functioning of financial markets, or they may not take the form of credit rationing. A heated debate (both theoretical and empirical) has been raging in the macroeconomics community about the relevance of credit rationing.

One implication of credit rationing for monetary intervention aimed at smoothing out business fluctuations is that a policy that mainly varies the supply of investment funds may not have the desired impact on investment. For example, it may not help much to increase the supply of funds in a recession if a large fraction of firms is credit rationed. A policy aimed at subsidizing investment more directly may be more effective. Another implication of credit rationing is that credit markets may amplify rather than dampen business fluctuations. An unexpectedly large drop in aggregate consumption may be only partially offset by an increase in investment, for example, if credit rationing is prevalent. Finally, credit rationing may have important implications for economic and financial development. Informational asymmetries are if anything worse in less developed economies. These economies may therefore have a handicap in catching up with the more developed and informationally transparent economies.[7]

2.2.2 Optimal Income Taxation

A central theme of this chapter is the trade-off between allocative efficiency and distribution in the presence of adverse selection. This trade-off has first been highlighted by Mirrlees (1971) in the context of redistributive taxation, where distributional concerns are weighed against incentive efficiency when lump-sum taxation is not feasible. One central reason why lump-sum taxation may not be feasible is that the size of the lump-sum tax is limited by the lowest income. To be able to raise higher tax revenues, the tax must inevitably be based on unobservable variables like an individual's earnings potential or productivity.

7. See Banerjee (2003) for a survey of these issues.

More formally, consider the following setup where "income" q is produced with "effort" e according to the production function: $q = \theta e$. Here, θ is an individual productivity parameter that can take two values, θ_L and θ_H, with $\theta_L < \theta_H$. Suppose that a proportion β of individuals has low productivity θ_L and a proportion $(1 - \beta)$ high productivity θ_H. All individuals have the same utility function given by

$$u[q - t - \psi(e)]$$

where t is the net tax (subsidy) the individual has to pay (receive) from the government and $\psi(e)$ is an increasing and convex cost function. The government's budget constraint is given by

$$0 \le \beta t_L + (1-\beta) t_H \tag{2.8}$$

where t_L is the (possibly negative) tax levied on low-productivity individuals and t_H the tax on high-productivity individuals. In the absence of adverse selection, a utilitarian government, maximizing the sum of individual utilities, solves the following problem:[8]

$$\underset{\substack{t_H, t_L \\ e_H, e_L}}{\text{Max}} \{\beta u[\theta_L e_L - t_L - \psi(e_L)] + (1-\beta) u[\theta_H e_H - t_H - \psi(e_H)]\} \tag{2.9}$$

subject to condition (2.8). Note that we assume here that the government can impose its tax scheme on individuals, so that it does not face any individual rationality constraints. At the optimum, condition (2.8) must be binding, and the first-order conditions yield

$$u'_L \equiv u'[\theta_L e_L - t_L - \psi(e_L)] = u'[\theta_H e_H - t_H - \psi(e_H)] \equiv u'_H$$

$$\psi'(e_L) = \theta_L$$

$$\psi'(e_H) = \theta_H$$

In words, the utilitarian optimum is attained when marginal utilities are equalized across individuals (that is, $u'_L = u'_H$). When all individuals have the same utility function, then they all reach identical utility levels if $u(\cdot)$ is concave. The other condition for an optimum is that the marginal cost of effort is equal to its marginal productivity for each type.

8. Note that when the government can observe θ_i it can also tell how much effort an individual has supplied when observing the individual output $q = \theta_i e$. Hence, by setting type-contingent output targets $q_i = \theta_i e_i$, the government can effectively control individual effort.

