Social Mobility and Redistributive Politics
Author(s): Thomas Piketty
Source: *The Quarterly Journal of Economics*, Vol. 110, No. 3 (Aug., 1995), pp. 551-584
Published by: Oxford University Press
Stable URL: http://www.jstor.org/stable/2946692
Accessed: 22-11-2016 17:09 UTC

THE
QUARTERLY JOURNAL
OF ECONOMICS

Vol. CX August 1995 Issue 3

SOCIAL MOBILITY AND REDISTRIBUTIVE POLITICS*

THOMAS PIKETTY

Just like economists, voters have conflicting views about redistributive taxation because they estimate its incentive costs differently. We model rational agents as trying to learn from their dynastic income mobility experience the relative importance of effort and predetermined factors in the generation of income inequality and therefore the magnitude of these incentive costs. In the long run, "left-wing dynasties" believing less in individual effort and voting for more redistribution coexist with "right-wing dynasties." This allows us to explain why individual mobility experience and not only current income matters for political attiitudes and how persistent differences in perceptions about social mobility can generate persistent differences in redistribution across countries.

I. INTRODUCTION

This paper develops a rational-learning theory[1] of redistributive politics seeking to explain important stylized facts concerning the effect of social mobility on both individual political attitudes and aggregate political outcomes.

The idea that social mobility plays a crucial role in shaping political attitudes (in particular toward redistribution) has a long history in the social sciences. De Tocqueville [1835] first stressed

*I am grateful to seminar participants at the Massachusetts Institute of Technology, Harvard University (Department of Economics and Kennedy School of Government), Columbia University, Boston University, Princeton University, Tarragona University, Brown University, New York University, Yale University, University of Minnesota, Northwestern University, Madrid, and in particular to Daron Acemoglu, Roland Benabou, Abhijit Banerjee, and two referees for their many comments and encouragement.

1. That is, as we understand it, a theory precisely describing not only the values and preferences individuals are promoting and the institutions aggregating their actions, but also the information sets they are exposed to and the way they learn from them. This differs from standard rational-choice theories, as well as from most sociological "explanations" of the effect of one's mobility experience on one's political attitudes.

the idea that the difference in attitudes toward redistribution between Europe and the United States could be explained by presumed differences in mobility rates. Since then, many authors have followed this line to explain the absence of any strong socialist movement in the United States.[2] On the other hand, comparative empirical studies of social mobility rates have long demonstrated the absence of any significant difference between industrial nations.[3] Lipset and Bendix [1959] and Lipset [1966, 1977, 1992] have repeatedly suggested that persistent differences between European and U. S. redistributive politics may be due to persistent differences in popular beliefs about social mobility.[4]

But social mobility is known to have crucial effects at the individual level as well. Although current income is positively correlated with voting attitudes toward redistribution (higher-income groups vote less for left-wing redistributive policies), the correlation is much less than one, and most of all the residual is strongly correlated with past income: upwardly or downwardly mobile voters always exhibit an intermediate position between stable low-income and high-income voters;[5] that is, Table I summarizes the typical voting patterns observed across time and industrial democracies with a remarkable degree of stability.[6] That is, seven out of ten lower-class voters born in the lower class typically vote for left-wing parties, against less than one-half of lower-class voters born in the middle class. Similar qualitative results are obtained in survey studies trying to isolate the specific redistributive component of political attitudes.[7] From this matrix it would appear that parents' income class determines one's political atti-

2. Among which Marx [1852], Sombart [1906], and Petersen [1953].

3. See, e.g., Lipset and Bendix [1959] and Erikson and Goldthorpe [1985, 1992].

4. "What explains the contrast in the political values and allegiances of American workers with those of other democratic nations? (. . .) the belief system concerning class rigidities stemming from varying historical experiences (· · ·) seems much more important than slight variations in rates of mobility" [Lipset 1992, pp. xx–xxi].

5. A few studies found that upwardly mobile agents are on average more right-wing than stable middle-class (mostly in the United States). However, later studies have shown that this was nonrobust (see Thomson [1971]), and this thesis has apparently been abandoned.

6. See, e.g., Abramson [1973], Thomson [1971], Boy [1980], and Cherkaoui [1992]. This sociology/political science literature usually cuts the society in half: lower-class, manual occupations; and middle-class, nonmanual occupations. Although this is highly rudimentary, more sophisticated studies with more than two income groups confirm the basic findings (see Turner [1992]), which casts serious doubts on the simple measurement error explanation for these findings. Table I does not show up simply because upper-half agents whose parents were in the upper half are in fact richer than other upper-half agents.

7. Otherwise, one could argue that not only redistribution is involved when voting for some political party. The point is that the same picture survives when

TABLE I
PERCENTAGE OF VOTES FOR LEFT-WING PARTIES AS A FUNCTION OF INDIVIDUAL MOBILITY EXPERIENCE

		Respondent's income	
		Low income	High income
Parents income	Low income	72%	38%
	High income	49%	24%

(Average matrix for six countries: Germany 1953, Britain 1962, United States, 1953, Finland 1949, France 1966, Norway 1957. Standard deviation = 5.78%.
Source. Cherkaoui [1992, p. 189].

tudes almost as much as one's current income, whereas straight economic rationality should imply that only current income and not past family income[8] should determine one's interests in redistribution, as in the standard public-choice models of redistributive politics.[9]

Our primary objective is to provide a common framework to account for these various stylized facts and, by doing so, to develop a new conceptual framework to think about redistributive politics. The basic idea of our theory is that voters may develop conflicting views about redistribution not because they are maximizing different objective functions but rather because through their various mobility experiences they (rationally) happen to learn and to believe different things concerning the incentive costs of redistributive taxation for society as a whole. That is, our modeling exercise consists of describing rational agents as having a priori the same distributive goals and as trying to learn from their income trajec-

survey studies directly ask the agents what they think about inequality and redistribution. See the studies edited by Turner [1992].

8. Unless one assumes that there are strong "dynastic permanent-income" effects. That is, one could reconcile Table I with a simple model of selfish, forward-looking, and well-informed voters only by assuming that ability exhibits sufficient memory along dynastic histories, so that kids' income prospects depend sufficiently on the grandparents' achievements for a given parental income. We feel that such an alternative explanation would eventually have to deal with the formation of beliefs about such dynastic transmission processes, which would bring it very close to the theory developed in this paper.

9. See, e.g., Mueller [1989] for the standard economic models of redistributive politics. Aside from the stylized facts mentioned above (which by nature these theories cannot accommodate), the basic prediction according to which a lower median-income/mean-income ratio should result in higher redistribution does not seem to be particularly consistent with the evidence (see, e.g., Perotti [1994] and subsequent references). See Piketty [1993] for an alternative viewpoint on the political economy of redistribution with perfectly informed, selfish voters.

tory not only the mobility matrix of their society but mostly how responsive individual probabilities of promotions and achievements are to individual effort (as opposed to predetermined factors that are beyond one's control), so as to evaluate the incentive costs of redistributive taxation. However, completely learning the relative role of effort in the generation of inequality would require a lot of costly experimentation that each single generation is not willing to undertake, which implies that in the long run different income histories lead dynasties to converge toward different beliefs regarding society's mobility parameters and therefore different beliefs concerning the socially optimal redistribution rate.

The key point is that in the long run the same reasons lead some dynasties to support higher redistributive taxation and at the same time to supply less effort, while some other dynasties support lower redistribution and at the same time work harder to be successful. Namely, in the long run some dynasties believe (maybe rightly) that predetermined factors are more important than individual effort in shaping individual achievements, while some others believe (maybe rightly) that individual effort is the key to success and social rigidities are second-order.[10] This implies that in steady state there are more "left-wing dynasties" in the lower class and more "right-wing dynasties" in the middle class (regardless of which dynasties have the "right" beliefs, if any), although everybody started with the same distributive goal. Moreover, upwardly and downwardly mobile groups include intermediate fractions of left-wing and right-wing dynasties as compared with stable lower-class and upper-class agents, which leads exactly to the voting patterns depicted in Table I.

The multiplicity of steady states explains at the same time why different countries can remain in different redistributive equilibria, although the underlying structural parameters of mobility are essentially the same. This is particularly likely if a country exhibited for some time in the past a significantly different experience of social mobility before joining the "common" pattern. The "canonical" application is the United States, whose nineteenth century mobility and class structure differed significantly from that of Europe before the two countries converged in the twentieth century.[11]

10. In fact, there is a whole continuum in between these two extreme dynasties.

11. Note that this provides a more rigorous explanation for this persistence phenomenon than the sociologists' "explanation" referred to above. Our theory

Four different pieces of evidence lead us to think that this theory has some relevance. First, when asked what they think about inequality and redistribution and why they vote the way they do, it appears that people from different social backgrounds share a wide consensus about abstract principles of distributive justice (ability per se is usually considered as an irrelevant basis for desert unless it is seen as being a result of previous efforts. People can deserve unequal rewards only on the basis of features—such as effort—that are subject to voluntary control), but that they differ substantially on practical assessments concerning the key to personal success (the poor emphasizing structural factors; the rich, personal qualities such as effort and ambition) (see Rytina, Form, and Pease [1970]; Kluegel and Smith [1986, Chapters 3–4]; and Miller [1992]). In some sense, this paper chooses to take seriously people's justification of their attitudes toward redistribution, instead of describing them as egoists and liars from the beginning.[12]

Next, voting patterns indeed exhibit an amazingly high rate of dynastic reproduction. Abramson [1973] shows Italian data where more than 80 percent of voters with left-wing parents voted for left-wing parties, irrespective of their social class and their mobility experience. This gives strong support to our theory,[13] which says that in the long run individual mobility experience has a substantial but completely indirect effect on individual political attitudes. That is, conditioning individual political attitudes on parents' political attitudes almost completely cancels the effect of individual social mobility on voting behavior depicted in Table I.[14] Our model makes transparent this distinction between the direct, "learning" effect and the indirect, "sampling" effect of mobility on political attitudes. Also, note that the idea that a common cause leads some

shows why it is possible to persist with a wrong estimate of social rigidities and the role of effort. Although agents may eventually learn the right matrix of actual mobility rates, they have a biased estimate of how this matrix would respond to changes in individual effort decisions (and this is the functional form that one needs to know to assess the incentive costs of redistribution).

12. One could obviously argue that people are basically egoistic and ex post "find" some beliefs to justify their behavior. But then one has to explain why income is not perfectly correlated with one's vote (see Table I). Methodologically, it makes sense to assume that agents lie in survey studies only if this is necessary to account for the actions and facts under consideration, which is not the case here.

13. It is hard to reconcile these very high rates of dynastic reproduction with the basic voting patterns of Table I without a theory giving a common reason why some dynasties vote for more redistribution and at the same time have lower rates of upward mobility.

