

Can Income Redistribution Be Privatized?

Erzo F.P. Luttmer*

Abstract

Income redistribution can be viewed as a process that screens individuals based on their innate ability and then transfers income conditional on ability. In this paper, I develop a mechanism that privatizes income redistribution. This mechanism implements the income distribution desired by the government by inducing profit-maximizing firms to solve the screening problem for the government. I discuss several likely reasons why private firms may solve the screening problem more efficiently than the government, thereby reducing the cost of redistribution.

* J.F. Kennedy School of Government, Harvard University, Cambridge, MA 02138, and National Bureau of Economic Research. E-mail: erzo_luttmer@harvard.edu. I would like to thank Douglas Bernheim, David Cutler, Mark Duggan, Matt Ellman, Martin Feldstein, Maitreesh Ghatak, Ed Glaeser, Larry Katz, Ellen Meara, Antonio Rangel, Andrei Shleifer, Kent Smetters and seminar participants at Boston College, Boston University, Chicago, Harvard, Northwestern and Stanford for suggestions and helpful discussions. Financial support from the Olin Fellowship is gratefully acknowledged. All errors are my own.

1. Introduction

Markets, rather than governments, produce most goods in developed economies. Economists generally prefer market production because market incentives encourage innovation and efficient production. There may be reasons why it is desirable to leave the production of certain goods to the government. These reasons have been widely studied and include adverse selection and issues of contractibility and control (see, e.g., Rothschild and Stiglitz, 1976, or Grossman and Hart, 1986). Nevertheless, market incentives increasingly govern sectors with historically strong government involvement, such as health care, transportation, telecommunications and prison services. The advantages and drawbacks of using market incentives in these sectors have received tremendous attention (e.g., see Vickers and Yarrow, 1988; Donahue, 1989; Wilson, 1989; or Hart, Shleifer and Vishny, 1997). In this light, economists devote surprisingly little attention to the possibilities of using market incentives in the production of “income redistribution,” a good on which society spends between 5-35% of GDP.¹

To identify a role for market incentives in income redistribution, one must first consider the production process of redistribution, or the fundamental reason why redistribution is costly. The Mirrlees (1971) optimal non-linear taxation framework answers this question: redistribution is costly because individuals have private information about their levels of ability. If the government could observe ability levels, it could redistribute without deadweight loss by using lump-sum taxes based on ability but not on income. However, in a world with private information, information about ability levels can be extracted only by distorting individuals’ behavior, and these distortions make

¹ The cost of income redistribution is the deadweight loss from taxation, because the government could collect lump-sum taxes without deadweight loss if redistribution were no concern. Suppose the *average* cost of public funds is half the *marginal* cost of public funds. With estimates of the marginal cost of public funds ranging between 0.3 (Ballard, Shoven and Whalley, 1985) and 2.0 (Feldstein, 1999) and government revenue of about 35% of GDP, the cost of redistribution ranges between 5% and 35% of GDP.

redistribution costly. Mirrlees determined which tax schedule extracts private information about abilities and implements the desired redistribution with the least distortions.²

In this paper, I describe a mechanism that introduces market incentives in the production of redistribution. In this mechanism, the government specifies the desired amount of redistribution, but the actual redistribution is performed by profit-maximizing firms. In a basic optimal income tax model, the mechanism exactly replicates the allocation implemented by the second-best non-linear income tax schedule. In this mechanism, however, the allocation follows from decisions taken by firms maximizing their profits rather than from a social planner optimizing a social welfare function. In practice, this mechanism may increase the efficiency of redistribution over income taxation for three reasons. Firms have incentives for innovation, can observe workers more easily and may use more flexible screening methods.³

The mechanism, formally described in section 2, has three key features. First, all redistribution takes place through implicit cross-subsidies within firms. To finance public expenditure, the government can use lump-sum taxation, which is feasible because the redistribution performed within firms enables every person to pay this tax. Because this tax is invariant, workers and firms receive the full returns to labor supply, including returns to hours, effort and investments in human capital. Second, all working-age individuals are matched to firms and firms compete in the labor market for groups or random allocations of workers, driving expected profits to zero. Third, workers receive an

² Piketty (1993) presents a mechanism that can implement any income redistribution without deadweight loss. The intuition behind his mechanism is that everyone gets severely punished if the distribution of self-reported abilities does not correspond exactly to the true distribution of ability. This gives everyone an incentive to truthfully reveal their ability. However, the mechanism is not very robust, as one deviator can cause everyone to be severely punished.

³ Cooper (1998) presents a model in which non-altruistic communities voluntarily redistribute human capital expenditure towards poorer communities. In her model, private redistribution is driven by a production function with spillover effects across human capital levels.

outside option that protects them from exploitation by firms. By letting the value of this outside option depend on the worker's ability, a wide range of income distributions can be implemented.

This mechanism can be seen as a way to privatize income redistribution.⁴ Firms that redistribute more efficiently than other firms earn supernormal profits, which provides incentives for finding innovative ways of redistributing with fewer distortions. Others may wish to view the mechanism as a form of private workfare. Under the mechanism, all people are guaranteed a job but, unlike standard workfare, these guaranteed jobs are in the private sector.

In many ways, the distinction between privatized redistribution and government redistribution is similar to that between markets and central planning in the production of ordinary goods. Fundamental welfare theory shows that a market allocation with appropriate lump-sum transfers can replicate any Pareto-efficient allocation made by a social planner, yet most economists are convinced that, in practice, markets lead to more efficient allocations than central planning. This conviction stems less from formal theory than from a combination of empirical evidence on the performance of centrally planned economies and general notions about the importance of incentives for innovation, agency problems in government and the benefits of Schumpeterian competition. Similarly, while section 3 shows that in theory the allocations under second-best non-linear taxation and privatized redistribution are identical, in practice many differences are likely to arise. A formal treatment of these differences lies beyond the scope of this paper. Instead, I discuss in section 4 some likely reasons why firms might in practice redistribute more efficiently than the government, such as their incentives for innovation, their ability to observe

⁴ I use the term privatization with the emphasis on redistribution being performed in decentralized fashion within firms that operate in a competitive market. Whether these firms are privately owned or government owned is not relevant for the mechanisms discussed. Hence, I discuss a fundamentally different aspect of privatization than is explored in the recent literature that emphasizes incomplete contracts and residual rights of control (such as in Hart, Shleifer and Vishny, 1997).

workers on a daily basis and their ability to use more flexible screening methods. I also discuss a number of practical issues that one must address before the privatized redistribution mechanism could be implemented. While none of these issues seem insurmountable, some of them may reduce the efficiency of privatized redistribution.

