Federal Reserve Bank
of Minneapolis

Spring 2002

Quarterly Review

Inequality and Fairness

Christopher Phelan*
Senior Economist
Research Department
Federal Reserve Bank of Minneapolis

While conceding that normative issues like *fairness* or *justice* may be appropriate considerations when discussing social policies informally, economists have a strong tradition of shying away from such considerations when they evaluate policies formally. Instead, economists traditionally see their expertise as in objectively evaluating whether a social policy is *efficient*. Perhaps the best illustration of this tendency to avoid normative issues is the familiar Edgeworth box diagram. In an endowment economy with two goods and two agents, efficiency implies that each agent's marginal rate of substitution between the two goods must equal the other agent's. If both agents like both goods, then this equality can occur when the allocation of the goods is quite unequal, such as when either agent has almost all of the aggregate endowment of both goods. Economists tend to be willing to criticize allocations in which marginal rates of substitution are not equated (since then both agents can be made better off by an allocation in which these rates are equated). But economists tend to be unwilling, at least formally, to choose among the many efficient allocations.[1]

Nevertheless, starting with the work of Edward Green (1987), economists have developed a large literature on dynamic contracting theory which has implications for inequality of consumption across households and, I argue, on whether such inequality is fair.[2] In that literature, society faces a trade-off between insurance and incentives, and inequality results as the product of efficient mechanisms intended to induce good behavior. I argue here that

the *efficiency* notion of these studies corresponds closely (but not perfectly) to the *fairness* notion of political philosopher John Rawls (1971). If my argument is accepted, then dynamic contracting theory has something to say about fairness as well as efficiency. It says that fairness and inequality are not incompatible: a fair allocation can imply inequality. But more surprisingly, the theory says that a fair allocation can imply extreme inequality of both opportunity and result.

To make this argument, I present a model which allows the consideration of equality of opportunity and result. The model economy has a fixed amount of land that is a necessary input to household production; that is, a household with little or no land has little or no opportunity to produce.[3] Thus, in the model, one household having

---

[1] There are, of course, major exceptions to this general tendency. Work by Nobel laureates James Mirrlees and Amartya Sen and many other mainstream economists focuses on normative, distributional issues. Further, many economists will say inequality is "getting worse" rather than simply "increasing." Nevertheless, most economists agree that such language, while not necessarily incorrect, goes against a tradition in the profession of not making value judgments.

[2] Examples of this literature include the work of Stephen Spear and Sanjay Srivastava (1987), Jonathan Thomas and Tim Worrall (1990), and Robert Townsend and me (1991). The model in this study borrows essential characteristics from the work of Andrew Atkeson and Robert Lucas (1992) and Aubhik Khan and B. Ravikumar (2001).

[3] In Khan and Ravikumar 2001, capital plays a role similar to that which land plays here. The difference is that here land is in fixed supply.

more land than another implies inequality of opportunity while one household having more consumption than another implies inequality of result. This fixed amount of land is also useful for deriving my most extreme conclusion: in the fair allocation, incentive issues cause the land allocation of almost all households to converge to zero, implying eventual infinite inequality of both opportunity and result.

Rawls' Fairness vs. Economists' Efficiency

To get to that extreme conclusion, I must establish a connection between Rawls' notion of fairness and economists' notion of efficiency in models of dynamic contracting. Rawls' notion of fairness generally corresponds to economists' notion of ex ante efficiency.

To see this, consider a world in which, by pure luck, some are born with valuable skills (the *skilled*) and some, with less valuable skills (the *unskilled*). Again, in evaluating alternative social policies, economists tend to avoid considering what is the best transfer from the skilled to the unskilled because that involves making a value judgment: best from whose perspective? Instead, economists tend to evaluate social policies in terms of ex post efficiency. Economists are willing to call social policy A "inferior to" social policy B if both the skilled and the unskilled weakly prefer social policy B (and at least one group strictly prefers policy B). But economists are traditionally silent on the ranking of these policies if one group prefers policy A and the other, policy B (as is true if policy A is a transfer from the skilled to the unskilled and policy B, a zero transfer).

Rawls (1971), however, is quite willing to consider the best transfer from the skilled to the unskilled. His method is to step back behind "a veil of ignorance" to a world in which everyone knows the two groups will exist, but no one knows which group anyone will be in. By definition, many things may be uncertain from behind the veil, but everyone is in the same informational position. In such a world, there is no question of "best from whose perspective?" There is only one perspective—that of individuals who see all possible outcomes as possibly happening to them. This common perspective, then, gives a way to rank social policies that economists are traditionally loathe to rank. Social policy A is better than social policy B if and only if it is better from the behind-the-veil perspective, and the best social policy is the one that maximizes welfare from that perspective.

Still, in some types of economic problems, notably, those involving insurance, Rawls' notion of fairness corresponds to the economist's notion of ex ante efficiency. Economists have little reticence about arguing that when households pool risk regarding home fires and as a result transfers go from households which do not have fires to households which do, this is efficient risk-sharing. In reality, people actually do make insurance contracts before they know which group they are going to be in. Formally, Rawls' criterion amounts to simply pretending that the skilled vs. unskilled scenario (in which everyone knows which group they are in) is the same as the fire vs. no fire scenario (in which no one knows, at the point of contracting, which group they are in) and choosing the transfer from the skilled to the unskilled to correspond with the efficient insurance arrangement that people would have chosen if they had had an opportunity to choose.