14. See also Kelley [1992] for some very detailed evidence showing that the effect of social origins is mostly indirect, i.e., goes through the parents' political preference and not the class per se.

agents to support redistribution and to supply less effort is similar to the old view that highly politicized workers do not try to use chances of social ascent as much as workers with less class consciousness (see Kaelble [1985, p. 61]).[15]

Finally, the view that there exist wide and persistent disagreements about the incentive costs of taxation is supported by the strong lack of consensus among economists when they attempt to quantify these costs. Everybody agrees that a 90 percent marginal income tax rate may well discourage labor supply and that a 10 percent rate leaves room for more taxation, but the consensus is not long preserved if we try to go farther.[16] This is hardly surprising since economists face the same basic limitations as the agents described in this paper. The only way to know the optimal redistribution rate for sure would be to try it for a while, and this entails substantial social costs. The difference (hopefully) is that most agents base their assessment on their limited personal experience (so that their eventual beliefs are to a large extent forecastable), whereas scholars may perform more sophisticated cognitive processes than those of the agents, or have more time to find more information.[17]

Another application of the ideas developed in this paper is worth mentioning. Forgetting completely about the redistributive taxation aspects of the theory, our learning model predicts that income inequality for a given, homogeneous cohort should grow with age. When people are young and start with the same beliefs, they put forth the same effort, and the only inequality comes from the shocks. As time passes, people who have received bad shocks

15. This example illustrates that left-wing dynasties may very well expend a lot of effort for other objectives which are not related to social ascent (such as trade-union activism or teaching).

16. See MaCurdy [1992] and subsequent references for recent developments of the long-standing controversy between economists about the work disincentives effects of taxes. Note that this empirical literature does not actually offer any estimate of the elasticity of the transition probabilities between different income levels with respect to changes in the after-tax income distribution, since it typically uses cross-sectional data on working hours, wage rates, and total income to estimate the elasticity of working hours with respect to the net wage rate and nonwage income, thereby completely neglecting the incentive costs of redistribution via lower individual effort to be promoted to a higher wage rate (whereas one would expect that this is the source of most of the incentive costs, if any, as opposed to working-hours responses within a fixed wage-rate category). In any case, note that an econometrician estimating the role of effort for these transition probabilities would also need to estimate agents' beliefs about the role of effort in order to get unbiased estimates of the true parameters (otherwise the econometrician will always tend to confirm what agents believe).

17. Section VI shows how a sophisticated outside observer can use our theory and cross-country evidence to make some (limited) progress in assessing these incentive costs.

may get (rationally) discouraged and supply less effort; whereas more successful agents keep putting forth more effort. Eventually a lot of persistent inequality has been created simply because of endogenous beliefs dynamics. This provides a new explanation for this widely noted phenomenon.[18]

The rest of this paper is organized as follows. Section II sets up a simple model of income inequality and redistribution. Section III describes the learning dynamics when dynasties learn only from their own income trajectory. Section IV analyzes the long-run steady states of this learning process and proves the main result; i.e., that voting patterns always look like Table I in the long run. Section V introduces the possibility of learning from other dynasties and shows how this restricts in interesting ways the degree of heterogeneity that one ought to observe in any single country while preserving the heterogeneity of long-run beliefs. Section VI attempts to make some outside observer's welfare comparisons of the various steady states. Section VII gives concluding comments.

II. A Model of Income Inequality and Redistribution

In order to isolate the heterogeneity of voting behavior stemming from the endogenous heterogeneity of beliefs about incentive costs, we consider a model of redistribution where different income groups do not a priori have different distributive objectives when they vote over redistributive policies.[19] This may arise simply because redistribution is of a pure social-insurance nature (each agent faces equal chances at the beginning of each period), or more generally because all agents share the same principles of distributive justice, although they may have different

18. See Deaton and Paxton [1994] for recent evidence that the variance of labor earnings (and not only the variance of total income or consumption) grows with age for a given cohort. Our proposed explanation differs from the usual explanation (the true inequality of ability between agents is the same at all ages, but it takes time for employers to learn these ability differentials. See, e.g., Baker, Gibbs, and Holmstrom [1994, pp. 925–28] and references cited therein) in that in our model there is no true inequality between agents, and everything comes from endogenous discouragement and encouragement effects. One way to distinguish empirically between the two theories would be to observe the effect on earnings dynamics of personal events that affect individual beliefs about the role of effort without affecting employers' perceptions about ability.

19. As we repeatedly stress throughout the paper, a model where voting heterogeneity comes entirely from heterogeneous, well-informed economic interests can hardly explain the voting patterns of Table I. This does not preclude real-world individual concerns for redistribution to be some complex combination of selfish and social values (as long as this is consistent with Table I and the observed rates of political reproduction).

material interests in redistribution; we choose to focus on the latter case.

We assume a discrete infinite horizon, $t = 1, 2, \ldots$, and we consider an economy made up of a continuum of agents $i = [0;1]$. For convenience we shall think of each period as a generation, and of each generation as having exactly one offspring each period.[20]

During each period t each dynasty i can obtain one of two possible pretax incomes $y_{it} = y_0$ or y_1, with $y_1 > y_0 > 0$. We note L_t (respectively, $H_t = 1 - L_t$) the mass of agents born at time t in low-income families (respectively, in high-income families); L_t is the mass of dynasties obtaining income y_0 at period $t - 1$ (L_t= m(i subject to $y_{it-1} = y_0$), where $m(.)$ is the Lebesgue measure over I). Agents obtain income y_0 or y_1 depending on luck, how much effort e one spent, and social origins (i.e., parents' income). The material welfare of agent i at period t is given by[21]

$$U_{it} = y_{it} - C(e_{it})$$

with

$$C(e) = e^2/2a, \qquad a > 0.$$

More precisely, the probability that an agent with social origins y_0 (respectively, y_1) and with effort supply e obtains income y_1 is given by

$$\text{proba}(y_{it} = y_1 | e_{it} = e, y_{it-1} = y_0) = \pi_0 + \theta e$$

(respectively,

$$\text{proba}(y_{it} = y_1 | e_{it} = e, y_{it-1} = y_1) = \pi_1 + \theta e).$$

In equilibrium these probabilities will be strictly between zero and one, so that there is positive intergenerational mobility in this economy. We assume, however, that $0 < \pi_0 < \pi_1$ to reflect the fact that children from high-income families have access to better opportunities (on average). $\theta > 0$ measures the extent to which individual achievement is responsive to individual effort.

It is irrelevant for our purposes where the parameters (π_0, π_1, θ) come from (and in particular whether poor kids have lower opportunities because of genetic endowment, parental environ-

20. Although nothing would be changed if lifetimes last several periods, as we shall see later on.

21. We assume a to be small enough so that the transition probabilities defined below will always be between 0 and 1. We choose this simple functional form for $C(e)$ for the sake of notational simplicity only.

ment, or schooling), because we only consider public policies that are purely redistributive; i.e., which take as given these parameters and simply try to make low achievement less painful by redistributing from y_1 to y_0.[22] In this simple two-income world where we assume that both effort and social origins are not publicly verifiable (second-best) optimal redistributive policies simply take the form of a tax rate $\tau \in [0;1]$.[23] Income is taxed at rate τ, and all tax revenue is redistributed in a lump-sum way, so that the after-tax income $y_{0\tau}$ (respectively, $y_{1\tau}$) of someone obtaining pretax income y_0 (respectively, y_1) is $(1-\tau)y_0 + \tau Y$ (respectively, $(1-\tau)y_1 + \tau Y$), where Y is aggregate income at the corresponding period.[24]

The timing of actions for each generation is as follows: after they choose their effort level e_{it} and their income shock $y_{it} = y_0$ or y_1 is realized, agents vote over the redistributive policy τ_{t+1} to be applied next period. Tax policies are chosen one period in advance to avoid time-consistency issues (the relevant tax rate is known prior to effort-taking) and to ensure that at the time of the vote agents know their own income group, so that there are four types of voters (as in Table I): the stable lower class, noted SL_t (those whose parents' income was y_0 and whose income is also y_0), the downwardly mobile, noted DM_t (those whose parents' income was y_1 and who have gone down to y_0), the upwardly mobile, noted UM_t (those whose parents' income was y_0 and who have moved up to y_0), and the stable high-income (or middle class, noted SH_t (those whose parents' income was y_1 and whose income is also y_1).

We assume that when voting over redistribution these differ-

22. Our analysis can readily be extended to a world with a larger set of policy tools (e.g., schooling and parental aid policies aimed at reducing $\pi_1 - \pi_0$) in case agents have common beliefs regarding how these policies affect the parameters (π_0,π_1,θ), whatever they may be. For given beliefs about (π_0,π_1,θ) agents will favor the same, socially optimal policy package, and these beliefs will be determined through individual income histories in the same way as in Section III, IV, and V below. However, if, as one would expect, agents have different beliefs regarding how different policies can affect some given parameters (e.g., agents who experienced different schooling or parental histories view differently the relative efficiency of pure redistributive taxation versus schooling subsidies), then the theory has to be substantially enriched in order to account for these endogenous variations in beliefs. We leave this for future research.

23. If redistributive transfers could be made conditional either on individual social origins or on individual effort, then one could redistribute without affecting individual incentives to expend a lot of effort, and the size of socially optimal transfers would just depend on $\pi_1 - \pi_0$ (and not on θ). Note also that once such conditional transfers are not feasible (second-best) optimal policies charges the same flat tax rate on everybody's income. It is useless to try to charge lower rates against lower lump-sum transfers to high-social-origins agents by inducing them to self-select (this is because effort matters the same way for all agents).

24. Using the notation introduced above, aggregate income at period t Y_t is equal to $L_{t+1}y_0 + H_{t+1}y_1$.

ent agents share the same social welfare function. To fix ideas, we assume that they all think that unequal opportunities that are beyond one's control (i.e., $\pi_0 < \pi_1$) are a bad thing, and that the state should try to correct this as long as this is in the interest of the most disadvantaged, i.e., should try to maximize the expected welfare of lower-class children by redistributing income from y_1 to y_0.[25] That is, when voting at period t, all agents are maximizing $V_{t+1} = \int_{i\in Lt+1} U_{it+1}di$. Note that there is nothing contradictory between maximizing a "social" objective function when voting and maximizing private welfare when choosing one's effort level. In the latter case, no positive-mass effect is imposed on the aggregate.[26] This is the traditional distinction between private and social values (see, e.g., Arrow [1963, p. 18]).

The important point is that every voter is going to balance the social benefits of equalizing opportunities with the incentive costs of taxation. That is, setting a tax rate τ_{t+1} will lead period $t + 1$'s agents to choose an effort level $e(\tau_{t+1},\theta)$ maximizing their own expected welfare and therefore to reduce their effort if obtaining a high income shock is less rewarded (regardless of their social origins):

$$\begin{aligned} e(\tau_{t+1},\theta) &= \text{argmax}_{e\geq 0}\, Ey - C(e) \\ &= \text{argmax}_{e\geq 0}(\pi_0 + \theta e)[(1 - \tau_{t+1})y_1 + \tau_{t+1}Y_{t+1}] \\ &\quad + (1 - \pi_0 - \theta e)[(1 - \tau_{t+1})y_0 + \tau_{t+1}Y_{t+1}] - C(e) \\ &= \text{argmax}_{e\geq 0}\, \theta e(1 - \tau_{t+1})(y_1 - y_0) - C(e); \end{aligned}$$

that is,

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

25. This Rawls-type social objective seems to be broadly consistent with what people express in survey studies (see above). Those readers who feel unhappy with this social objective can replace it by another social welfare function (such as the utilitarian sum of utilities, assuming risk aversion), without changing the substance of what follows (see below).

26. As to why people go and vote despite their negligible importance, we have nothing original to say. One Kant-like theory would be that they act the way they want everybody to act (as if they were the dictator). For a more conventional theory assume, for example, that the continuum economy we described so far is in fact a large finite economy with some positive probability of being the decisive voter. The economy must be sufficiently large so that agents' social concerns do not show up when choosing effort levels. If the agents' total utility is $W_t = Bu_{it} + Cv_{t+1}$, then for any $\in > 0$ there exists b/c sufficiently small and n sufficiently large such that individual effort and voting decisions are $\in$-close to those we are considering.

