Social Arrangements and Economic Behavior *

Andrew Postlewaite
University of Pennsylvania

July 17, 2001

Abstract

The social organization of society can have an important effect on the economic organization of that society. By this, we mean that there are not markets to mediate the determination of all things people legitimately care about from an economic point of view. This "incompleteness" of markets will typically lead to an indeterminacy in the social organization, that is, in who associates with whom, who marries whom, and so on. We discuss these points in detail and illustrate them with several formal models.

*Support from the National Science Foundation is gratefully acknowledged. I want to thank Ichiro Obara for careful proofreading and helpful comments.

1 Introduction

The aim of this paper is to argue that the social organization of society can have an important effect on the economic organization of that society. When we say that the social organization of society affects the economic organization of that society, we mean generally that there are not markets to mediate the determination of all things people legitimately care about from an economic point of view. This "incompleteness" of markets will typically lead to an indeterminacy in the social organization, that is, in who associates with whom, who marries whom, and so on. After informally discussing these points in detail, we illustrate them in section 3 with several formal models.

2 Social arrangements: A Lancastrian approach[1]

We begin with what can be called a "Lancastrian" point of view. People care ultimately about a few very basic things: they want to eat and procreate, they want to be safe and secure in their physical and social environments, they want to be protected from the elements and they want these things for their children. In our economic models we typically consider a more detailed description of preferences, describing them with utility functions whose arguments are particular items of clothing, food, jewelry, vacation trips, and so on. From a Lancastrian point of view, these arguments of the utility functions as we typically model them are essentially inputs into a production process that transforms these items into the basic goods mentioned above. Groceries and restaurant meals are all converted into satisfaction from hunger, clothing into protection from the elements, etc.

Why does it make a difference whether we take as the primitive of our models the more basic Lancastrian utility function or the more common utility function that takes as arguments more detailed descriptions of the things we actually purchase? There is certainly an advantage in using the latter, in that the additional detail allows us to address particular issues like the burden of taxation or distortions caused by market intervention. Further, if the process that converts the standard consumption goods into Lancastrian basic goods or needs is fixed, the more detailed utility function

[1] The material in this section is largely derived from a talk given at the Meeting of the Society for Economic Dynamics in Alghero, Sardinia July, 1999.

is simply the concatenation of this conversion function with the Lancastrian utility function.

But a central point of this paper is that the mapping from standard goods into satisfaction of basic needs is *not* a primitive of the environment, exogenous and fixed, but is to a large extent a social construction. Some part of the mapping is fixed: the thermal and water resistant properties of a coat determine completely the degree to which it satisfies the need for protection from the elements. But expenditure on clothing isn't driven solely by the extent to which is satisfies the desire for protection from the elements. In the U.S., single men that earn over $40,000 per year spend twice as much on clothing as men with average earnings. It's doubtful that the richer men are that much warmer or drier than the average, despite this additional expenditure.

What seems clear is that the purpose of the bulk of the expenditure on clothing is to influence others' opinions of us. Whether an individual is trying to impress someone of the opposite sex or to reassure the members of the community that he or she properly "belongs", the basic need that is being satisfied by most clothing purchases is not protection from the elements. The same can be said for a great deal of consumption: people don't buy Rolex watches to tell time nor do they buy Ferrari's to get to the grocery store. And there would be far less money spent on vacations if there were a law against telling people where you went when you got back home.

2.1 The importance of social arrangements

When we say that the mapping from consumption goods to the satisfaction of basic needs is to a large extent a social construction, we don't mean that the point of some particular consumption is to satisfy basic needs that are social in nature, such as the desire for mates or friends. We take these needs to be substantively no different from the need for food and protection from elements. What we mean by the mapping being a social construction is that the relationship between the inputs - the clothes you buy - and the basic needs they are meant to satisfy - influencing potential mates and friends - is not fixed and exogenously given in the way that the protection from the elements one obtains from clothes is. There is a fundamental difference between the way clothes satisfy the basic need for protection from the elements and the way clothes satisfy basic social needs. The shoes you buy will keep your feet equally warm and dry whether the people around you know the designer's name or not, but the degree to which you impress these people is *not* independent of their awareness of the designer.