I depart here from Rawls' notion of fairness regarding the preferences of those behind the veil. I assume that preferences regarding risk and time are the same for those in the world and those behind the veil and take forms commonly assumed by economists. Rawls argues that individuals behind the veil will be more risk averse than individuals observed in the world and will not discount the future. (In fact, Rawls argues that behind the veil, individuals will be infinitely risk averse, or care only about maximizing welfare in their own worst-case scenario.) This is how Rawls (1971, pp. 152ff) derives his familiar *maximin* criterion: policy A is preferred to policy B from behind the veil if and only if the welfare of the worst outcome for any individual under policy A exceeds the welfare of the worst outcome under policy B. Here, I show that if more reasonable levels of risk aversion and some discounting are allowed, then what is chosen from behind the veil can change drastically.

A Static Model

I start by considering a simple contracting model, an economy with only one time period. Even in this simple model, we shall see that fairness implies some inequality.

This economy has a continuum of identical households and a single divisible unit of land. The single consumption good c is produced by households using land not allocated to private use. Specifically, if z is the household's land allocation and $\theta \in \{0, \bar{\theta}\}$ is the fraction of the household's land allocated to production of the consumption good, then the household's production is

$$(1) \quad y = z\theta\mu$$

where $\mu \in \{0,1\}$ is an independent random variable (across

households) and $0 < \bar{\theta} < 1$. Let $0 < \pi_L < 1$ denote the probability of realization $\mu = 0$ and $\pi_H = 1 - \pi_L$. In words, equation (1) says that a household's output y equals $z\bar{\theta}$ if fraction $\bar{\theta}$ of the household's land is allocated to production (as opposed to none of it being so allocated) and the random variable μ is drawn to be one instead of zero. Otherwise (because either $\theta = 0$ or $\mu = 0$), a household's output equals zero. Let $p_L(\theta) = 1 - (\pi_H \theta / \bar{\theta})$ and $p_H(\theta) = \pi_H \theta / \bar{\theta}$ denote the endogenous probabilities of low and high outputs.

In this model, households care about their consumption and the amount of land allocated to their private use. A household's preference-ordering over certain pairs of consumption and land allocated to consumption $(c, z(1-\theta))$ is determined by the Cobb-Douglas specification $c^\alpha [z(1-\theta)]^{1-\alpha}$, where $0 < \alpha < 1$. As I argued above, to be most consistent with Rawls' concept of fairness, preferences over risk should be defined from two perspectives: in the world and from behind the veil. This would allow households to be more risk averse regarding allocations from behind the veil as opposed to in the world. However, defining preferences from both perspectives is technically difficult and conceptually problematic. So here I focus solely on the situation in which preferences toward risk from behind the veil correspond to those in the world. In the world, a household's preference-ordering over lotteries is determined by the expected value of the constant relative risk aversion specification

$$(2) \qquad U(c, z(1-\theta)) = (c^\alpha [z(1-\theta)]^{1-\alpha})^{1-\sigma} / (1-\sigma)$$

where U is the household's utility and $0 < \sigma < 1$.[4]

By Rawls' criterion, for an allocation to be fair, it must be *symmetric*—each household must be treated identically from the behind-the-veil perspective. If the allocation is deterministic (involves no lotteries), then symmetry implies that the land allocation and division, z and θ, are common across households. This implies that an allocation is simply a vector (z, θ, c_L, c_H) denoting each household's land allocation, land division, and consumption given a low (zero) and high $(z\bar{\theta})$ output realization. Since consumption occurs after production, a household's consumption can depend on its observable production in a deterministic symmetric allocation. However, a symmetric allocation could, in principle, be random. If a lottery is involved, some households could receive large land allocations while others receive small ones. As long as each household faces the same odds, though, such a random allocation can be considered

symmetric. For now, I restrict consideration to deterministic allocations. Later, I will show that random allocations will never be chosen from behind the veil.

Suppose θ and μ are publicly observable. Then an ex ante efficient allocation $(z^*, \theta^*, c_L^*, c_H^*)$ solves

$$(3) \qquad \max_{(z, \theta, c_L, c_H)} p_L(\theta) U(c_L, z(1-\theta)) + p_H(\theta) U(c_H, z(1-\theta))$$

subject to the resource constraints

$$(4) \qquad z \leq 1$$

and

$$(5) \qquad p_L(\theta) c_L + p_H(\theta) c_H - p_H(\theta) z\theta \leq 0.$$

In words, an ex ante efficient allocation maximizes ex ante utility subject to resource constraints on land (4) and the consumption good (5).

It is nearly immediate that a solution to this problem implies that $z = 1$, $\theta = \bar{\theta}$, and $c_L = c_H = \pi_H \bar{\theta}$, or that consumption is a constant over all realizations of μ and equal to expected output.[5] Having the consumption of risk-averse households depend on the realization of their output (which depends only on their realization of μ) does not serve the purpose of maximizing households' utility. It only makes households less happy. Thus, a fair allocation in this economy has not only (by assumption) *equality of opportunity*—every household has the same income-producing possibilities as every other since all have the same land allocation and the same land division—but also *equality of result*—income is redistributed so that all households have the same consumption.