Taking this into account, the tax rate τ_{t+1} maximizing the expected welfare of lower-class children at period $t + 1$ is given by

$$\begin{aligned}\tau_{t+1}(\pi_1 - \pi_0,\theta) &= \operatorname{argmax}_{\tau\geq 0} V_{t+1} = \int_{i\in Lt+1} U_{it+1}di \\ &= \operatorname{argmax}_{\tau\geq 0} (\pi_0 + \theta e(\tau,\theta))(1 - \tau)y_1 \\ &\quad + (1 - \pi_0 - \theta e(\tau,\theta))(1 - \tau)y_0 \\ &\quad + \tau[(\pi_0 L_{t+1} + \pi_1 H_{t+1} + \theta e(\tau,\theta))(y_1 - y_0) + y_0] \\ &\quad - C(e(\tau,\theta));\end{aligned}$$

that is,

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

Unsurprisingly, the socially optimal tax rate is an increasing function of $(\pi_1 - \pi_0)$ and a decreasing function of θ. The larger the inequality of opportunity $\pi_1 - \pi_0$, the more it needs to be corrected, and the higher the income elasticity θ with respect to effort, the more severe the moral-hazard incentive problem.[27] Note that these properties do not depend on the particular social welfare function that we chose for illustrative purposes.[28]

If the parameters (π_0,π_1,θ) were known to everybody with certainty, then by assumption everybody would agree regarding the level of socially optimal redistribution. Political attitudes would not vary with income, and one can easily show that starting from any initial condition $(L_0,H_0 = 1 - L_0,\tau_0)$, the economy would converge toward a unique steady-state distribution $(L_\infty,H_\infty = 1 - L_\infty,\tau_\infty)$.

III. Dynastic Learning

Now, assume that agents initially have different beliefs about society's structural parameters (π_0,π_1,θ). That is, all agree that opportunities are to some extent unequal and that incentives are to some extent important, but they disagree on the relative quantitative importance of the two. Some think that the "deterministic" difference in opportunities $\pi_1 - \pi_0$ is small as compared with the importance θ of individual effort in shaping individual achieve-

27. Note also that no public intervention is required if opportunities are equal, i.e., $\pi_1 - \pi_0 = 0$ (this is because we assumed no risk aversion), and that more equalization of opportunities is less costly when the society is richer (i.e., H_{t+1} is larger).

28. In particular, the same properties would hold if one maximizes any (weighted-)utilitarian social welfare function (assuming positive risk aversion, otherwise the optimal utilitarian tax rate is always zero).

ments, so that they want very little state intervention so as not to offset individual incentives. Whereas some others agree that incentives are a problem, but that overall θ is sufficiently small as compared with $\pi_1 - \pi_0$ (that is, structural and predetermined factors outweigh individual factors) so that the state can play a substantial role in raising revenue to equalize opportunities without that much harm.

The questions we want to investigate are the following: assume that there is some "true," stationary set of parameters $(\pi_0^*,\pi_1^*,\theta^*)$. What happens in the long run if agents start with uncertain beliefs about these parameters? What do the long-run voting patterns look like? What role is played by social mobility in this learning and voting process?

To address these questions, we must first specify how agents learn about society's mobility parameters. We consider a learning process where each dynasty learns only from its own experience and not from others' beliefs nor from the aggregate income distribution. That is, each dynasty i believes that others' dynasties are right-wing or left-wing for all sorts of complex reasons that i does not understand and therefore treats their beliefs as completely exogenous and uninformative (we assume Bayesian rationality but not common knowledge of Bayesian rationality). In Section V we will analyze how learning from others can reduce the long-run heterogeneity of beliefs without canceling it completely.

The initial state of the economy is $(L_0,H_0,\tau_0,(\mu_{i0})_{i\in I})$, where $\mu_{i0}(.)$ is the initial prior of dynasty i. We allow μ_{i0} to be any probability measure defined over the set of all logically possible (π_0,π_1,θ).[29] At period $t = 0$, each dynasty i chooses an effort level $e_{i0}(\mu_{i0},\tau_0)$ maximizing its expected private welfare, rationally updates its beliefs μ_{i0} given its income achievement $y_{i0} = y_0$ or y_1, takes part in the voting process over τ_1 by supporting what i believes to be the socially optimal policy $\tau_{i1}(\mu_{i1}(.))$ given its posterior beliefs μ_{i1}, and finally transmits its posterior to its offspring—and so on for the next generation.

First, note that as far as effort-taking is concerned, only the averages of the probability measures $\mu_{it}(.)$ are relevant. By linearity, agents just choose the optimum effort level associated with the

29. For notational simplicity, we will restrict ourselves in what follows to beliefs with finite support. By logically possible we mean that all parameters (π_0,π_1,θ) receiving a positive weight define mobility probabilities strictly between 0 and 1, given the effort cost (as summarized by the parameter a). This allows us to forget completely about corner solutions.

certainty-equivalent beliefs; that is,

$$e_{it}(\tau_t,\mu_{it}) = e(\tau_t,\theta(\mu_{it}))$$

with

$$\theta(\mu_{it}) = \sum_{\mathrm{supp}(\mu it)} \theta\mu_{it}(\pi_0,\pi_1,\theta).$$

Individual voting decisions are potentially more complicated. One way individual agents with beliefs μ_{it} could determine their voting attitude is by computing their best estimate of the true social optimum $\tau_{it}(\mu_{it})$ as if everybody shared their beliefs, i.e., by ignoring the fact that other agents have other beliefs and therefore may react to taxes in a way that they view as suboptimal (in one way or another). This would lead to a most-preferred tax rate schedule $\tau_{it}(\mu_{it})$ which would be increasing with one's best estimate of the inequality of opportunities $\pi_1(\mu_{it}) - \pi_0(\mu_{it})$ and decreasing with one's best estimate of the role of effort $\theta(\mu_{it})$.[30]

However, even if dynasty i treats others' beliefs as exogenous and uninformative, the maximization of V_t implies that i should take into account the way other agents are responding to taxes.[31] That is, if dynasty i can observe that the average beliefs of other agents are $\theta_t = \int_{j\neq i}\theta(\mu_{jt})\mathrm{d}j$, then i's best estimate of the policy maximizing V_t is given by

$$\begin{aligned}\tau_{it}(\mu_{it}) = \mathrm{argmax}_{\tau\geq 0} \sum_{\mathrm{supp}(\mu it)} [(\pi_0 + \theta e(\tau,\theta_t))(1-\tau)y_1 \\ + (1-\pi_0-\theta e(\tau,\theta_t))(1-\tau)y_0 + \tau[(\pi_0 L_t + \pi_1 H_t \\ + \theta e(\tau,\theta_t))(y_1 - y_0) + y_0] - C(e(\tau,\theta_t))]\mu_{it}(\pi_0,\pi_1,\theta);\end{aligned}$$

30. However, note that τ_{it}, may increase with the variance of μ_{it} for given averages, since $\tau(\pi_1 - \pi_0,\theta)$ is proportional to $1/\theta^2$.

31. With common knowledge of Bayesian rationality and regardless of the possibility of learning from others, agents may not take others' beliefs entirely as given if voters did not care only about V_{t+1} (the average welfare of the poor kids next period) but rather about something like $\Sigma_{s\geq t+1}\delta^{s-t-1}V_s$. In that case, they would not choose their most-preferred policy by taking as given others' reactions to taxes, but rather they would try to influence others' future learning processes through strategic policy choices, thereby trading off less justice in the short-run against more truth and therefore more justice in the longer run. For example, if you believe that poor kids have beliefs that are inefficiently biased against effort, you may want to support even less redistribution than you view as socially optimal, so as to get them to learn what you view as the truth. In the same way, low-effort, predetermined factors advocates may want to tax even more the yuppies putting forth what they view as inefficiently high effort. In other words, the political conflict will be intensified by such attempts to influence others' beliefs. This may actually bring individual most-preferred policies closer to those computed under the assumption that everybody shared one's beliefs. This seems to capture well the idea that individuals are somehow responsible for their wrong beliefs, and we leave this important issue for future research.

that is,

$$\tau_{it}(\mu_{it}) = \frac{H_t\pi_1(\mu_{it}) - \pi_0(\pi_{it})}{a(y_1 - y_0)\theta_t^2} + \frac{1 - \theta(\mu_{it})}{\theta_t}.$$

In other words, for given average beliefs θ_t individual most-preferred tax rates still depend positively on individual estimates of the inequality of opportunities $\pi_1(\mu_{it}) - \pi_0(\mu_{it})$ and negatively on individual estimates of the role of effort $\theta(\mu_{it})$.[32] In what follows, we will simply summarize most-preferred policy schedules $\tau_{it}(\mu_{it})$ by a function $\tau(\pi_0(\mu_{it}) - \pi_1(\mu_{it}),\theta(\mu_{it}))$ that is increasing in its first argument and decreasing in its second argument. In order to concentrate on the endogenous path of voters' perceptions, we adopt the crudest possible view of the political process: parties are assumed to be purely opportunistic. Individual preferences over tax rates are single-peaked, and both parties advocate the median most-preferred tax rate. Thus, $\tau_{t+1} = \text{med}((\tau_{it+1}(\mu_{it}))_{i\in i})$.[33]

The dynastic Bayesian updating process is perfectly standard. Consider, for example, the learning effects of an upwardly mobile trajectory, i.e., an agent $i \in UM_t$ with initial beliefs μ_{it}. For any $(\pi_0,\pi_1,\theta) \in \text{supp}(\mu_{it})$, we have

$$\mu_{it+1}(\pi_0,\pi_1,\theta) = \mu_{it}(\pi_0,\pi_1,\theta)\frac{\pi_0 + \theta e(\theta(\mu_{it}),\tau_t)}{\sum_{\text{supp}(\mu_{it})}(\pi_0' + \theta' e(\theta(\mu_{it}),\tau_t))\mu_{it}(\pi_0',\pi_1',\theta')}.$$

That is, i will put more weight on those parameters predicting more upward mobility than the prior. Note that Bayes' rule puts few restrictions on short-run learning from one's own experience. One's effort level and political attitudes can go in every direction following, say, an upwardly mobile trajectory, depending on how initial beliefs determine the interpretation of the event. For example, an upwardly mobile trajectory need not increase one's estimate of the role of effort. In the example above, $\theta(\mu_{it+1})$ can be

32. Note also that θ_t has an ambiguous effect on $\tau_{it}(\mu_{it})$. A higher θ_t induces i to support less equalization of opportunities because i knows that other agents will overreact to tax increases, thereby increasing the cost of redistribution ($(H_t(\pi_1(\mu_{it}) - \pi_0(\mu_{it}))/a(y_1 - y_0)\theta_t^2$ is a decreasing function of θ_t). On the other hand, a higher θ_t induces i to tax more (respectively, subsidize less) what he views as inefficiently high (respectively, low) effort-taking, regardless of any redistributive objective ($(1 - \theta(\mu_{it})/\theta_t$ is an increasing function of θ_t). Note that the second effect dominates if $\pi_1(\mu_{it}) - \pi_0(\mu_{it}) = 0$; i.e., if i believes that there is no need for pure redistribution. Also see footnotes 31 and 39.

