

Galiani, Sebastian

Working Paper

Social Mobility: What is it and why does it matter?

Documento de Trabajo, No. 101

Provided in Cooperation with:

Centro de Estudios Distributivos, Laborales y Sociales (CEDLAS), Universidad Nacional de La Plata

Suggested Citation: Galiani, Sebastian (2010) : Social Mobility: What is it and why does it matter?, Documento de Trabajo, No. 101, Universidad Nacional de La Plata, Centro de Estudios Distributivos, Laborales y Sociales (CEDLAS), La Plata

This Version is available at:

<https://hdl.handle.net/10419/127621>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

C | E | D | L | A | S

Centro de Estudios
Distributivos, Laborales y Sociales

Maestría en Economía
Universidad Nacional de La Plata



Sebastian Galiani

Social Mobility: What is it and why does it matter?

Documento de Trabajo Nro. 101
Junio, 2010

ISSN 1853-0168

Social Mobility: What is it and why does it matter?¹

Sebastian Galiani
Washington University in St. Louis

**April 2007
Draft Version**

¹ Document elaborated for the 7th Social Equity Forum, IADB. I thank the excellent research assistance of Maria Eugenia Garibotti.

INTRODUCTION

The definition of social mobility is the object of some discussion, and although there is a common thread that runs through all of these discussions, the actual definition varies from study to study. There is agreement that social mobility refers to “*movements* by specific entities between periods in *socioeconomic status indicators*” (Behrman, 2000) and that it aims to quantify “the movement of given [entities] *through the distribution of economic well-being over time*, establishing how *dependent* one’s current economic position is on one’s past position, and relating people’s mobility experiences” to the overall conditions of the economy in which they operate (Fields, 2000). Differences arise, however, when an attempt is made to endow these definitions with empirical content (i.e., when an effort is made to determine what variable should be used to measure mobility, what exactly should be considered “movement” in a distribution, or what time spans should be used to evaluate mobility). In the following discussion, we briefly comment on some of the conceptual issues that have been raised in the literature on mobility.

Among the multiple considerations concerning the definition of mobility, in this paper we define social mobility as a situation in which the relative economic status of an agent is not dependent on starting conditions such as parental income or family background. Therefore, analyzing the determinants of mobility involves exploring the channels through which offspring’s income is correlated to its parents’, such as inherited bequest, education, formal rules, skills, opportunities, working spirit, among many others.

As parental linkage is a source of differences in income among individuals, there is a deep relation between social mobility and inequality. They are jointly determined, and the most prevalent theoretical association between mobility and inequality is negative; since structural conditions that lead to low mobility also tend to favor unequal outcomes.

Assessing inequality, leads us to investigate analytical frameworks to analyze the sources of differences in incomes: from effort, education and ability to the beliefs about the nature of the income generating process. As long as these determinants are at least partially related to the parental background, we will find a strong link between social mobility and: equal opportunities, meritocracy, human capital accumulation, politico-economic considerations and beliefs, which we will account for in this survey.

Furthermore, when initial conditions prevent individuals from allocating resources optimally, if social mobility is limited, economic efficiency is reduced. For example, this happens whenever liquidity constraints prevent individuals from acquiring an efficient level of human capital. Therefore, it is of great relevance to understand that social mobility and efficiency are also determined together.

Finally, throughout this survey, we map the key parameters affecting mobility into proposed policy actions, from each model studied. We find that policies that break the dependence of outcomes on initial conditions, such as universal public education or early intervention programs may be successful in raising income and social mobility. Furthermore, due to the relationship between mobility and other social and economic dimension such as efficiency and inequality; those policies can have broader short and long term effects.

This survey is divided into 8 sections. In the first one, some preliminary concepts related to the definition of mobility are briefly explained. Section 2 comprises the review of a range of models of income mobility that relates it to human capital, credit constraints, education and technology, among others. The third section deals with the politico-economic considerations regarding inequality and mobility. In Section 4, it is established the relationship between social mobility and both equal opportunity and meritocracy. The next section introduces a key concept: efficiency, in the context of the models that assess social mobility. In section 6, there is a discussion about the effects of social mobility from an economic efficiency and also from political outcomes points of view. Section 7 includes models that consider the perception of mobility. To conclude, in section 8, there is a summary of the policy implications of the surveyed models, showing the different policies that have the potential to improve social mobility in the short and long terms.

1. PRELIMINARY CONCEPTS

Socioeconomic Status Indicators

When studying mobility, we are interested in using an indicator that can capture some element of economic well-being. This entails taking variables into consideration which measure long-term status rather than short-term fluctuations. (We do not, after all, want to confuse social mobility with economic insecurity.) When trying to encapsulate this general idea of concrete, operational content, we may want to consider consumption (which is, presumably, closely linked to *permanent* income), educational attainment, asset holdings (wealth), or some composite measure of socioeconomic level.

Data availability considerations, unfortunately, quickly limit the scope of the indicators used in actual studies, as mobility research calls for long panels, or, at the least, information on parents and offspring. It is already quite hard to find high-quality datasets that provide this kind of information on income and labor earnings, which are some of the most commonly measured socioeconomic indicators. Reliable long-term information on consumption, or socioeconomic status, is practically nonexistent, especially in developing countries. Data on intergenerational educational attainment, on the other hand, is easier to find, partly because retrospective questions (questions on parental educational attainment, for instance) are bound to be reliably answered. As a result, most of the studies now available focus on income and educational attainment.

Time Period

There are two main types of mobility that we may want to study, depending on the length of time we allow for changes to take place. *Intragenerational* mobility refers to movements in the indicator of choice that occur in a relatively short time span – typically within an individual’s adulthood. *Intergenerational* mobility, on the other hand, refers to changes that are observed from one generation to the next. Thus, *intragenerational* mobility studies follow individuals throughout their lives, while *intergenerational* studies focus on entire dynasties, tracking their movements from one generation to the next.

Movements

The main source of controversy regarding social mobility is arguably the issue of how to define “movement” within a distribution. Different concerns give rise to very different ideas about how to quantify changes in an individual’s economic status. A very preliminary observation is that no meaningful “movement” can ever occur unless there is variability in the values of the indicator to be studied. In other words, income mobility would be an empty concept if income inequality did not matter.

Fields (2000) provides a useful categorization of different mobility concepts, as well as stating some of the implicit assumptions and value judgments underlying each one. He shows that the kind of normative (or subjective) choice that is made in studies of social mobility goes beyond simply deciding whether it is a “good” or “bad” feature. Following

Fields (2000), imagine the following societies numbered from 1 to 5, formed by 3 individuals (these are somewhat extreme examples, but serve to point out the sources of disagreement):

Table 1: Hypothetical Changes in Income Distributions

	1	2	3	4	5
Andrea	20	5	20	40	5
Brian	10	10	30	20	1
Chris	30	10	20	60	15

Any relative measure of mobility would indicate that, in going from society 1 to society 2, there has been perfect mobility. The correlation between the two, indeed, is zero. But note that average wealth has gone down and, as a result, so have income levels for all but one of the individuals. The transition from society 2 to society 3, where the correlation is 0.5 – mobility is much lower – but all incomes have gone up, is the mirror image of the first. In this sense, it could be argued that the movement from 2 to 3 is more desirable than from 1 to 2.

The main point here is that relative measures of mobility do not take changes in the average income level into account. While this is seen as a flaw by researchers who would say that an environment where all incomes are growing is “mobile” or “dynamic”, even if social ranks are not moving, we believe that mobility concerns arise mainly because of relative considerations. In other words, societies tend to favor situations where individuals’ relative status is considerably independent from their parents’ degree of success, and this is best captured by measures of relative mobility.

One distinction should be noted with regard to measures of relative mobility. *Positional movement* takes place when an individual changes her position in the overall distribution. Related to this is *share movement*, which focuses on the change in an individual’s share of the overall total.² Both are instances of what Behrman (2000) calls “relative” or “exchange” mobility.

A process whereby all members double their incomes would be perfectly immobile under both measures (this would correspond to going from society 1 to society 4). Also, a process in which total income decreases and inequality increases while ranks remain the same would show no positional mobility, but shares would have changed (this is the case in going from distribution 4 to 5). Rankings can also change drastically with very little change in income, as is the case in the shift from society 2 to society 5.³

² Although we are trying to avoid referring to “income” in order to drive home the point that we are interested in a broader notion of “social” mobility, it is clear that there must be some quantitative element in order to talk meaningfully about “shares”.

³ This kind of objection is valid to some extent, and we mention it because some of the researchers in this area actually do raise such objections. But the fact is that this kind of example does not seem to be in line with what we see – or would expect to see – in real life settings. And relative movements do have the advantage that they capture the persistence of inequality in a way that absolute measures do not precisely because they abstract away from changes in mean income, or even in its distribution. The main point to be

The aim is to isolate the distributional component from mobility. We believe that mobility concerns are entirely related to ranks; thus, there may be social mobility even if the overall distribution becomes more unequal. In any case, the lesson to be learned from these examples is that a complete picture of the desirability of an income process cannot be derived from any one summary measure because changes in average income levels, changes in the overall distribution, and the degree to which relative performance in one generation depends on relative performance in the past must all be taken into account.

Absolute movement measures, in contrast, focus on quantifying the total change, regardless of whether any exchange mobility has taken place. We will not go into detailed definitions of these instruments here, however, since they do not seem to be relevant to the type of concern that drives this study.

Although we may be content with *quantifying* movement – whatever definition we decide upon – it is usually the case that we also want to capture other aspects of movement. In particular, studies of social mobility are generally concerned with gauging how dependent final outcomes are on initial conditions. In other words, income mobility may exist because individual incomes follow an identical random process, independent of initial income, or there may be some sort of dependence – usually captured by a correlation coefficient – between income realizations for the same individual (or dynasty) across time. Most studies of mobility equate time dependence with immobility and, generally speaking, this is the approach that is favored in this paper.

We believe that the concept of mobility is quite distinct from that of economic growth and is closely connected to notions of efficiency, fairness, and political conflict which follow paths that are largely unconnected to overall growth. This is why we choose to focus on measures of relative, rather than absolute, mobility. As will be made clear in this paper, the time span to be chosen will depend on the context of the analysis, while the issue of time dependence will invariably play an important role: *increasing social mobility implies breaking the dependence of individual economic outcomes on initial conditions.*

Normative Analysis of the Correlation Coefficient

A low correlation coefficient indicates that mobility is high and that parental income is weakly correlated with individual income, whereas a high correlation coefficient signals low mobility. However, authors such as Swift (2005) and Feldman et al. (2000) claim that a zero correlation is not a morally desirable goal because any serious attempt to disconnect the life-income of parents and their offspring would severely compromise important values of family life and privacy. Thus, they argue, instead of pursuing the objective of zero intergenerational correlation, we should ask ourselves which mechanisms of intergenerational transmission are unfair and should then design our

made here is that, just as reporting on inequality does not seem complete without studying mobility, measuring mobility will not give us a perfect picture of how “desirable” a given society is. Although this is undeniably the case, it is also true of any other one-dimensional measure.

policies accordingly. Following the same line of reasoning, Jencks and Tach (2005) claim that if we want to know whether opportunity is becoming more or less equal, we need to track intergenerational linkages that violate certain norms associated with meritocracies and developmental opportunity. Under their definition of developmental opportunity, society must either make families and communities more alike or find ways to offset the adverse effect of growing up in a disadvantaged family or community.

Implicit in the reasoning of these authors is the idea that some sources of advantages are acceptable or even desirable while others not. For instance, many people agree that ethnicity should not be a source of advantage, but that effort should. Family contacts, intelligence, or other genetic or environmental traits are more debatable. Therefore, if acceptable sources of advantages are somehow inherited by children of advantaged parents, we would observe a positive correlation between parental and individual incomes even if no unacceptable source of advantages remained.

However, even if all sources of unacceptable advantages were eliminated and the correlation between parental and individual income were fully explained by acceptable advantages, we could still argue that policy interventions to help the disadvantaged would reduce the correlation. Normative evaluation of such policies would ultimately depend on the social preferences of the evaluator and his preference for mobility *per se*.

Relation with Other Distributional Features

Concerns about static aspects of a society's welfare distribution – such as inequality, poverty or polarization – are, as a rule, intimately connected with any discussion on social mobility. Most theoretical studies focus on the joint determination of mobility and inequality, since concepts such as poverty and polarization are hard to incorporate into a model. We will therefore examine the relationships between inequality and mobility.

The most common characterization points to inequality as being “a snapshot” of the income-generating process at any point in time, while mobility is seen as tracking income distribution over time based on individual trajectories. When considered from this standpoint, the absence of mobility accentuates all issues associated with inequality. It also underscores the fact that a comparison between the prevailing income distributions in two societies – two “snapshots” – may reverse when mobility is incorporated into the analysis.

This reversal would only happen if higher mobility were somehow associated with higher inequality and vice versa, so that superficially “equal” societies might actually have more persistent inequality than highly unequal societies. As we will see, this is far from obvious, and theoretical models as well as empirical evidence provide mixed results. In fact, the most prevalent theoretical association between mobility and inequality is negative, since structural conditions that lead to low mobility also tend to favor unequal outcomes.

2. MODELS OF INCOME MOBILITY

Becker and Tomes (1979) were one of the first to model income mobility and inequality. In their economy, there is a continuum of dynasties. Each generation values its own consumption as well as its offspring's income:

$$U_t = U_t(Z_t, I_{t+1}) \quad (1)$$

where U_t is parents' utility, Z_t stands for parents' consumption, and I_{t+1} is the adult wealth of their children.

This wealth is embodied in a stock of human and nonhuman capital that has three sources: a direct investment from each generation's parents (y); endowed "luck" that is correlated across contiguous generations of the same dynasty (e) (which can be interpreted as ability, intelligence, social background or values, and family connections); and sheer market luck (u).

$$I_{t+1} = w_{t+1}y_{t+1} + w_{t+1}e_{t+1} + w_{t+1}u_{t+1} \quad (2)$$

In other words, we think of individuals in this model as helping their children out in two ways. They can decide to leave them "capital" by investing in their education or providing them with financial capital. This is a deliberate decision on the part of the parents. In addition, they pass on family characteristics that were assigned by "luck" to each dynasty. Parents cannot alter such characteristics or prevent them from being transmitted. A natural interpretation would involve genetic characteristics, but there are also family connections, "norms" or "customs" within each household, as well as other non-biological traits, that are largely the result of each family's history and that cannot be radically changed by any individual dynasty. A poor parent cannot build family connections out of thin air.

Of course, there is a final element of "luck" in determining individual income; people will typically differ in the extent to which they take full advantage of their endowment, and their resourcefulness or drive may lead them to find new opportunities. Serendipity may also intervene. This is the way in which the term u should be interpreted.

As can be seen from equation (2), each part of this capital stock has the same (constant) market return, w .⁴ Parents are assumed to be risk-neutral and to choose their bequests in order to maximize their own expected utility, subject to a budget constraint (where r is the intergenerational rate of return),

$$Z_t + \frac{w_{t+1}y_{t+1}}{1+r_t} = I_t \quad (3)$$

⁴ Thus, both inequality and mobility refer to the distribution of capital stock or of all sources of income.

and a law of motion for endowed luck that depends on the family endowment and the average level of endowed luck in the economy.

$$e_{t+1} = (1 - h + f)\bar{e}_t + he_t + v_{t+1} \quad (4)$$

In this law of motion, h represents the fraction of family endowment that is inherited, and f is the aggregate rate of growth, while v is a random shock. We thus allow for slow changes in family “luck”. It may be the case that a rich family loses part of its connections, or that someone from a family that does not value “hard work” may still turn out to be hardworking and pass this trait on to his children. In any case, as we will see, h is a very important parameter of the model, as it measures the extent to which family “luck” is equalized across dynasties. When h is close to 0, family endowments are basically identical, with some random variation. When it is close to 1, dynasties differ permanently.

