, Thomas, "Social Mobility and Redistributive Politics," *Quarterly Journal of Economics*, 110 (3), August 1995, pp 551-84.

Uchitelle, Louis, "Insecurity forever," *New York Times*, 11/20/94.

Table 1
Numerical Results

Parameters:	ϕ=.7 v=1	ϕ=.7 v=2	ϕ=.7 v=3	ϕ=.52 v=3
$F(n^*(z,0))$	.182	.378	.633	.240
$F(n^*(z,1))$	.982	1.00	1.00	1.00
$[w^*(z,1)]/[w^*(z,0)]$	1.368	1.169	1.103	1.334
$d(w^*(z,1)/d\sigma_w$	.036	.283	.430	.520
$d(w_{S=1}/w_{S=0})/d\sigma_w$	.003	.225	.301	.466
$d(w^*(z,1)/w^*(z,0))/d\sigma_w$	.131	.514	.671	.742
$d$var.log.wage$/d\sigma_w$	.001	.003	-.001	.011
$h(vE(T\|S=0)-n^*(z,0))$	.168	.322	.417	.248
$h(vE(T\|S=1)-n^*(z,1))$	.368			
$d\Omega_0/dn^*(z,0)$	.118	.120	.073	-.006
$d\Omega_1/dn^*(z,1)$	.799			
$d\Omega/d\sigma_w$	-.018	.012	.012	-.0005
$dQ/d\sigma_w$	-.043	.042	.097	.022
$dW/d\sigma_w$	.004	.205	.311	.280