33. We certainly do not suggest that this is a good positive theory of political parties. For example, one may want to model political parties as having beliefs on their own and trying to influence voters' perceptions (see Roemer [1994] for research along these lines).

smaller than $\theta(\mu_{it})$ if μ_{it} is such that i puts forth very little effort and views a high-income shock as indication of a high π_0 and a low θ. It follows that unless we make unjustified restrictions on the initial priors (μ_{i0}) the direct learning effects of social mobility will not deliver the voting patterns depicted in Table I. We now turn to the long-run evolution of beliefs, where these ambiguous, direct learning effects are counterbalanced by indirect, sampling effects that always push toward Table I.

IV. STEADY-STATE POLITICAL ATTITUDES

The first property of this dynamic process of learning and voting is that it converges; that is, in the long run beliefs about society's mobility parameters and the resulting equilibrium tax rate are stationary.[34] This is a direct consequence of the martingale convergence theorem.

PROPOSITION 1. Whatever the initial condition $(L_0,H_0,\tau_0,(\mu_{i0})_{i\in I})$, and for every dynasty $i \in I$, the belief $\mu_{it}(.)$ converges with probability 1 toward some stationary belief $\mu_{i\infty}(.)$ as t goes to ∞. The equilibrium tax rate τ_t converges toward some tax rate τ_∞.

Proof of Proposition 1. For any given tax rate sequence $(\tau_t)_{t\geq 0}$, the stochastic process $(\mu_{it}(.))_{t\geq 0}$ is defined by a standard, fully rational process of Bayesian updating, and as such has the martingale property. What dynasty i expect its offspring to know next period is exactly what dynasty i knows today; otherwise dynasty i would know it right away (see, e.g., Aghion et al. [1991]). Thus, the martingale convergence theorem applies, and the society converges toward some stationary set of beliefs $(\mu_{i\infty}(.))_{i\in I}$. It follows that the equilibrium tax rate, as a continuous function of these beliefs, converges.

QED

Now, the interesting question is whether every dynasty necessarily adopts the same beliefs in steady state, and whether the long-run tax rate is necessarily equal to the true socially optimal tax rate. That is, we want to know how initial conditions $(\mu_{i0})_{i\in I}$ map into long-run beliefs $(\mu_{i\infty})_{i\infty I}$. First, we must rule out "doctrinaire" initial conditions if we want to say anything of interest,

34. Obviously, this would not be true if society's mobility parameters are not stationary, which may well be the case in practice. We leave this for future research.

since Bayes' rule does not allow us to learn anything that was ruled out by the prior. If each dynasty i starts with beliefs μ_{i0} concentrated on a single point (say, $\mu_{i0} = 1_{(\pi 0i,\pi 1i,\theta i)}$), then these beliefs will trivially persist in the long run, and any voting pattern can be steady state. To exclude these degenerate situations, we assume that initial priors (μ_{i0}) put a strictly positive weight on the true parameters ($\forall i \in I, \mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$), so that at least they are "given a chance" to learn the truth. This can be viewed as a stability condition in the sense that all steady states that do not originate from such initial conditions would not survive any small perturbation of individual beliefs putting any small positive weight on the true parameters.

However, this does not guarantee that agents will converge to the truth in the long run. For any tax rate $\tau \in [0;1]$ define $S(\tau)$ as the set of beliefs $\mu(.)$ such that

$$\text{(a)} \qquad \forall(\pi_0,\pi_1,\theta) \in \mathrm{supp}(\mu),\ \pi_0 + \theta e(\mu,\tau) = \pi_0^* + \theta^* e(\mu,\tau)$$

$$\pi_1 + \theta e(\mu,\tau) = \pi_1^* + \theta^* e(\mu,\tau)$$

$$\text{(b)} \qquad (\pi_0^*,\pi_1^*,\theta^*) \in \mathrm{supp}(\mu).$$

Conditions (a) and (b) say that when the tax rate is τ beliefs in $S(\tau)$ generate effort decisions $e(\mu,\tau)$ that lead to expected probabilities of upward mobility, downward mobility, etc., which are the same across all points of their support and coincide with the true probabilities. Therefore, a dynasty starting with such beliefs will never modify these beliefs, whatever income trajectory may be observed. Most of all, such beliefs are stable. Because they lead to expectations that entail no contradiction with experience, no perturbation in the direction of the truth will be recognized with certainty (see the Appendix). Conversely, beliefs that do not verify these properties are unstable. An agent starting with beliefs predicting mobility rates higher or lower than the true mobility probabilities will easily recognize its mistake if its beliefs are perturbated in the direction of the truth.

Now, the point is that there are many beliefs that lead to no contradiction between expectation and experience. Define $\Delta(\tau)$ as the set of all (π_0,π_1,θ) such that

$$\pi_0 + \theta e(\tau,\theta) = \pi_0^* + \theta^* e(\tau,\theta)$$

$$\pi_1 + \theta e(\tau,\theta) = \pi_1^* + \theta^* e(\tau,\theta);$$

that is,

$$\Delta(\tau) = ((\pi_0(\theta),\pi_1(\theta) = \pi_1^* - \pi_0^* + \pi_0(\theta),\theta)_{\theta\geq 0})$$

with

$$\pi_0 = \pi_0^* + (\theta^* - \theta)e(\tau,\theta).$$

(Figure I represents the locus $\Delta(\tau)$ in the (θ,π_0) space. Parameters in $\Delta(\tau)$ are indistinguishable from the true parameters $(\pi_0^*,\pi_1^*,\theta^*)$ in the sense that if one believes these are the true parameters one will take an effort level leading to expectations about income mobility that exactly coincides with experience. One can see on Figure I why there are many such parameters. It is difficult to realize that one puts too much weight on effort if at the same time one puts too little weight on predetermined factors. By definition, all beliefs $\mu \in S(\tau)$ have their averages $(\pi_0(\mu),\pi_1(\mu),\theta(\mu)) \in \Delta(\tau)$. Conversely, all parameters $(\pi_0,\pi_1,\theta) \in \Delta(\tau)$ define many corresponding beliefs $\mu \in S(\tau)$ whose average $(\pi_0(\mu),\pi_1(\mu),\theta(\mu))$; all probability measure μ whose support is on the line passing through (π_0,π_1,θ) and $(\pi_0^*,\pi_1^*,\theta^*)$, putting a positive weight on $(\pi_0^*,\pi_1^*,\theta^*)$; and whose average is exactly (π_0,π_1,θ) belong to $S(\mu)$ (see Figure I).

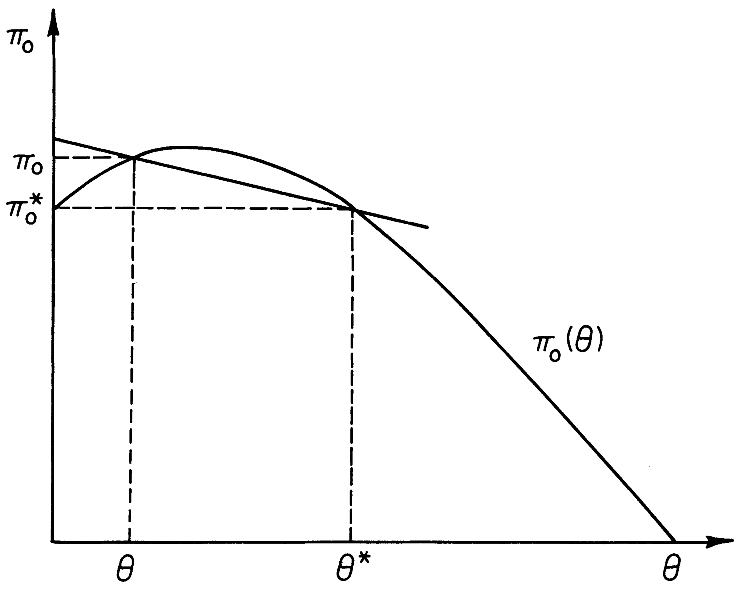

FIGURE I

Note that these many beliefs of $S(\mu)$ averaging to the same point of $\Delta(\tau)$ are essentially the same as far as anything "material" is concerned. They lead to the same effort levels, mobility rates, and political attitudes (see below).

Based upon the discussion above, we have the following results.

PROPOSITION 2. (1) Whatever the initial condition $(\mu_{i0})_{i\in I}$ subject to $\forall\, i \in I\ \mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$, the long-run steady state is such that

(i) $\forall i \in I,\ \mu_{i\infty} \in S(\tau_\infty)$
(ii) τ_∞ is the median of $(\tau_{i\infty}(\mu_{i\infty}))_{i\in I}$.

(2) Conversely, for any beliefs distribution and tax rate $((\mu_{i\infty})_{i\in I},\tau_\in)$ verifying (i) and (ii), there exists some initial condition $(\mu_{i0})_{i\in I}$ subject to $\forall i \in I\ \mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$ converging toward $((\mu_{i\infty})_{i\in I},\tau_\infty)$.

Proof of Proposition 2. See the Appendix.

The fact that rational Bayesian learning does not converge to the truth in this model should not come as a surprise to anybody familiar with the costly-experimentation literature pioneered by Rothschild [1974]. Whenever one is trying to learn some optimum action by using signals whose informativeness depends on the current action, the only way to guarantee complete learning is to take all possible actions during sufficiently long periods, which is privately optimal only if one is sufficiently patient. In our model, agents are trying to learn the functional form relating effort to mobility probabilities, and the only way to learn everything about such a functional form would be for several generations to "sacrifice" their life by trying to supply no effort or to work like mad in order to see what happens to their socioeconomic status. We ruled this out by assuming that each generation chooses its effort level by maximizing its own private welfare U_{it}, thereby making active experimentation strategies unattractive. However, note that exactly the same pattern of long-run mistakes would prevail if each generation was choosing its effort level by maximizing some dynastic utility $\Sigma_{s\geq t}\delta^{s-t}U_{is}$ so long as the discount factor is not sufficiently close to 1 (i.e., as long as $0 \leq \delta \leq \delta_0$ for some $\delta_0 < 1$).[35]

35. McLennan [1984] and Easley and Kiefer [1988, pp. 1060–62] prove this property in a monopolistic pricing model where the monopolist is initially uncertain about two possible linear relationships between price and probability of consumer purchase. As in our model (and unlike in Rothschild [1974]) the action space is

Since not everything is learned in the long run, social mobility trajectories and actual long-run beliefs and political attitudes are jointly determined. What is more interesting is that although long-run beliefs can be wrong in many different ways every single steady state will necessarily exhibit the voting patterns represented in Table I. What Proposition 2 really tells us is that for any stable steady state all dynasties can be ranked along a one-dimensional scale, namely, their position on the curve $\Delta(\tau_\infty)$. That is, in the long run all dynasties believe that the predetermined opportunity difference $\pi_1 - \pi_0$ between lower-class and middle-class children is (on average) $\pi_1^* - \pi_0^*$ (the true opportunity difference), but they have different estimates of θ, i.e., of how much individual effort can undo the effects of social rigidities. All dynasties are mobile, so that one can find proponents of all redistributive policies in every income group. But the point is that because the same beliefs lead some dynasties to supply less effort and to support more redistribution, in steady state there are more left-wing voters among the lower class (irrespective of who has got the right belief, if any), and the political composition of socially mobile agents is strictly intermediary between that of the stable.