Families fully anticipate inherited endowment (presumably, any relevant “random” shocks to ability, personality, or the like can be observed before making the optimal investments), but not random luck. From standard expected utility maximization, and normalizing r and w to equal 1, the income of children in the $t+1$ generation of the i th family can be expressed as:

$$I_{t+1}^i = \alpha a + \beta I_t^i + \alpha h e_t^i + \alpha v_{t+1}^i + \alpha u_{t+1}^i \quad (5)$$

Using Table 2, equation (5) says that income in a given generation basically depends on parental wealth and endowments, although an additional effect is generated through endowed and market luck. Crucially, parental preferences determine how strong this dependence is. This model captures the fact that anticipated luck affects how much capital parents leave to their offspring, which is why both market and endowment luck are not fully appropriated by their children (they only get to enjoy a proportion α).

Table 2: Key Variables in Becker and Tomes (1979)

Variable	Interpretation/Definition	Depends on
α	Proportion of parents' wealth that goes to children's income	- Parents' utility function – specifically, how much children's income affects parental welfare - Intergenerational rate of return.
β	Actual value of the propensity to invest in children	$\beta = \alpha(1 + r)$
a	Baseline endowed ability	$\bar{e}(1 - h)$

All families are assumed to be identical. This means income inequality depends on the past history of market luck. Assuming that parameters are such that the persistence of initial conditions is limited ($\beta, h < 1$), income in generation $t+1$ is:

$$I_{t+1}^i = \frac{\alpha a}{(1-\beta)(1-h)} + \alpha \sum_{k=0}^{\infty} \beta^k u_{t+1-k}^i + \alpha \sum_{k=0}^{\infty} \left(\frac{\beta^{k+1} - h^{k+1}}{\beta - h} \right) v_{t+1-k}^i \quad (6)$$

and the (squared) coefficient of the variation in income – a measure of inequality⁵ – is:

$$CV_I^2 = \frac{1-\beta}{1+\beta} CV_u^2 + \frac{(1+h\beta)(1-\beta)}{(1-h\beta)(1+\beta)} CV_e^2 \quad (7)$$

That is, inequality has two components: one that comes from market luck (u), and one that comes from endowed advantages (e).⁶ Clearly, if any of these components were less variable, income inequality would go down. The additional insight derived from this model comes through the effects of β and h .

Note that the coefficient for endowed family luck is larger than the coefficient for market luck. The difference between them grows as β and h become larger. That is, as we increase the *actual value of the propensity of each generation to invest in the next* and the degree of *family inheritability* of characteristics that affect income, we give *greater weight in overall inequality to family-specific advantages*.

However, whereas an increase in h unambiguously increases inequality, an increase in β – say, as the result of a change in preferences or through an increase in market returns – **reduces the coefficient of variation in income**. This is the result of the fact that, while it increases the variance in income, it increases the level of income even more.

We next turn to income mobility and then analyze how β and h affect the trade-off between inequality and mobility. The preferred measure of income mobility in this article is the effect of an increase in income for generation t on the incomes of subsequent generations. This gives us an idea of how quickly “temporary” increases in social standing fade away. It measures how long it takes to go “from rags to riches and back”. In other words, it captures our preferred concept of mobility as the degree of dependence of current outcomes on past performance. Lower persistence is equivalent to higher mobility.

⁵ The coefficient of variation of a random variable is defined as the ratio of its standard deviation to its mean. That is, $CV(x) = \sigma_x / \mu_x$. In this case, v has zero mean, so that its standard deviation is normalized by average endowed luck, or \bar{e} .

⁶ The last term actually reflects exogenous variation in endowed luck, since $\sigma_e^2 = \frac{\sigma_v^2}{(1-h^2)}$.

Table 3: Effects of One-Time Changes in Income

Source of Change	δI_{t+m}^i	Behavioral and Parametric Assumptions	As m Goes to Infinity
Market or endowed luck	$\beta^m \delta I_t^i$	h negligible $\beta < 1$	Declines monotonically
Endowment	$h^m \delta I_t^i$	No optimization – investments are given	Declines monotonically
Endowment	$\bar{I} h(1-\beta) \frac{\beta^m - h^m}{\beta - h} \frac{\delta e_t^i}{\bar{e}}$	Optimal behavior $\beta + h > 1$	Rises up to a peak in a generation that depends on $\beta + h$, then declines monotonically
		Optimal behavior $\beta + h < 1$	Declines monotonically

This effect will depend on the source of the increase in income and, as the authors show, its magnitude hinges on the interplay between optimizing behavior and inheritability. Indeed, if we call δI_t a change in income for generation t , and we trace changes in the m th generation, we can envision the different scenarios described in Table 3.

When the degree of inheritability of endowed luck is close to zero, a look at equation 5 tells us that each generation will receive a fraction β of any change in the income of the previous generation, regardless of the source of this income. This change will decrease as the distance between generations increases, given our assumption that $\beta < 1$. In other words, when families do not provide sizeable (dis)advantages to their offspring, mobility is high.

If inheritability was high, but investments were not influenced by changes in endowments – that is, if families did not optimize – we would see a similar response to a change in endowment. However, if we take into account the fact that families do optimize, we can see that, for some parameter values, a change in family endowment at time t can compound over time before it fades away (implying that a “lucky” generation will make its descendants luckier than average for a relatively long time). This is exactly the result that we would associate with a **lack of social mobility**.

We can now see how income inequality and social mobility may be negatively or positively correlated, depending on what is driving their movements. When *intergenerational inheritability of advantage* increases, income inequality goes up and social mobility goes down. Thus, differences between families are exacerbated *and* made more persistent.

This is intuitive: if family connections or race become more important in determining individual income, we would expect differences between families to become more pronounced and permanent.

When the *propensity to invest in children* goes up, mobility again is reduced, because a one-time shock to income takes more time to fade away. However, in the long run, income differences are reduced, so income inequality goes down.

We can summarize these findings in the following table.

Table 4: Relationship between Inequality and Mobility when Parameter Values Change

Parameter	Interpretation	Effect of an Increase in Parameter	
		on Inequality	on Mobility
β	Parental propensity to invest in children	Reduces Inequality	Reduces Mobility
h	Inheritability of family advantage	Increases Inequality	Reduces Mobility

Clearly, β is not a parameter that can be manipulated by public policy. It is quite impractical and, arguably, even unethical for the state to try to convince parents to care less (or more) for their children’s well-being. However, it is not implausible for alterations in h to fall within the scope of public policy. Note that a decrease in h is equivalent to a homogenization of endowments. That is, a decline in h makes endowments more likely to equal the average. Hence, a policy that, for example, calls for investments to be made in raising the quality of public education while making it more homogeneous⁷ would probably result in a decrease of h . Note that this reduction would also tend to reduce inequality.

Another way of changing h is by promoting meritocratic employment policies, possibly through encouragement of market competition among potential employers. Why would market competition help attain a meritocracy? Advantages in the labor market that are not related to higher productivity – such as race or family connections – can only persist if companies that engage in these hiring practices are shielded from their consequences. Firms in a competitive environment, in contrast, will come under pressure to adopt practices that favor productivity over personal loyalties or racial prejudice.

The main claim underlying these suggestions is that policies that **break the dependence of outcomes on initial conditions** are unambiguously desirable.

Becker and Tomes (1979) abstract from the determination of the return to human capital, equating it to physical capital and operating within a stationary economy. Although the basic insight – that mobility and inequality are affected by the way in which parents optimize when making their decisions, as well as by inheritance of personal qualities – is

⁷ And, potentially, more similar to education in private schools.

not in dispute, some unbundling of the process of human capital accumulation as an additional force driving inequality and mobility remains to be done. We will comment on two lines of work. One focuses on purely economic forces, while the second adds politico-economic considerations.

Building on Becker and Tomes: Human Capital and Credit Constraints

Hassler, Rodríguez Mora and Zeira (2003) and Owen and Weil (1998) show how inequality and mobility may be jointly determined in a general equilibrium model with overlapping generations of workers. In both cases, they draw a distinction between skilled and unskilled workers. They analyze mobility by determining the odds that the child of an uneducated worker will become skilled (upward mobility), or the odds that the child of a skilled worker will not receive an education (downward mobility).⁸

The key issue in both models is the lack of access to a credit market, since such access would allow “high-ability” individuals to invest in their education. This is a very practical concern, inasmuch as it has the consequence of making educational decisions (and, therefore, income) strongly dependent on parental background. In other words, it is a reinterpretation of the parameter h in the Becker and Tomes (1979) model.

Both models assume that **ability is not genetically determined**. This is why credit market failures end up accounting for most of inefficient immobility. The key point to bear in mind, though, is that this is a simplifying assumption which serves to highlight what aspects *other* than genetic endowments may affect intergenerational mobility. These effects would persist even if ability were, to some degree, genetically inherited.

Owen and Weil: Liquidity Constraints and Multiple Equilibria

Owen and Weil (1998) model the joint determination of aggregate output, income distribution, mobility, and returns to education in general equilibrium. Skilled and unskilled labors are complements in production, and changes in returns to skill stem from changes in the aggregate supply of each kind of labor.⁹ In other words, a large supply of unskilled workers increases the skill premium, and vice versa.

Agents differ across two dimensions: they receive different parental transfers, and they are born with different ability levels. Ability is independent across generations and is defined as the amount of efficiency units supplied to the labor market, regardless of skill level (which, instead, affects the wage level at which those units are rewarded). Thus, ability is not genetically determined and can be likened to “industriousness” (i.e., how hardworking a person is).

⁸ Note that these two probabilities may move in opposite directions: a policy that increases overall educational achievement may raise upward mobility and reduce downward mobility. Thus, we once again run into the problem of clearly defining what we want to include in the definition of mobility, or must ask ourselves whether bundling all these phenomena into one concept even makes sense.

⁹ Underlying this result, there is a neoclassical production function and perfect capital mobility, so the marginal productivity of capital is constant.

The timing is as follows. Individuals can be said to live through two stages. In the first, they receive transfers from their parents, then choose an educational level, and work. In the second, they consume and leave a bequest to their children.

An individual can acquire skills by buying education at a fixed cost \bar{e} . If we call q_t^i an individual's ability level and w_t^s the wage for skill level s , then we find that there is an ability threshold above which it pays to become educated.¹⁰ In an efficient outcome, individuals whose ability level exceeds this threshold will obtain an education.

$$q_t^* = \frac{\bar{e}}{w_t^e - w_t^u} \quad (8)$$

However, there are no credit markets available to finance educational decisions. This means that parental bequests (x) are the only source of funding. Thus, individual resources are given by:

$$\begin{aligned} x_{i,t} - \bar{e} + q_{i,t} w_t^e & \text{ if } x_t^i \geq \bar{e} \text{ and } q_{i,t} \geq q^* \\ x_{i,t} + q_{i,t} w_t^u & \text{ otherwise} \end{aligned} \quad (9)$$

In other words, agents will receive an education only if their bequests enable them to afford it.

Resources are split between individual consumption and bequests to children, with γ being the weight that is given by parents to bequests.

$$U(x_{t+1}, c_{t+1}) = \gamma \ln(x_{t+1}) + (1 - \gamma) \ln(c_{t+1}) \quad (10)$$

Clearly, in choosing a given bequest level, parents are determining the expected value of their children's education level. For instance, a bequest lower than \bar{e} ensures that children will be unskilled workers.

The optimal decisions of each family define, for each distribution of labor supply and each level of wages, what the transfers and resulting education levels will be for the next generation. The authors look for a *steady state* of this model – a relative wage level and skill distribution such that:

- Families expect wages to stay the same and therefore choose transfers in such a way that the education distribution remains unchanged; and
- Given the education distribution, this wage level is the outcome of market equilibrium.

¹⁰ Equation (8) also shows that, in equilibrium, skilled wages must be higher than unskilled wages.

In other words, the starting point for the situation is such that the economy will remain in the same state indefinitely.

Owen and Weil find that when individuals face liquidity constraints, there are multiple steady-state equilibria which exhibit a **positive** association between equality and mobility. That is, whenever there is **high** mobility, there is **low** inequality.

There can, of course, also be a situation in which there is no mobility at all. In such a situation, a handful of skilled workers have high wages that enable them to educate all of their children while wages for unskilled workers are so low that education is not affordable, even for the most industrious families. In fact, when the cost of education is high enough, this is the only kind of steady state there is.¹¹

When the cost of education is below a certain ceiling, there is at least one **high-income, high-mobility** equilibrium. In such a situation, the workforce is highly educated, which lowers the equilibrium skill wage premium and hence the chances that an unskilled worker will find herself limited by her borrowing constraint. This not only reduces inequality; it also makes it easier for a high-ability child of an unskilled worker to receive an education and lowers the incentive for a low-ability child of a skilled worker to acquire skills. In other words, it raises both upward and downward mobility.

In addition, in these equilibria the allocation of education is more sensitive to actual ability and less so to parental background than it is in a more unequal economy where borrowing constraints are disproportionately greater for the children of low-income parents. This is, in fact, why income is higher, as the most industrious workers are in high-productivity positions.

In low-income societies, these conclusions are reversed: the stock of skills is small, the wage differential is high (so is inequality), and mobility is very low. As a consequence, the allocation of resources is more inefficient.¹²

Thus, credit market imperfections make current skill levels dependent on past skill levels, as embodied in parental income. The inefficient steady states arise from the fact that this dependence is not based on a *productivity* difference (as would be the case if, for instance, ability were inheritable).

The key aspect of this model is that these different steady states coexist as possible outcomes **for the same economy**. In other words, two economies with the same parameters (access to the same technology, equal weight given to children in the utility function) may end up with widely different mobility and inequality levels. If it were possible to change the situation in one period by means of a single (large-scale) policy

¹¹ There can also be a continuum of this sort of immobile steady state. Starting from any one of these states, we can slightly increase the proportion of educated individuals and reduce the wage differential accordingly, and we will have another steady state with no mobility.

¹² Note that upward and downward mobility move together when going from one steady state to another. This is because we are changing the nature of the equilibrium, and high-mobility equilibria are more efficient in that they are more sensitive to the actual ability of each individual.

intervention, it would become self-sustainable. In contrast, the interventions required in economies where it is a question of changing the fundamentals, such as Becker and Tomes, are generally long-term policy measures.

What would a “one-time” intervention involve? One possibility would be a large-scale redistribution of income that could be accomplished either by giving funds directly to parents, by giving them “vouchers” for education, or by using income tax revenues to subsidize public education – anything that would temporarily break the dependence of education on background, thereby increasing the supply of educated workers and moving the system toward a low-inequality, high-income steady state.