What if a household's θ and μ are not publicly observable, but instead are observed by only that household? In particular, suppose that a household can surreptitiously divert land intended for production to private use.[6] Now θ can no longer be perfectly inferred from observed output. If a household has positive output, it must have chosen θ

---

[4]Requiring that $\sigma < 1$ puts a limit on the risk aversion of households. If $\sigma \geq 1$ (where, as usual, $\sigma = 1$ implies that $U(c, z(1-\theta)) = \log(c^\alpha [z(1-\theta)]^{1-\alpha})$, then exposition becomes more difficult, but my main results still hold.

[5]If $\theta = 0$, then output equals zero and, thus, $c_L = c_H = 0$. This gives a utility of zero (or $-\infty$ if $\sigma \geq 1$). But $c_L = c_H = \pi_H \bar{\theta}$ gives positive utility (or finite utility if $\sigma \geq 1$). To see that $c_L = c_H$, suppose that (c_L, c_H) satisfies the resource constraint, but $c_L \neq c_H$. Setting $c_L = c_H = \pi_H \bar{\theta}$ also satisfies the resource constraint with equality and increases the objective function value since the objective function is concave.

[6]For simplicity, assume that this ability to divert is one-way. A household can divert land intended for production to its private use, but cannot divert land intended for its private use to production.

$= \bar{\theta}$, but if a household has zero output, this does not imply that $\theta = 0$ since zero output can also occur due to $\mu = 0$. Thus, for an allocation to be feasible, when it specifies a land division θ, actually choosing θ must be individually rational.

Specifically, an allocation (z,θ,c_L,c_H) is incentive-compatible if and only if

$$(6) \quad p_L(\theta)U(c_L,z(1-\theta)) + p_H(\theta)U(c_H,z(1-\theta)) \geq U(c_L,z).$$

The left side of this incentive constraint is the expected utility of setting θ to the value specified by the allocation, and the right side is the certain utility of setting $\theta = 0$, which gives the household the low output with certainty. (If the allocation specifies $\theta = 0$, then the left and right sides of (6) are automatically equal.)

Imposing the incentive constraint implies that $c_H > c_L$. Thus, there is inequality of result.[7] If an allocation specifies that $\theta = \bar{\theta}$ and $c_L = c_H > 0$, then the allocation is not incentive-compatible since the household is better off setting $\theta = 0$ and receiving the added utility associated with using all of its land for personal use. (That is, $U(c,z) > U(c,z(1-\bar{\theta}))$ if $c > 0$.) Thus, equality of result is feasible only if $\theta = 0$. Then, since production equals zero, $c_L = c_H = 0$. But this allocation gives a utility of zero. Society can always do better by setting $\theta = \bar{\theta}$ and setting c_H sufficiently greater than c_L to satisfy the incentive constraint. Since when $z = 1$ two constraints on allocations (expressions (5) and (6)) remain and, with $\theta = \bar{\theta}$, two choice variables (c_L and c_H) remain, maximizing expected utility essentially amounts to requiring the resource and incentive constraints to hold with equality. (Making c_H/c_L greater than necessary to provide incentives only decreases expected utility.) Thus, in this model, inequality of result makes production possible.

But what of equality of opportunity? So far this has been simply assumed by setting z and θ common for all households. To prove that equality of opportunity is in fact optimal, I examine the dual of the primal programming problem defined by maximizing utility subject to resource and incentive constraints. (For an elaboration of this method, see the 1992 work of Andrew Atkeson and Robert Lucas and my 1994 work.)

Suppose that an allocation (z,θ,c_L,c_H) maximizes (3) subject to the resource constraints (4) and (5) and the incentive constraint (6). Further, suppose that the allocation delivers a maximized utility w^* and a shadow price of

land q, where q is the relative value of the Lagrange multipliers on constraints (4) and (5). This allocation has a resource cost (in terms of the consumption good) of $pz + p_L(\theta)c_L + p_H(\theta)c_H - p_H(\theta)z\theta$. Standard arguments imply that the utility-maximizing allocation (z,θ,c_L,c_H) must also minimize the resource cost of providing w^* subject to the incentive constraint (6). Otherwise, the surplus resources could be redistributed to households to get higher levels of utility.

This cost-minimization problem can be stated for arbitrary w and q as

$$(7) \quad V(w|q) = \min_{(z,\theta,c_L,c_H)} qz + p_L(\theta)c_L + p_H(\theta)c_H - p_H(\theta)z\theta$$

subject to the incentive constraint (6) and a condition that the allocation actually deliver w utils, or

$$(8) \quad w = p_L(\theta)U(c_L,z(1-\theta)) + p_H(\theta)U(c_H,z(1-\theta)).$$

In proving Lemma 1 (below), I show that the functional form assumption that $U(c,z(1-\theta)) = [1/(1-\sigma)] \times (c^\alpha[z(1-\theta)]^{1-\alpha})^{1-\sigma}$ implies that $V(w|q) = V(1|q)w^{1/(1-\sigma)}$. Since $V(w|q)$ is a convex function of w (since $1/(1-\sigma) > 1$), equality of opportunity is implied. That is, while a symmetric lottery could conceivably be set up in which some households are given an allocation which delivers expected utility w_1 with probability ρ and others, expected utility w_2 with probability $1 - \rho$, the convexity of $V(w|q)$ implies that this setup requires more resources than does giving both groups an expected utility of $\rho w_1 + (1-\rho)w_2$. Such a lottery is, thus, inefficient.