The results shown in this paper suggest that in theory redistribution can be privatized. There is a long way to go, however, before mechanisms for privatized redistribution would approach practical feasibility. Market incentives plausibly can increase the efficiency of redistribution, but this question cannot be answered conclusively by theory alone. This paper argues therefore that forms of privatized redistribution should not be ruled out when looking for ways redistribute more efficiently.

2. A Mechanism for Privatized Redistribution

A. Description of the Privatized Redistribution Mechanism

Under the privatized redistribution mechanism, the government does not engage in redistribution through the tax system. Government revenue is raised by a lump-sum tax on working-age individuals, which is feasible because the privatized redistribution mechanism ensures that everyone is able to pay this tax. To attain the desired income distribution, the government defines the structure of the labor market such that firms and workers, each pursuing their own self-interest, make decisions that result in this income distribution. The structure of the labor market is defined by:

1) The Employment Relationship. The employment relationship cannot be ended without agreement of *both* the worker and the employer. In particular, the employer is not allowed to

fire the worker. While the worker cannot unilaterally end his employment relationship with the original or primary employer, he is free to sell his labor services to 3rd parties. In this case, the primary employer will receive a fraction τ (determined by the government) of the labor earnings of the worker.

2) The Outside Option. The worker's outside option consists of two parts. First, the primary employer must pay a basic salary, \underline{C} , to the worker (whether or not the worker provides any labor services to the primary employer). As explained below, the level of \underline{C} is determined in the labor market. Second, the worker may perform labor services for 3rd parties and retain a fraction $1-\tau$ of those earnings, with the remainder going to the primary employer.⁵

3) Clearing of Labor Market. For a given τ , employers reveal their demand for workers as a function of the basic salary. The government acts as the auctioneer to determine the market-clearing basic salary, \underline{C} , at which *all* working-age individuals can be matched to an employer, and allocates to each employer the number of workers that she demands at the market-clearing level of the basic salary. The allocation is *random* to prevent employers from trying to select only high-ability workers. Since all workers are matched to firms, there is no unemployment.

Given this structure of the labor market, employers offer each of their workers an incentive schedule, which defines salary as a function of labor supply. Workers select their preferred labor supply from the incentive schedule, supply this amount of labor, and are paid accordingly.

Alternatively, they can exercise their outside options. As explained below, it is optimal for firms to offer the incentive schedules that implicitly redistribute from higher-ability workers to lower-ability workers.

⁵ The fact that, if exercised, this second component of the outside option would distort the labor supply downwards (due to the implicit tax rate of τ) does not matter, since this outside option is not exercised in equilibrium.

B. Informal Discussion of the Operation of Privatized Redistribution

Before turning to a more formal model of the redistribution mechanism in section 3, I discuss the basic intuition behind the mechanism. For ease of exposition, assume for the moment that (i) firms can fully observe worker ability and (ii) that τ equals one, which means that the second outside option of working for a third party is never attractive (since 100% of the earnings must be given to the primary employer). Both assumptions will be relaxed later in this subsection.

The basic intuition is relatively simple. Since a worker cannot not leave the employment relationship unilaterally, an employer will extract the maximal surplus (positive or negative) out of the worker *subject to the constraint* that the worker must at least attain the utility level associated with the basic salary. Hence, under perfect information, the employer will drive down every worker's utility to the level of the basic salary and have all workers perform the first-best level of labor supply because this level maximizes total surplus. However, since firms must compete for allocations of workers, the basic salary will be driven up until firms no longer earn supernormal profits. Workers differ in their productivity but their utility levels are equalized to the level associated with the basic salary. Hence, firms extract positive surpluses from higher-ability workers that are used to finance the negative surpluses extracted from lower-ability workers.⁶ Privatized redistribution thus takes place through implicit cross-subsidies within firms. These cross-subsidies are driven by two factors. First, all workers have the same outside option. Second, on the labor market, firms compete for *randomly-allocated* workers rather than for specific individuals. Hence, competition will reduce any supernormal profits associated with employing an "average" worker, but surpluses associated with employing higher-ability

⁶ I use the term "higher-ability workers" to denote workers who are sufficiently productive that a firm earns a profit on them. Conversely, "lower-ability workers" are workers who are implicitly subsidized by a firm.

individuals will not be competed away.

This basic intuition remains valid when employers have imperfect information about workers' abilities. Instead of having each worker perform the first-best labor supply, an employer now offers an incentive schedule in which labor supply for lower-ability workers is distorted downwards in order to extract more surplus from higher-ability individuals. Moreover, higher-ability individuals may capture informational rents, thus attaining a level of utility in excess of the level associated with the basic salary. With better information about worker ability, both the labor supply distortion of the lower-ability workers and the informational rent of the higher-ability workers can be reduced, increasing profits. Hence, firms have incentives to improve their assessments of workers' abilities.

When τ equals unity, ability is not rewarded in utility terms as long as employers can perfectly observe abilities. If workers' abilities are imperfectly observable, higher-ability workers will earn informational rents. However, in an effort to reduce these informational rents, it is optimal for firms to distort the labor supply of lower-ability individuals to the point that the marginal cost of this distortion fully dissipates the marginal reduction in informational rents paid to higher-ability types. In other words, this mechanism implicitly places a social value of zero on utility of higher-ability workers. When the government reduces τ , it strengthens the workers' outside option since workers retain a fraction $(1-\tau)$ of their labor earnings at 3rd parties. Because higher-ability workers are more productive, the value of this outside option increases with ability, which induces employers to offer incentive schedules that partly reward ability. Moreover, since a reduction in τ limits the surplus that can be extracted from higher-ability workers, it also limits need for firms to distort the labor supply of the lower-ability workers in order to maintain incentive compatibility.

Hence, by choosing τ , the government can select a point on the schedule that trades off the amount of redistribution with the costs of redistribution.

To make privatized redistribution feasible, many practical issues would need to be addressed including opportunities for employment changes, bankruptcy, retirement rules and the age at which individuals are matched to firms. While the aim of this paper is limited to showing that privatized redistribution is possible in theory, I discuss these and other practical issues in section 4 and argue that they are unlikely to pose insurmountable problems for privatized redistribution.