To see that, note that those dynasties $i \in I$ who have converged toward a higher $\theta = \theta(\mu_{i\infty})$ vote for less redistributive policies ($\tau(\pi_1 - \pi_0,\theta)$ is decreasing with θ) and supply more effort ($e(\tau_\infty,\theta)$ is increasing with θ) so that a higher fraction of them $H_\infty(\theta)$ has a high income in steady state. Indeed, $H_\infty(\theta)$ is given by the condition that the mass going out of the high-income class is equal to the mass coming in:

$$(\pi_1^* + \theta^* e(\tau_\infty,\theta))H_\infty(\theta) + (\pi_0^* + \theta^* e(\tau_\infty,\theta))L_\infty(\theta) = (1 - \pi_1^* - \theta^* e(\tau_\infty,\theta))H_\infty(\theta);$$

that is,

$$H_\infty(\theta) = \left(\frac{1 - 2(\pi_1^* - \pi_0^*)}{1 - 2\pi_1^* + \pi_0^* - \theta^* e(\tau_\infty,\theta)}\right) - 1$$

$$\text{so that } H_\infty'(\theta) > 0 \qquad (\text{as long as } H_\infty(\theta) < 1).$$

It follows that a higher fraction of lower tax rates supporters has a high income. In the same way, lower tax rates supporters have a higher probability of being upwardly mobile than stable in the

continuous (price instead of effort), and their continuity result (i.e., the incomplete learning for $\delta = 0$ survives for δ small but positive) can be directly applied to our setting.

lower class, but a lower probability of being upwardly mobile than stable in the middle class. Indeed, the steady-state fractions of θ-dynasties who are upwardly mobile $UM_\infty(\theta)$, downwardly-mobile $DM_\infty(\theta)$, stable at high-income $SH_\infty(\theta)$, and stable at low-income $SL_\infty(\theta)$ are given by

$$
\begin{aligned}
UM_\infty(\theta) &= (\pi_0^* + \theta^* e(\tau_\infty,\theta))L_\infty(\theta)\\
DM_\infty(\theta) &= (1 - \pi_1^* - \theta^* e(\tau_\infty,\theta))H_\infty(\theta)\\
SH_\infty(\theta) &= (\pi_1^* + \theta^* e(\tau_\infty,\theta))H_\infty(\theta)\\
SL_\infty(\theta) &= (1 - \pi_0^* - \theta^* e(\tau_\infty,\theta))L_\infty(\theta).
\end{aligned}
$$

It follows that the fraction of θ-dynasties who are mobile as compared with the fraction of θ-dynasties who are stable at high income (respectively, low-income) decreases (respectively increases) with θ. Therefore, the mobile as a whole have a political orientation which is intermediate between those of the stable.

PROPOSITION 3. In any stable steady state described by Proposition 2, the voting patterns mimic those presented in Table I. That is, for any two redistributive policies τ, τ', with $\tau > \tau^3$,

$$H_\infty(\tau,\tau') < L_\infty(\tau,\tau')$$

$$SH_\infty(\tau,\tau') < UM_\infty(\tau,\tau'), DM_\infty(\tau,\tau') < SL_\infty(\tau,\tau'),$$

where $X(\tau,\tau')$ is the fraction of class X preferring τ to τ'.

Proof of Proposition 3. Because preferences over tax rates are single-peaked, there exists τ'', with $\tau > \tau'' > \tau'$, such that dynasties $i \in I$ preferring τ to τ' are those whose most-preferred tax rates $\tau(\theta(\mu_{i\infty}))$ is above τ''; i.e., those whose $\theta(\mu_{i\infty})$ is below some θ''. Since the fraction of θ-dynasties $H_\infty(\theta)$ obtaining a high income in steady state increases with θ, the fraction of the high-income class whose $\theta(\mu_{i\infty})$ is below some θ'' is lower than that of the low-income class. Similarly, because $SH_\infty(\theta)/DM_\infty(\theta)$ and $SH_\infty(\theta)/UM_\infty(\theta)$ increase with θ, $SH_\infty(\tau,\tau') < UM_\infty(\tau,\tau')$, and $SH_\infty(\tau,\tau') < DM_\infty(\tau,\tau')$, and conversely with SL.

QED

Thus, in the long run social origins have an effect on political attitudes only because they are informative about which type of dynasty one belongs to. This is what we referred to in the introduction as the indirect, sampling effect of social mobility on political attitudes, as opposed to the direct learning effect whose direction can be ambiguous. Prior to convergence, however, one cannot completely distinguish between the indirect and the direct effect. Many lower-class agents are in the lower class because their ideology does not push them to work hard to be promoted, but also their poor economic performance confirms their initial ideology (and conversely for the right-wing ideology).

V. Robustness of Long-Run Heterogeneity

The analysis of the possible long-run beliefs when dynasties learn only from their own income trajectory should be regarded as a benchmark case based upon minimal learning opportunities. In practice, agents can learn in many different ways, for example, by looking at achievements and beliefs around them, through direct communication and debate, by designing large-scale econometric studies, and so forth. A complete analysis of these learning schemes and associated limits to beliefs heterogeneity is far beyond the scope of this paper.[36] This section will simply address the issue of learning by looking at cross-section aggregates in order to show that this alters only partially the long-run heterogeneity of beliefs described by Proposition 2, thereby providing some minimal robustness property.

We now assume common knowledge of Bayesian rationality. Each dynasty now believes that other dynasties are rational Bayesian updaters and tries to infer as much as possible from their political attitudes and the observed income distribution. We make minimal observational assumptions: each dynasty can only observe the aggregate income distribution (L_t,H_t) as well as the actual tax rate τ_t (that is, the median point of the cross-section distribution of most-preferred tax rates but not the entire distribution).[37]

In order to make inferences from these observations, each dynasty i must be endowed at $t = 0$ with some prior beliefs $\zeta_{i0}((\mu_{j0}')_{j\neq i})$ defined over all possible distributions of others' priors $(\mu_{j0}')_{j\neq i}$. We assume that the mapping $\mu_{i0} \rightarrow \zeta_{i0}(.) = \zeta(.,\mu_{i0})$ is

36. We can nonetheless indicate some straightforward consequences of enlarging learning opportunities. First, if the economy is partitioned into small (zero-mass) disjoint sets of friends and relatives who can perfectly observe one another, then the steady-state results of Proposition 2 still holds. But things would probably be different if these small sets are overlapping (see Bala and Goyal [1994] for an analysis of such issues in a different context). Next, if all dynasties start with common priors and can costlessly communicate with one another, then it is clear that they cannot disagree forever (see Geanakoplos and Polemarchakis [1982].

37. Proposition 4 below could easily be extended to allow for the observation of mobility rates and of Table I for given τ,τ' (see the proof of Proposition 4 and footnote 41).

common knowledge, which completes the specification of the information structure.

Given some initial condition $((\mu_{i0},\zeta_{i0})_{i\in I},\tau_0)$, this collective learning process follows some trajectory $((\mu_{it},\zeta_{it})_{i\in I},\tau_t)_{t\geq 0}$ and converges toward some long-run steady state $((\mu_{i\infty},\zeta_{i\infty})_{i\in I},\tau_\infty)$, for the same reasons as in the pure dynastic-learning case. For the same reasons as in Section III, we are only interested in long-run steady states originating from initial conditions putting positive weight on the truth (both on the true parameters and on the true distribution of others' priors).[38]

First, note that the long-run heterogeneity of beliefs will survive only if everybody consistently misperceives the average effort $e_\infty = \int_{i\in I} e_{i\infty}(\mu_{i\infty})di$ of other agents. For example, a dynasty i underestimating the role of effort as compared with the average agent and as compared to the truth $(\theta(\mu_{i\infty}) < \min(\int_{j\in I}\theta(\mu_{j\infty})dj,\theta^*))$ will have to overestimate the average effort e_∞ in order to make the observation of aggregate income H_∞ consistent with its beliefs (and conversely for agents on the other side).[39] But the point is that such a permanent misperception of the average effort of others can be perfectly rational if dynasty i is sufficiently uncertain about the initial distribution of priors $(\mu_{j0})_{j\in I}$, given that dynasty i only observes the median most-preferred policy (from which almost nothing can be inferred about the average effort).

However, such misperceptions can be sustained only if the median and average beliefs i has to attribute to others are "rationalizable" from i's viewpoint. Steady states where agents disagree too much will always be ruled out by this process of learning from others, regardless of the uncertainty about initial priors. To see that, consider a dynasty i whose long-run beliefs $(\mu_{i\infty} \in S(\tau_\infty)$ put positive weight on a point $(\pi_{0i},\pi_{1i},\theta_i) \in \Delta(\tau_\infty)$ and

38. Note that if we do not restrict the ζ_{i0}s to put positive weight on the true initial distribution of priors, then new steady states (as compared with Proposition 2's characterization) would actually appear. For example, even if μ_{i0} puts positive weight on the true parameters, $\mu_{i\infty}$ will put zero weight on the true parameters in case ζ_{i0} puts probability 1 on an average effort which is not consistent with the true parameters and the observed income distribution (this type of dramatic inference is driven by the fact that the aggregate income distribution is a deterministic function of the effort profile).

39. Thus, in the steady states described in Proposition 4 below, dynasties with lower $\theta(\mu_{i\infty})$ will tend to have higher estimates of the average $\theta(\mu_{j\infty})$, which may reverse the ordering of the schedule of most-preferred tax rates described in Section III. Over some range, dynasties with lower $\theta(\mu_{i\infty})$ may support less redistribution because they believe more than others that other agents will overreact to tax increases. We do not know whether this is capturing something of interest, and in any case, this nonmonotonicity would probably disappear if voters were trying to influence others' beliefs (see Section III and footnote 31).

who is able to compute the set of (averages of) possible long-run beliefs $\Delta_{(\pi 0i,\pi 1i,\theta i)}(\tau_\infty)$ that other agents can have if the true parameters are $(\pi_{0i},\pi_{1i},\theta_i)$. One can obtain the equation of this set by replacing $(\pi_0^*,\pi_1^*,\theta^*)$ by $(\pi_{0i},\pi_{1i},\theta_i)$ in the equation defining the true set of possible long-run beliefs $\Delta(\tau_\infty)$. On Figure II we represent two such sets $\Delta_{(\pi 0i,\pi 1i,\theta i)}(\tau_\infty)$: one for a dynasty underestimating effort $(\theta < \theta^*)$, and one for a dynasty overestimating the role of effort $(\theta' > \theta^*)$. One can see on Figure II that this rules out steady states where there are at the same time some very left-wing agents and a majority of very right-wing agents. This is because a national left-wing dynasty $(\theta_i = \theta)$ believes that the maximum long-run "mistake" is $\theta^-(\pi_{0i},\pi_{1i},\theta_i)$, so that such a dynasty could not rationalize the existence of dynasties whose most-preferred tax rate is below the tax rate $\tau(\pi_1^* - \pi_0^*,\theta^-(\pi_{0u},\pi_{1i},\theta_i))$ associated with beliefs $\theta^-(\theta_i)$. So the median most-preferred tax rate τ_∞ has to be higher than $\tau(\pi_1^* - \pi_0^*,\theta^-(\pi_{0i}^*,\pi_{1i}^*,\theta_i^*))$ for the most left-wing dynasty i^*, which will be the case if everybody is relatively left-wing or if everybody is relatively right-wing (for example, if everybody is at the right of θ^* or if everybody is at the left of θ^*). For the same