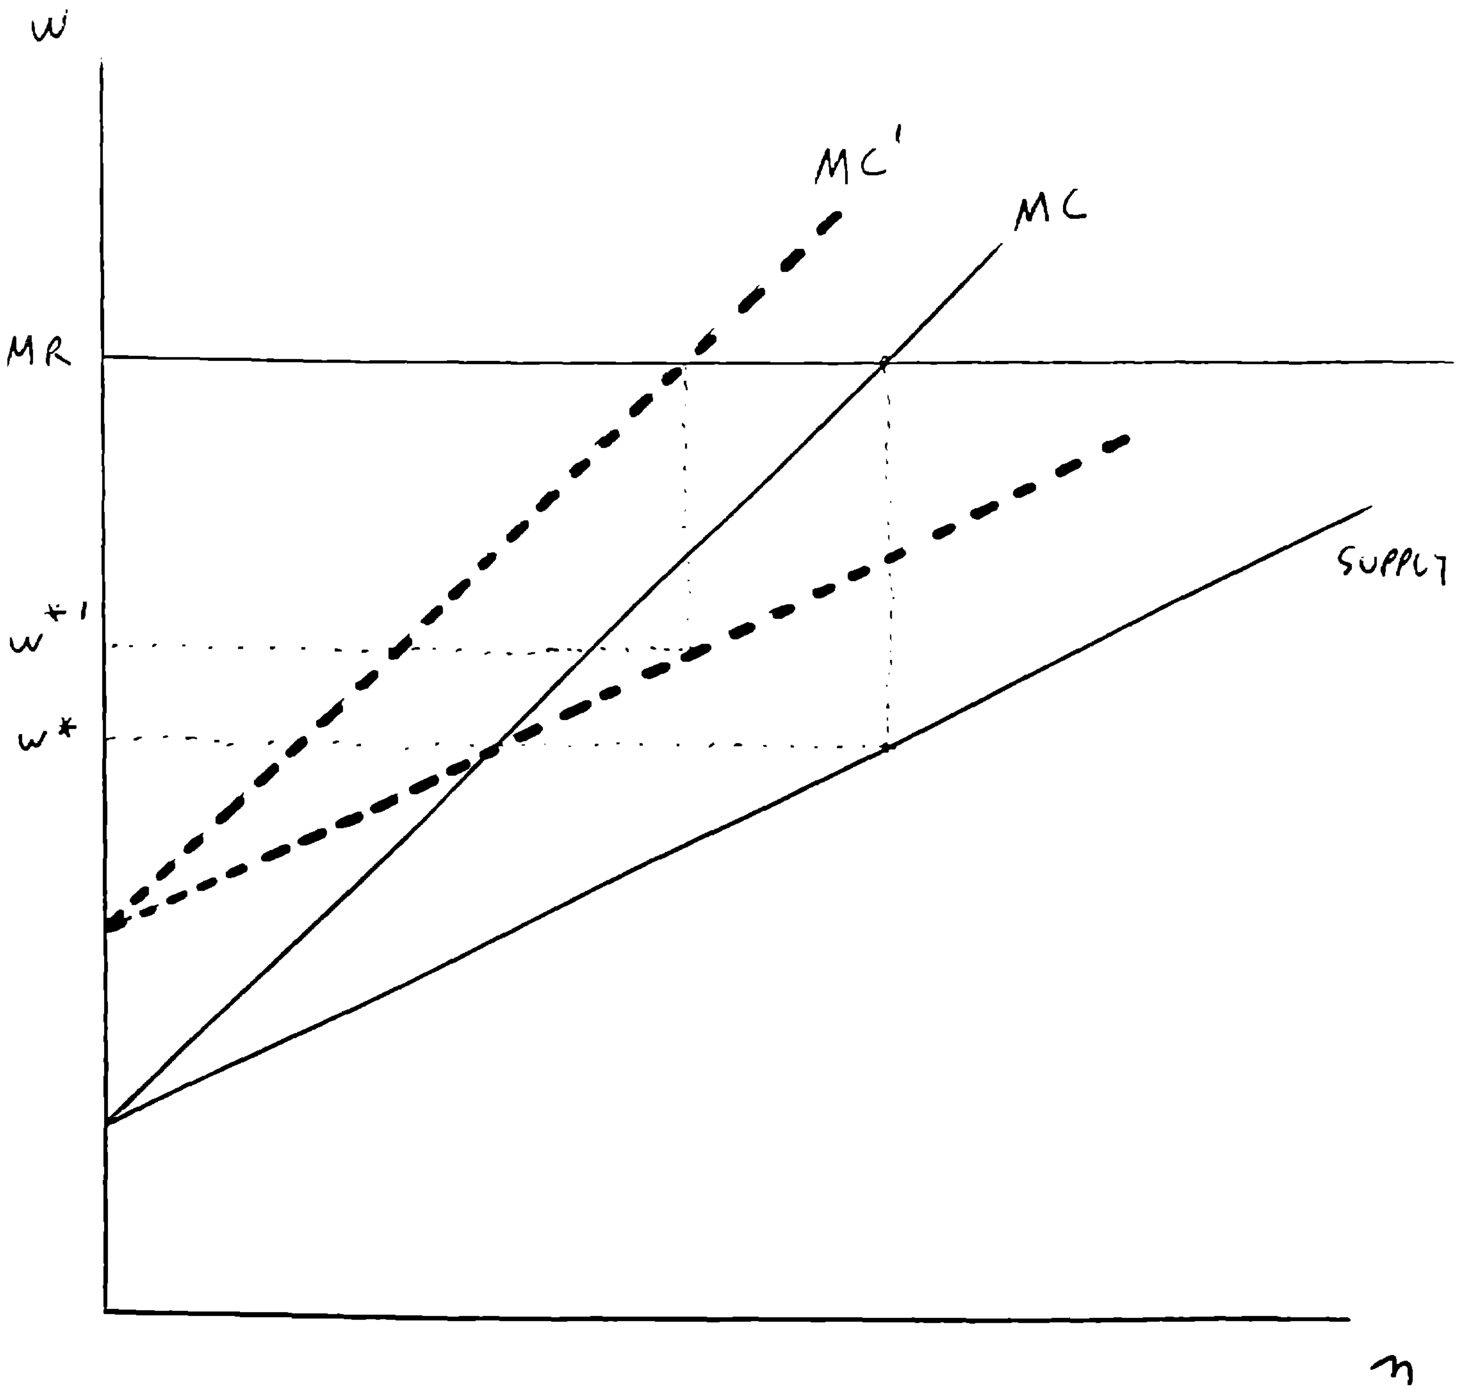

FIGURE 1
WAGE DETERMINATION

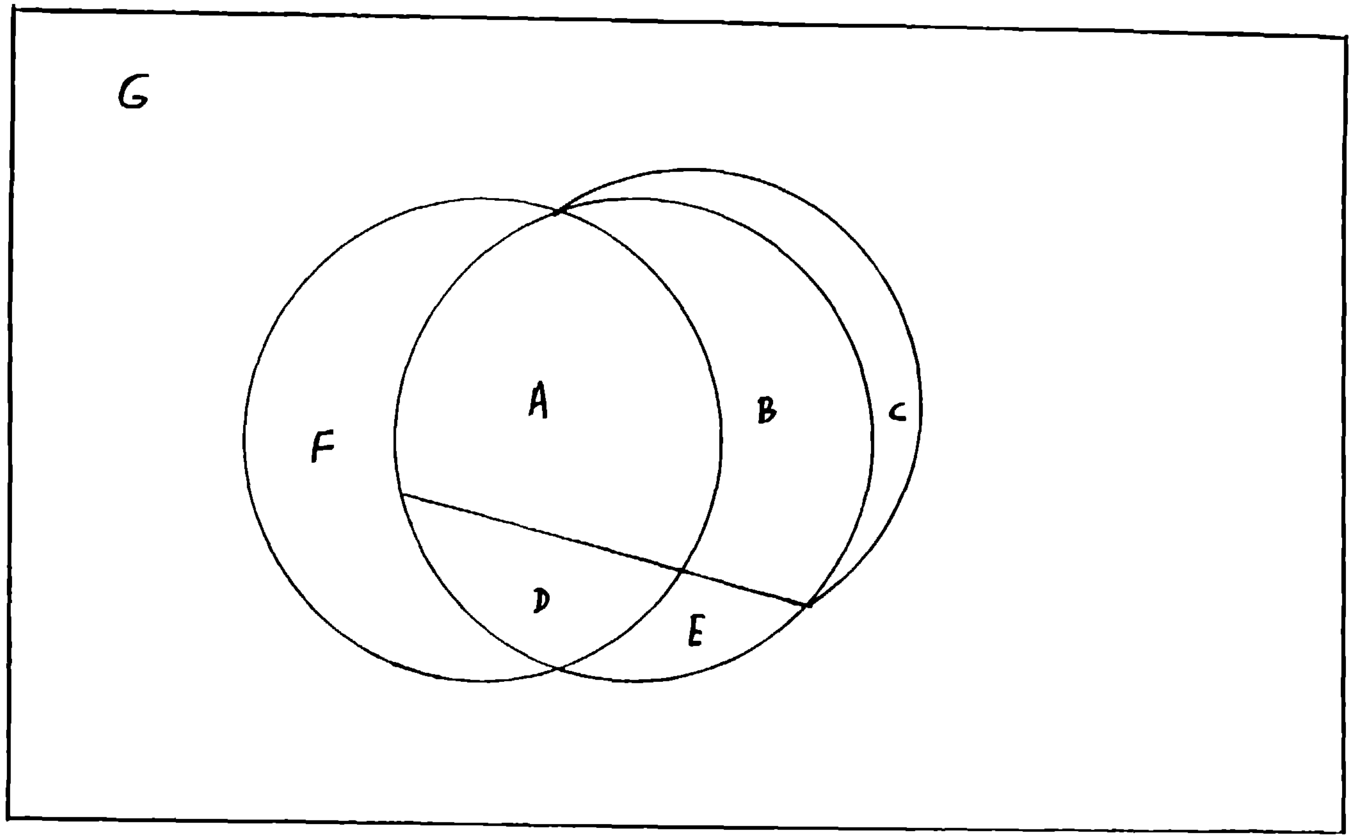

FIGURE 2

THE SETS WHERE $T=1$, $S=1$ AND $R=1$