C. Key Features of Privatized Redistribution

The key insight behind the privatized redistribution mechanism is that competition between firms for specific workers, which causes salaries to be linked to individual abilities, can be eliminated without reducing competition between firms for workers in general, which is necessary to prevent firms from capturing rents. This can be achieved by having firms compete for large groups or random allocations of workers rather than for specific individuals. This insight has been used before, in particular in Diamond's (1992) proposals for reform in the health insurance market. In that proposal, the government would divide the entire population into large groups and there would be competition between insurance companies to insure these groups. As a result, selection problems (making insurance prohibitively expensive for individuals with poor health risks) would be eliminated, the costs associated with marketing, administration and underwriting would be reduced but the benefits of market competition among insurance companies would be preserved.

By definition, all redistribution mechanisms entail some loss of freedom in order to take

resources from the rich. In the regular income tax system, individuals lose the freedom to keep all their earnings. In the privatized redistribution mechanism, individuals are tied to a firm, but are protected by the outside options of the basic salary and of working elsewhere and retaining a fraction $(1-\tau)$ of those earnings. In fact, if τ is made a function of earnings such that it replicates the income tax structure, individuals would lose no more freedom under the privatized redistribution mechanism than under income taxation.⁷ In the privatized redistribution mechanism, the parameter τ determines the extent of redistribution. If τ equals zero, workers can keep their full labor earnings if they exercise the outside option of working for a third party. In this case, the primary employer has no power to extract surplus from higher-ability workers and the basic salary will be zero. In other words, a value of τ equal to zero replicates a laissez-faire society. A value of τ equal to unity will equalize the outside option across all workers (to the basic salary), and would yield a flat utility distribution under perfect information. Under imperfect information, higher-ability workers would still be better off due to informational rents. Hence, by varying τ between zero and one, a wide range of distributional preferences can be accommodated.

3. Model

A. Basic Framework

In this subsection, I present a simple income tax model as a benchmark for privatized

⁷ In some sense, this mechanism is the complement to a mechanism proposed by Kremer (1997) that allows the most productive individuals to buy themselves out of distortionary taxation. Hence, in Kremer's proposal, individuals have the outside option of leaving the regular tax system, whereas in the privatized redistribution mechanism, individuals have the outside option of entering a regular tax-system.

redistribution, modeled later in this section.⁸ This model follows the standard Mirrlees framework in which individuals differ in their levels of ability, which are unobservable. In this framework, the government can redistribute from high-ability individuals to low-ability individuals only if it provides incentives to high-ability individuals not to pretend to have a low level of ability. It can provide these incentives by distorting the labor supply of low-ability individuals downwards because such a distortion is less costly for them than for high-ability individuals. The distortion makes redistribution costly, but is necessary to discriminate between individuals of different abilities.

For simplicity, I use a one-period model with two types of persons, low-ability workers (fraction α) and high-ability workers (fraction $1-\alpha$). High-ability workers have a constant marginal product of w_H and low-ability workers have a constant marginal product of w_L , where $0 < w_L < w_H$. The two types have the same utility function, $U(C,L) = C - v(L)$ where C is consumption, L is labor supply and v is an increasing and convex function with $v(0)=0$ and $v'(0)=0$.⁹ Let Y denote a worker's total production, i.e. $Y = wL$. Finally, I assume that the government optimizes an additive social welfare function of the form $\sum_i \varphi(U_i)$, where i indexes individuals and φ is an increasing and concave function.

If the government could observe each person's ability, it would be able to implement the first-best allocation. In this case, the government maximizes the social welfare function subject to the resource constraint.

⁸ This model is a simple variant of optimal income tax models discussed extensively in the literature (for a discussion, see Atkinson and Stiglitz, 1980). No original results are presented in this subsection, it merely serves as a benchmark for the privatized redistribution mechanism.

⁹ Labor supply is broader than just hours. It includes any way a worker can affect his output, including components such as effort and investments in human capital.

$$\underset{\{C_L, C_H, Y_L, Y_H\}}{\text{Max}} \quad \alpha \varphi(C_L - v(Y_L/w_L)) + (1-\alpha) \varphi(C_H - v(Y_H/w_H)) \quad (1a)$$

$$\text{s.t.} \quad \alpha(C_L - Y_L) + (1-\alpha)(C_H - Y_H) \leq 0 \quad (1b)$$

The conditions for the first-best allocation are:

$$1 - v'(Y_H^{FB}/w_H)/w_H = 0 \quad (2a)$$


$$1 - v'(Y_L^{FB}/w_L)/w_L = 0 \quad (2b)$$

$$C_H^{FB} = v(Y_H^{FB}/w_H) + \text{Surplus}^{FB} \quad (2c)$$

$$C_L^{FB} = v(Y_L^{FB}/w_L) + \text{Surplus}^{FB} \quad (2d)$$

where $\text{Surplus}^{FB} \left(\equiv \alpha Y_L^{FB} + (1-\alpha)Y_H^{FB} - \alpha v(Y_L^{FB}/w_L) - (1-\alpha)v(Y_H^{FB}/w_H) \right)$ is the first-best level of surplus, or total production net of total disutility of labor. The superscript FB indicates that the values apply to the first-best allocation. Conditions (2a) and (2b) show that labor supply is efficient for both types. Conditions (2c) and (2d) show that the two types achieve the same level

Figure 1: First-Best Allocation



of utility. To implement this allocation, a high-ability worker needs to pay a lump-sum tax of $(Y_H^{FB} - C_H^{FB})$ and a low-ability worker receives a transfer of $(C_L^{FB} - Y_L^{FB})$.

The first-best allocation is depicted graphically in figure 1. The indifference curve of the low-ability worker (U_L) is steeper than the indifference

curve of the high-ability worker (U_H) because a low-ability worker has to increase labor supply more (and, hence, be compensated more) than a high-ability worker for a given increase in production. In the first-best allocation, low-ability workers receive bundle A and high-ability workers bundle D . In this allocation, both types provide the efficient labor supply (the indifference curves are tangent to the 45° line) and both types achieve the same utility (the indifference curves intersect at $Y=0$ where $v(0)=0$ for both types). The tax revenue DE from high-ability workers is used to finance the subsidy AB to low-ability workers.