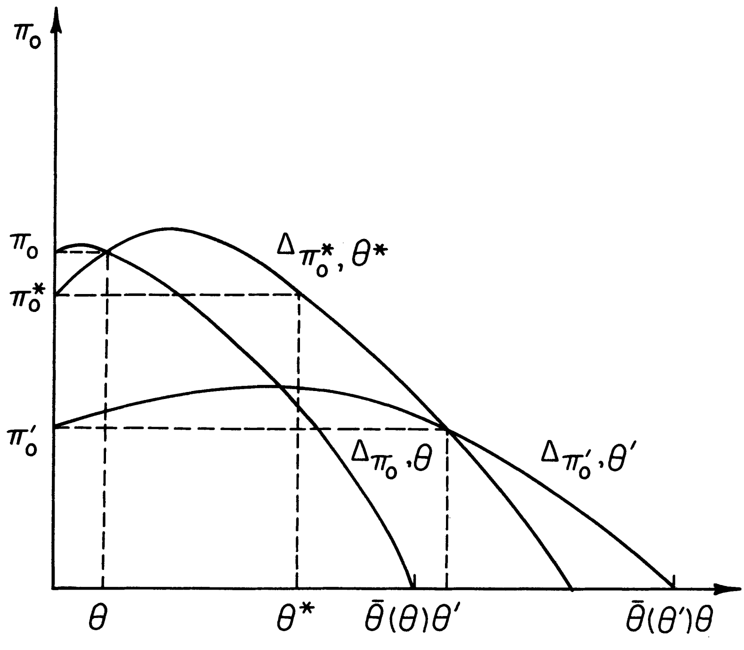

FIGURE II

reasons, the average dynasty must not be too far away from the most left-wing dynasty i^*, so that the effort level the latter has to attribute to the former be lower than $e(\theta^-(\pi_{0i^*},\pi_{1i^*},\theta_{i^*}))$, the maximum effort i^* is ready to attribute to rational Bayesian updaters.

To summarize, learning from others rules out steady states where there is too much diversity, which seems to capture well one reason why each single country tends to be politically homogeneous. Rational agents cannot disagree too much forever.[40] However, note that this narrowing of the set of steady states is biased against extreme left-wing dynasties, as opposed to extreme right-wing dynasties who can move on with their extreme beliefs without ever worrying about the existence of agents in the other corner!

Based upon the intuition given above, one can prove that these are actually the only steady states that are ruled out by the process of learning from others.

PROPOSITION 4. (1) Whatever the initial condition $((\mu_{i0},\zeta_{i0})_{i\in I},\tau_0)$ subject to $\forall i \in I$ $\mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$ and $\zeta_{i0}((\mu_{j0})_{j\neq i}) > 0$, the long-run steady-state is such that,

(i), (ii) of Proposition 2

(iii) $\forall\, i \in I$, $\forall\, (\pi_{0i},\pi_{1u},\theta_i) \in \text{supp}(\mu_{i\infty})$, $\pi_{0i}L_\infty + \pi_{1i}H_\infty + \theta_i e(\theta^-(\pi_{0i},\pi_{1i},\theta_i)) > H_\infty$ and $\tau(\pi_1^* - \pi_0^*,\theta^-(\pi_{0i},\pi_{1i},\theta_i)) < \tau_\infty$.

(2) Conversely, for any beliefs distribution and tax rate $((\mu_{i\infty})_{i\in I},\tau_\infty)$ verifying (i), (ii), and (iii), there exists some initial condition $(\mu_{i0},\zeta_{i0})_{i\in I}$ subject to $\forall\, i \in I \mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$ and $\zeta_{i0}((\mu_{j0})j \neq i) > 0$ converging toward $((\mu_{i\infty})_{i\in I},\tau_\infty)$.

Proof of Proposition 4. See the Appendix.

The incomplete learning result of Proposition 4 relies heavily on the uncertainty about the initial distribution of priors $(\mu_{i0})_{i\in I}$ We believe, however, that this captures well the actual difficulty of learning something from other agents with different political views. It is difficult to sort out the informational content from the prior, in particular because the strength of the latter is in general unknown.

40. In this model there is actually another force pushing toward homogeneity, although in a less drastic way. Namely, lower redistribution implies higher effort (for given beliefs) and therefore less opportunity to learn that low effort is actually the optimal choice (and conversely for high redistribution). That is, by influencing individual experimentation, the majority tends to attract the rest of the economy in its direction. We leave a more precise analysis for future research.

However, note that we have not proved whether these steady states are stable against aggregate shocks. If, following a shock to many agents' beliefs, the tax rate shifts to another value, then the history of tax rates along the adjustment process may reveal so much that agents do not converge back to the initial steady state.[41]

In sum, what we learn from Proposition 4 is that even if we assume agents have a very sophisticated cognitive ability (they know the right model and are able to compute its dynamic properties), learning from others does not change dramatically the amount of long-run heterogeneity of beliefs that obtain under pure dynastic learning. Learning from others will of course homogenize long-run beliefs for given initial conditions, but there is no general presumption that the heterogeneity will shrink completely.

VI. Some Welfare Analysis

Now consider what a very rational agent could learn by making inferences from cross-country evidence based on this model. Imagine an outside observer knowing the model and looking at the pieces of international evidence that we have on inequality, mobility, and redistribution in western democracies. Assume also that this outside observer is ready to assume that these countries have the same structural parameters $(\pi_0^*\pi_1^*,\theta^*)$. The first piece of international evidence is that important and fairly stable differences in levels of redistribution are being observed. Typically, there tend to be many fewer redistributive transfers in the United States than in Western Europe and especially in Scandinavia.[42] From this one can infer that these countries are in different steady-state equilibria of the model (this is confirmed by the observation that working hours, i.e., some limited signal of effort, tend to be longer

41. A general analysis of the exact conditions required for complete learning (which we do not offer in this paper) would most likely boil down to dimensionality comparisons, as is usually the case for questions of information transmission. That is, complete learning will obtain when the dimension of observable signals is higher than the dimension of what one is learning about. In our model the unknown parameters are three-dimensional, dynastic income history is two-dimensional (because of noise), and this is why pure dynastic learning fails. Inferences from the aggregate distribution and the median most-preferred add two more dimensions of signals, and this is why we have to introduce two dimensions of uncertainty about the initial distribution of priors (the median and the average) to preserve incomplete learning. The history of taxes and aggregate distributions adds some $(0,1, \ldots, t \ldots)$-dimensional signals, which in principle could be matched by the infinite-dimensional uncertainty about the priors or, more deeply, by a larger dimensionality of the set of unknown parameters $(\pi_0^*,\pi_1^*,\theta^*)$ (with more than two income levels, with nonlinear effects of effort, and so forth).

42. It is hard to give a global quantification of this multidimensional phenomenon. Mueller [1989, p. 326] presents some data showing that the size of transfers as a fraction of GNP is twice as large in Western Europe as in the United States.

in the United States). Of course, if the outside observer looking at these countries knows the true parameters, he can easily say which country redistributes too much, which country works too much, which agents have a "wrong" ideology, and so on. He knows that the "truly optimal" rate of redistribution τ^*, effort level e^*, and GNP $L^*y_0 + H^*y_1$ are given by the true parameters $(\pi_0^*,\pi_1^*,\theta^*)$.

But if the outside observer does not know a priori which beliefs are the right ones (just as us), what can he say if he wants to compare the actual welfare of these different dynasties and countries? The answer may first seem to be not much. Indeed, one can find steady states where the agents spending the highest amount of effort are in fact not working enough (given the true returns to effort), and others where the agents spending the lowest amount of effort work too much. Maybe there is too much redistribution in the United States, and maybe there is too little in Sweden.

In a desperate need to refine his beliefs, the observer may compare the GNPs of these different countries. The theory predicts that a country with less redistribution should have a higher GNP (whatever the true parameters), and that this should be all the more so if the incentive problem is more severe (that is, if the true social optimum is relatively little redistribution). Here the evidence is not very conclusive. EC countries tend to have a somewhat lower GNP per capita than the United States, but this is less so for Scandinavia. Coming down to less and less secure grounds, the observer may want to compare mobility rates. The theory again predicts that countries with less redistributive taxation should have higher mobility rates, and again this should show up particularly strongly if individual incentives play the key role postulated by these countries. The striking observation here is that all quantitative studies that have tried to compare mobility rates across developed countries have concluded that these rates were amazingly similar (see the references in Section I). The observer may choose to conclude that since the rigid, redistribution-intensive societies of Western Europe are as mobile as the United States, there is little reason to believe in such a strong need to preserve individual incentives. This is a very unsecured inference process, in particular because higher transfers may not only alleviate inequality but also make mobility easier (unlike in our model), but this may be the best one can do to refine arbitrary priors, and we believe that this is the kind of instrumental comparison on which a number of observers "decide" on which side of the Atlantic we are closer to the social optimum.

In theory, one can say more than that by looking in more detail at the class composition of the electorates supporting different redistributive policies (say, different parties). For example, if there is a lot of class polarization (i.e., very high partisan voting in each social class), this suggests that (at least) one class is very far from its socially optimal welfare level. In the same way, very different most-preferred policies (i.e., main political parties advocating very different rates) suggest that (at least) some dynasties have got it all wrong. Assume that we observe very different policy proposals, but very little class polarization. This suggests that very different effort levels do not have a major effect on individual achievements, and therefore that the truly socially optimal policy involves a lot of redistribution and that those working the most should slow down. Similarly, substantial class polarization around comparable policy proposals indicates that individual factors are the key to success and that the social optimum involves little redistribution. This analysis of class polarization of electorates versus polarization of the political spectrum can also be conducted at the cross-country level. One would have to look at these data in more detail, but there does not seem to be striking dissimilarity across western countries from which information could be inferred. In any case, these are again very approximate ways to infer some information, but these may be the best ones available given what we want to learn.

VII. Concluding Comments

This paper has two main objectives. First, providing some theoretical foundations to understand better the political economy of redistribution and particularly some important stylized facts concerning the effect of social mobility on political attitudes toward redistribution (namely, the fact that voters with identical incomes but different social origins vote differently). This gives a richer picture of redistributive politics than the standard public-choice model (which cannot account for these stylized facts), and this shows the importance of belief systems for the generation and dynamics of inequality.[43]

43. In effect, our learning story provides a mechanism generating persistent inequality across dynasties that are intrinsically equal. Other mechanisms producing persistent inequality out of self-fulfilling beliefs include the well-known statistical discrimination model (Arrow [1973]; Phelps [1972]; also see Coate and Loury [1993] for recent developments), as well as a model proposed recently by Roemer and Wets [1994], where agents learn about a convex relationship between human capital investment and income through linear extrapolation of the (human capital investment, income) pair of their social neighborhood, so that poor kids

We believe that our theory also provides a tractable framework for analyzing the fluctuations of redistributive politics, e.g., one that can be used to look at the effects of changes in the pretax distribution on redistributive policies. For example, we have not analyzed how shocks to the fundamentals $(\pi_0^*,\pi_1^*,\theta^*)$ determine transitions between steady states. This looks like an important extension if one accepts the view that historical changes in attitudes toward redistribution are to a large extent driven by changing perceptions about the incentive costs of redistribution, as opposed to changing strategic positions of the decisive income groups (see Piketty [1995] for such an extension).