The proof of Lemma 1 also shows that this problem exhibits a *scaling* property that will prove useful for consideration of dynamic versions of this model. That is, $V(w|q) = V(1|q)w^{1/(1-\sigma)}$ precisely because the best way to treat a group owed w utils is proportional to the best way to treat a group owed 1 util, scaling z, c_L, and c_H up and down proportionally to satisfy the promise-keeping constraint.[8]

---

[7] A qualification: If σ is above a critical value (greater than one), then adequate incentives to allow $\theta = \bar{\theta}$ cannot be provided without setting $c_L = 0$, which ensures that expected utility equals minus infinity. Then the economy has many optimal allocations, all of which imply this worst-case expected utility.

[8] If $\sigma > 1$, then utility is negative and $V(w|q) = V(-1|q)(-w)^{1/1-\sigma}$, which is also convex.

LEMMA 1. *Scaling*

The cost-minimization function $V(w|q)$ takes the form $V(w|q) = V(1|q)w^{1/(1-\sigma)}$. Further, if $(z(w),\theta(w),c_L(w),c_H(w))$ denotes the cost-minimizing allocation as a function of w, then $z(w) = z(1)w^{1/(1-\sigma)}$, $\theta(w) = \theta(1)$, $c_L(w) = c_L(1)w^{1/(1-\sigma)}$, and $c_H(w) = c_H(1)w^{1/(1-\sigma)}$.

Proof. Consider $w > 0$ and the programming problem defined by minimizing (7) subject to (6) and (8). Without loss, we can choose $\hat{z} \equiv z/w^{1/(1-\sigma)}$, $\hat{c}_L \equiv c_L/w^{1/(1-\sigma)}$, and $\hat{c}_H \equiv c_H/w^{1/(1-\sigma)}$ instead of choosing (z,c_L,c_H) directly. (The actual (z,c_L,c_H) values are simply implied by $(\hat{z},\hat{c}_L,\hat{c}_H)$.) Imposing the definitions of $(\hat{z},\hat{c}_L,\hat{c}_H)$ into the programming problem and simplifying delivers

$$(9) \qquad V(w|q) = w^{1/(1-\sigma)}\min_{(\theta,\hat{z},\hat{c}_L,\hat{c}_H)} q\hat{z}$$
$$+ p_L(\theta)\hat{c}_L + p_H(\theta)\hat{c}_H - p_H(\theta)\hat{z}\theta$$

where the minimization is subject to

$$(10) \qquad 1 = p_L(\theta)[(\hat{c}_L^{\alpha}[\hat{z}(1-\theta)]^{1-\alpha})^{1-\sigma}/(1-\sigma)]$$
$$+ p_H(\theta)[(\hat{c}_H^{\alpha}[\hat{z}(1-\theta)]^{1-\alpha})^{1-\sigma}/(1-\sigma)]$$

and

$$(11) \qquad p_L(\theta)[(\hat{c}_L^{\alpha}[\hat{z}(1-\theta)]^{1-\alpha})^{1-\sigma}/(1-\sigma)]$$
$$+ p_H(\theta)[(\hat{c}_H^{\alpha}[\hat{z}(1-\theta)]^{1-\alpha})^{1-\sigma}/(1-\sigma)]$$
$$\geq (\hat{c}_L^{\alpha}\hat{z}^{1-\alpha})^{1-\sigma}/(1-\sigma).$$

This minimization problem is stated independently of w and is, in fact, the minimization problem associated with $w = 1$.   Q.E.D.

To recap, when incentives must be provided in this static model, equality of opportunity is attained, but equality of result is abandoned. Through no fault of their own, the households with low output realizations have lower consumption levels than those with high output realizations. Since the allocation is incentive-compatible, all households indeed do the right thing and set $\theta = \bar{\theta}$. The low output households simply have the bad luck of realizing $\mu = 0$ instead of $\mu = 1$. Inequality of result in this model occurs not to punish the dishonest—when the allocation is implemented, no household is dishonest—but instead to ensure honesty. Why is this fair? It is what all households would have agreed to if they could have chosen in advance.

A Dynamic Model With Two Periods . . .

Now I transform the static model into a model in which time passes. In this dynamic economy, the efficient (fair) outcome involves inequality of both opportunity and result.

Consider a two-period version of the economy just analyzed, with periods $t = \{0,1\}$. Per period preferences are identical in this economy, and households place weight 1 on period 0 utility and weight $\beta < 1$ on period 1 utility. The consumption good is assumed to be nonstorable.