The first-best allocation is not attainable, however, if the government cannot observe ability. In this case, the government can only base its taxes and transfers on observed income. When ability is not observable, the government needs to provide an incentive to each type to truthfully reveal its ability. As figure 1 shows, a high-ability worker would prefer the allocation of a low-ability worker, bundle A , over his own allocation, bundle D , but not vice versa. Hence, the incentive constraint will bind only for the high-ability type when ability is unobservable. The second-best allocation is obtained by maximizing social welfare subject to the resource constraint and the incentive constraint for the high-ability type:

$$\underset{\{C_L, C_H, Y_L, Y_H\}}{\text{Max}} \quad \alpha \varphi(C_L - v(Y_L/w_L)) + (1-\alpha) \varphi(C_H - v(Y_H/w_H)) \quad (3a)$$

$$\text{s.t.} \quad \alpha(C_L - Y_L) + (1-\alpha)(C_H - Y_H) \leq 0 \quad (3b)$$

$$C_L - v(Y_L/w_H) \leq C_H - v(Y_H/w_H) \quad (3c)$$

The first-order conditions for the second-best allocation are given by:

$$\left(1 - v'(Y_H^{SB}/w_H)/w_H\right) = 0 \quad (4a)$$

$$\left(1 - v'(Y_L^{SB}/w_L)/w_L\right) \alpha = \left[v'(Y_L^{SB}/w_L)/w_L - v'(Y_L^{SB}/w_H)/w_H \right] \times \left[\frac{\alpha(1-\alpha)(\varphi'_L - \varphi'_H)}{\alpha\varphi'_L + (1-\alpha)\varphi'_H} \right] \quad (4b)$$

$$C_H^{SB} = v(Y_H^{SB}/w_H) + \alpha \text{ InfoRent}^{SB} + \text{Surplus}^{SB} \quad (4c)$$

$$C_L^{SB} = v(Y_L^{SB}/w_L) - (1-\alpha)InfoRent^{SB} + Surplus^{SB} \quad (4d)$$

where the superscript SB indicates that the values apply to a second best allocation, where

$\phi'_L \left(\equiv \phi'(C_L^{SB} - v(Y_L^{SB}/w_L)) \right)$ and $\phi'_H \left(\equiv \phi'(C_H^{SB} - v(Y_H^{SB}/w_H)) \right)$ denote the marginal social valuation of income for respectively the low-ability and high-ability type, and where total surplus is denoted by $Surplus^{SB} \left(\equiv \alpha Y_L^{SB} + (1-\alpha)Y_H^{SB} - \alpha v(Y_L^{SB}/w_L) - (1-\alpha)v(Y_H^{SB}/w_H) \right)$. The high-ability type captures an informational rent, $InfoRent^{SB} \equiv v(Y_L^{SB}/w_L) - v(Y_H^{SB}/w_H)$, equal to the difference in utility that the high- and low-ability type attain from the bundle for the low-ability person. Since $v(\cdot)$ is an increasing and convex function and $w_L < w_H$, this informational rent decreases as Y_L^{SB} is reduced. In other words, by reducing the labor supply of the low-ability type, the informational rent is reduced and the distribution becomes more equal. Condition (4a) shows that the labor supply of the high-ability type is efficient in the second-best allocation. The labor supply of the low-ability type, however, is distorted downwards as condition (4b) shows. This downward distortion is optimal in the second-best allocation because it eases the incentive constraint for the high-ability type, thus allowing more redistribution to the low-ability type. The left-hand side of (4b) shows that marginal cost of reducing Y_L^{SB} is equal to the lost production minus the decrease in disutility of labor, both multiplied by α , the fraction of low-ability workers. The right hand side shows that the marginal benefit of decreasing Y_L^{SB} is equal to the marginal reduction in informational rents (the first term between square brackets) times the social valuation of this reduction (the second term between square brackets). Thus, Y_L^{SB} will be distorted downwards more when (i) informational rents are more sensitive to distortions in Y_L^{SB} , (ii) there is a larger difference between the social marginal valuation of income of the high- and low ability type, $\phi'_L - \phi'_H$, or (iii) the fraction α of low-ability types is lower. The latter result can be seen by dividing condition

(4b) by α , but the intuitive explanation is that if the fraction of low-ability types is high, it is relatively costly to distort the labor supply of the low-ability type (because there are many of them), and relatively cheap to pay an informational rent to the high-ability type (because there are few of them). Hence, the downward distortion of Y_L^{SB} is decreasing in α and the informational rent is increasing in α .

Compared to the first-best allocation, social welfare in the second-best allocation is lower for two reasons. First, the total surplus has decreased because the labor supply of the low-ability type is distorted below its efficient level:

$$DWL \text{ from labor distortion} = \alpha \left(Y_L^{FB} - Y_L^{SB} - v(Y_L^{FB}/w_L) + v(Y_L^{SB}/w_L) \right) = \alpha \int_{Y_L^{SB}/w_L}^{Y_L^{FB}/w_L} (w_L - v'(L)) dL \quad (5a)$$

Second, the allocation of the surplus is no longer efficient because high-ability types receive informational rents. This reduces social welfare by:

$$Misallocation \text{ Cost} = \int_0^{InfoRent^{SB}} \left(\alpha \varphi'(Surplus^{SB} - (1-\alpha)x) - (1-\alpha) \varphi'(Surplus^{SB} + \alpha x) \right) dx \quad (5b)$$

For future reference, consider the outcome for extremely redistributive preferences, such as a Rawlsian social welfare function that places a marginal social valuation of income of 1 on the least well-off individual (the low-ability type) and zero on anybody else. Hence, φ'_L equals 1 and φ'_H equals 0. In this case, conditions (4a), (4c) and (4d) would all remain applicable, but equation (4b) would simplify to:

$$\left(1 - v'(Y_L^{SB}/w_L)/w_L \right) \alpha = \left[v'(Y_L^{SB}/w_L)/w_L - v'(Y_L^{SB}/w_H)/w_H \right] \times [1-\alpha] \quad (4b')$$

The expression (5a) for DWL would remain applicable but the expression for the misallocation cost would simplify to the amount by which the informational rent reduces the utility of the low-ability type:

mechanism is likely to be more efficient in practice than income taxation.