Next, this paper suggests that instead of always looking at politics as a game designed to aggregate well-informed, conflicting interests, it may sometimes be valuable to consider that the main difference between voters is not their differing interests and objective functions but rather the information and ideas about policies that they have been exposed to during their social life. Not only is majority rule ill-suited to aggregate conflicting interests (see, e.g., majority cycles), but differing beliefs and ideas about government intervention in the economy are pervasive (not only among economists). The point is that although people can have different beliefs about the best-possible policy, these beliefs are not arbitrary. Agents are naturally exposed to different pieces of information depending on their economic positions. We hope that this general approach can be tractable and rewarding enough to solve interesting political-economy questions in the future.[44]

APPENDIX

Proof of Proposition 2

We first prove part (1) of Proposition 2.

Consider $((\mu_{i\infty})_{i\in I},\tau_\infty)$ some steady-state originating from some initial condition $((\mu_{i0})_{i\in I},\tau_0)$ subject to $\forall\, i \in I\ \mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > 0$, and

invest less and remain poor while rich kids invest more and remain rich (a problem with this mechanism is that agents cannot rationally account for the observed inequality around them—e.g., the poor cannot rationalize the high investment they must attribute to the rich without becoming rich themselves, unlike in our model).

44. For example, consider the interesting model of Saint-Paul [1993], where agents with different employment status vote over firing costs. One could extend Saint-Paul's theory to allow for the existence of some socially optimal, possibly positive firing costs depending on how much employers internalize the human-capital social costs of firing. In such a case, it may be reasonable to expect that various employment histories lead to various informational exposures regarding employers' excessive propensity to fire, leading to different political attitudes and possibly important positive and normative implications.

consider some dynasty i. We have to prove that $\mu_{i\infty}$ is in $S(\tau_\infty)$. We first prove that $\mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*) > 0$. Pick $(\pi_0,\pi_1,\theta) \neq (\pi_0^*,\pi_1^*,\theta^*)$ subject to $\mu_{i0}(\pi_0,\pi_1,\theta) > 0$, and define for any $t \geq 0$ the likelihood ratio $l_t = \mu_{it}(\pi_0,\pi_1,\theta)/\mu_{it}(\pi_0^*,\pi_1^*,\theta^*)$.

Consider all periods t subject to $i \in L_t$. If $i \in UM_t$, then Bayes' rule implies that

$$\mu_{it+1}(\pi_0,\pi_1,\theta) = \mu_{it}(\pi_0,\pi_1,\theta)\frac{\pi_0 + \theta e(\mu_{it},\tau_t)}{\sum_{\mathrm{supp}(\mu it)}(\pi'_0 + \theta' e(\mu_{it},\tau_t))\mu_{it}(\pi'_0,\pi'_1,\theta')}$$

$$\mu_{it+1}(\pi_0^*,\pi_1^*,\theta^*) = \mu_{it}(\pi_0^*,\pi_1^*,\theta^*)\frac{\pi_0^* + \theta^* e(\mu_{it},\tau_t)}{\sum_{\mathrm{supp}(\mu it)}(\pi'_0 + \theta' e(\mu_{it},\tau_t))\mu_{it}(\pi'_0,\pi'_1,\theta')}$$

so that

$$l_{t+1} = l_t\frac{\pi_0 + \theta e(\mu_{it},\tau_t)}{\pi_0^* + \theta^* e(\mu_{it},\tau_t)}.$$

Conditional on the true parameters $(\pi_0^*,\pi_1^*,\theta^*)$, the event $i \in UM_t$ will occur with probability $[\pi_0^* + \theta^* e(\mu_{it},\tau_t)]$. If $i \in SL_t$, Bayes' rule implies that

$$\begin{aligned}\mu_{it+1}(\pi_0,\pi_1,\theta) &= \mu_{it}(\pi_0,\pi_1,\theta)\\ &\quad\times\frac{1 - \pi_0 - \theta e(\mu_{it},\tau_t)}{\sum_{\mathrm{supp}(\mu it)}(1 - \pi'_0 - \theta' e(\mu_{it},\tau_t))\mu_{it}(\pi'_0,\pi'_1,\theta')}\\ \mu_{it+1}(\pi_0^*,\pi_1^*,\theta^*) &= \mu_{it}(\pi_0^*,\pi_1^*,\theta^*)\\ &\quad\times\frac{1 - \pi_0^* - \theta^* e(\mu_{it},\tau_t)}{\sum_{\mathrm{supp}(\mu it)}(1 - \pi'_0 - \theta' e(\mu_{it},\tau_t))\mu_{it}(\pi'_0,\pi'_1,\theta')}\end{aligned}$$

so that

$$l_{t+1} = l_t\frac{1 - \pi_0 - \theta e(\mu_{it},\tau_t)}{1 - \pi_0^* - \theta^* e(\mu_{it},\tau_t)}.$$

Conditional on the true parameters $(\pi_0^*,\pi_1^*,\theta^*)$, the event $i \in SL_t$ will occur with probability $[1 - \pi_0^* - \theta^* e(\mu_{it},\tau_t)]$. It follows that if t is subject to $i \in L_t$, then conditional on the true parameters $(\pi_0^*,\pi_1^*,\theta^*)$, we have

$$\begin{aligned}E(l_{t+1}|l_t) &= [\pi_0^* + \theta^* e(\mu_{it},\tau_t)]l_t\frac{\pi_0 + \theta e(\mu_{it},\tau_t)}{\pi_0^* + \theta^* e(\mu_{it},\tau_t)}\\ &\quad + [1 - \pi_0^* - \theta^* e(\mu_{it},\tau_t)]l_t\frac{1 - \pi_0 - \theta e(\mu_{it},\tau_t)}{1 - \pi_0^* - \theta^* e(\mu_{it},\tau_t)},\end{aligned}$$

that is,

$$E(lt+1|lt) = lt.$$

One can show in the same way that if t is subject to $i \in H_t$, $E(l_{t+1}|l_t) = l_t$. That is, the likelihood ratio l_t follows a martingale (conditional on the true parameters) (this is actually a general property of Bayesian learning processes; see Smith and Sorensen [1994]). It follows that l_t converges with probability 1, and most of all that l_t cannot converge to $+\infty$ with positive probability (since $E(l_t|l_0) \neq l_0 < +\infty \ \forall\ t$). It follows that if $\mu_{i\infty}(\pi_0,\pi_1,\theta) > 0$, then $\mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*) > 0$ (otherwise $l_\infty = +\infty$). Since there must exist such a (π_0,π_1,θ) (otherwise $\mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*) = 1 > 0$), it follows that $\mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*) > 0$.

Finally, note that by definition of stationarity $\mu_{i\infty}$ must be such that whatever dynasty i observes it does not change the weight of any point of its support; that is, it must be that

$$\forall\ (\pi_0,\pi_1,\theta),(\pi'_0,\pi'_1,\theta') \in \text{supp}\ (\mu_{i\infty})\pi_0 + \theta e(\mu_{i\infty},\tau_\infty) = \pi'_0 + \theta' e(\mu_{i\infty},\tau_\infty)$$
$$\pi_1 + \theta e(\mu_{i\infty},\tau_\infty) = \pi'_1 + \theta' e(\mu_{i\infty},\tau_\infty).$$

But since $(\pi_0^*,\pi_1^*,\theta^*) \in \text{supp}\ (\mu_{i\infty})$, it follows that

$$\forall\ (\pi_0,\pi_1,\theta) \in \text{supp}\ (\mu_{i\infty})\pi_0 + \theta e(\mu_{i\infty},\tau_\infty) = \pi_0^* + \theta^* e(\mu_{i\infty},\tau_\infty)$$
$$\pi_1 + \theta e(\mu_{i\infty},\tau_\infty) = \pi_1^* + \theta^* e(\mu_{i\infty},\tau_\infty).$$

That is, $\mu_{i\infty} \in S(\tau_\infty)$. This ends the proof of part (1) of Proposition 2.

Part (2) of Proposition 2 is trivial. Take any $((\mu_{i\infty})_{i\in I,\tau_\infty})$ such that $\forall\ i\mu_{i\infty} \in S(\tau_\infty)$ and $\tau_\infty = \text{med}\ (\tau_{i\infty}(\mu_{i\infty}))_{i\in I})$. If $((\mu_{i0})_{i\in I},\tau_0) = ((\mu_{i\infty}(_{i\in I},\tau_\infty)$, then $\forall\ i \in I, t \geq 0$, $\mu_{it} = \mu_{i0}$.

QED

Stability of the Steady States of Proposition 2

Moreover, there are many other initial conditions leading to a given steady state $((\mu_{i\infty})_{i\in I},\tau_\infty)$ satisfying (i) and (ii). For example, consider some dynasty i subject to supp $(\mu_{i\infty}) = [(\pi_0^*,\pi_1^*,\theta^*),(\pi_0,\pi_1,\theta)]$. Then if μ_{i0} is subject to supp (μ_{i0}) = supp $(\mu_{i\infty})$, and $\mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) < \mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*)$, then $\mu_{it} \to \mu_{i\infty}$ with probability 1. To see that, note that if $\mu_{it}(\pi_0^*,\pi_1^*,\theta^*) < \mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*)$, then $\mu_{it+1}(\pi_0^*,\pi_1^*,\theta^*) < \mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*)$ and that $\mu_{it}(\pi_0^*,\pi_1^*,\theta^*) \to 0$ with probability 0 (for the same reasons as in the proof of part (1) of Proposition 2). If $\mu_{i0}(\pi_0^*,\pi_1^*,\theta^*) > \mu_{i\infty}(\pi_0^*,\pi_1^*,\theta^*)$, then $\mu_{it} \to \mu_{i\infty}$ with probability p and to 1 with probability $1 - p$ subject to $\mu_{i0}(p\mu_{i\infty} + (1-p)) + (1 - \mu_{i0})\mu_{i\infty} = \mu_{i0}$ (this comes directly from the martingale condition; see McLennan [1984, pp. 343–44]. The unconditional expectation of the long-run beliefs equals initial beliefs; i.e., $p = 1 - (\mu_{i0} - \mu_{i\infty})/\mu_{i0}(1 - \mu_{i\infty})$. Figure III represents the transition functions for the probability $\mu_{it}(\pi_0^*,\pi_1^*,\theta^*)$ (for the case $\theta^* > \theta$) and characterizes entirely the dynamics of beliefs. If the initial beliefs are at the left of $\mu_{i\infty}$, then long-run beliefs will be attracted with probability 1 by $\mu_{i\infty}$. Whereas if the initial beliefs are at the right of $\mu_{i\infty}$, then they will be attracted with positive probability $p(\mu_{i0})$ by μ_{i0} and with positive probability $1 - p(\mu_{i0})$ by the truth (unsurprisingly, $p(\mu_{i0})$ tends to 1 as μ_{i0} tends to $\mu_{i\infty}$ and tends to 0 as μ_{i0} tends to 1). One can prove similar stability properties when # supp $(\mu_{i\infty}) > 2$. Thus, the steady state $((\mu_{i\infty})_{i\in I},\tau_\infty)$ is stable against perturbations of individual beliefs.