Other policy recommendations come from “outside” the model. The focus on human capital acquisition allows Owen and Weil to consider different policy experiments, all of them related to changes in the education system. Given that inefficiencies result from the presence of liquidity constraints, a (permanent) program of education loans is the first option considered. These authors find that such a program reduces inequality and raises average income, effectively moving the economy to a high-education, low-inequality steady state. The second policy they analyze is a *meritocratic* public education system, where high-ability individuals get the education that maximizes net output. Again, this kind of policy effectively does away with the inefficiencies created by the borrowing constraint. Whether this kind of policy is materially feasible in low-income, low-mobility countries depends on functional forms adopted in the model.¹³

Why would the implementation of a meritocratic schooling system be an effective policy? The answer is that such a system would tend to wipe out the differences in schooling that arise solely by virtue of parental income: in other words, for the same reason that we argued that an increase in schooling quality would reduce h and therefore increase mobility. Note that this is not an abstract consideration. Countries such as China, South Korea and Japan have a strong public education sector which uses testing and assessments early on as a basis for assigning students to schools and channeling them into careers based on their cognitive ability. To a lesser extent, European countries such as Germany and France have a strong, high-quality public education sector as well. These countries also have greater social mobility than Latin America.

Table 5: Types of Policy Intervention in Owen and Weil’s Model

Policy	Type of Intervention	Effect within the Model
Massive Redistribution	One-Time	Keeping fundamentals such as education cost the same, movement toward a better steady state.
Education Loans	Permanent	Change in fundamentals: slackens liquidity constraint. Can potentially improve upon the best pre-intervention steady state.
Meritocratic Public Schooling	Permanent	Maximizes net output, breaking the link between parental income and education.

¹³ It is even more doubtful whether it would be politically feasible, since in a poor country such a measure might entail seizing the entire labor income of an adult generation in order to finance education for the young.

So far, the effects of differences in background are unambiguously inefficient. However, this changes when skilled parents transmit advantages because they are, in some sense, more efficient than their unskilled counterparts.

Hassler, Rodriguez Mora and Zeira: The Roles of Education and Technology

Moving on, Hassler, Rodriguez Mora and Zeira (2003) present a similar setup in which workers may be skilled or unskilled and must finance their education without access to a loan market. In their setting, however, different agents face different costs of education. In order to become skilled, a child needs a certain amount of schooling time, and this is affected by two factors: innate ability (which is independent across generations and reduces the time needed), and parental background. High-ability individuals need to spend less time obtaining formal schooling in order to acquire skills, and the same is true for children of skilled parents, who need less outside help in order to achieve the same goals, presumably because they have a better understanding of what needs to be done and how best to do it. This is an additional, and plausible, pathway through which parental background affects children's well-being.

There are two alternative technologies in this economy: a constant-returns-to-scale production sector, where each skilled worker produces output according to her productivity, and a sector which uses unskilled workers and natural resources as inputs (with such resources being assumed to be equally divided among unskilled workers). The difference between skilled and unskilled workers lies in their productivity, a , which can take two values, $a_s > a_n$, according to the skill level.

$$\begin{aligned} y_s &= a_s \\ y_n &= a_n x^\alpha \end{aligned} \tag{11}$$

In an equilibrium, skilled workers will earn a_s , while unskilled workers will earn $a_n x^\alpha = a_n \left(\frac{X}{N} \right)^\alpha$, where X is the aggregate stock of natural resources and N the aggregate supply of unskilled labor. The income of unskilled workers decreases when the aggregate supply of unskilled workers rises (since $\alpha \in (0,1)$).

Teachers are hired to provide this schooling time (i.e., to “produce education”) and are paid the skilled-wage rate. Each skilled worker who is employed as a teacher produces a certain amount of units of education: $h < 1$. This links innate ability and parental background to education costs: the less outside time needed for the acquisition of skills, the lower the cost of skill acquisition.

The economy consists of overlapping generations of individuals whose lives can be divided into two periods. In the first, they acquire skills. In the second, they work,

consume, and invest in education for their children. They derive utility from their own consumption and from the well-being of their offspring:¹⁴

$$V_{par} = \ln c_{par} + \beta E(V_{off}) \quad (12)$$

where E is the expectations operator, and V is generational utility.

The process used to model education starts out by specifying innate ability. This is indicated by the amount of time that a person would need to be educated in order to become skilled, if born to a skilled parent. It is labeled *inaptitude* and denoted by e (and is assumed to be uniformly distributed between 0 and 1). The educational barrier faced by children of unskilled parents is introduced by saying that an individual with inaptitude e needs be units of education in order to become skilled if born to an unskilled household, with $b > 1$. Note that when background is introduced in this way, its effect is not “inefficient,” since it actually **requires more resources** to educate children from poor families. This differs from the effect of liquidity constraints, which are also present in the model and which limit educational investment even if it would be efficient to invest in an individual’s education, given her innate ability.

This model has a unique steady state whose equilibrium we will now examine. Under this model, parents will choose to invest in their children’s education if:

$$\begin{aligned} \ln(y - i) + \beta V_s &\geq \ln y + \beta V_n \\ \text{where } i &= \begin{cases} (e/h)y_s & \text{if parent is skilled} \\ (be/h)y_s & \text{if parent is unskilled} \end{cases} \end{aligned} \quad (13)$$

Because of the structure of preferences, both kinds of parents will choose to spend a maximum fraction of their income in education, which we will call m ¹⁵. This fraction is chosen optimally, given the value of education (i.e., the difference in expected welfare between skilled and unskilled individuals).

In turn, m defines two threshold inaptitude levels, such that parents will only invest in education for lower values of e . These thresholds satisfy:

$$\begin{aligned} \frac{e_s}{h} y_s = m y_s &\Rightarrow e_s = hm \\ \frac{be_n}{h} y_s = m y_n &\Rightarrow e_n = \frac{y_n}{y_s} \frac{hm}{b} = \frac{hm}{Ib} \end{aligned} \quad (14)$$

¹⁴ It should be noted that they do not care about the amount of resources spent on their children in and of itself, but rather about the result of their expenditure.

¹⁵ M is obtained as the solution to $-\ln(1 - m) = \beta(V_s - V_n)$.

Where I is the measure of income inequality chosen by the authors, given by the ratio between the incomes of skilled and unskilled workers:

$$I = \frac{y_s}{y_n} \quad (15)$$

Note that, given the distributional assumptions of e , these thresholds give the probability that the son of a skilled (or unskilled) worker will be skilled and thus serve as indicators of mobility. As expected, the probability that the child of an unskilled worker will become skilled is lower than it is for the son of a skilled worker.

Now, we want to solve for the gains to education as a function of the fraction of income that is spent in equilibrium in order to solve for the optimal m , given the outcomes in the labor market:

$$V_s - V_n = \ln I + h \left(1 - \frac{1}{bI} \right) \left[-\ln(1-m) - m \right] \quad (16)$$

What does equation (16) tell us? The gains from education are increasing in inequality. This effect is propagated through a direct channel in the form of what could be termed an income effect, but it is also generated by the fact that it increases the likelihood that children from skilled backgrounds will be more able to afford an education than children from unskilled backgrounds. This amplifies the advantage of being born to a skilled parent. We can also see that gains from education rise when the share of income that is allotted to education increases.

In order to close the model, we need to know the share of individuals who will be unskilled in equilibrium, which will in turn determine relative wages. In other words, we need to take into account what happens in the labor market. For every ratio of skilled to unskilled labor supply, there is a corresponding inequality level, which in turn influences the proportion of income that goes into education. In a steady state, this ratio will be such that the optimal education demands generated will keep the proportions of skilled and unskilled unchanged in the following generation.

Since the model defines a unique steady-state equilibrium, comparative statics come from changes in the model parameters. The key finding is that there is no unique correlation between inequality and upward mobility (measured by the probability that an individual born to an unskilled family will become skilled, e_n). As we have previously shown, this

level is given by $e_n = \frac{hM(I)}{Ib}$, where we have incorporated the fact that the share of income going into education is endogenous.

There are two effects. Through $M(I)$, inequality tends to raise mobility by increasing individuals' *incentives* for investment in education. However, there is also a negative

effect, which is created by the difficulty of paying for a teacher when the wage differential is high. This is the *distance* effect. At low levels of inequality, the incentive effect dominates. When inequality is high, the distance effect prevails.

Table 6: Effects of an Increase in Inequality

	When Initial Inequality is	
	Low	High
Higher Inequality	Increases Mobility	Reduces Mobility

Of course, inequality is endogenous and depends on the labor market’s structural conditions. Hence, Hassler, Rodriguez Mora and Zeira first analyze exogenous changes in the production sector. **Skill-biased technical change** raises the wage differential for any given workforce skill composition. This kind of change increases the incentive to acquire education, thereby raising the proportion of income that goes into education. It also increases inequality, and this latter effect acquires greater weight when inequality is already high. In other words, skill-biased technical change increases inequality the most in economies which were already unequal. The change in mobility mirrors what we said above: for low levels of initial inequality, skill-biased technical change increases upward mobility; for high levels of inequality, it reduces mobility.

The second set of exogenous changes that these authors analyze relates to the educational sector, that is, changes in h and b . An increase in h , i.e., in the overall productivity of education, has two general equilibrium effects: it increases upward mobility, and it reduces inequality. A reduction in the educational barrier faced by children from unskilled backgrounds, as measured by b , has a less clear effect. By making investment by unskilled parents more productive, it reduces the fraction of income invested in education across the board. This implies greater downward mobility, as children from high-skill backgrounds become more likely to forgo an education. However, it can still be shown that inequality goes down as mobility goes up.

Table 7: Effects of Policies in the Hassler, Rodriguez Mora and Zeira Model

Policy/Change	Description	Implemented in	Effect on	
			Inequality	Mobility
Increase in a_s	Skilled-biased technical change	Labor market	Increase	Increase / Decrease
Reduction in h	Increase in productivity of educational system		Decrease	Increase
Reduction in b	Reduction in educational barrier faced by children of unskilled parents	Education sector	Decrease	Increase
Public education p	State financing of p units of education for every child		Decrease (for most parameter values)	Increase

The last change that is analyzed is the introduction of **public education**, which amounts to a certain level of education, p , being purchased by the government (financed by taxes) and given to all children, with parents being free to supplement that education with additional outlays. These authors further assume that there is a proportional tax on income, T , and that h equals 1. With these additions to the original model, the education thresholds are modified and turn into:

$$e_s = p + m(1-T) \quad e_n = \frac{1}{b} \left[p + \frac{m}{I}(1-T) \right] \quad (17)$$

The main point is that skilled parents make more effective use of public education. Carrying out a general equilibrium analysis similar to the one used in the original model, including the relationship between taxes and the level of public education, the authors conclude that an increase in public education raises upward mobility and, unless both β and b are too close to 1, reduces inequality.

Summarizing the findings in Table 7, we see that, here again, public interventions that operate through the educational sector tend to induce a virtuous cycle of reduced inequality and increased mobility. However, it is possible that higher inequality will be met by higher mobility, particularly when the factors driving the increases come from the labor market and are the result of changes in technology.

Although these three models emphasize different aspects of the determination of income inequality and intergenerational mobility, they share a number of common aspects. The main intuition is that **the more parental background determines individual outcomes, the more likely it is that high inequality will be associated with lower mobility**. In Becker and Tomes (1979), reducing the degree of inheritability lowers inequality and raises mobility, which is effectively the same as reducing the educational barriers faced by children of unskilled parents in the Hassler, Rodríguez Mora and Zeira (HMZ) (2003) model or relaxing the tightness of the borrowing constraint in the Owen and Weil (1998) model. In contrast, when we **increase incentives to invest in children** starting from a relatively equal setup, we may find **that inequality increases together with mobility**. This is what happens in the Becker and Tomes model, as well as in the HMZ setup, when skill-biased technical change increases the incentive to education.

These models both suggest a **natural constraint on how much mobility can be changed** and point to policies that may affect it. If ability is intrinsically inheritable, as is the family endowment in the Becker and Tomes model, then it is a source of permanent differences between dynasties and a brake on social mobility. However, this also suggests that efforts to “level the playing field” through the use of such measures as compulsory and universal public education, tutoring, or early intervention programs may be successful in raising income and social mobility. Given the evidence that there is a great deal of scope for early interventions to improve cognitive abilities, the assumption that ability is by and large independent across generations but that parental background can give children an advantage is probably the most relevant for policy purposes.

3. INEQUALITY AND MOBILITY: POLITICO-ECONOMIC CONSIDERATIONS

So far, we have abstracted from the actual workings of governments. A vast amount of literature discusses how politico-economic considerations affect the distribution of income, as well as social mobility, when considered intertemporally. We will now discuss these kinds of results.

At this point, we are mainly interested in the association between inequality and mobility, rather than in the politico-economic consequences of mobility per se (which we deal with in a separate section). The main findings discussed in the literature are summarized in Benabou (1996).

This author sets up an overlapping generations model in which generations within a dynasty are not altruistic toward each other. The utility of each generation is given by:

$$U_t^i = \ln c_t^i + \rho \ln d_t^i \quad (18)$$

where c is consumption when young and d is consumption when old. People are endowed with resources w distributed independently and identically across dynasties. These resources can be invested in capital and used to generate second-period income, according to the technology used:

$$y_t^i = r(k_t^i)^\beta (w_t^i)^{1-\beta} \quad (19)$$

where k is the amount of the investment and w_t is the average level of resources in the community. Given that there are imperfect capital markets, k is limited by individual resources – that is, $k_t^i \leq w_t^i$.

The linkage between generations comes through this resource level, which can be interpreted as the endowment of basic skills: skills possessed by young agents are derived from the interaction between parents' incomes as adults – determined by their own skill level – and innate ability, which is independent across generations. Thus:

$$w_{t+1}^i = \varepsilon_{t+1}^i y_t^i \quad (20)$$

There is a political element in the model which comes in the form of income taxation and redistribution before individual investment decisions are made. For convenience, the author adopts a log-linear specification, whereby post-redistribution income is given by:

$$\hat{y}_t^i = (y_t^i)^{1-\tau} (\tilde{y}_t)^\tau \quad (21)$$

where \tilde{y}_t is determined by budget balance. The tax scheme is progressive when $\tau \in (0,1]$ and regressive when $\tau < 0$. The tax rate results from the workings of the political

process. In order to incorporate the possibility of systems or societies with different levels of wealth bias, Benabou introduces the variable p , which indicates the position in the wealth distribution of the pivotal voter. Higher p implies higher wealth bias.

When a system displays a pro-poor bias, in the sense that the pivotal voter is poorer than the mean, the economy converges toward a unique steady state. The stronger the populist bias, the higher redistribution and mobility are and the lower after-tax inequality is. Thus, we see an **inverse relation between inequality and mobility, as greater equality is associated with higher mobility**, which in this case is induced by political considerations.

When a system exhibits a pro-rich bias, however, **multiple steady states** are possible. Equilibria with low redistribution have **high (after-tax) inequality and lower mobility**. They are also **less efficient than equilibria with higher redistribution**. Thus, we again find that inequality and mobility are inversely correlated. Both are jointly determined by the nature of the political system, which in turn affects the cross-sectional wealth distribution, and by the degree to which it persists from generation to generation.

4. SOCIAL MOBILITY, EQUALITY OF OPPORTUNITY AND MERITOCRACY

A sizeable portion of the literature on social mobility links it with both equality of opportunity and meritocracy, but without – for the most part – specifying the precise conceptual link between these notions.¹⁶ In some cases, these two concepts are taken to be almost exact synonyms, while in others they are considered to be related ideas that sometimes have a causal link. This ambiguity is attributable partly to a certain disregard for semantic discussions by the economic community and partly to the intrinsic difficulty of defining “equality of opportunity” and “meritocracy”. In the next sections, we will attempt to shed some light on these concepts.

Equal Opportunity and Social Mobility

Numerous authors have given a nuanced treatment to the definition of equality of opportunity and its relationship to notions of fairness and justice.¹⁷ We do not attempt to provide a full survey, but rather simply a definition that draws on different sources and is intuitively appealing.