If a household's θ_t and μ_t are again observable, then in this dynamic model, society can do no better than simply to repeat the equal land, equal consumption solution to the static problem. But if a household's θ_t and μ_t are observed by only that household, things become more complicated. In particular, the definition of a *symmetric allocation* is more complicated. For instance, land allocation is no longer trivial since a household's allocation in period 1 can depend on its realized output in period 0. A symmetric allocation in the two-period economy is a vector $(z_0,\theta_0, c_L,c_H,z_L,z_H,\theta_L,\theta_H,c_{LL},c_{LH},c_{HL},c_{HH})$. Here, c_L and c_H represent the period 0 consumption allocation for each output realization, (z_L,z_H) and (θ_L,θ_H) represent the period 1 land allocation and land division for each period 0 output realization, and $(c_{LL},c_{LH},c_{HL},c_{HH})$ represents the period 1 consumption allocation for each history of output realizations. For such an allocation to be incentive-compatible, choosing θ_L must be individually rational (in period 1) for households with low first-period realizations, or

$$(12) \qquad p_L(\theta_L)U(c_{LL},z_L(1-\theta_L)) + p_H(\theta_L)U(c_{LH},z_L(1-\theta_L))$$
$$\geq U(c_{LL},z_L).$$

Similarly, choosing θ_H must be individually rational (in period 1) for households with high first-period realizations, or

$$(13) \qquad p_L(\theta_H)U(c_{HL},z_H(1-\theta_H)) + p_H(\theta_H)U(c_{HH},z_H(1-\theta_H))$$
$$\geq U(c_{HL},z_H).$$

Finally, choosing θ_0 must be individually rational (in period 0) for all households, or

$$(14) \qquad \sum_{i \in \{L,H\}} p_i(\theta_0) \times$$
$$\left[U(c_i,z_0(1-\theta_0)) + \beta\sum_{j \in \{L,H\}} p_j(\theta_i)U(c_{ij},z_i(1-\theta_i))\right]$$
$$\geq U(c_L,z_0) + \beta\sum_{j \in \{L,H\}} p_j(\theta_L)U(c_{Lj},z_L(1-\theta_L)).$$

For such an allocation to be feasible, resource constraints on land and consumption must hold in each period, or

(15) $\quad z_0 \leq 1$

(16) $\quad p_L(\theta_0)z_L + p_H(\theta_0)z_H \leq 1$

(17) $\quad p_L(\theta_0)c_L + p_H(\theta_0)c_H - p_H(\theta_0)z_0\theta_0 \leq 0$

(18) $\quad p_L(\theta_0)[p_L(\theta_L)c_{LL} + p_H(\theta_L)c_{LH}]$
$\qquad\qquad + p_H(\theta_0)[p_L(\theta_H)c_{HL} + p_H(\theta_H)c_{HH}]$
$\qquad\qquad - p_L(\theta_0)p_H(\theta_L)z_L\theta_L$
$\qquad\qquad - p_H(\theta_0)p_H(\theta_H)z_H\theta_H \leq 0.$

An allocation $(z_0^*, \theta_0^*, c_L^*, c_H^*, z_L^*, z_H^*, \theta_L^*, \theta_H^*, c_{LL}^*, c_{LH}^*, c_{HL}^*, c_{HH}^*)$ is ex ante efficient (or *fair* by Rawls' criterion) if it solves

(19) $\quad \max_{(z_0, \theta_0, c_i, z_i, \theta_i, c_{ij})} \sum_{i \in \{L,H\}} p_i(\theta_0) \times$
$\qquad \big[U(c_i, z_0(1-\theta_0)) + \beta \sum_{j \in \{L,H\}} p_j(\theta_i) U(c_{ij}, z_i(1-\theta_i)) \big]$

subject to conditions (12) through (18).

What can be said about a solution to this problem? Thanks to the assumed functional form, quite a lot. We know that a solution involves *memory*. Specifically, the land a household is allocated in period 1 (which determines the household's opportunity to produce in period 1) depends on the household's output in period 0. We also know that the solution again *scales*. In the solution, all households receive the same efficient land division θ, and each household's consumption in period 1 for each period 1 output is proportional to the household's land in period 1, or $c_{HH}/c_{LH} = c_{HL}/c_{LL} = z_H/z_L$. Households that have a good outcome in the first period are simply given proportionally more of everything in the second period, including the ability to produce. In the second period, then, the economy has equality of neither opportunity nor result.

Explaining this outcome again requires examining the problem of minimizing the resource cost of providing w lifetime utils, subject to incentive constraints. In the two-period model, the relevant shadow prices are those of land in each period (denoted by q_0 and q_1) in terms of period t consumption and that of period 1 consumption in terms of period 0 consumption (denoted by δ). The cost-minimization problem is, then,

(20) $\quad \min_{(z_0, \theta_0, c_i, z_i, \theta_i, c_{ij})} q_0 z_0 + \sum_{i \in \{L,H\}} p_i(\theta_0) \times$
$\qquad \big[c_i - z_0\theta_0\mu_i + \delta(q_1 z_i + \sum_{j \in \{L,H\}} p_j(\theta_i)(c_{ij} - z_i\theta_i\mu_j)) \big]$

subject to the incentive constraints (12), (13), and (14) but not the resource constraints on land or consumption.