Under adverse selection, the following set of incentive constraints must be imposed on the government's optimization problem:

$$\theta_L e_L - t_L - \psi(e_L) \geq \theta_H e_H - t_H - \psi\left(\frac{\theta_H e_H}{\theta_L}\right) \tag{2.10}$$

$$\theta_H e_H - t_H - \psi(e_H) \geq \theta_L e_L - t_L - \psi\left(\frac{\theta_L e_L}{\theta_H}\right) \tag{2.11}$$

Indeed, when the government cannot identify each productivity type, it can only offer everybody an income-contingent tax, where t_i has to be paid if income $q_i = \theta_i e_i$ is produced. This would allow an individual with productivity θ_j to pay t_i by producing q_i at the cost of an effort level $\theta_i e_i / e_j$. As always, the incentive constraints ensure that this behavior is not attractive for either type of individual.

When $u(\cdot)$ is concave, the complete information optimum is such that

$$q_L - t_L - \psi(e_L) = q_H - t_H - \psi(e_H)$$

This allocation, however, violates incentive constraint (2.11): high-productivity individuals would then prefer to choose (q_L, t_L) instead of (q_H, t_H).[9] Therefore, condition (2.11) must be binding in a second-best optimum, and the same is of course true of the budget constraint (2.8). Using these two constraints to eliminate the tax levels from the maximand and taking the first-order conditions with respect to e_H and e_L then yields

$$\psi'(e_H) = \theta_H \tag{2.12}$$

and

$$\psi'(e_L) = \theta_L - (1-\beta)\gamma\left[\psi'(e_L) - \frac{\theta_L}{\theta_H}\psi'\left(\frac{\theta_L e_L}{\theta_H}\right)\right] \tag{2.13}$$

where $\gamma = (u_L' - u_H')/[\beta u_L' + (1-\beta)u_H']$.

That is, γ represents the difference between the marginal utilities of low- and high-ability individuals as a percentage of average marginal utilities.

Condition (2.12) tells us, as is by now familiar, that the second-best allocation for type θ_H is efficient (efficiency at the top). Condition (2.13) tells

9. Note that this conclusion is robust to changes in the utility function: it would remain true, for example, if effort appeared additively outside the utility function {e.g., if we assume a payoff $u[q-t] - \psi(e)$, with $\psi(\cdot)$ being linear or convex in effort}.

us that type θ_L underprovides effort. Indeed, first-best efficiency requires that $\psi'(e_L) = \theta_L$. But since $\theta_L < \theta_H$, we have $0 < \psi'(e_L)(1-\theta_L/\theta_H) < \psi'(e_L) - \theta_L/\theta_H\, \psi'(\theta_L e_L/\theta_H)$. Moreover, for any strictly increasing, concave utility function $u(\cdot)$, γ is strictly positive, since then $u_L' > u_H' > 0$. Therefore, the second term of the right-hand side (RHS) of condition (2.13) is strictly positive, resulting in underprovision of effort. The reason why it is second-best efficient to underprovide effort here is that a lower e_L limits the welfare difference between high- and low-productivity individuals, which is given by

$$[\theta_H e_H - t_H - \psi(e_H)] - [\theta_L e_L - t_L - \psi(e_L)] = \psi(e_L) - \psi\left(\frac{\theta_L e_L}{\theta_H}\right) \tag{2.14}$$

when the incentive constraint (2.11) is binding. This brings about a first-order gain in the utilitarian welfare function that exceeds the second-order loss from a reduction in productivity of the low type. A reduction in e_L brings about a reduction in inequality of welfare because high-productivity individuals have the option to produce output q_L while saving a proportion $(\theta_H - \theta_L)/\theta_H$ of the effort e_L that low-productivity individuals have to exert.