¹⁶ In fact, Benabou and Ok (2001b) argue that since equal opportunity is usually seen as the reason why mobility is desirable, the degree of mobility in the income process should be ranked according to how strongly opportunities are equalized.

¹⁷ See, among others, Rawls (1971), Roemer (1998, 2000), Dworkin (1981) and, more recently, Hansson (2004), Sugden (2004) and Hild and Voorhoeve (2004). All of these authors provide numerous additional references.

As we mentioned in the discussion on mobility, in every society there are certain generally desirable social outcomes with regard to income, status, and the like. Egalitarian philosophies of justice suggest that some sort of equality should prevail in the distribution of these outcomes. Equality in the *opportunity* to achieve these goals, as opposed to equality in the *actual distribution* of outcomes, seems to reflect the most widespread view of what constitutes a fair situation. It is what the “equal playing field” metaphor refers to, and it seems to have almost universal appeal.¹⁸

What exactly, then, is equality of opportunity? Drawing from the common elements in Dworkin (1981), Roemer (1998, 2004) and Sen (2000), we can sketch out a very simple model that helps illustrate this concept. Say that an individual’s outcomes depend on her effort and her circumstances. *Effort* relates to all actions willfully undertaken by the individual that potentially affect her attainment of the desired outcome. *Circumstances*, on the other hand, refer to aspects of the environment that affect the likelihood of attaining the goal but that are largely beyond the individual’s power to change. In a society which values equality of opportunity, individuals will be held responsible for differences in outcomes stemming from differential effort, but not for those that result from two agents applying the same effort level starting from different circumstances.

This brief description highlights the common element and the sources of diverging opinions among proponents of equality of opportunity. As for the common element, note that equality of outcomes *per se* is not deemed desirable a priori. Instead, differences in effort (chosen autonomously by each individual) make differential outcomes not only acceptable, but desirable. Differences arise, though, when the time comes to decide what constitutes “circumstance” and what constitutes “effort” and, hence, in determining how far the equalization of outcomes should go.

One extreme view is to define “equal opportunity” as “non-discrimination”. Under this definition, guaranteeing that people are judged by their qualifications only – for instance, when applying for a job – is enough to achieve equal opportunity. In this case, the relevant circumstances are taken to be “irrelevant” factors that are generally beyond the individual’s ability to change: gender, or race, or other such characteristics. If, on the other hand, only college-educated individuals are given a chance to compete for a given job, this is not considered to be unfair. The implicit assumption is that everyone has had the opportunity to acquire the relevant skills and that those who did not do so should be accountable for their choice.

The opposite extreme is represented by egalitarians such as, for instance, John Roemer (1998), who departs from this view in two significant ways. First, he argues that the set of relevant circumstances should include innate ability, family background (in terms of connections and of attitudes toward effort, availability of role models, and the like), peer group effects, and all factors that affect individual *preferences* but that come into play before an individual can be reasonably expected to notice their importance. To make the

¹⁸ This is the theme of the 2005 *World Development Report* (World Bank, 2005); a large proportion of adults in different countries agreed with the statement: “It’s fair if people have more money and wealth, but only if there are equal opportunities” (cited in Jencks and Tach, 2005).

point explicit, he would argue that whether an individual was raised in an affluent suburb where parents expect their children to go to college or whether he was born into a poor family where scholastic performance was not deemed important are factors that should be regarded as individual circumstances.

The second departure is that he proposes that effort should not be quantified independently from circumstances. That is, he argues that the effort expended by an individual who completes a two-year post-secondary degree despite having been raised in an environment where the median educational level is that of a high school dropout is greater than the effort made by a college graduate who comes from a well-to-do setting in which the median educational achievement is the completion of a college degree. In effect, he suggests circumstances should be used to define groups that serve as a benchmark for gauging individual effort. He also calls for a broad definition of what factors actually constitute “circumstances”.¹⁹

Most opinions lie somewhere in the wide expanse existing between these extremes. Most people would accept the proposition that family connections and privileges should not matter in determining lifetime income and status. Possibly fewer would agree with the idea that differences in innate cognitive or physical ability should play a significant role.²⁰ It is fair to say that relatively few would want to compensate individuals for differences in their preferences that lead them to exert less effort than similar individuals, regardless of where those preferences come from.

At this point, we should see a clear difference between meritocracy and equality of opportunity as we have defined it here, although we may want to count as “merits” aspects that are effectively circumstances. For instance, we may find it plausible that the work ethic of immigrant children is not really of their own doing, yet still believe it is morally defensible to reward them for their deeds. This underscores the fact that equality of opportunity may have the wide support it enjoys partly because of the looseness with which the concept is usually employed.

We will now explore the connection between equal opportunity and social mobility. Both terms are frequently used interchangeably, yet our previous discussion should have established the fact that perfect mobility will only be equivalent to perfect equality of opportunity if we adopt an extreme view of what circumstances should be compensated for when equalizing opportunity.

By our preferred measures, perfect social mobility implies no correlation between parental and offspring outcomes. How could this come about in an actual economic

¹⁹ Roemer (1998) goes as far as to argue that for the children of Asian immigrants in the US, median educational achievement is higher because of their heritage, which encourages them to excel at school; thus, they should not enjoy a higher status as adults than children from African-American families who were equally poor and unconnected to begin with, but whose environment made it unimportant to go to school.

²⁰ Agreement would probably be greater in the case of people with disabilities.

situation? Under perfect social mobility, it should be the case that all differences in status that come from being born into a given family are wiped out.

As the models we have described make clear, achieving perfect social mobility means that the differences in income-generating assets must be fully compensated for. However, these do not just include financial capital and family connections. They also encompass, as we have seen, factors such as intelligence, social environment, attitudes about effort and the like, which are clearly linked to the attainment of social outcomes and which, to a large degree, are inherited through a mix of genetic and cultural transmission mechanisms. We discuss this further below, but we would like to point here to the work of Bowles and Gintis (2002), who document the extent to which factors other than intelligence are transmitted through the family and affect individual earning capacity.²¹

Several models show that the degree of social mobility implied by equality of opportunity depends on how we define “circumstances”. Roemer (2004) drives this point home by stating that parents influence children’s outcomes through the “provision of social connections; formation of beliefs and skills in children through family culture and investment; genetic transmission of ability; formation of preferences and aspirations in children”.²² He then formally shows that in order to achieve zero intergenerational correlation in outcomes, a policy should be devised to equalize all of these sources of differential achievement. As is apparent, it is far from obvious that all of these factors will be considered as “circumstances” by a majority of the population or by policy-makers.

This point bears repeating: achieving perfect mobility involves embracing policies that attempt to compensate for facts such as parental upbringing, or even innate intelligence. In addition, these policies – such as, for example, one geared toward providing more educational resources to children with low IQs – must be permanent if they are to have lasting results. This combination is quite likely to make them thoroughly unpopular and underscores the fact that it may well not be feasible to achieve perfect mobility and that, instead, the question should be recast to focus on how far in that direction we may be able to go.

A different view of the channel through which equality of opportunity may *not* imply higher mobility appeared in early discussions of the topic, particularly in a paper written by John Conlisk (1974) in reply to remarks by Richard Herrnstein, who claimed that equalizing opportunity would effectively create castes of individuals ranked by their inherited cognitive ability. The model in Conlisk’s paper shows the extent to which this is, indeed, true.

²¹ Other references are Feldman, Otto and Christiansen (2000) and Korenman and Winship (2000), who explore what cultural transmission mechanisms are and how they account for intergenerational correlations in several relevant traits.

²² A variety of models in the economic literature, including all those which we have summarized in this article, show that this is indeed the case: that differences in any of these aspects can generate persistent inequality across dynasties, even in the presence of complete markets. Piketty (2000) provides a survey. The case of beliefs and aspirations is especially interesting and will be examined in detail below.

Assume that income obeys a simple system of three linear equations. The income (Y) of a given individual depends on realized intelligence (I) and parental income – reflecting the return to ability and the degree to which parents transfer income-generating assets to their children. Intelligence depends both on the genetic “potential” (G) and on parental income, presumably because, as we have already argued in Section 2, parents with higher incomes provide a more stimulating environment for their children or may have an advantage in education. Lastly, genetic potential is inherited from parents through a first-order autoregressive process. Of course, there is also an element of “chance” in each of these functional dependences (the usual error term).

$$\begin{aligned}
 Y_t &= \alpha_0 + \alpha_1 I_t + \alpha_2 Y_{t-1} + u_{Yt} \\
 I_t &= \beta_0 + \beta_1 G_t + \beta_2 Y_{t-1} + u_{It} \\
 G_t &= \gamma_0 + \gamma_1 G_{t-1} + u_{Gt}
 \end{aligned}
 \tag{22}$$

Now, the author argues that fully equalizing opportunity would make some of the coefficients in these regressions equal to zero, in particular, α_2 and β_2 . If this were the case – if, say, we could somehow reduce the link between parental income and realized intelligence – social mobility would clearly increase in both the short- and the long-run. That is, we would see a drop in the intergenerational correlation of income. If we were to keep the discussion on this level, we would return to our previous point about the extent to which equalizing opportunity would imply setting the coefficients to zero (or reducing them to the fullest possible extent). However, the main point in Conlisk’s paper is a different one.

Equalization of opportunity, it can be argued, should include reductions in the variability of the error term. That is, if the error term captures aspects such as geographical location, quality of public education and other environmental variables, an equalization of these aspects would strengthen the link between effort and outcomes, thereby making individuals with similar characteristics more likely to enjoy similar outcomes. Conlisk shows that this reduction would effectively decrease social mobility by making inherited traits more determinant.

Note that we are once again reminded of the limits imposed by biological considerations. This does not mean that the current situation cannot be improved upon, but it certainly should be borne in mind.

Almost any notion of justice calls for individuals to enjoy similar opportunities to achieve their goals – for people to interact on a “level playing field”. We have shown the extent to which this equality of opportunity would imply social mobility. As we have seen, where the equalization of opportunity helps to break down the family transmission of relevant traits, it will increase mobility. However, these two concepts are not completely equivalent, in that only the advocates of a very extreme position on the issue would contend that a satisfactory equalization of opportunity implies perfect intergenerational mobility. This extreme position would probably be untenable from a political standpoint and would greatly expand the scope of state action.

Meritocracy and Social Mobility

A concept related to equal opportunity is that of “meritocracy”. Again, as commonly used by economists, meritocracies are assumed to go hand in hand with social mobility. In this section, we provide a brief conceptualization of the term and discuss its relation to mobility and equal opportunity.²³

As a very first approximation, the idea of a meritocracy describes a society where income and social status reward “merit”; that is, they accrue to those who “deserve” it. As Sen (2000) puts it, however, this leaves the concept essentially underdefined, for the definition of what constitutes merit is contingent on the definition of what a good society is and of what helps to bring it about.

What is merit, then? Meritorious actions are presumably those that have desirable consequences, possibly from an ex-ante point of view, and are therefore to be encouraged. A well-planned investment designed to generate a positive impact may, of course fail, as it inevitably entails a degree of risk –as any other investment–, but it would presumably be regarded as having more merit than one with a negative expected value. In this sense, a meritocratic society should tend to encourage wise investments. Work effort in productive activities is another example of an action that would be rewarded under this concept.

This line of thought would seem to be in step with what most people consider fair. “Good”, productive actions should be rewarded, and meritocracy is a reasonable goal for society. Note that if we think that some level of equality is a desirable end, then a meritocratic society should also reward actions that tend to reduce inequality (including, potentially, affirmative action, education programs, and the like).

However, the customary approach taken to this subject in economics does not follow this path but instead sees “merit” as being an intrinsic characteristic of an individual. What is arguably the initial characterization of this idea, by sociologist Michael Young in his book *The Rise of Meritocracy, 1870 – 2033*, defines merit as the combination of ability and effort. Although the initial definition of merit can be expanded to accommodate some measure of effort, since it probably leads to better outcomes, the element of ability is a different proposition. Once ability is taken into consideration, then the task of tracing the relationships among meritocracy, social mobility and equal opportunity requires us to specify how ability comes to be and how it is transmitted across generations.^{24,25}

²³ Much of the discussion in this section is based on the collected papers presented by Arrow, Bowles and Durlauf (2000).

²⁴ Ability should be broadly defined as anything that is rewarded in the labor market. It is generally taken, though, to be roughly equivalent to cognitive ability or IQ.

²⁵ In an article published shortly before his death, Michael Young states:

“Ability of a conventional kind, which used to be distributed between the classes more or less at random, has become much more highly concentrated by the engine of education.
(...)”

This idea of meritocracy could, however, fail to induce social mobility in the long run for two linked but distinct reasons.

On the one hand, we may think, as Herrnstein and Murray do in *The Bell Curve*, that intelligence is for the most part genetically inherited. If this were indeed the case, then, as we have argued in the previous section, a meritocracy would typically be associated with low mobility. People with high cognitive ability have a high return to education, which makes them acquire skills and enjoy high wages. In turn, their offspring will also have high wages because of their high cognitive ability. This would put mobility at odds with the idea of meritocracy. In a way, it also goes against equal opportunity, insofar ability is part of an individual's circumstance.

However, there are strong reasons to believe that there is still a great deal of work left to do before differences in cognitive ability can be claimed to be solely determined by genetic endowments. Most of the work on intelligence is based on IQ measurements. Flynn (2000) and Feldman et al. (2000) argue that, while there is a high correlation of IQ scores between generations of the same family, it is quite probable that a large share of this correlation is due to environmental factors (in other words, that it is a result of "cultural" transmission). This is also supported by research on early child development, which suggests that the human brain is quite malleable during the early years.²⁶ If this is the case, it may be possible to devise policies that increase the reward both to merit and to social mobility. If people are put on roughly equivalent "cognitive starting points" by means of, for example, targeted early intervention programs, then educational attainment – and eventually labor-market outcomes – would mostly depend on effort and career choices. Since this reflects mainly individual preferences, we could expect mobility (in the sense of relative, or exchange, mobility) to increase. We will further explore this line of reasoning in Section 6 based on the survey conducted by Cunha, Heckman, Lochner and Masterov (2005).

The second reason why meritocracy could lead to reduced social mobility is to be found in the arguments put forward by Bowles and Gintis (2002), who contend that much of the intergenerational transmission of status is not explained either by genetic transmission of IQ or by the transmission of assets. Instead, they reason that personality traits such as fatalism, a work ethic and respect for authority are both important in the determination of

With an amazing battery of certificates and degrees at its disposal, education has put its seal of approval on a minority, and its seal of disapproval on the many who fail to shine from the time they are relegated to the bottom streams at the age of seven or before.

(...)

I expected that the poor and the disadvantaged would be done down, and in fact they have been. If branded at school they are more vulnerable for later unemployment.

They can easily become demoralized by being looked down on so woundingly by people who have done well for themselves.

It is hard indeed in a society that makes so much of merit to be judged as having none. No underclass has ever been left as morally naked as that."

²⁶ For research on this topic, see Young (2002).

labor-market outcomes and transmitted from generation to generation, presumably through parental role models, peer effects, and other cultural mechanisms.²⁷ These channels are much more difficult to neutralize, as they involve a change in the transmission of values. Unless these channels are modified, it is very unlikely that increased rewards to merit will raise long-run social mobility. Quite to the contrary, they may lead to the outcome dreaded by Michael Young (1958): a system of merit-based castes. This might be considered desirable on moral grounds, but it runs counter to the attainment of a society with high mobility. It could also be inefficient from a purely economic standpoint, as we will argue in Section 6.