The usefulness of the dual problem (cost minimization) as opposed to the primal problem (utility maximization) stems from the fact that the continuation of a cost-minimizing allocation is itself a cost-minimizing allocation. That is, consider a household which realizes outcome $i \in \{L,H\}$ in the first period. This household's second-period allocation is the one-period allocation $\{z_i, \theta_i, c_{iL}, c_{iH}\}$. Let w_i denote the one-period expected utility associated with $\{z_i, \theta_i, c_{iL}, c_{iH}\}$. This allocation can be shown to solve the static cost-minimization problem outlined above.[9] Thus, the result that $c_{HH}/c_{LH} = c_{HL}/c_{LL} = z_H/z_L$ has already been proven in Lemma 1.

Note that at this point only scaling has been shown, not memory. From what has been shown so far, it is still possible that $c_{HH}/c_{LH} = c_{HL}/c_{LL} = z_H/z_L = 1$, or that a household's first-period outcome does not affect its second-period allocation. To see that $z_H/z_L > 1$ (or households with good outcomes in the first period get better second-period allocations), we must examine the two-period problem imposing the result that $\{z_i, \theta_i, c_{iL}, c_{iH}\}$ solves the one-period problem. That is, for each $i \in \{L,H\}$, instead of directly choosing $\{z_i, \theta_i, c_{iL}, c_{iH}\}$, choose w_i with the understanding that w_i implies $\{z_i, \theta_i, c_{iL}, c_{iH}\}$ through the static minimization problem. Then the cost-minimization problem is

(21) $\quad \min_{(z_0, \theta_0, c_i, w_i)} q_0 z_0$
$\qquad + \sum_{i \in \{L,H\}} p_i(\theta_0)[c_i - z_0\theta_0\mu_i + \delta V(1 | q_1) w_i^{1/(1-\sigma)}]$

subject to the first-period incentive constraint

(22) $\quad p_L(\theta_0)\big[U(c_L, z_0(1-\theta_0)) + \beta w_L' \big]$
$\qquad + p_H(\theta_0)\big[U(c_H, z_0(1-\theta_0)) + \beta w_H' \big]$
$\qquad\qquad\qquad\qquad \geq U(c_L, z_0) + \beta w_L'$

and the first-period promise-keeping condition

[9] This is established by arguments from Spear and Srivastava (1987) and Townsend and me (1991), which rely on the work of Dilip Abreu, David Pearce, and Ennio Stacchetti (1986, 1990).

$$(23) \cdot \quad w = p_L(\theta_0)\Big[U(c_L, z_0(1-\theta_0)) + \beta w_L'\Big]$$
$$+ p_H(\theta_0)\Big[U(c_H, z_0(1-\theta_0)) + \beta w_H'\Big].$$

The proof of Lemma 2 (below) shows that $w_H > w_L$, or that the efficient allocation for this economy has memory.

LEMMA 2. *Memory*
Second-period utility following a high first-period output realization is greater than second-period utility following a low first-period output realization: $w_H > w_L$.

Proof. Suppose that $(z_0^*, \theta_0^*, c_i^*, w_i^*)$ solves the cost-minimization problem (21) with $w_H = w_L = w^*$. Since otherwise first-period production equals zero, $\theta_0^* = \bar{\theta}$. Next, consider the following perturbation: $w_H = w^* + x/\pi_H$, $w_L = w^* - x/\pi_L$, where x is arbitrary. With this setup, a choice of $x = 0$ should be optimal. Note that as x varies, the effect on ex ante utility is zero; thus, the promise-keeping constraint is not affected, and neither is the left side of the incentive constraint. But the right side of the incentive constraint is decreased as x increases; as x increases, the incentive constraint is loosened. Since without the incentive constraint a solution involves $c_L = c_H$, the incentive constraint is binding, so loosening it has positive value. The derivative of the objective function with respect to x equals

$$(24) \quad \delta V(1|q_1)[(w^* + x/\pi_H)^{\sigma/(1-\sigma)} - (w^* - x/\pi_L)^{\sigma/(1-\sigma)}]$$

which at $x = 0$ equals zero. Thus, $x = 0$ is not optimal. A similar argument rules out $w_H < w_L$.   Q.E.D.

To recap, in the two-period model, incentive issues ensure that both equality of result and equality of opportunity (in the second period) are abandoned. In this dynamic model, as in the static model, in the first period households with low output realizations—again, through no fault of their own—have lower consumption levels than those with high output realizations. Further, these low output realizations follow them into the second period. The households with low output realizations receive smaller land allocations and, for each second-period output realization, proportionally less of the consumption good than households with high output realizations. Further still, this dependence of the second-period allocation on first-period outcomes is not necessary to ensure that $\theta = \bar{\theta}$. From behind the veil, the households could have chosen to retain equality of opportunity in the second period and still have $\theta = \bar{\theta}$ simply by repeating the solution to the static model. They choose not to do this simply because they can do better ex

ante by introducing memory. Again, this is fair because it is what all households would have agreed to in advance.

. . . And With Infinite Periods
Now I extend the time line of the dynamic two-period model to infinity: periods $t = \{0,1,...,\infty\}$. In this infinite-period economy, per period preferences are identical, and households place weight β^t on period t utility. We shall see that the theoretical results from the last model hold here as well, but that the extension of time adds extreme implications on inequality as the economy ages. In particular, as $t \to \infty$, inequality goes to infinity as almost all land and consumption go to a vanishingly small fraction of the population. Thus, in the ex ante efficient allocation, the economy not only has inequality of opportunity and result; eventually, it has as much of both kinds of inequality as is possible.