$T=1$: $A \cup D \cup F$

$S=1$: $A \cup B \cup D \cup E$

$R=1$: $A \cup B \cup C$

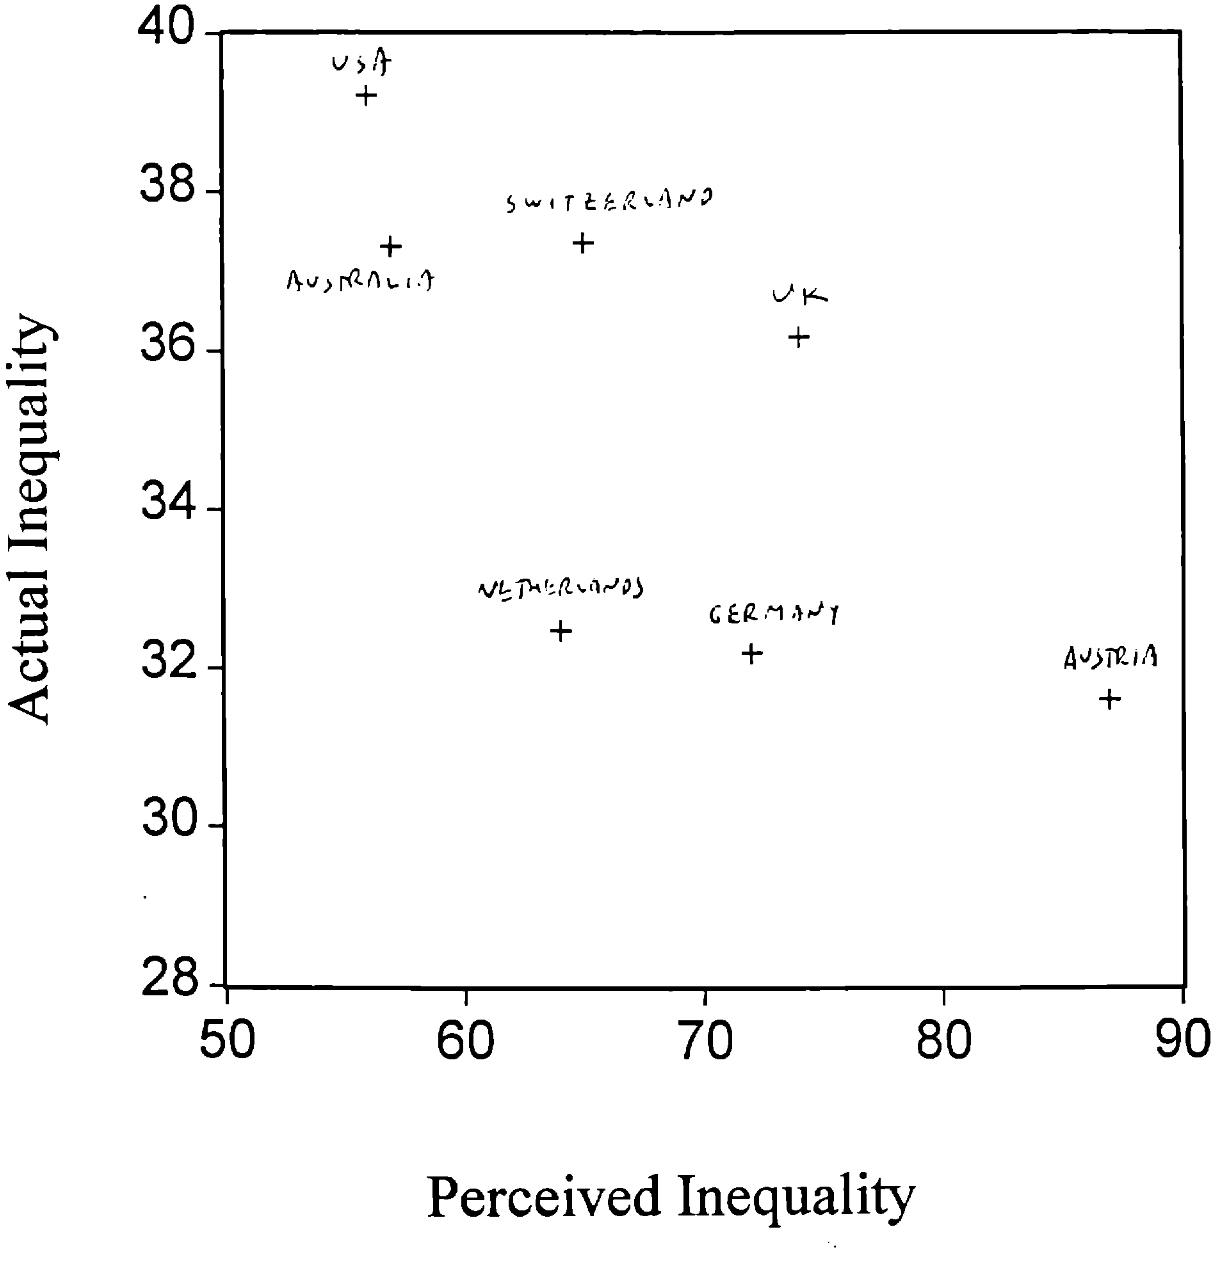
40
38
36
34
32
30
28
50
60
70
80
90
Actual Inequality
Perceived Inequality
USA
SWITZERLAND
AUSTRALIA
UK
NETHERLANDS
GERMANY
AUSTRIA

Figure 3