In the privatized redistribution mechanism, firms maximize expected profits subject to three constraints. First, a firm must offer a low-ability worker a contract that is at least as good as the outside option. Second, it must offer a high-ability worker a contract that is at least as good as the outside option. Finally, it must induce high-ability workers to choose the contract intended for them over the contract intended for low-ability workers. The outside option consists of two components, (i) the basic salary, \underline{C} , and (ii) the possibility of working for a 3rd party and keeping a fraction $(1-\tau)$ of those earnings. Let the utility associated with this second component be denoted by $U_L(\tau)$ and $U_H(\tau)$ respectively for the low- and high-ability type.¹⁰ Since an allocated worker is of low ability with probability α , expected profit is given by $\alpha(Y_L - C_L) + (1-\alpha)(Y_H - C_H)$, and firms solve:

$$\underset{\{C_L, C_H, Y_L, Y_H\}}{\text{Max}} \quad \alpha(Y_L - C_L) + (1-\alpha)(Y_H - C_H) \quad (6a)$$

$$\text{s.t.} \quad C_L - v(Y_L/w_L) \geq \underline{C} + U_L(\tau) \quad (6b)$$

$$C_H - v(Y_H/w_H) \geq \underline{C} + U_H(\tau) \quad (6c)$$

$$C_H - v(Y_H/w_H) \geq C_L - v(Y_L/w_H) \quad (6d)$$

Constraint (6b), the outside option for the low type, is always binding. Only cases where the incentive constraint (6d) is binding are interesting. Otherwise we would be considering an anomaly of the two-type model where costless redistribution is possible. The outside option for the high-type, constraint (6c), may or may not be binding depending on the values of τ and α . For sufficiently high values of τ and α , constraint (6c) is not binding, and the operative conditions for the optimum are:

¹⁰ $U_L(\tau)$ equals the maximum of $(1-\tau)Y - v(Y/w_L)$ with $Y \geq 0$ being the choice variable. $U_H(\tau)$ is found analogously.

$$\left(1 - v'(Y_H^P/w_H)/w_H\right) = 0 \quad (7a)$$

$$\left(1 - v'(Y_L^P/w_L)/w_L\right) \alpha = \left[v'(Y_L^P/w_L)/w_L - v'(Y_L^P/w_H)/w_H\right] \times [1 - \alpha] \quad (7b)$$

$$C_H^P = v(Y_H^P/w_H) + \underline{C} + U_L(\tau) + \left(v(Y_L^P/w_L) - v(Y_L^P/w_H)\right) \quad (7c)$$

$$C_L^P = v(Y_L^P/w_L) + \underline{C} + U_L(\tau) \quad (7d)$$

where the superscript P indicates that the values apply to the privatized redistribution mechanism.

Condition (7a), which is identical to condition (4a), indicates that the labor supply of the high-

ability type is not distorted: $Y_H^P = Y_H^{SB} = Y_H^{FB}$. Condition (7b), which is identical to condition (5b'),

shows that the labor supply by the low-ability type is distorted downward and is equal to the value

that a Rawlsian social planner would choose: $Y_L^P = Y_L^{SB}$. Since the levels of labor supply in the

privatized redistribution mechanism are equal to those in the second-best, it follows that $Surplus^P$

equals $Surplus^{SB}$ and that $InfoRent^P$ equals $InfoRent^{SB}$. To show that the consumption levels in the

current problem are also equal to those chosen by the Rawlsian social planner, we need to

eliminate $\underline{C} + U_L(\tau)$ from conditions (7c) and (7d). We can solve for $\underline{C} + U_L(\tau)$ by setting

expected profit to zero, as is required for an equilibrium in the labor market. Expected profit is

given by:


$$\begin{aligned} \pi(\alpha) &= \alpha(Y_L^P - C_L^P) + (1-\alpha)(Y_H^P - C_H^P) \\ &= \alpha(Y_L^P - v(Y_L^P/w_L) - \underline{C} - U_L(\tau)) + (1-\alpha)(Y_H^P - v(Y_H^P/w_H) - \underline{C} - U_L(\tau) - (v(Y_L^P/w_L) - v(Y_L^P/w_H))) \\ &= -\underline{C} - U_L(\tau) + \left(\alpha(Y_L^P - v(Y_L^P/w_L)) + (1-\alpha)(Y_H^P - v(Y_H^P/w_H))\right) - (1-\alpha)(v(Y_L^P/w_L) - v(Y_L^P/w_H)) \\ &= -\underline{C} - U_L(\tau) + \quad \quad \quad \mathit{Surplus}^P \quad \quad \quad - (1-\alpha) \mathit{InfoRent}^P \end{aligned} \quad (8)$$

Setting profits to zero shows that the outside option for the low-ability type, $\underline{C} + U_L(\tau)$, is equal to

the surplus minus the total informational rent paid to the high-ability types. Substituting this into

equations (7c) and (7d) shows that the consumption levels are the same as obtained by the

Figure 3: Privatized Redistribution



becomes more attractive and at some point becomes binding. Similarly, as α falls, informational rents decrease until the outside option for the high-ability type becomes binding. The case with the binding outside option for the high-ability type is shown graphically in figure 3. The indifference curves of the low- and high-ability types are tangent to the outside option line, indicating that the outside option is binding for both. Moreover, the indifference curve of the high-ability type runs through the allocation for the low-ability type, F , indicating that the incentive constraint is binding. Due to this incentive constraint, the labor supply of the low-ability type is distorted downwards, below the efficient level, A . Algebraically, the operative conditions for the optimum are:

$$\left(1 - v'(Y_H^P/w_H)/w_H\right) = 0 \quad (9a)$$

$$v(Y_L^P/w_L) - v(Y_L^P/w_H) = U_H(\tau) - U_L(\tau) \quad (9b)$$

$$C_H^P = v(Y_H^P/w_H) + C + U_H(\tau) \quad (9c)$$

$$C_L^P = v(Y_L^P/w_L) + C + U_L(\tau) \quad (9d)$$

Rawlsian social planner.

For sufficiently low values of τ and α , the outside option for high-ability workers (constraint 6c) becomes binding. As τ , the tax rate collected by the primary employer on earnings from 3rd parties, is decreased, the outside option for the high-ability type

Equation (9a) is equal to equation (4a) and establishes that the labor supply of the high-ability type is efficient, like in the second-best optimum. With constraint (6c) binding, the equation determining the labor supply of the low-ability type, (9b), is markedly different from corresponding equation, (4b) of the second-best optimum. Rather than being determined by the trade-off between the cost of the downward distortion and the benefit of reducing the informational rent, the labor supply of the low-ability type is now determined by setting the informational rent (left-hand side of 9b) equal to the utility difference in the outside options. This utility difference is determined by the slope $(1-\tau)$ of the outside option line. Hence, to replicate the second-best solution, the policymaker needs to set τ such that the solution for Y_L^P from equation (9b) is equal to the value for Y_L^{SB} found in equation (4b).¹¹ Since the levels of labor supply in the privatized redistribution mechanism are equal to those in the second-best, it follows that $Surplus^P$ equals $Surplus^{SB}$ and that $InfoRent^P$ equals $InfoRent^{SB}$. Finally, to show that the consumption levels determined by equations (9c) and (9d) are also equal to those found in the second-best optimum, we first need to solve for the equilibrium level of \underline{C} , which is done by setting expected profits to zero:

¹¹ To prove that such a $\tau \in [0,1]$ exists, note that :

- (i) The right-hand side of (9b), $U_H(\tau) - U_L(\tau)$ is a continuously decreasing function of τ ,
 - (ii) $U_H(1) - U_L(1) = 0$, because at a tax rate of 100% it is optimal for both types not to work and receive zero utility, and
 - (iii) $U_H(0) - U_L(0) = Y_H^{FB} - v(Y_H^{FB}/w_H) - Y_L^{FB} + v(Y_L^{FB}/w_L) \geq Y_L^{FB} - v(Y_L^{FB}/w_H) - Y_L^{FB} + v(Y_L^{FB}/w_L) = v(Y_L^{FB}/w_L) - v(Y_L^{FB}/w_H) > v(Y_L^{SB}/w_L) - v(Y_L^{SB}/w_H)$, where the first inequality follows from revealed preference of the high type and the second inequality follows from the fact that the informational rent, $v(Y/w_L) - v(Y/w_H)$, is increasing in Y and that $Y_L^{FB} > Y_L^{SB}$.
- From (i), (ii) and (iii), it follows that there exists a $\tau \in [0,1]$ such that $v(Y_L^{SB}/w_L) - v(Y_L^{SB}/w_H) = U_H(\tau) - U_L(\tau)$

$$\begin{aligned}
\pi(\alpha) &= \alpha(Y_L^P - C_L^P) + (1-\alpha)(Y_H^P - C_H^P) \\
&= \alpha(Y_L^P - v(Y_L^P/w_L) - \underline{C} - U_L(\tau)) + (1-\alpha)(Y_H^P - v(Y_H^P/w_H) - \underline{C} - U_H(\tau)) \\
&= -\underline{C} - U_L(\tau) - (1-\alpha)(U_H(\tau) - U_L(\tau)) + (\alpha(Y_L^P - v(Y_L^P/w_L)) + (1-\alpha)(Y_H^P - v(Y_H^P/w_H))) \\
&= -\underline{C} - U_L(\tau) - (1-\alpha) \mathit{InfoRent}^P + \mathit{Surplus}^P = 0
\end{aligned} \tag{10}$$

Substituting the solution found for \underline{C} from (10) into equations (9c) and (9d), shows that the consumption levels in the privatized redistribution mechanism are equal to those found in the second-best. Thus, the privatized redistribution mechanism replicates the allocation under non-linear income taxation both when the outside option of the high type is binding and when it is not.

4. Would Privatized Redistribution Be More Efficient Than Income Taxation?

While privatized redistribution and optimal income taxation are equivalent in the stylized model above, important differences are likely to arise in practice. Competitive pressures can spur innovation in the way firms screen their workers on ability, thus making privatized redistribution more efficient. On the other hand, the privatized redistribution mechanism raises many practical issues, such as how to deal with bankruptcy, match-specificities or worker relocations . It seems that none of these issues are insurmountable problems, but some of the solutions may entail costs that reduce the efficiency of privatized redistribution. A rigorous analysis of whether in practice the benefits of privatized redistribution would outweigh the costs is beyond the scope of this paper, but the next two subsections offer some considerations relevant for evaluating these costs and benefits.

A. Possible Benefits of Privatized Redistribution

The main benefit from privatized redistribution is that it decentralizes the decision of *what mechanism* to use to screen workers on ability without decentralizing the decision of how much to redistribute to which types of workers. This decentralization creates competition which provides incentives for innovations that reduce the cost of redistribution. This incentive can be seen by examining the expression found in equations (8) and (10) for the expected profit of a firm:

$$\pi(\alpha) = - (\underline{C} + U_L(\tau)) + \text{Surplus}^P - (1-\alpha) \text{InfoRent}^P \quad (11)$$

Equation (11) shows that the profits of a particular firm depend on three terms: (i) the cost of the outside option for low-ability workers, (ii) the surplus created by the firm, (iii) the informational rents paid by the firm to high-ability workers. The first term is exogenous to the firm. The surplus depends on the distortions created by the screening technique used by the firm. Hence, any reduction in the deadweight loss of these distortions will increase the firm's profits dollar-for-dollar. If the outside option for the high-ability type is not binding, the informational rent depends on the firm's choice of screening technique, and the informational rent will reduce the profits one-for-one. Hence, the firm has the correct incentive to reduce deadweight loss and implicitly values any payment of informational rents to the high-ability type as a pure loss, which corresponds to the Rawlsian definition of misallocation costs found in equation (5b'). If the outside option of the high-ability type is binding, the informational rent is determined by τ (see eq. 9b) and is thus exogenous to the firm. Hence, in this case, the firm's objective is to maximize surplus, subject to the distributional requirements defined by the outside option line.

While it is difficult to predict what innovations in screening techniques would be prompted by competition between firms, firms enjoy three advantages relative to the government in developing

more efficient screening techniques. First, as employers, they can observe more about a worker than the government, and use this information to assess the worker's ability. Second, their incentive schedule, unlike a tax system, need not be rule-based, but may rely on discretion. Hence, they can use information that is hard to quantify to decide what kind of incentive schedule to offer a particular worker. Finally, they may have more instruments than the government to induce high-ability individuals to reveal their ability. For example, they could have low-ability workers perform tasks that low-ability workers dislike less than high-ability workers.

Instead of modeling these advantages separately, I assume that they are reflected in a signal that a firm obtains about each of its workers. The precision of the signal may vary across firms. Firm j receives a signal about the worker's ability that is completely informative with probability p_j and completely uninformative with probability $(1-p_j)$. Hence, conditional on getting a negative signal, the probability for firm j that the worker is of low ability is $\alpha_j^N = p_j + (1-p_j)\alpha$. Similarly, conditional on a positive signal, the probability that the worker is of low ability is: $\alpha_j^P = (1-p_j)\alpha$. Independently of the precision of the signal, firms receive a negative signal with probability α and a positive signal with probability $(1-\alpha)$.

First, consider the case where constraint (6c), the outside option for the high-ability type, is never binding ($\tau=1$). In this case, the firm will offer a different incentive schedule to workers about whom it receives a negative signal than to workers about whom it receives a positive signal. If the firm receives a negative signal and worker is thus likely to be of low ability, the firm offers a schedule that creates a relatively small labor distortion for low-ability workers and a relatively large informational rent for high-ability workers. If the firm receives a positive signal, it offers a schedule that creates a relatively large labor distortion and a relatively small informational rent.