To conclude this somewhat normative discussion, we find that a transition to a society where merit determines social rewards would not necessarily increase social mobility, unless we accept the proposition that all the determinants of merit – including preferences and possibly aspirations – are independent of parental origin. This echoes our previous section on equal opportunity, with the cautionary note that the link between social mobility and meritocracy, which is usually assumed away in most economic work on mobility, is actually dependent on a series of assumptions about the inheritability of income-generating traits.

In other words, while we may think that both meritocracy and mobility in a given society are desirable, it may well be the case that policies that advance one of these causes work against policies designed to attain the other. As an example, opponents of affirmative action as a means of redressing past injustices usually claim that it runs counter to the establishment of a meritocracy, yet affirmative action may serve to increase mobility.

5. MERITOCRACY, OPPORTUNITY, MOBILITY AND EFFICIENCY

As a means of transitioning to our discussion of mobility and its economic consequences and correlates, we survey a set of models that point to its interrelationship within an economic context, as well as its efficiency consequences, while keeping the normative concern present.

The Size of the Pie

Benabou (2000) provides a general framework for an explicit examination of the links among these four notions. In a model very similar to those we have discussed, he starts by assuming that individual pre-tax income reflects both social background and personal characteristics. As a first approximation, these personal characteristics are assumed to be distributed independently from social background and to represent “ability”. At this point, he equates equality of opportunity with meritocracy (under the common heading “meritocracy in opportunities”) and identifies it with the share of total income variance that comes from variability in individual traits.

Granting that the assumption of independence between individual traits and social background is not quite realistic, he redefines **meritocracy in opportunities** as the share of total variance in pre-tax income that is “unexpected” given social background. That is, he measures how much of current income is not dependent on parental status and the like, but instead on other autonomous factors. Insofar as these factors reflect variability in personal choice, effort, and so forth, this definition arguably comes closer to an actual meritocracy than the first one.

A related idea in this context is that, in a meritocracy, individual outcomes should also differ – in the sense that after-tax incomes should not be completely equalized. In other words, the existence of a meritocracy should imply some sort of incentive system whereby good actions get higher rewards. Thus, it should also be the case that “luck” or talent is effectively rewarded even after taking into account any fiscal redistribution – what Benabou terms “**inequality of outcomes**”.

Table 8: Aspects of Meritocracy

Societal Outcome	Definition	Environment
Meritocracy in opportunities	Degree to which market outcomes are unrelated to social background, explained by personal choices or characteristics	Market, pre-tax
Inequality of outcomes	Degree to which effective income is determined by individual talent and abilities	After-tax

Both of these societal attributes are desirable, but they are not logically connected; it is quite possible that a society may be very meritocratic in its market opportunities, but that it fully redistributes income, so that effective outcomes are equalized. At the same time, whether inequality of outcomes is desirable clearly depends on whether opportunities have been allocated fairly. This consideration leads to the formulation of a two-dimensional measure of meritocracy – a “meritocracy utility function” – which expresses an ordering of all possible combinations of equal opportunity and unequal outcomes.²⁸ This makes the relationship between these two notions explicit. In particular, this measurement should reflect the fact that, as opportunities become more equal, inequality in outcomes is more acceptable (even desirable) and that as outcome inequality rises, the value of additional equality of opportunity increases.

What about the rewards to effort? In an argument that echoes Roemer (1998), Benabou claims that differences in the level of effort “must ultimately reflect different (perceived) returns to effort, hence differences in either background or ability”. Thus, the degree to which increased rewards to effort lead to improvements in terms of an overall

²⁸ As a technical note, the function is assumed to be increasing in both arguments, with a positive cross-derivative, and quasi-concave. The value of the function is minimal when either dimension tends to zero: an *aristocracy*, if there is no meritocracy in opportunity, or a *mediocracy* if all post-tax incomes are perfectly equalized.

meritocracy would depend on the extent to which attitudes toward effort are shaped by family background. If the influence of this factor is high, then equality of opportunity will be reduced (although an increase in rewards to effort will certainly raise the inequality of outcomes).

Table 9: Effect of Rewards to Effort in terms of Meritocracy

Outcome	Effect
Meritocracy in Opportunities	Depends on factors affecting costs and returns to effort: if these factors depend on parental characteristics, an increase in rewards to effort would reduce the meritocracy of opportunity
Inequality of Outcomes	As long as there are any differences in effort level, inequality of outcomes will increase with rewards to effort

When considering policies from an intertemporal, long-run point of view, it becomes apparent that tradeoffs are inevitable; most policies will affect both dimensions in different directions. Consider a policy that redistributes income: it reduces inequality of outcomes, but it increases equality of opportunity in the following generation, as deviations from mean market income are less likely to come from differences in parental background. To understand these effects more fully, the author presents a model of infinitely-lived dynasties.

Each generation is born with a certain level of human capital, which depends both on random luck (independent across generations) and on the individuals' parents, through their realized level of human capital and through an educational choice as well. That is, we assume that skills are hereditary to some level, but that they are also affected by education. Each generation must choose three quantities: labor supply (which is costly in terms of utility but increases pre-tax income), own consumption,²⁹ and education for the next generation. There are no capital markets, so education and consumption must be financed by current income. Current income depends on the level of the labor supply and on human capital. Income is taxed and redistributed, which means that taxes can relax the liquidity constraint for families with small human-capital endowments.

Because of these credit-market imperfections, income inequality is associated with lower output growth. However, starting from a given level of inequality, the tax rate that maximizes current growth is lower than the rate that maximizes the long-run level of income.³⁰ Intuitively, we know that redistributing today reduces the current incentive to work but also induces growth further in the future, since it reduces inequality in all subsequent periods by shifting income to liquidity-constrained households.

²⁹ An overlapping generational interpretation would include "children's" consumption.

³⁰ Throughout the discussion, it is assumed that the tax rate is the same in every period.

How do equality of opportunities and inequality of outcomes appear in this model? The former is a measure of social mobility, and it decreases when the relative importance of education and parental human capital in the determination of current human capital are high, as well as with the importance of skills in the labor market. It increases, on the other hand, with the level of redistribution. Inequality of outcomes, by contrast, increases with the importance of skills and decreases with redistribution.

Table 10: Impact of Parameters in terms of Meritocracy

Parameter	Effect of an Increase in the Parameter
Impact of education on current human capital	Reduces meritocracy in opportunities in the presence of liquidity constraints
Impact of parental human capital on current human capital	Reduces meritocracy in opportunities
Market return to acquired skills	Reduces meritocracy in opportunities Increases inequality of outcomes
Extent of redistribution	Increases meritocracy in opportunities Reduces inequality of outcomes

What is their impact on growth and long-run output? On the one hand, inequality of outcomes increases incentives to work, thereby raising overall efficiency. Equality of opportunity is also efficient, as it tends to reduce the growth losses that operate through income inequality, i.e. through differential access to education. However, the extent to which this aspect of meritocracy goes hand in hand with efficiency is limited, as reducing disparities in access to education also decreases the efficiency gains to be derived from a level of education that is more productive for the offspring of highly skilled individuals.

In summary, meritocracy and efficiency are closely related when capital markets are imperfect, but their correlation depends on how parental background affects the productivity of investments that raise income.³¹ In addition, meritocracy has two distinct aspects, equality of opportunity and inequality of outcomes, and a normative ordering should incorporate both dimensions. By and large, social mobility, equality of opportunity and meritocracy in general tend to increase efficiency in economies with limited access to education finance.

³¹ This suggests the possibility that breaking the link between parental background and the productivity of education enhances efficiency and mobility. Cunha et al. (2005) argue that this is what early childhood interventions do.

Inequality of Opportunity and the Need to Provide Incentives

A related approach is developed in Phelan (2006). He argues that in a dynamic contracting model with unobservable effort where agents demand incentives not to shirk, it is efficient to have both unequal opportunity *and* social mobility. In other words, agents born into families with high realized output should enjoy higher income (unequal opportunity) but, over time, all dynasties should have a positive probability of changing positions in the income ranking (social mobility).

In this model, there is a continuum of dynasties, each composed of a generation that lives for one period. Production depends on effort: if a household exerts effort $a \in \{a_0, \dots, a_N\}$, then output $q \in \{q_0, \dots, q_N\}$ occurs with a probability $p(q|a) > 0$. Intuitively, low outcomes are more likely when effort is low, while high outcomes are more likely when effort is high. Households receive an “allocation” of effort (i.e., they are told what effort level to exert), but this effort level is not observable. In other words, households can claim to exert a certain level of effort, but may actually shirk instead.

Household utility in one period is given by:

$$U(c_t, a_t) = u(c_t) - v(a_t) \quad (23)$$

where c is consumption in one period and U is increasing in consumption and strictly decreasing in effort. A household cares about the expected value of the average discounted payoffs, $(1 - \beta) \sum_{t=0}^{\infty} \beta^t U(c_t, a_t)$.

As we have said, this is a dynamic contracting model. A planner designing this dynamic social contract must choose the best among all feasible allocations. What are “allocations” in this context? Formally, allocations specify an initial distribution of promises of lifetime utilities, Ψ_0 , a sequence of effort allocations for each level of promised expected utility, $\{a_t(w_t)\}_{t=0}^{\infty}$, a sequence of consumption allocations for each level of observed output and promised expected utility, $\{c_t(w_t, q_t)\}_{t=0}^{\infty}$, and a sequence of “promises” or “continuation values” for every realization of output and initial promise, $\{w_{t+1}(q_t, w_t)\}_{t=0}^{\infty}$. In other words, it starts by specifying how “rich” (in terms of utility) a family can expect to be over its members’ lifetimes and then provides a rule specifying how much effort a generation must exert. Next, for each possible realization of outcome, it states how much will be consumed at that time and how much will be “inherited” by the generations to follow in the form of “promised” utility.

For an allocation to be feasible, it must:

- Deliver on its promises: The expected utility arising from following the allocation starting at time t with a promised utility w_t must equal w_t , so that the outcome does not depend on people being systematically wrong about their prospects.³²
- Provide the right incentives: If an allocation calls for a certain level of effort, that effort should be optimal for the household in expected terms, so that its members will freely choose it from among all possible alternatives.³³
- Respect the resource constraint: The amounts consumed in each realization of aggregate output must not exceed that output.³⁴

There are many different allocations that satisfy these three conditions, which means we must have some criterion we can use to choose among them. Phelan first recasts the problem as one of several individuals who are linked to each other by some sort of altruism. Their individual welfare therefore depends both on the existing level of utility when they are alive and on the promised utility to their descendants. In other words:

$$w_t = \sum_{q_t} P(q_t | a_t(w_t)) \left[(1-\beta)U(c_t(w_t, q_t), a_t(w_t)) + \beta w_{t+1}(w_t, q_t) \right] \quad (24)$$

This is a reinterpretation of the previous model, with identical formal results. The advantage of thinking through the problem in this way is that it suggests a criterion which can be used to rank feasible allocations. The author argues that the correct welfare criterion is the average welfare of all generations, across both time and dynasties. This departs from the traditional criterion, which maximizes a weighted average of $(1-\beta) \sum_{t=0}^{\infty} \beta^t U(c_t, a_t)$, the “per generation” utility of a dynasty at time zero, where the average is over dynasties. The argument backing up this change is that the traditional formulation, which has appeared in several previous papers, puts no weight on the welfare of subsequent generations.³⁵

The author shows that the optimal allocation under this criterion exhibits both **unequal opportunities** and **social mobility**.

³² Formally, $w_t = \sum_{q_t} P(q_t | a_t(w_t)) \left[(1-\beta)(u(c_t(w_t, q_t)) - v(a_t(w_t))) + \beta w_{t+1}(w_t, q_t) \right]$.

³³ Formally, $w_t \geq \sum_{q_t} P(q_t | \hat{a}) \left[(1-\beta)(u(c_t(w_t, q_t)) - v(\hat{a})) + \beta w_{t+1}(w_t, q_t) \right]$.

³⁴ Formally, $0 \geq \int_V \sum_{q_t} P(q_t | a_t(w_t)) [c(w_t, q_t) - q_t] d\Psi_t(w_t)$ where $V = \{u(c) - v(a_0) | c \in C\}$ and C is the consumption set.

³⁵ Formally, a planner must choose a dynamic allocation that maximizes $\bar{v} = \liminf_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T \bar{v}_t$,

where $\bar{v}_t = \int_V w_t d\Psi_t(w_t)$.

Unequal opportunities are reflected in the fact that dynastic utility during every period tends toward a limiting yet variable distribution (at every point in time, some people will be born with lower expected welfare than others). These disadvantaged individuals will be those born into dynasties with low realized output in the past. This is a consequence of the need to provide incentives for agents to expend effort. In other words, since incentive constraints are binding, a way to relax them – i.e. to provide incentives in a more efficient way – is to make the utility of generations down the line (which current generations care about, as they are altruistic) dependent on current output. In other words, it pays to have parents who are able to provide lifetime advantages to their children.

Social mobility is associated with the fact that an individual is born with a given expected dynastic utility does not preclude him from moving up or down in this distribution as periods go by. In fact, this is bound to happen. In particular, there are no castes (i.e., income groups whose composition does not change over time). This effectively means that, eventually, all dynasties will have a chance of changing their ranking, no matter where their starting point is.

What is remarkable about this result is how it relates to the veil of ignorance under which allocations are to be judged, according to philosophers whose postulates can be traced back to John Rawls. This is an optimal allocation in the sense that it would be chosen by an individual who does not know which dynasty or generation she will be born into. Even so, and despite the fact that all agents are ex-ante equal, the optimal arrangement exhibits inequality of opportunity. In addition, it features social mobility as a characteristic of an optimal arrangement, thereby pointing to its efficiency.

This result treats all agents as perfectly identical in terms of the income opportunities they face, at least in the initial period. This should be borne in mind, since no information is provided as to what should be done if these opportunities were to differ by dynasty.

Intelligence, Mobility and Efficiency

So far, all the models we have discussed relate to *earnings* mobility for workers. A different case in which meritocracy may relate to higher efficiency and mobility is modeled by Hassler and Mora (2000), who consider the social mechanisms that allocate agents to occupational categories: entrepreneurs vs. workers. They define “intelligence” as “what you use when you don’t know what to do”, borrowing from Jean Piaget, and explore the link between intelligence and entrepreneurship in high- and low-growth environments.

In their model, individuals are born with a given level of intelligence, which is less than perfectly correlated with their parents’, and with a stock of social assets that is fully determined by their background. They interpret meritocracy as the degree to which social positions are influenced by innate intelligence as opposed to background (note how this relates closely to the “equality of opportunity” aspect of meritocracy defined by Benabou,

2000). In this model, a higher degree of meritocracy implies higher mobility, since innate traits are less correlated across generations than social background is.

Individuals, then, must choose whether to become entrepreneurs or workers. The latter earn a fixed economy-wide competitive wage, w_t , which is fully known at the time of the decision, whereas entrepreneurial income depends on the ability to forecast the value of a random variable, interpreted as the “correct” decision.³⁶ Entrepreneurs hire a certain number of workers, l_t , and choose a as close as possible to a stochastic process x_t . Profits are given by:

$$\Pi = e^{-(x_t - a)^2} (2e^{r_t} l_t^{1/2} - w_t l_t) \quad (25)$$

For any rational choice of l , profits are maximized by correctly predicting x . Individuals have a certain belief as to the true value of x_t , which they assume to be normally distributed with mean $\mu(j)$ and variance $1/P(j)$.