This result comes from the combination of an infinite number of periods and three factors of the two-period model:

- Land is in fixed supply.

- Allocations scale. (Lemma 1 holds.)

- Households with high output realizations receive better future allocations than do households with low output realizations. (Lemma 2 holds.)

As argued above, scaling occurs regardless of the shadow prices of land and consumption in each period. These shadow prices affect how to treat a household owed 1 util, but don't affect the fact that the best way to treat all other households is proportional to this. Suppose the shadow prices were such that the cost-minimizing way to treat a household owed 1 util has the household's expected land allocation higher or lower in period $t + 1$ than in period t. Then this situation would hold for all households in period t. If all households have their land allocation increasing or decreasing in expectation, then the aggregate land allocation is increasing or decreasing with certainty. But this is impossible since land is in fixed supply. Thus, shadow prices must be such that every household's expected land allocation in period $t + 1$ equals its actual land allocation in period t. That is, land allocation is a *martingale,* a stochastic process with the property that its expectation in period $t + 1$ equals its value in period t. Since a household's conditional consumption in each period is proportional to its land allocation, the implications of land being a martingale hold for consumption as well.

To deliver the result that, eventually, almost all households get arbitrarily little land (and consumption), with almost all land (and consumption) going to a vanishingly small fraction of the population, we need more than scaling and land in fixed supply. Land tomorrow must also be used as an incentive device today. Note that for the infinite-period economy, except for the distribution of promised utilities, every period is exactly the same as every other period; no period is closer to the end of time. But scaling implies that the distribution of utilities is essentially irrelevant. (For a detailed justification, see my 1994 work.) Thus, the shadow price of land never changes, and neither does the price of future consumption. So the best way to treat a household owed 1 util never changes, and Lemma 2 ($w_H > w_L$) continues to hold. Suppose a household owed 1 util gets z units of land in period t and gets z_L or z_H in period $t + 1$ if it has a low or high realization in period t. We know that $\pi_L z_L + \pi_H z_H = z$ and that z_L and z_H are the same for all t. Further, scaling implies that a household which has $2z$ units of land in period t gets $2z_H$ or $2z_L$ in period $t + 1$.

How do these facts imply eventual extreme inequality? Suppose that $\pi_L = \pi_H = 1/2$, $z_L = (1/2)z$, and $z_H = (3/2)z$, or that every day a household either loses half its land holdings or increases them by 50 percent. (The numbers do not matter. The following reasoning holds for all (π_L, π_H) adding to 1 and $z_H > z_L$ with expected value z.) If a household starting with z_0 has a high realization and then a low realization (or a low and then a high), its land holdings after two periods are $(3/4)z_0$. If high then low or low then high occurs again, the household's land holdings after four periods are $(9/16)z_0$. (Every two periods, the household loses a quarter of its land.) Over time, if the household keeps receiving exactly half high and half low realizations, then its land holdings converge to zero. But note that the law of large numbers implies that as $t \rightarrow \infty$, almost all households have arbitrarily close to half their realizations high and half low. Thus, almost all households have their land holdings converge to zero. Eventually, expected utility averaged across households is arbitrarily low.

Why is such an allocation agreed to from behind the veil? The key to this is that the standard from behind the veil is expected *discounted* utility ($\beta < 1$). Before the world begins, how it will eventually look is of little importance to those who exponentially discount the future. Nevertheless, in each period, making the future depend on the current realization always loosens the incentive constraint, and thus, in each period, adding at least a little more inequality makes sense.

Reinterpreting Time

From the perspective of the beginning of time, in the efficient allocation of the infinite-period model, all households receive the same allocation. This is, in fact, implicit in the definition of a *fair allocation*. Thus, from this perspective, the model has equality of opportunity. But from the perspective of any later period, the model no longer has that equality. A household's past stays with it forever and directly affects its ability to produce through its land allocation.

This distinction becomes more important depending on how time is interpreted in the model. If a time period is interpreted as one period in the long life of a single individual, then this model could perhaps be considered as having equality of opportunity, accompanied by ever-increasing inequality of result. However, a time period could also be interpreted as an entire lifetime of a member of a dynastic household that cares about its descendants (although discounting by β). The allocation in the infinite-period model is exactly the allocation the first generation would choose when setting up society. That society cannot be considered as having equality of opportunity for any but the first generation.

But is this allocation, under Rawls' criterion, truly fair? Specifically, from behind the veil of ignorance, the allocation in the infinite-period model is the one agreed to only if the first generation is the only one whose preferences matter. From the perspective of someone born into a later generation, inequality of opportunity (dependence of a later generation's allocation on its ancestors' outcomes) is simply a welfare-reducing lottery.

To handle this case, suppose the veil also hides which generation an agent will be born into. Then the behind-the-veil agent would put greater weight than β^t on the utility of the tth generation. For instance, consider a two-generation economy ($t \in \{0,1\}$) in which parents (those born in period 0) put weight 2/3 on their own utility and weight 1/3 on their children's. (This corresponds to $\beta = 1/2$.) However, the children in this economy (those born in period 1) care only for themselves. Given this, if from behind the veil an agent sees a 50 percent chance of being a parent or a child, the agent puts weight 1/3 on first-generation utility and weight 2/3 on second-generation utility. Formally, the problem appears as one in which the discount rate of the agents from behind the veil is different from the discount rate of the households actually in society. (For discussion of the appropriate societal rate of discount, see the 2001 work of Andrew Caplin and John Leahy.)