Hence, an incentive schedule that depends on the signal increases the firm's expected profit by reducing the expected deadweight loss and the expected payments of informational rents.

For values of τ below 1, the strength of the signal will determine whether constraint (6c), the outside option for the high-ability type, will be binding. By the same logic as above, it is optimal for a firm receiving a negative signal to reduce the labor distortion and increase the informational rent. If the signal is sufficiently strong and negative (the posterior value of α sufficiently high), it is optimal for the firm to set the informational rent at a very high level, where constraint (6c) will no longer bind. When (6c) ceases to bind, any further negative information about the worker's ability will reduce the optimal labor distortion of the low-ability. Conversely, if the signal is sufficiently strong and positive, indicating that the worker is very likely to be of high ability, it is optimal for the firm to increase the labor supply distortion of the low-ability type and to reduce the informational rent going to the high ability type. As the signal becomes stronger, the informational rent will be reduced to such an extent that the outside option for high-ability type becomes binding. At this point, any further increase in the signal's strength will no longer increase the labor distortion nor decrease the informational rent. Hence, the outside option for the high type effectively puts an upper bound on the labor distortion of the low-ability type and a lower-bound on the informational rent going to the high-ability type.

Algebraically, the incentive schedule offered to workers with a negative signal is defined by equations (9a-d) if constraint (6c) is binding and by equations (7a-d) otherwise where α is replaced by α_j^N . Let the expected profit conditional on a negative signal be denoted by $\pi(\alpha_j^N)$. Similarly, the firm offers workers with a positive signal the incentive schedule defined by equations (10a-d) if constraint (6c) is binding and by equations (7a-d) otherwise where α is

replaced by α_j^P . The expected profit, conditional on a positive signal, is $\pi(\alpha_j^P)$. Because profits are a convex function of α , the unconditional expected profit is:¹²

$$\pi|_{p_j} = \alpha \pi(\alpha_j^N) + (1-\alpha) \pi(\alpha_j^P) \geq \pi(\alpha \alpha_j^N + (1-\alpha) \alpha_j^P) = \pi(\alpha) \quad (12)$$

where the inequality is strict if constraint (6c) is not always binding. Hence, the presence of a signal of the worker's ability weakly increases the firm's expected profit. By a similar argument, an increase in the precision p_j of the signal weakly increases firm j 's expected profit. This provides the incentive for firms to invest in technologies that increase the precision of the signal.¹³ While an increase in signal precision leads to supernormal profits for an inframarginal firm, once the signals of other firms also become more precise, the competition for workers will increase and drive up the basic salary until expected profits again equal zero. This increase in the basic salary implies an improvement in social welfare over the level attained in the second-best allocation.

B. Practical Issues and Possible Drawbacks of Privatization

There are many design issues and practical questions that need to be addressed when considering the practical feasibility of the privatized redistribution mechanism. The model abstracted from these issues because they do not alter the basic idea behind privatized redistribution. However, in practice, they are important and some of them may entail costs that

¹² This is most easily seen using a standard argument in economics. Suppose, we increase α slightly. If the firm does *not* adjust its incentive schedule (defined by Y_L , Y_H , C_L and C_H), profits would decrease linearly, as the profit function shows: $\pi(\alpha) = \alpha(Y_L - C_L) + (1-\alpha)(Y_H - C_H)$. However, because the firm can adjust the incentive schedule, profits must weakly exceed the level predicted by the linear decrease. Next, suppose we decrease α slightly. If the firm does *not* adjust its incentive schedule, profits would increase linearly. However, because the firm can adjust the incentive schedule, profits exceed the level predicted by the linear increase. Hence, profits are a convex function of α .

¹³ These incentives may be weaker or stronger than socially optimal for the same reasons as why the level of research and development in traditional markets may not be socially optimal. See Tirole (1988) for an overview of this literature.

reduce the efficiency of privatized redistribution. Below, I discuss the most prominent practical issues.

1. Match-specificities. One obvious concern with the privatized redistribution mechanism is whether the randomly-allocated workers have the attributes and skills sought by the employers. Matching on location can be accommodated by making the allocations of workers region-specific (with adjustments for any regional differences in the average ability levels of workers). This still leaves the problem that an employer may need particular skills that a randomly-allocated worker is unlikely to possess. This problem can be solved in two ways, but both ways probably entail some transaction costs. First, the primary employer of a worker may detail the worker to a third company in return for a periodic payment. The primary employer remains responsible for the worker's salary and any other employer obligations. Second, the worker may transfer to a new employer who needs his skills or where he prefers to work. Since the employment relationship cannot be ended unilaterally, the worker needs to obtain permission from his current employer. Suppose the employment change takes place immediately after the worker has selected the preferred labor supply from the incentive schedule but before he has performed any work. At this point, the worker has revealed his type because he has chosen his labor supply. Because the employer makes a profit of $(Y_H^P - C_H^P)$ on a high-ability worker, she will only agree to let a high-ability worker go in exchange for a transfer sum that is at least as large as this profit. Similarly, she makes a loss of $(C_L^P - Y_L^P)$ on low-ability workers, and is willing to pay a transfer sum of at most this amount to a prospective employer. If the worker is at least as productive at the prospective employer as at the current one, a prospective employer can hire a high-ability worker for a transfer sum of $(Y_H^P - C_H^P)$ and a salary of C_H^P without making a loss. Similarly, after

receiving a transfer sum of $(C_L^P - Y_L^P)$, a prospective employer can pay a low-ability worker a salary of C_L^P .¹⁴ Hence, transfer payments exist that allow workers to change employment when such a change leads to a better match.

2. Small Firms and Risk-Aversion. Firms requesting random allocations of workers face uncertainty with regard to the average ability of their allocation. If firms are risk-averse and insurance markets are imperfect, this uncertainty is costly. In practice, smaller or more risk-averse firms would probably not request random allocations of workers but instead hire workers from larger companies that do request random allocations of workers. Of course, there are likely transaction costs associated with hiring workers from other companies.

3. Age at which individuals are matched to firms. In the absence of early matching, an individual would only receive part of the returns to educational investments, with the remainder going to the future employer. This would result in underinvestment in education (just like income taxes reduce educational investments). When individuals are matched early, however, the employer can share in educational expenses and has an incentive to do so. On the other hand, an employer is solely interested in educational investments that raise a worker's economic productivity and not in types of education to which society may attach "merit value" (e.g. enabling political or civic engagement).