The authors show that regardless of their beliefs, all entrepreneurs will demand $l_t = \left(\frac{e^{r_t}}{w_t}\right)^2$ workers. The best action a , on the other hand, will depend on those beliefs and, in this case, will be $a = \mu(j)$. Given this, the **expected utility**³⁷ of agent j if she becomes an entrepreneur depends crucially on the precision of her belief and equals:³⁸

$$\begin{aligned} V^e(j) &= E \ln \Pi \\ &= 2r_t - \ln w_t - \frac{1}{P(j)} \end{aligned} \quad (26)$$

If the agent chooses to become a worker instead, utility is no longer stochastic. Instead, it is fully known before the decision and equals $\ln w_t$. Note that this implies that expected profits must be higher than wages, and we can therefore think of entrepreneurs as “rich”.

Agents choose the occupation that offers a higher expected utility, which is determined by a threshold degree of accuracy in their beliefs about the true state of the world. Agents with a high degree of accuracy end up being entrepreneurs, while agents with a lower level of accuracy choose to become workers. In other words, an agent will choose to be an entrepreneur if:

³⁶ This decision changes from period to period; in the real world, it would be a multidimensional object indicating the “best way” to run a firm.

³⁷ We are assuming a log utility function.

³⁸ “Precision” is defined as the inverse of the variance of a random variable.

$$\begin{aligned}
V^e(j) &\geq V^w(j) \\
2r_t - \ln w_t - \frac{1}{P(j)} &\geq \ln w_t \\
P(j) &\geq \frac{1}{2(r_t - \ln w_t)}
\end{aligned} \tag{27}$$

This shows that the process that determines such beliefs will be the driving force behind the “class structure” in this model. Entrepreneurial activities will be the domain of people with access to high-quality information. Workers will be those with higher variance in their beliefs about the way the world functions.

Technology is a pair $[r, x_r]$ that follows an exogenous stochastic process in continuous time (specifically, a Brownian motion). For the initial analysis, technological growth is exogenous and follows:

$$x_{r+g} = x_r + \int_0^g \sqrt{\sigma} dz \tag{28}$$

where dz is a standard Wiener increment. In particular, by setting the exogenous rate of technological growth between periods to g , we obtain:

$$x_{t+1} = x_t + \sqrt{g\sigma} \varepsilon_t \tag{29}$$

with ε_t a standard normal, independent over time. Note that $g\sigma$ can be interpreted as the amount of information needed to keep up with the new technology and suggests that, at times of high rates of technological growth, the “best way” to run a firm may vary widely from one period to the next. We can now turn to the key issue: how beliefs about x are formed.

Each individual has two sources of private information, neither of which he can trade or transfer.

First, he obtains an unbiased signal about x_{t-1} from his parents. Depending on social background, the signal will vary and will be more or less precise: children of entrepreneurs will get to know the exact value of x_{t-1} , whereas children of worker families will have an unbiased signal of x_{t-1} with identical variance $\gamma > 0$.

A second source is the individual’s cognitive ability, or intelligence.³⁹ Each individual is born with an unbiased signal of x . Individuals differ in the precision of their signal, which

³⁹ We omit the lengthy discussion presented in the original paper about what the politically (and biologically) correct term is for the ability to make good decisions in uncertain environments with less than perfect information.

may take on one of two values, the higher of which corresponds to highly intelligent individuals. Specifically, the precision for low-intelligence individuals is 1, while for high-intelligence individuals it is $\alpha > 1$. Intelligence is independent across individuals and along generations. Individuals are rational Bayesian agents, and they update their prior signal based on both sources. The resulting precision levels are:

Table 11

		Family Background	
		Entrepreneur	Worker
Intelligence	High	$\frac{1}{g\sigma} + \alpha$	$\frac{1}{\gamma + g\sigma} + \alpha$
	Low	$\frac{1}{g\sigma} + 1$	$\frac{1}{\gamma + g\sigma} + 1$

There are two sources of increased precision. Knowledge of the past realization is useful in constructing a forecast. Individuals from all groups use this knowledge, but those from working-class backgrounds start off with worse information. This is captured in Table 11 by the fact that $\frac{1}{g\sigma} > \frac{1}{g\sigma + \gamma}$. Both groups have to deal with the fact that the accuracy of this forecast will depend on how fast the environment is changing. A second source of increased precision is individual intelligence.

Our interest in this “precision” stems from its role in determining the class structure. We should keep in mind that, in an equilibrium, there should always be both entrepreneurs and workers. This means that agents with the lowest signals (in this case, low-intelligence individuals from working-class backgrounds) will be workers. By the same token, high-intelligence individuals from entrepreneurial families have the best information and thus will be entrepreneurs. The threshold level will be somewhere in between, and the ordering between high-intelligence working-class individuals and low-intelligence entrepreneurial-class persons depends on the model’s parameters.

For meaningful possibilities of social mobility and meritocracy, therefore, the key inequality is whether a highly intelligent agent from a working-class background will have higher precision (and, thus, a higher probability of choosing to be an entrepreneur) than a low-intelligence individual from an entrepreneurial family. For high enough values of $g\sigma$, which measures the pace of change in the economic environment, this will indeed be the case, as the value of past information becomes negligible in comparison to understanding current conditions.

In fact, when these authors study the steady state of this model,⁴⁰ they find that the parameters that explain whether there is intergenerational mobility depend on $g\sigma$ and on the distribution of intelligence (specifically on the proportion of highly intelligent individuals that are born in each generation, q).

⁴⁰ In the steady state, the proportion of entrepreneurs and workers is kept constant.

Table 12: Factors affecting Social Mobility

Model Parameter	Interpretation	Effects on Social Mobility	Pathway
$g\sigma$	Pace of change in the economic environment, complexity	Higher $g\sigma$ leads to higher mobility	Intelligence becomes more decisive in the occupational decision. When $\frac{1}{\gamma + g\sigma} + \alpha > \frac{1}{g\sigma} + 1$, smart children from working-class families become entrepreneurs, “replacing” low-intelligence children from entrepreneurial families.
q	Proportion of highly intelligent individuals in the population	In high-complexity environments, a very large or very small q leads to lower mobility	When q is too large, not all intelligent individuals can be entrepreneurs in an equilibrium. Thus, some of the intelligent persons from working-class households will remain workers. When q is too small, some of the less intelligent agents from entrepreneurial families will become entrepreneurs. Thus, extreme levels of q induce a high degree of dependence for individual outcomes on initial family conditions.

When the environment is not very complex ($g\sigma$ is low), there can be no social mobility in a steady state. Why? For a society to be in a steady state, the proportion of entrepreneurs must remain the same. Yet, in order for there to be some degree of social mobility, some children from working-class families must become entrepreneurs. As we have argued, these children must be those with high intelligence.

In low-complexity environments $\frac{1}{\gamma + g\sigma} + \alpha < \frac{1}{g\sigma} + 1$, if high-intelligence individuals from working-class families choose to become entrepreneurs (a necessary condition for mobility), so will low-intelligence individuals from entrepreneurial families. But if this were the case, then the share of entrepreneurs would increase (it would include all former entrepreneurs, plus at least some of the highly intelligent individuals from working-class families), thus creating a situation that departs from a steady state. Consequently, **any steady state in a low-complexity society will involve social immobility.**

What happens in high-complexity environments? This depends on q . In general, in these environments it is relatively hard to be a successful entrepreneur, and inherited information has little value. This leads intelligent individuals to choose to become entrepreneurs. Will the set of highly intelligent individuals coincide with the set of entrepreneurs? This would indicate both mobility and meritocracy, but the answer depends on the parameters. For instance, q could be so large that if all highly intelligent individuals were to become entrepreneurs, their income relative to wages would be too low, and some might choose to be workers instead (and the ones who would make that choice would be from working-class families). Alternatively, q could be low, making wages too low and inducing (some of) the low-intelligence agents from entrepreneurial families to become entrepreneurs themselves.

Therefore, in highly complex societies, there will always be a greater degree of meritocracy, in the sense that **entrepreneurs will be more intelligent than the population as a whole**. However, the degree of mobility will depend on the proportion of intelligent individuals. A large proportion of intelligent individuals reduces downward mobility, thus increasing the extent of dependence between past and current outcomes.

These authors add one last twist: technological growth may be partly related to investments undertaken by entrepreneurs. In this case, individuals who choose to be entrepreneurs can also choose whether to innovate; in other words, they choose their own level of g . Higher levels of g increase both the mean and the variance of expected profits and are best used by intelligent individuals. The authors show that the level of innovation will not depend on social background, but rather on intelligence alone. As before, whenever a low-intelligence individual chooses to become an entrepreneur, so will a high-intelligence agent from the same background. Also, whenever an individual from a working-class background finds it profitable to become an entrepreneur, so will an individual from an entrepreneurial background.

The end result is that, in this case, there will in general be two steady states. In one of them, there is no mobility and little innovation, since innovation is driven by intelligent individuals only. In the other, intelligent individuals from both classes become entrepreneurs in any generation, and so technology (and, therefore, output) grows, thus increasing the degree of meritocracy and making mobility possible.

The policy implications of this model are not absolutely clear. We may think that, starting from a rigid society, opening up to competition and technological adoption will tend to reduce the “informational advantages” that are inherited by children of entrepreneurial families. This model abstracts from failures in the credit market, so there is no role for policies in tackling this issue. However, if it were a case of potential entrepreneurs from working-class families facing some sort of credit constraint, there could be room for incentives for research and innovation that would more than likely be used by intelligent, as opposed to well-off, individuals.

Finally, a note of caution is warranted regarding these policy prescriptions, which amount to making the economic climate more “challenging” –for instance, by opening up to

international trade and foreign technological progress— so that “inherited” knowledge is not useful in determining outcomes. This, however, may make for a tough transition, as entrepreneurs make initially bad decisions that drive down real wages for the economy as a whole.

It is not clear how to fit education into the model. One interesting way in which it might work with the model would be if it increases the precision of the private, innate signal of each individual. In other words, if it increases intelligence. There does not seem to be any sort of consensus on this point, however, particularly since there is little evidence that education, other than at a very early age, actually changes cognitive ability. Also, if this were the channel through which education acted, it would likely do relatively little to reduce the barrier between the rich and the poor; consequently, starting from a situation with no mobility, it would not erase the advantage of a low-intelligence individual from an entrepreneurial family, as he would now also be more intelligent.

However, it may be that education acts through another channel. It might provide more information about past environments, therefore directly reducing the advantage enjoyed by children in entrepreneurial families. This second role could potentially move a society away from a steady state of perfect immobility by providing intelligent agents from working-class families with an extra edge.

Thus, in short, although we may think equality of opportunity or meritocracy are desirable attributes in a society, they bear no necessary connection to economic efficiency. In many cases, they are positively related, but in these cases the mechanism seems to go from economic efficiency to opportunity or meritocracy and mobility. This is particularly true when economies have imperfect market arrangements that create room for efficiency-enhancing policies – namely, when imperfect capital markets preclude individuals from fully realizing their potential human capital and skills.

6. EFFECTS OF SOCIAL MOBILITY

We will now elaborate upon our line of reasoning in arguing that mobility may be a necessary feature of an efficient economic system. As we have suggested, it is hard to argue, from a purely economic point of view, that mobility should be a goal in and of itself, as opposed to being the consequence of efficiency-enhancing policies involving human capital accumulation. However, when we add a political dimension to the analysis, we find that mobility may have consequences per se, as it adds an extra dimension to distributional conflict. Lastly, we will see that the relevant effect does not derive from the actual income process, but rather stems from the perception that agents have of it, which informs their position in the political arena.

Economic Efficiency

At this point, we have briefly gone through most of the main body of the economic literature relating to social mobility and economic performance. We will now summarize the key findings, without going into the specific details of each particular model.

The main conclusion, as set forth in Becker and Tomes (1979), Owen and Weil (1998), and Hassler, Rodriguez Mora and Zeira (2002), is that whenever liquidity constraints prevent individuals from acquiring an efficient level of human capital, social mobility is limited and economic efficiency is reduced. In such cases, redistribution helps ease the liquidity constraint and increases long-term mobility.

When markets are complete, however, social mobility may be low (equivalently, the effects of different initial conditions may be persistent), but efficiency considerations do not justify interventions – only equity concerns provide a reason to redistribute. A number of these models can be found in Piketty (2000). He identifies two main reasons for “persistent inequality” with complete markets.

The first is the intergenerational transmission of financial assets in a model where income is determined by labor supply and returns on wealth:

$$y_{it} = v_t a_{it} + r_t w_{it} \quad (30)$$

where a is ability that is rewarded in the labor market at rate v , and rw is capital income – r is the interest rate and w the stock of inherited wealth. Without any behavioral assumption, income inequality will be greater than the inequality of labor earnings, unless a (which is taken to be exogenous) and w are negatively correlated. This is a feature observed in empirical studies.

If we assume further that bequeathed wealth is an increasing function of income, then one can see that income disparities are persistent even if labor income is fully equalized (with a set at 1 for everyone). This is indeed what motivates our study of mobility. The further insight, which is also borne out by available empirical evidence, is that when assets can be bequeathed, income is likely to be more closely correlated intergenerationally than labor earnings.

$$y_{i,t+1} = v_t + r_{t+1} S(y_{it}) \quad (31)$$

As Piketty illustrates, initial differences in asset holdings will only be equalized in the long run if savings are a concave fraction of income, labor earnings are the same for every individual, and fertility behavior is the same for every family. A departure from any of these assumptions (by positing that families endogenously determine their bequests, for instance, or that there is a degree of heterogeneity in family size which affects the potential for leaving bequests) leads to persistent wealth inequalities in the

long run. In the case of differential fertility behavior, for example if dynasty i has $(1+n_i)$ children, then the transition for individual wealth is:

$$w_{i,t+1} = \frac{S(v_t + rw_t) + (1-\delta)w_t}{1+n_i} \quad (32)$$

Consequently, different dynasties may end up having different steady states of wealth. This will be the case if poor families tend to have more children.

Piketty argues that, within this framework, inheritance taxation is the most effective policy instrument. Of course, in any but the simplest model, such taxation would have to be a permanent feature of the economy.⁴¹ The distortionary consequences of this type of taxation clearly depend on the driving forces behind inheritance. If bequests are mainly accidental by-products of precautionary saving, then taxing them has no obvious cost. However, if they are the product of intergenerational altruism, the distortionary effect is more likely to be negative. The strength of this conclusion, however, depends on particular assumptions as to the precise reason why bequests are left. It makes a difference, for instance, whether it is bequests per se that affect parental utility, or whether there is a “dynastic utility function” that each generation maximizes. In fact, inheritance taxation can potentially increase long-run wealth disparities. These results lead Piketty to quote Mulligan (1997): “much more research (...) [is] necessary to arrive at a strong conclusion regarding the unimportance of taxes for intergenerational mobility”.

A second (efficient) reason for persistent inequality is family inheritance of ability and tastes. Piketty uses a very simple reduced-form model, similar to the one employed by Becker and Tomes (1979), and stresses the fact that intergenerational transmission of ability works at least as much through family environment and learning as it does through genetic transmission. As we have repeatedly argued, this intergenerational transmission reduces mobility. Piketty considers the question of what sort of policies could break down this transmission. He concludes that if capital markets are indeed perfect and this persistence of ability is therefore efficient, then any public intervention would compound the problem, as it would make high-ability families even better off by covering part of the cost of human-capital acquisition, while overall levels would not be changed. Again, even if one were to argue that redistribution in this context is desirable on purely ethical grounds (based on an appeal for equality of opportunity, for instance), the consequences of a redistributive policy are far from obvious, since we do not know much about the magnitude of the incentive effects of earnings taxes.