If the discount rate of agents behind the veil differs from that of the households in society, the analysis of this model becomes considerably more complicated. In particular, scaling does not necessarily hold. Nevertheless, for a simplified version of the two-period model, we can directly compute examples.

For example, with $\alpha = 0.5$, $\sigma = 0.5$, $\bar{\theta} = 0.5$, and $\pi_L = \pi_H = 0.5$, the static model implies that $c_L = 0.11$ and $c_H = 0.39$ (which average to $\pi_H \bar{\theta} = 0.25$). If agents from behind the veil put weight 1/3 on period 0 utility and weight 2/3 on period 1 utility (the case in which $\beta = 1/2$ and parents care about children but children do not care about parents), then the resulting allocation is not a repetition of the static model's. Instead, $c_L = 0.13$, $c_H = 0.37$, $z_L = 0.85$, $z_H = 1.15$, $c_{LL} = 0.09$, $c_{LH} = 0.33$, $c_{HL} = 0.12$, and $c_{HH} = 0.45$. That is, the first generation's consumption is less variable here than in the static model while the second generation's consumption and land allocation depend on its parents' output realization. In this example, agents from behind the veil prefer inequality of opportunity for the second generation even when the second generation matters directly as opposed to mattering only through the altruism of its parents. In fact, for every interior weighting scheme and altruism level attempted, the computed solution always entails some dependence of second-period consumption and land on first-period outcomes. Inequality of opportunity is robust.[10]

Conclusion

I have used Rawls' (1971) behind-the-veil of ignorance device as a fairness criterion to evaluate social policies and applied it to a contracting model in which the terms *equality of opportunity* and *equality of result* are well defined. My results suggest that fairness and inequality—even extreme inequality—are compatible. In a static world, when incentives must be provided, fairness implies equality of opportunity but inequality of result. In a dynamic world of long-lived individuals, fairness implies not only inequality of result; it also implies, eventually, infinite inequality of result. If each period of this dynamic model is interpreted as a generation, then eventual infinite inequality holds for opportunity as well, subject to the condition that fairness is from the perspective of the first generation, that later generations matter only through the first generation's altruism. A computed example shows that if the preferences of later generations are explicitly taken into account, inequality of opportunity still occurs, although perhaps not at extreme levels.

---

[10]In related work (Phelan 2002), I prove that in the two-period moral hazard problem of William Rogerson (1985), efficiency requires that the consumption of children depend on the outcomes of their parents as long as parents care about their children and that the parents' incentive constraint bind in the social planner's problem (which occurs quite generally). Recent work by Scott Freeman and Michael Sadler (2002), however, argues in a similar model that if parents are free to choose bequests, children's consumption depends too much on parental outcomes compared to what would be chosen by a social planner that maximizes steady-state expected lifetime utility.

References

Abreu, Dilip; Pearce, David; and Stacchetti, Ennio. 1986. Optimal cartel equilibrium with imperfect monitoring. *Journal of Economic Theory* 39 (June): 251–69.

_____. 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58 (September): 1041–63.

Atkeson, Andrew, and Lucas, Robert E., Jr. 1992. On efficient distribution with private information. *Review of Economic Studies* 59 (July): 427–53.

Caplin, Andrew, and Leahy, John. 2001. The social discount rate. Discussion Paper 137. Institute for Empirical Economics (Federal Reserve Bank of Minneapolis).

Freeman, Scott, and Sadler, Michael A. 2002. On the optimality of inheritance. Manuscript. University of Texas.

Green, Edward J. 1987. Lending and the smoothing of uninsurable income. In *Contractual arrangements for intertemporal trade*, ed. Edward C. Prescott and Neil Wallace, pp. 3–25. Minnesota Studies in Macroeconomics, Vol. 1. Minneapolis: University of Minnesota Press.

Khan, Aubhik, and Ravikumar, B. 2001. Enduring relationships in an economy with capital. Manuscript. Federal Reserve Bank of Philadelphia.

Phelan, Christopher. 1994. Incentives and aggregate shocks. *Review of Economic Studies* 61 (October): 681–700.

_____. 2002. Inequality of opportunity. Manuscript. Research Department, Federal Reserve Bank of Minneapolis.

Phelan, Christopher, and Townsend, Robert M. 1991. Computing multi-period, information-constrained optima. *Review of Economic Studies* 58 (October): 853–81.

Rawls, John. 1971. *A theory of justice.* Cambridge, Mass.: Belknap Press of Harvard University Press.

Rogerson, William P. 1985. Repeated moral hazard. *Econometrica* 53 (January): 69–76.

Spear, Stephen E., and Srivastava, Sanjay. 1987. On repeated moral hazard with discounting. *Review of Economic Studies* 54 (October): 599–617.

Thomas, Jonathan, and Worrall, Tim. 1990. Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory* 51 (August): 367–90.