4. Retirement. There are several ways in which retirement could be incorporated in the mechanism. One option is that the government stipulates that firms must raise the basic salary at retirement to a level that is deemed socially acceptable for retired persons. Retirees are still

¹⁴ The current employer also requires from the new employer that the salary remains C_L^P because this prevents that the low-ability type will perform the first-best labor supply at the new employer. This is necessary to ensure that the incentive constraint (6d) is not violated.

allowed to work within the privatized redistribution mechanism, but the higher level of the basic salary ensures that working is not a financial necessity. The obligation of firms to pay a higher basic salary at retirement creates the need to extract more surplus out of workers during working life, lowering the basic salary during working life. Hence, in effect, it is a forced savings mechanism administered by firms. More generally, by making the outside option age-dependent, redistribution over the life-cycle can be incorporated in the mechanism.

5. Morbid incentives. Since lower-ability individuals are costly to a firm (because the firm implicitly subsidizes lower-ability workers), the firm has a perverse financial interest in their death. Clearly, this is undesirable. To eliminate this interest, a firm whose employee dies (independent of cause) should pay the government a levy that exceeds the cost of implicit subsidies to the lowest-ability worker.

6. Bankruptcy. To ensure that primary employers honor their financial obligations towards their workers, they would need regulations similar in spirit to current banking regulations. In particular, these regulations should prevent primary employers from practices such as transferring their higher-ability workers to other companies, siphoning off the associated transfer fees, and then declaring bankruptcy because they are not able to fulfill their obligations towards lower-ability employees. For example, these regulations could include a requirement that primary employers post bonds that are forfeited in case of bankruptcy.

7. Bargaining power. If employers do not have all the bargaining power, workers will capture some fraction of the surplus created by working rather than exercising an outside option. Since this surplus is greater for higher-ability workers than for lower-ability one, this shifts the salary distribution in favor of higher-ability workers. The government could correct this shift by

increasing τ . Average salaries, however, are not affected since the zero-profit condition will cause the basic salary to drop by exactly the average increase in surplus going to workers. Firms have incentives to increase their bargaining power because this would earn them supernormal profits.

8. Multiple periods. In a multiple period setting, a ratchet effect may arise: if a high-ability worker reveals his type in the first period, the firm may expropriate his informational rents in the following periods. The analysis of the ratchet effect is complicated (for an overview, see e.g. chapter 9 of Laffont and Tirole, 1993). Whether the ratchet effect would be important in practice is unclear. Ratchet effects are also present in multi-period models of optimal income taxation, which has not hampered the operation of income tax systems. In addition, firms that are best able to overcome the ratchet effect, will earn higher profits and would be more likely to prevail.

9. Political economy considerations. Political economy reasons against privatized redistribution may include the inherent uncertainties associated with a radical change and dangers of rent-seeking activities (e.g. if firms lobby to reduce the basic salary below the market-clearing level). People might suffer a utility loss from feeling indentured to their employer (despite the outside options). On the other hand, privatized redistribution may protect citizens against arbitrary redistribution by a government sensitive to special interests.

5. Conclusions

In this paper, I present a mechanism that, in theory, privatizes income redistribution. There are three main reasons why privatized redistribution may be more efficient than redistribution through the government's tax and transfer system. First, firms are likely to have more information than the government about their workers, which limits the need to distort labor supply to extract

information about workers' abilities. Second, firms can devise more flexible incentive schedules than the one implicit in the government's tax and transfer system, because firms can use more instruments, such as the type of work performed, and because firms' decisions about which incentive schedules to offer a particular worker need not be codified but may rely on discretion. Finally, and perhaps most importantly, under privatized redistribution, firms have incentives to develop and invest in techniques that allow them to redistribute more efficiently. However, many issues would arise in the implementation of privatized redistribution. Some of these practical issues may entail costs that reduce the benefits of privatized redistribution.

While the mechanism for privatizing redistribution presented in the paper is stylized, I argue that the ideas underlying them may have practical applications. In particular, they may stimulate new approaches to deal with the persistent unemployment problems in Europe or to raise the wages of the working poor in the U.S. without creating unemployment. Even if the practical feasibility of privatized income redistribution is uncertain, this paper makes the case that income redistribution need not be the exclusive domain of the government and may be provided more efficiently by markets.

References

- Atkinson, Anthony B. and Joseph E. Stiglitz (1980), *Lectures on Public Economics*, London: McGraw-Hill.
- Ballard, Charles, John Shoven and John Whalley (1985), "General Equilibrium Computations of the Marginal Welfare Cost of Taxes in the United States," *American Economic Review*, 71(1), pp. 128-138.
- Cooper, Suzanne J. (1998), "A Positive Theory of Income Redistribution," *Journal of Economic Growth*, 3(2), pp. 171-195.
- Diamond, Peter (1992), "Organizing the Health Insurance Market," *Econometrica*, 60(6), pp.1233-1254.
- Donahue, John D. (1989), *The Privatization Decision*, New York: Basic Books.
- Feldstein, Martin S. (1999), "Tax Avoidance and the Deadweight Loss of the Income Tax," *Review of Economics and Statistics*, 81 (4), pp. 674-680.
- Grossman, Sanford J. and Oliver D. Hart (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration", *Journal of Political Economy*, 94(4), pp. 691-719.
- Hart, Oliver, Andrei Shleifer and Robert W. Vishny (1997), "The Proper Scope of Government: Theory and an Application to Prisons," *Quarterly Journal of Economics*, 112 (4), pp. 1127-1161.
- Kremer, Michael (1997), "Pareto Improving Tax Reform," mimeographed, MIT.
- Laffont, Jean Jacques and Jean Tirole (1993), *A Theory of Incentives in Regulation Procurement*, Cambridge, MA: MIT Press
- Mirrlees, James A. (1971), "An Exploration in the Theory of Optimal Income Taxation", *Review of Economic Studies*, 38, pp. 175-208.
- Piketty, Thomas (1993), "Implementation of First-Best Allocations via Generalized Tax Schedules", *Journal of Economic Theory*, 61, pp. 23-41.
- Rothschild, Michael and Joseph E. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90(4), pp. 630-649.
- Tirole, Jean (1988), *The Theory of Industrial Organization*, Cambridge, MA: MIT Press.
- Vickers, John and George Yarrow (1988), *Privatization: An Economic Analysis*, Cambridge, MA: MIT Press.
- Williamson, Oliver E. (1975), *Markets and Hierarchies, Analysis and Antitrust Implications: a Study in the Economics of Internal Organization*, New York : Free Press.
- Wilson, James Q. (1989), *Bureaucracy: What Government Agencies Do and Why They Do It*, New York: Basic Books.