A similar concern is raised by persistence in social outcomes that is explained by differences in tastes, given the fact that individuals measure themselves against reference groups and low-income individuals have low-income reference groups. In this and other

⁴¹ For a non-permanent tax rate, we would need a situation where a one-time equalizing tax leaves the economy in a steady state. But, of course, this means that equal wealth was a steady state to begin with, and this is true only in models with few sources of heterogeneity in economic behavior.

similar models, where individuals are somehow responsible for their own misfortune, it is quite hard to come up with acceptable policy interventions. In a sense, it implies that perfect social mobility is not efficient.

A related connection between efficiency and mobility is given when parental background acts as an input in the production of human capital and the acquisition of skills. That is, higher-income families make better use of the resources at their disposal to educate their children. This is a feature of Hassler, Rodríguez Mora and Zeira (2003), Raut (1996), and Benabou (2000). The relevant tradeoff is then between the fact that educating the children of the rich is relatively cheaper and the potential output that is lost because high-ability children born into poor families do not receive an education that fully develops their ability. This is true under imperfect capital markets, but also under no liquidity constraints: if ability is correlated across generations, then reducing the educational advantage of high-income families helps make low-income/high-ability dynasties into high-income/high-ability families. This tradeoff is at the heart of Benabou's argument.

Hassler and Mora (2000) make a related point, noting that the advantage enjoyed by rich entrepreneurial parents is fully inefficient, in that society would be better off in the long run if these advantages were not strong determinants of economic position. However, it remains true, even in this model, that the children of parents with a strong informational advantage will be entrepreneurs. The downside is that this kind of informational advantage is most likely to be relevant in situations marked by low growth and low levels of long-run output.

Overall, then, efficiency does not seem to be determined by mobility – rather, it seems to be a jointly determined result. The relation is at its strongest when there are liquidity constraints that prevent agents from making efficient investments in their stock of human capital. However, there are constraints on what level of social mobility is efficient, as, even in contexts of perfect availability of financing opportunities, there would be some intergenerational persistence, and in many cases the appeal to policies that increase social mobility stems from the belief that it is in itself a desirable feature of the income process.

Political Outcomes

So far, we have taken an approach in which redistribution is analyzed insofar as it relates to efficiency, without focusing on the degree to which it is more or less likely that redistribution will come about. This aspect is the main contribution of the models we will analyze in this section.

The first finding, which has been widely cited, points to a direct connection between income mobility and redistribution when individuals care about their own interests only. This is the formalization by Benabou and Ok (2001a) of the “Prospect of Upward Mobility” (POUM) hypothesis. This is a customary answer to the question as to why low-income voters may oppose substantial redistribution and is based on the idea that they may “take into account the fact that they, or their children, may move up in the income

distribution and therefore be hurt by such policies”. Three assumptions underlie this hypothesis:

- Policies are persistent, so that what is decided today is expected to remain in place for a number of periods.
- Agents are not too risk-averse, since otherwise they would favor the income-smoothing possibilities provided by redistribution.
- Individual families who are currently poorer than average (and therefore would be among the beneficiaries of redistribution) expect to become richer than average.

The usual assumption has been that this last condition cannot exist in situations where agents have rational expectations. This paper shows that this is not the case.

The setting is one where income is strictly exogenous and follows some given stochastic process. Agents are allowed to vote on a flat, proportional tax rate, the proceeds of which accrue to each citizen in equal proportion. In a static setting, we obtain the traditional median-voter result, where, if the voter with the median income is poorer than the mean, her preferred tax rate will be the policy outcome of electoral competition between two parties. In this context, where taxation has no deadweight loss, the preferred tax “rate” is actually complete expropriation.

This result breaks down when a temporal dimension is given to the problem. The simplest way to see the underlying mechanism is to analyze what happens when income is deterministic. The transition function $y_{t+1} = f(y_t)$ specifies how (pre-tax) income changes from one period to the next. Transition functions are assumed to be increasing.

The first case to be analyzed focuses on a two-period scenario where individuals vote at date 0 on a tax scheme to be implemented at date 1. If we call F_t the cumulative income distribution at date t and μ_t the average income at date t , an individual with time-zero income y will prefer a laissez-faire approach (a zero tax-rate, r_0) to full redistribution (r_1) if her date-1 income will be higher than the date-1 average.

$$f(y) > \int f dF_0 = \mu_1 \tag{33}$$

From Jensen’s inequality, if f is strictly concave, the agent with a mean income in period 0 will oppose redistribution.

$$f(\mu_0) = f\left(\int y dF_0\right) > \int f dF_0 = \mu_1 \tag{34}$$

At the same time, the individual with the lowest initial income will surely favor redistribution. This underlies the result that there is a critical income level **below the average income** such that individuals with higher income oppose redistribution and those with lower income favor redistribution. The main result, though, is that, for any income

distribution F_0 , a laissez-faire policy will be preferred by the median voter under some concave transition function f^* . In fact, this will be the case for all functions that are more concave than f^{*42} .

A second step, which highlights the role of persistence in government policies, assumes that agents choose at time 0 a tax scheme that will be in place during periods 0 through T . Agents care about their present discounted income stream. Given a transition function f and initial income y , an individual will prefer a laissez-faire policy if:

$$\sum_{t=0}^T \delta^t f^t(y) > \sum_{t=0}^T \delta^t \mu_t \quad (35)$$

where f^t is the result of applying f , t successive times. The results are similar. In the first place, there is, as before, an income level y^* such that people with higher incomes will favor a laissez-faire approach and those with lower incomes will want redistribution. This support for a laissez-faire policy changes depending on two features of the model: the transition function, and the time horizon of the policy change (which includes the number of time periods and the discount factor). Indeed, longer temporal horizons and higher discount factors promote support for a laissez-faire response. Once again, this effect operates through the convexity of transition functions: as the time horizon grows, transition functions are composed, generating f^1, f^2, \dots, f^T , and these functions are increasingly concave.

These observations carry over to the case where income is stochastic. In this case, the analog of the concavity of the transition function is the existence of a stochastic process for individual income which is concave *in expectation*. That is, the expected value of income tomorrow, given today's income, is a concave function of the latter. In this case, again, for every distribution we can find a stochastic income process such that the median voter opposes redistribution. This case of the POUM hypothesis posits agents with completely rational expectations and is consistent with very unequal *realized* income distributions in each period.

Although there is no model of the underlying income-generating process, this result captures one of the essential consequences of social mobility in the political arena. Social mobility can be seen as implying concave transition functions (this idea is explored in Benabou and Ok, 2001b). It reduces political conflict between the pivotal voter and the rich, thereby reducing equilibrium redistribution.⁴³

⁴² f is "more concave" than g if it is obtained from g by applying an increasing, strictly concave transformation.

⁴³ In a related result, Harms and Zink (2003) develop a model where this result (the median voter preferring no redistribution) follows from an explicit economic choice problem, where individuals vote on a tax on their earnings and choose their skill level. If the median voter can finance his education and there is a skill premium, he may expect to lose from redistribution when he enters the labor market. In that case, the prospect of increasing the distance between the median voter and the poor – even if there is no real exchange mobility – reduces his support for redistribution.

As we have already mentioned, Benabou (1996) summarizes several models in which political variables affect mobility. Basically, in the presence of imperfect capital markets, redistribution is economically desirable. Whether it can be expected to be the outcome of the political process, however, depends on what group in society has the most political weight. That is, he allows for a departure from the one person = one vote paradigm. This departure is summarized by the distance, in terms of wealth, between the median agent (who we would expect to be pivotal in a traditional setting) and the pivotal voter (whose preferences are represented by the political system).

Benabou finds that if political power is held disproportionately by a wealthy minority, then redistribution is inefficiently low (in the sense that increasing redistribution unambiguously increases efficiency), as is mobility in the asset distribution. Moreover, there are multiple steady states, with intertemporal efficiency being lower in less redistributive, more inegalitarian steady states with lower mobility. On the other hand, if the political system has some degree of pro-poor bias, there will be only one steady state, with high redistribution.

The main idea is, therefore, that the degree of distributional conflict in a society affects its equilibrium level of redistribution. Increasing social mobility reduces this conflict and thus curbs the demand for redistribution. In contexts where redistribution is inefficient, this would be a desirable result. In environments where redistribution is efficient, however, it is often the case that mobility is not exogenous to the political outcome, but rather is endogenously determined within the model. In this case, it is far from obvious that reducing redistribution is in any way desirable.

7. THE PERCEPTION OF MOBILITY

In the models we discuss in this section, individual income depends on parental background, effort, and luck. With no market imperfections, redistributive schemes based on income taxes are inefficient, as they distort the labor/leisure choice. Individuals' *beliefs* about the income process affect their preferences regarding the amount of redistribution. At the same time, the existence of a redistributive scheme changes individual incentives to exert effort. Low transfers provide an incentive for work, which in turn lowers the impact of luck. High transfers, in contrast, lower the incentive to exert effort, with the consequence that luck plays a more important role in individual outcomes. As a consequence, individual beliefs are self-fulfilling.

We first draw attention to the model developed by Piketty (1995). His intention is to understand why it is that, although by and large poor people support redistribution more than the rich do, family background affects these preferences. Indeed, most of the available data indicate that poor people who have experienced downward mobility (that is, their parents were relatively rich) *oppose* redistribution more than the average poor individual does, whereas high-income people from a poor background tend to support redistribution more than the average rich person does. The point of the paper is to show how this can be brought about, not by a difference in redistributive aims (all agents will

have the same goals), but rather by differences in income trajectories which inform agents' beliefs about the world.

There is a continuum of dynasties. At each point in time there is one generation that votes on taxes for the following period, works, and pays the taxes specified by the preceding generation. There are two possible pre-tax income levels: y_1, y_0 (with $y_1 > y_0$), which are randomly assigned in the population, with a probability distribution that depends on two factors: family background (children of parents who obtained a high income in the previous period have a higher probability of having a high income than do children of poor parents), and effort, where all individuals share the same cost of effort, C .

$$\begin{aligned}
\Pr(y_{it} = y_1 | e_{it} = e, y_{i,t-1} = y_0) &= \pi_0 + \theta e \\
\Pr(y_{it} = y_1 | e_{it} = e, y_{i,t-1} = y_1) &= \pi_1 + \theta e \\
\theta > 0, \pi_0 < \pi_1 & \\
U_{it} &= y_{it} - C(e_{it}) \\
C(e) &= e^2 / 2a, \quad a > 0
\end{aligned} \tag{36}$$

There are two relevant parameters in this model: the advantage derived from being born in a rich family, $\pi_1 - \pi_0$, and the degree to which effort affects the probability distribution over outcomes, θ . These are unknown to the agents in the model.

Agents vote for their preferred tax rate after they choose their effort level and after their income is realized. Voters can be divided into four groups based on their background and current situation:

Table 13: Voting Groups

		Parents' Income	
		y_0	y_1
Agent's Income	y_0	Stable Low Income	Downward Mobile
	y_1	Upward Mobile	Stable High Income

The tax redistribution scheme operates by taking income from those with high realizations and making transfers to those with low realizations. All agents share the same objective function: they would like to minimize the effective disadvantage suffered by children from poor families – that is, they want to maximize their expected welfare.

Agents choose their effort levels, and this choice is distorted by setting a high tax rate. In fact:

$$e(\tau_{t+1}, \theta) = a\theta(1 - \tau_{t+1})(y_1 - y_0) \tag{37}$$

Taking this into account, the optimal tax rate for an individual is:

$$\tau_{t+1}(\pi_1 - \pi_0, \theta) = H_{t+1}(\pi_1 - \pi_0) / a(y_1 - y_0) \theta^2 \quad (38)$$

where H is the proportion of individuals who receive y_1 .

The tradeoff which agents face when choosing their preferred tax rate is that high taxes discourage effort and lower total output. If all agents shared the same beliefs, however, all would choose the same rate: if they believe income differences are mainly driven by background, they will favor high redistribution; if, on the other hand, they believe effort is the main driving force, they will oppose redistribution due to its incentive costs, since it will affect equilibrium effort and thus the long-term equilibrium income distribution for children from poor families. This is what is spelled out in equation (38).

Note that the only source of disagreement in this economy is about the fundamentals of the model: there is no conflict of interest in the strict sense of the word. How, then, are beliefs about the true state of the world formed? Piketty explores the implications of purely individual rational learning – that is, starting from a given prior and observing the experience of their parents, agents update their beliefs using Bayes’ rule. In turn, these posterior beliefs guide their choice of action.⁴⁴

The crux of the argument has to do with the fact that the restrictions placed by Bayes’ rule on individual learning are very lax. Starting from a prior, Bayesian updating will put more weight on posterior beliefs which are consistent with actual experience. In other words, if an individual experiences high mobility, posterior beliefs will be such that upward mobility, given the agents’ choices and priors, is most likely. However, there is no necessary link between this and the actual parameters of the model; interpretations will vary depending on beliefs, as we show in Table 14.

Table 14: Interpretation of Upwardly Mobile Trajectory

Prior Beliefs	Action Taken	Bayesian Update
High $\pi_1 - \pi_0$, Low θ	Low Effort	Large role of chance
Low $\pi_1 - \pi_0$, High θ	High Effort	Large role of effort

Piketty shows that there is a continuum of possible dynastic types in a steady state. The two ends of this spectrum correspond, on the one hand, to those who believe effort is more important and who therefore oppose redistribution, exert a high level of effort and have high average incomes (the “right-wing” dynasties) and, on the other, those who believe luck plays a more central role and who therefore support redistribution, exert a

⁴⁴ The optimal actions depend both on individual beliefs about the true model and on beliefs about what other people believe, since this will inform each individual’s response to tax rates. Piketty shows that, at the extremes, where people assume that everybody else shares their beliefs, or where they know the exact distribution of beliefs, the preferred tax rate still increases as an individual’s belief in the importance of luck rises, and decreases as individual’s belief in the significance of effort increases, which is what Piketty needs to arrive at the main result.

low level of effort and have low average incomes (the “left-wing” dynasties). A one-time experience of downward mobility despite a high level of effort, or of upward mobility despite a low level of effort, will not fully counterbalance these long-run dynastic beliefs, since every agent is aware of the “sampling effects” (i.e., of the fact that there is, after all, some degree of random luck).

In this model, therefore, conflicting views about social mobility and the incentive effects of redistribution are perfectly compatible with the absence of any deep conflict in terms of beliefs about what a fair social outcome is. It shows how history and personal mobility experience can lead to divergent views about the world and, in particular, how all the views can be consistent with one “true” state of the world.⁴⁵ When agents do not have an incentive to engage in widespread experimentation, this outcome is indeed likely. A second result is that there is no necessary relation between the optimal level of redistribution for the true parameter and the level which obtains in the long-run equilibrium. This model therefore underscores how fragile normative recommendations are and thus, how sensitive they are to the introduction of reasonable forms of uncertainty about the true model.

A related sort of connection between beliefs and mobility is introduced in the work of Benabou and Tirole (2006). They start from the fact that, as mentioned earlier, the extent to which people believe luck is an important determinant of individual outcomes is highly correlated with their preferred redistribution levels. These authors subsequently present a model with two equilibria: one in which people believe the world is fair and that effort is consequently an important determinant of success in life, and they therefore oppose high levels of redistribution; and one in which people believe the world is mostly unfair and hence support redistribution and exert low levels of effort. This is compatible with the actual world being either “fair” or “unfair”; thus, in one of the equilibria, people are consistently wrong.