How can we reinterpret conditions (2.12), (2.13), and (2.14) in terms of existing features of the income tax code? One can think of governments first setting *marginal tax rates* as well as *uniform tax rebates*, while individuals respond by choosing effort. In this perspective,

- Equation (2.12) implies that the marginal tax rate at output $q_H = \theta_H e_H$ should be *zero*, in order to induce efficient effort.
- By contrast, the marginal tax rate at output $q_L = \theta_L e_L$ should be *positive, and equal to*

$$\left\{(1-\beta)\gamma\left[\psi'(e_L) - \frac{\theta_L}{\theta_H}\psi'\left(\frac{\theta_L e_L}{\theta_H}\right)\right]\right\}\Big/\theta_L \tag{2.15}$$

where γ and e_L are the second-best values; indeed, this marginal tax rate is needed to induce the (inefficiently low) choice of e_L given by equation (2.13).
- In between these two output levels, the marginal tax rate should be positive and such that the additional tax payment coming from high-productivity individuals will be sufficient to reduce the ex post utility difference between the two types to that given by condition (2.11) [recall that,

from budget constraint (2.8), all marginal tax revenues are returned to individuals, in the form of uniform tax rebates].

This framework thus seems to plead for positive marginal tax rates at low income levels, and for zero marginal tax rates at high income levels. This result appears paradoxical, since we started from a government objective that was tilted toward redistribution. In fact, there is of course positive redistribution here, through the uniform tax rebate. And the positive marginal tax rate at low incomes is there to achieve higher redistribution, as explained earlier. The theory in fact offers a rationalization for the widespread practice of making various welfare benefits conditional on having low incomes: this results in very high effective marginal tax rates at low income levels. As for the zero-marginal-tax-rate-at-the-top result, it should not be overemphasized: If we allowed for a continuum of types (as in section 2.3), this zero marginal tax rate would appear only at the very maximum of the income distribution.

We close this subsection by analyzing the effect of changes in the government's information structure in this setup. Following Maskin and Riley (1985), we could indeed ask the following question: how would tax levels be affected if, instead of observing "income" $q = \theta \cdot e$, the government observed individual "input," that is, effort (or, equivalently, hours of work)? Would this new information structure result in more or less inequality, and more or less allocative distortion?

In comparison with the preceding analysis, only the incentive constraints (2.10) and (2.11) would be affected. They would become

$$\theta_L e_L - t_L - \psi(e_L) \geq \theta_L e_H - t_H - \psi(e_H) \tag{2.16}$$

$$\theta_H e_H - t_H - \psi(e_H) \geq \theta_H e_L - t_L - \psi(e_L) \tag{2.17}$$

since individuals now face effort-contingent taxes, instead of output-contingent taxes. As before, the problem is to prevent high-ability individuals from mimicking the choice of low-ability individuals. Consequently, condition (2.17) must be binding at the optimum, and the inequality in welfare between the two types is given by

$$[\theta_H e_H - t_H - \psi(e_H)] - [\theta_L e_L - t_L - \psi(e_L)] = (\theta_H - \theta_L) e_L \tag{2.18}$$

Thus, for a given e_L, welfare inequality is higher in equation (2.18) than in equation (2.14) whenever $\theta_H > \psi'(e_L)$. This is always the case at the optimum, since $\theta_H > \theta_L > \psi'(e_L^*)$. Intuitively, an effort-contingent tax makes it more attractive for the high type to mimic the behavior of the low type, since a given effort level allows the high type to enjoy comparatively more output than the low type, thanks to his higher productivity. Instead, in an income-contingent scheme, mimicking the low type really means sticking to a low income. It is true that this allows the high type to save on effort cost, but with a convex effort cost, the benefit of this lower effort is limited.

Solving the optimum tax problem with equation (2.17) instead of equation (2.11), one can show that efficient effort results again for the high type. Instead, the effort of the low type is further distorted downward, in comparison with the income tax case: when welfare inequality is higher, the marginal benefit from reducing inequality at a given allocative cost is also higher. In this setting, effort monitoring is thus inferior to income monitoring, since it leads both to more inequality and to more allocative inefficiency.

This last discussion is especially relevant when the principal is able to make decisions about which information structure to put in place. It would apply naturally, for example, in the context of the internal organization of a firm, where decisions have to be made about which information systems to set up for monitoring purposes. On this subject, see the general analysis of Maskin and Riley (1985) on the superiority of output- over input-monitoring schemes.