QED

Proof of Proposition 4

Part (1) of Proposition 4 is straightforward. Any steady state must verify (i) and (ii) for the same reasons as in Proposition 2, and it is clear from the discussion before the statement of Proposition 4

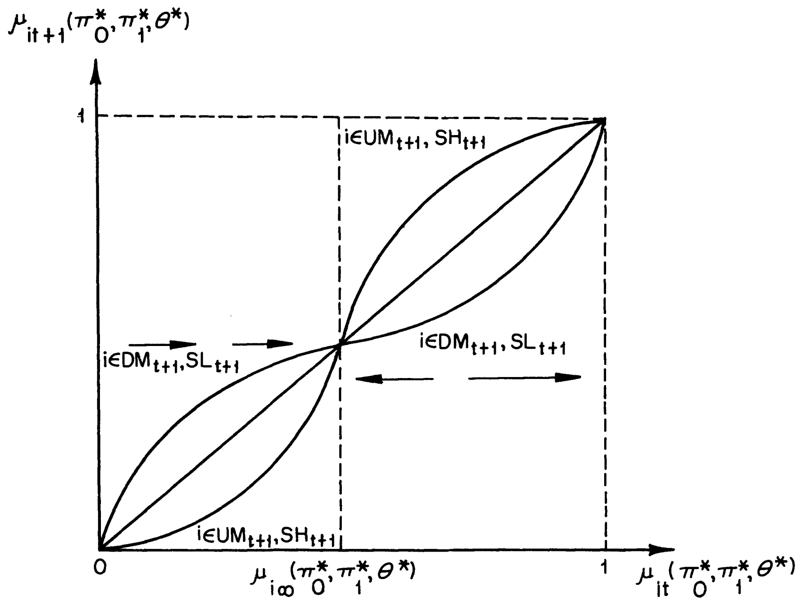

FIGURE III

that if (iii) is not satisfied for some dynasty i, then i cannot rationalize its observation of τ_∞ and (H_∞,L_∞).

We now prove part (2). Consider a candidate steady state $((\mu_{i\infty})_{i\in I},\tau_\infty)$ satisfying (i), (ii), and (iii). From Section III we know that the $(\mu_{i\infty})$ are such that nothing more can be learned from one's own income trajectory.

Consider $i \in I$ and $(\pi_{0i},\pi_{1i},\theta_i) \in \text{supp}\ (\mu_{i\infty})$. Let $e_\infty = \int_{i\in I} e(\theta_i)di$ be the average effort. Let $(L_\infty,H_\infty = 1 - L_\infty)$ be the aggregate income distribution $(\pi_0^*L_\infty + \pi_1^*H_\infty + \theta^*e_\infty = H_\infty)$. Let $e_\infty(\pi_{0i},\pi_{1i},\theta_i)$ be the only estimate of others' average effort that is consistent with the observation of aggregate income in case the true parameters are $(\pi_{0i},\pi_{1i},\theta_i)$:

$$\pi_{0i}L_\infty + \pi_{1i}H_\infty + \theta_i e_\infty(\pi_{0i},\pi_{1i},\theta_i) = H_\infty.$$

Let $\theta(\pi_{0i},\pi_{1i},\theta_i)$ be such that $e_\infty(\pi_{0i},\pi_{1i},\theta_i) = e(\theta(\pi_{0i},\pi_{1i},\theta_i))$. By condition (iii) we have $\theta(\pi_{0i},\pi_{1i},\theta_i) < \theta^-(\pi_{0i},\pi_{1i},\theta_i)$.

Let θ^m be such that $\tau(\pi_1^* - \pi_0^*,\theta^m) = \tau_\infty$. By condition (iii) we have $\theta^m < \theta^-(\pi_{0i},\pi_{1i},\theta_i)$.

Let $(\mu_j'^{(\pi 0i,\pi 1i,\theta i)})_{j\neq i}$ be any distribution of beliefs over $S_{(\pi 0i,\pi 1i,\theta i)}(\tau_\infty)$ such that the average av $((\theta(\mu_j'^{(\pi 0i,\pi 1i,\theta i)}))_{j\in I})$ is $\theta(\pi_{0i},\pi_{1i},\theta_i)$ and the median med $((\theta(\mu_j'^{(\pi 0i,\pi 1i,\theta i)}))_{j\in I})$ is θ^m (there are many such distributions, we are just fixing two dimensions of an infinite-dimensional variable).

Now consider $\zeta_{i0}(.)$ defined by

$$\text{supp}\ (\zeta_{i0}) = [(\mu_j'^{(\pi 0i,\pi 1i,\theta i)})_{j\neq i}]_{(\pi 0i,\pi 1i,\theta i)\in\text{supp}(\mu i\infty)}$$
$$\forall\ (\pi_{0i},\pi_{1i},\theta_i),(\pi'_{0i},\pi'_{1i},\theta'_i) \in \text{supp}\ (\mu_{i\infty}),$$
$$\mu_{i\infty}(\pi_{0i},\pi_{1i},\theta_i)\zeta_{i0}(\mu_j'^{(\pi 0i,\pi 1i,\theta i)})_{j\neq i} = \mu_{i\infty}(\pi'_{0i},\pi'_{1i},\theta'_i)\zeta_{i0}(\mu_j'^{(\pi 0i',\pi 1i',\theta i')})_{j\neq i}.$$

Assume that at time $t = 0$ the collective learning process starts with any initial condition $((\mu_{i0} = \mu_{i\infty},\zeta_{i0})_{i\in I},\tau_0 = \tau_\infty)$. Then we claim that $\forall\ t > 0\ ((\mu_{it},\zeta_{it})_{i\in I},\tau_t) = ((\mu_{i0},\zeta_{i0})_{i\in I},\tau_0)$. This is because for any dynasty i the probability of observing aggregates $(L_\infty,H_\infty = 1 - L_\infty)$ and τ_∞ is the same for all $(\pi_{0i},\pi_{1i},\theta_i) \in \text{supp}\ (\mu_{i\infty})$ (by construction of ζ_{i0}), so nobody learns anything from the aggregate observations. Since the $\mu_{i\infty}$ are such that nothing is learned from dynastic income history, it follows that this is a steady state. (One could also prove that these steady states are stable in the same way as in proposition 2.

QED

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REFERENCES

Abramson, P., "Intergenerational Social Mobility and Partisan Preference in Britain and Italy," *Comparative Political Studies,* VI (1973), 221–34.

Aghion, P., P. Bolton, C. Harris, and P. Julien, "Optimal Learning by Experimentation," *Review of Economic Studies,* LVIII (1991), 621–54.

Arrow, K., *Social Choice and Individual Values* (New York: John Wiley & Sons, 1963).

——, "The Theory of Discrimination," in *Discrimination in Labor Markets,* O. Ashenfelter and A. Rees, eds. (Princeton, NJ: Princeton University Press, 1973).

Baker, G., M. Gibbs, and B. Holmstrom, "The Wage Policy of a Firm," *Quarterly Journal of Economics,* CIX (1994), 921–57.

Bala, V., and S. Goyal, "Learning from Neighbors," mimeo, McGill University, 1994.

Boy, D., "Système Politique et Mobilité Sociale," *Revue Francaise de Science Politique,* XXX (1980), 925–58.

Cherkaou, M., "Mobilité," in *Traité de Sociologie,* R. Boudon, ed. (Paris: Presses Universitaires de France, 1992).

Coate, S., and G. Loury, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review,* LXXXIII (1993), 1220–40.

Deaton, A., and C. Paxton, "Intertemporal Choice and Inequality," *Journal of Political Economy,* CII (1994), 437–67.

Easley, D., and N. Kiefer, "Controlling a Stochastic Process with Unknown Parameters," *Econometrica,* LVI (1988), 1045–64.

Erikson, R., and J. Goldthorpe, "Are American Rates of Social Mobility Exceptionally High?" *European Sociological Review,* I (1985), 1–22.

Erikson, R., and J. Goldthorpe, *The Constant Flux: A Study of Class Mobility in Industrial Societies* (Oxford: Clarendon Press, 1992).

Geanakoplos, J., and H. Polemarchakis, "We Can't Disagree Forever," *Journal of Economic Theory,* XXVIII (1982), 192–200.

Kaelble, H., *Social Mobility in the 19th and 20th Centuries: Europe and America in Comparative Perspective* (New York: St. Martin's Press, 1985).

Kelley, J., "Social Mobility and Politics in the Anglo-Saxon Democracies," in *Social Mobility and Political Attitudes: Comparative Perspectives,* F. Turner ed. (New Brunswick, NJ: Transaction Publishers, 1992).

Kluegel, J., and E. Smith, *Beliefs About Inequality* (New York: Aldine de Gruytler, 1986).

Lipset, S. M., "Elections: the Expression of the Democratic Class Struggle," in *Class, Status and Power,* R. Bendix and S. M. Lipset, eds. (New York: Free Press, 1966).

——, "Why No Socialism in the US," in *Sources of Contemporary Radicalism,* S. Bialer and S. Blutzar, eds. (Boulder, CO: Westview Press, 1977).

——, "Foreword: The Political Consequences of Social Mobility," in *Social Mobility and Political Attitudes: Comparative Perspectives,* F. Turner, ed. (New Brunswick, NJ: Transaction Publishers, 1992).

Lipset, S. M., and R. Bendix, *Social Mobility in Industrial Society* (Berkeley: University of California Press, 1959).

MaCurdy, T., "Work Disincentive Effects of Taxes: A Reexamination of Some Evidence," *American Economic Review,* LXXXII (1992), 243–49.

Marx, K., *Der Achtzehnte Brumaire des Louis Napoleon* (New York: J. Weydemeyer, 1852).

McLennan, A., "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control,* VII (1984), 331–47.

Miller, D., "Distributive Justice: What the People Think," *Ethics,* CII (1992), 555–93.

Mueller, D., *Public Choice II* (Cambridge: Cambridge University Press, 1989).

Perotti, R., "Income Distribution and Investment," *European Economic Review,* XXXVIII (1994), 827–35.

Petersen, W., "Is America Still the Land of Opportunity?" *Commentary,* XVII (1953), 477–84.

Phelps, E., "The Statistical Theory of Racism and Sexism," *American Economic Review,* LXII (1972), 659–61.

Piketty, T., "Dynamic Voting Equilibrium, Political Conservatism and Income Redistribution," Working Paper No. 93-10, DELTA, Paris, 1993.
——, "Redistributive Responses to Distributive Trends, mimeo, Massachusetts Institute of Technology, 1995.
Roemer, J., "The Strategic Role of Party Ideology When Voters Are Uncertain about How the Economy Works," *American Political Science Review,* LXXXVIII (1994), 327–35.
Roemer, J., and R. Wets, "Neighborhood Effects on Belief Formation and the Distribution of Education and Income," Working Paper No. 94-02, University of California, Davis, 1994.
Rothschild, M., "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory,* IX (1974), 185–202.
Rytina, J., W. Form, and J. Pease, "Income and Stratification Ideology: Beliefs about the American Opportunity Structure," *American Journal of Sociology,* LXXV (1970), 703–16.
Smith, L., and P. Sorensen, "Pathological Models of Observational Learning," mimeo, Massachusetts Institute of Technology, 1994.
Sombart, W., *Warum gibt es in den Vereinigten Staaten keinen Sozialismus?* (Tübingen: J. C. B. Mohr, 1906).
Thomson, K., "A Cross-National Analysis of Intergenerational Social Mobility and Political Orientation," *Comparative Political Studies,* IV (1971), 3–20.
Tocqueville, A. de, *De la Démocratie en Amérique (Paris, 1835), in Oeuvres Complètes* (Paris: Gallimard, 1961).
Turner, F., ed., *Social Mobility and Political Attitudes: Comparative Perspectives* (New Brunswick, NJ: Transaction Publishers, 1992).