This model starts from a setup that is very similar to the one found in Piketty (1995). The main difference lies in the condition that the first generation – in a two-period game – chooses an optimal belief, given the beliefs of others and the redistribution level that will follow in equilibrium. A second departure is that tax rates are voted upon selfishly.

The setup is the following. There is a continuum of agents, indexed by i . As in Piketty (1995), income is the result of an interaction among innate advantages, effort, and luck.

$$\begin{cases} \Pr(y_i = 1) = \pi^i + \theta e^i \\ \Pr(y_i = 0) = 1 - (\pi^i + \theta e^i) \end{cases} \quad (39)$$

As before, e is effort and π^i reflects parental background: π^i takes on a high value for a fraction $\varphi < 1/2$, and a low value for the rest.

⁴⁵ These are empirically important, as is demonstrated by the evidence presented in Graham and Pettinato (2002) and in Ravallion and Lokshin (2000) and by the references in Piketty (1995).

At date zero, parents receive a signal which noisily informs them about the true value of θ . They subsequently choose what to tell their children (or themselves, in a different interpretation). Once they or their children have the revised signal (which they know to be potentially distorted), at time 1, they vote on the tax rate that would maximize their own welfare and choose their optimal effort level (in the intergenerational interpretation, this is done by their descendents). Lastly, their income is realized in the following period, and redistribution and consumption take place.

An individual's preferred tax rate has three components:

- The larger the innate differences in income prospects are ($\pi^i - \bar{\pi}$, the difference with respect to the average), the higher the rate is.
- There is a POUM effect: Optimistic individuals who expect a higher return to effort also expect to move up the income distribution, and this lowers their preferred tax rate.
- If individual effort is distorted downward because of the lag between effort and consumption, the optimal tax will be lower for individuals who see this policy channel as a way of correcting that distortion.

The authors show that two distinct equilibria arise. In one, a large proportion of the population chooses to believe that the world is fair. Therefore, these people vote for a low tax rate. This justifies their belief that effort pays, since, in equilibrium, low-income individuals are more than proportionately likely to have exerted a low level of effort. In the other, a majority chooses to believe that the world is intrinsically unfair and votes for a high degree of redistribution. This reduces individual incentives to exert effort, with the result that, in equilibrium, people with high pre-tax incomes are disproportionately from high-income backgrounds. Note that both equilibria are possible regardless of the true state of the world. The story told in the paper is that of a true world in which effort does not pay, but in which people may choose to believe it does because they need to “believe in a just world”.

A similar multiple-equilibria model with differences in beliefs that lead to differences in preferences over redistribution is that of Alesina and Angeletos (2005). In their case, however, these are truly rational expectation equilibria, with self-fulfilling beliefs.

Agents live for two periods and engage in a productive activity in each of them. Taxes are set at some point in the middle of an agent's life, which is parameterized in the model. Pre-tax income is determined by individual ability (A), investment during the first period of life (k), and effort in the second period (e), as well as by a certain amount of “noise” (η). Both ability and noise are random and i.i.d. across agents. α reflects the point in life when taxes are set (it reflects what investments are sunk when the tax rate is decided).

$$y_i = A_i \left[\alpha k_i + (1 - \alpha) e_i \right] + \eta_i \quad (40)$$

The government imposes a flat tax rate on income and shares the proceeds equally among all agents.

People have an intrinsic demand for social justice and fairness (they would like to reduce the impact of sheer luck in individual income) but at the same time would like to reward individual effort and talent.⁴⁶ Accordingly, they consider a fair income to be:

$$\hat{y}_i = A_i [\alpha k_i + (1 - \alpha) e_i] \quad (41)$$

Thus, there is a level of “socially fair” income for given characteristics of individuals, namely, that which an individual would obtain with no taxes and no noise. Overall social fairness is given by the extent to which after-tax incomes differ from this ideal. Social *injustice* is a weighted average of the “variance decomposition” of income inequality.

$$\Omega = \tau^2 \text{Var}(\hat{y}_i) + (1 - \tau^2) \text{Var}(y_i - \hat{y}_i) \quad (42)$$

If income were exogenous, and minimizing Ω were the only goal, then the preferred level of redistribution would be:

$$\frac{1 - \tau}{\tau} = \frac{\text{Var}(\hat{y}_i)}{\text{Var}(y_i - \hat{y}_i)} \quad (43)$$

Which is decreasing in the “signal-to-noise” ratio: the higher the variability in the noise component, the more desirable redistribution would be.

However, this ratio is endogenous in equilibrium, since equilibrium effort depends on the expected tax rate. A high expected tax rate reduces the incentive to engage in productive effort, thereby increasing the importance of noise in realized income and thus providing a rationale for voting for high levels of redistribution. The opposite holds true for a low expected tax rate. The consequence is that both kinds of equilibrium are possible and compatible with rational expectations. Thus, concerns for fairness are compatible with both high and low levels of redistribution in equilibrium.

In all three papers, we see that beliefs about the true nature of the income process may drive economically powerful forces regardless of whether they are correct (and, in a sense, with multiple rational-expectations equilibria, there is no “outside world” that is independent from beliefs). This relates to the POUM hypothesis, in that it may well be that *beliefs* about the existence of social mobility – or equality of opportunity, as in the papers that we have mentioned – drive demands for redistribution. As we will discuss in the following section on the empirical evidence, the studies conducted thus far indicate that this is indeed the case.

⁴⁶ Note the relation with the concepts of meritocracy and of equalization of opportunities by reducing variability in outcomes as outlined by Conlisk (1974).

8. CONCLUSIONS: A roundup of policy implications

In this discussion we explored a series of definitions of social mobility and settled on one according to which situations with high mobility are considered to be those where the relative economic status of an agent is not dependent on starting conditions such as family background or connections.

As is detailed in Table 15, most of the models we have studied suggest there is room for addressing social immobility through some sort of education policy. A question remains, though: What are the limits on education's impact? In other words: How far can we possibly go in the direction of social mobility?

Table 15 adds a policy intervention that, while related to education, is not explicitly addressed by the models we have surveyed but that is implied by their treatment of "ability" and its intergenerational transmission: early childhood interventions. This is motivated by a strand of research that departs from the unidimensional conception of "ability" embraced by most economists and scrutinizes the process of skill formation. We provide an overview, based on the survey by Cunha, Heckman, Lochner and Masterov (CHLM) (2005).

This section suggests that a combination of different policies has the potential to improve social mobility in the short and long terms.

In the short run, policy should deal with the fact that some degree of immobility is unavoidable, given the long-lasting consequences of poor parental environments. However, efforts should be made to ensure that individuals from poor households who have managed to acquire skills have the opportunity to develop them further. The steps to be taken in this direction should include the creation of mechanisms for financing higher education and discouraging discrimination and nepotism in labor markets, possibly by increasing market competition and transparency in State hiring.

Table 15: Policy Implications

Policy	Effect	Model	Possible Limitations
Public School System	Reduces the need for educational finance. Makes education more uniform.	C74 HMZ03 OW98	Takes ability (and its transmission) as given
Early Childhood Intervention	Reduces intergenerational transmission of skills, or parental advantage in education, by directly addressing the mechanisms of ability formation.	B00 BT79 C74 CHLM05 HMZ03	Political feasibility
Education Loans	Provides financing to those who would profit from education but face liquidity constraints.	B00 HMZ03 OW98	Takes ability (and its transmission) as given
Inheritance Tax	Reduces income inequality and immobility that stems from financial bequests.	BT79 P95	Endogeneity of bequests and other investments in children limits the feasibility of full equalization
Promotion of Competitive Markets	Reduces the impact of noncognitive parental advantages (race, connections) by making actual productivity matter in hiring decisions. Makes entrepreneurial activities more complex and dynamic.	BT79 C74 HM00	Political feasibility Transition

B00: Benabou, 2000
BT79: Becker and Tomes, 1979
C74: Conlisk, 1974

CHLM05: Cunha, Heckman, Lochner and Masterov, 2005
HM00: Hassler and Mora, 2000

HMZ03: Hassler, Rodríguez Mora and Zeira, 2003
OW98: Owen and Weil, 1998
P95: Piketty, 1995

For long-run results, however, the most promising policies appear to be early intervention programs, which can break down the dependence between the development of socially productive skills and parental background. These initiatives hold out the promise of greatly improving both social mobility and economic efficiency at one and the same time. Efforts should be made to design interventions which take into account local conditions, along with carefully designed evaluations, so that lessons can be learned in terms of how such initiatives are operated, how they can be improved, and how they can be extended into large-scale programs.

References

- Alesina, Alberto, and George-Marios Angeletos. 2005. "Fairness and Redistribution." *American Economic Review* 95, no. 4: 960-960.
- Arrow, Kenneth, Bowles, Samuel, and Steven, eds. Durlauf. 2000. *Meritocracy and economic inequality* Princeton, N.J.: 2000.
- Becker, Gary S., and Nigel Tomes. 1979. "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility." *Journal of Political Economy* 87, no. 6: 1153-1189.
- Behrman, Jere R. 2000. "Social Mobility: Concepts and Measurement." In *New markets, new opportunities? Economic and social mobility in a changing world*, 69-100. Washington, D.C.: 2000.
- Benabou, Roland, and Efe A. Ok. 2001a. "Social Mobility and the Demand for Redistribution: The Poupou Hypothesis." *Quarterly Journal of Economics* 116, no. 2: 447-487.
- Benabou, Roland, and Efe A. Ok. 2001b. "Mobility as Progressivity: Ranking Income Processes According to Equality of Opportunity."
- Benabou, Roland, and Jean Tirole. 2006. "Belief in a Just World and Redistributive Politics." *Quarterly Journal of Economics* 121, no. 2: 699-746.
- Benabou, Roland. 1996. "Inequality and Growth." In *NBER macroeconomics annual 1996*, 11-74. Cambridge and London: Harvard University Press, 1996.
- Benabou, Roland. 2000. "Meritocracy, Redistribution, and the Size of the Pie." In *Meritocracy and economic inequality*, 317-339. Princeton, N.J.: Princeton University Press, 2000.
- Bowles, Samuel, and Herbert Gintis. 2002. "The Inheritance of Inequality." *Journal of Economic Perspectives* 16, no. 3: 3-30.
- Conlisk, John. 1974. "Can Equalization of Opportunity Reduce Social Mobility?." *American Economic Review* 64, no. 1: 80-90 .
- Cunha, Flavio, Heckman, James J., Lochner, Lance, and Dimitriy V. Masterov. 2005. "Interpreting the Evidence on Life Cycle Skill Formation." , NBER.
- Dworkin, Ronald. 1981. "What is Equality?." *Philosophy and Public Affairs*, 10:3, 185-246, 283-345.

- Feldman, Marcus W., Otto, Sarah P., and Freddy B. Christiansen. 2000. "Genes, Culture, and Inequality." In *Meritocracy and economic inequality*, 61-85. Princeton, N.J.: 2000.
- Fields, Gary S. 2000. "Income Mobility: Concepts and Measures." In *New markets, new opportunities? Economic and social mobility in a changing world*, 101-132. Washington, D.C.: 2000.
- Flynn, James (2000) IQ trends over time: intelligence, race, and meritocracy. In K. Arrow, S. Bowles, & S. Durlauf (eds.), *Meritocracy and economic inequality*, Princeton, NJ, Princeton University Press 35-60
- Graham, Carol, and Stefano Pettinato. 2002. "Frustrated Achievers: Winners, Losers and Subjective Well-Being in New Market Economies." *Journal of Development Studies* 38, no. 4: 100-140.
- Hansson, Sven Ove. 2004. "What Are Opportunities and Why Should They Be Equal?." *Social Choice and Welfare* 22, no. 2: 305-316.
- Harms, Philipp, and Stefan Zink. 2003. "Eating the Rich vs. Feeding the Poor: Borrowing Constraints and the Reluctance to Redistribute." *Public Choice* 116, no. 3-4: 351-366.
- Hassler, John, and Jose V. Rodriguez Mora. 2000. "Intelligence, Social Mobility, and Growth." *American Economic Review* 90, no. 4: 888-908.
- Hassler, John, Rodriguez Mora, Jose Vicente, and Joseph Zeira. 2003. "Inequality and Mobility." C.E.P.R. Discussion Papers, CEPR Discussion Papers: 2497.
- Herrnstein Richard and Murray Charles. 1994. "Bell Curve: Intelligence and Class Structure in American Life". New York: The Free Press.
- Hild, Matthias, and Alex Voorhoeve. 2004. "Equality of Opportunity and Opportunity Dominance." *Economics and Philosophy* 20, no. 1: 117-145.
- Jencks, Christopher and Laura Tach. 2005. "Would Equal Opportunity Mean More Mobility?" John F. Kennedy School of Government, Faculty Research Working Paper Series RWP05-037, May 2005.
- Korenman, Sanders, and Christopher Winship. 2000. "A Reanalysis of The Bell Curve: Intelligence, Family Background, and Schooling." In *Meritocracy and economic inequality*, 137-178. Princeton, N.J.:, 2000.
- Owen, Ann L., and David N. Weil. 1998. "Intergenerational Earnings Mobility, Inequality and Growth." *Journal of Monetary Economics* 41, no. 1: 71-104.
- Phelan, Christopher. 2006. "Opportunity and Social Mobility." *Review of Economic Studies* 73, no. 2: 487-504.

Piketty, Thomas. 1995. "Social Mobility and Redistributive Politics." *Quarterly Journal of Economics* 110, no. 3: 551-584.

Piketty, Thomas. 2000. "Theories of Persistent Inequality and Intergenerational Mobility." In *Handbook of income distribution. Volume 1*, 429-476. Handbooks in Economics, vol. 16., 2000.

Raut, Lakshmi K. 1996. "Signalling Equilibrium, Intergenerational Mobility and Long-Run Growth."

Ravallion, Martin, and Michael Lokshin. 2000. "Who Wants to Redistribute? The Tunnel Effect in 1990s Russia." *Journal of Public Economics* 76, no. 1: 87-104.

Rawls, John. 1971. "A Theory of Justice". Cambridge, Mass.: Belknap Press of Harvard University Press.

Roemer, John E. 1996. "Theories of distributive justice". Cambridge and London: Harvard University Press.

Roemer, John E. 1998. "Equality of Opportunity". Cambridge, MA: Harvard University Press.

Roemer, John E. 2000. "Equality of Opportunity." In *Meritocracy and economic inequality*, 17-32. Princeton, N.J.: 2000.

Roemer, John E. 2004. "Equal Opportunity and Intergenerational Mobility: Going beyond Intergenerational Income Transition Matrices." In *Generational income mobility in North America and Europe*, 48-57. Cambridge; New York and Melbourne: Cambridge University Press.

Sen, Amartya. 2000. "Merit and Justice." In *Meritocracy and economic inequality*, 5-16. Princeton, N.J.: 2000. .

Sugden, Robert. 2004. "Living with Unfairness: The Limits of Equal Opportunity in a Market Economy." *Social Choice and Welfare* 22, no. 1: 211-236.

Swift, Adam. 2005. "Justice, Luck, and the Family: The Intergenerational Transmission of Economic Advantage from a Normative Perspective." In *Unequal chances: Family background and economic success*, 256-276. New York: 2005.

Young, Michael. 1958. *The Rise of the Meritocracy, 1870-2033: An Essay on Education and Equality*.