One could argue that the relationship between the shoes you buy and the degree to which they satisfy basic social needs *is* fixed and immutable; there could be a universal standard of taste that is constant across people and time. We don't think that such a claim is viable, however. The average expenditure on clothing in Milan is a great deal more than that in Philadelphia, a fact difficult to explain either by the wealth of the inhabitants or the weather of the cities. Rather, we would argue that the social organizations of the cities are different. The "equilibrium" in Milan is such that vast (by Philadelphia standards at least) expenditures on clothing pay off, indeed such expenditures may be nearly essential. One could spend the same amount on clothing in Philadelphia of course, but the reality is that individuals there find that, on the margin, expenditures on DVD players and other high-end stereo equipment are more effective in achieving the social status that is being sought.

One could attempt to account for the differences in expenditure between Philadelphia and Milan in a number of ways. First, one could simply posit genetic differences - Italians have better taste than Americans. But the fact is that the descendants of Italian immigrants in Philadelphia don't dress noticeably differently from Philadelphians of other national origins. Different prices or taxes might explain some of the differences, but for the point that we want to make - that there are differences across groups in the linkage between consumption goods and the satisfaction of the basic needs they lead to - we can point to differences in consumption habits across groups within the U.S. that are as striking as those between Milanese and Philadelphians. Taxes, quotas or different government policies simply cannot explain all the differences we see.

To summarize, we are arguing that the preferences over goods and services that we typically take as given and exogenous in most economic models are in fact social constructions that are endogenous in a more primitive model. In any particular group, the link between what an individual consumes and the satisfaction of basic needs that consumption leads to ultimately depends on the consumption habits of the other people in that group. A consequence of this is that even if people are identical with respect to their preferences over basic goods or needs, there can be substantial variation across different social groups in their preferences over specific goods and services - the preferences we typically deal with - because of the variation across those groups in the way that consumption of goods and services satisfy basic wants or needs.

The variation across groups in this connection between consumption of goods and services and the satisfaction of basic wants would be of little economic consequence if the variation across groups was limited to the kind of examples above, that is, whether it's clothing or stereo equipment that is a more important status symbol in some group. That amounts to little more than detailing the difference in toys that people buy. But if groups differ in the extent to which such things as investment in physical or human capital result in the satisfaction of basic needs or wants, these can lead to substantial differences in material consumption across different groups. We see large differences across countries in consumption levels, and these differences seem not to be simply a consequence of differences in endowment. Since endowment differences don't seem to be the primary explanation of the large differences that we see, we should look to the institutions within the countries to understand better these differences. Parente and Prescott (1999) take this approach, arguing that much of the differences can be accounted for by differences in political institutions.

The discussion above about the potential importance of social arrangements suggests that differences in social arrangements can have an important impact on economic behavior. We should never expect markets to be the sole mechanism through which peoples' basic needs are satisfied. We are accustomed to dealing with the kinds of transactions that markets handle reasonably well, that is, those in which we produce, buy and sell food, clothing and so on. The transactions that underlie the satisfaction of such basic needs as the desire to eat and be protected from the elements can probably be efficiently mediated through markets. But when we think about the satisfaction of other basic needs - the desire for mates or the approbation of those in our community, markets are less likely to be the vehicle that mediate transactions. In most developed economies we don't see a smoothly running market for mates, at least for long-term mates. Consider also hypothetical markets to mediate transactions aimed at satisfying the desire for the approbation of your peers. We could auction off the Nobel award, but the mere fact that one gets it as a consequence of a market transaction eliminates a large part of its value.

There are basic needs the satisfaction of which aren't/can't be organized through markets as we generally understand markets. We use the term social arrangements to refer loosely to the ways that these basic wants or needs are satisfied within a group. What is of interest is the interaction of activities governed by markets and activities governed by social arrangements. If these two spheres of human activity were separate and isolated, we could conduct our analysis of human behavior in two disjoint enterprises, one dealing with the market sphere, the other with the social sphere - economics and sociology if you will. But the activities are *not* disjoint. The social arrangements that organize peoples' activities aimed at satisfying those basic needs that aren't satisfied through market activities affects those peoples' incentives in market decisions. There are societies that reward with status and prestige those individuals who devote themselves to the study of holy script, while others bestow these honors on those who control large corporations. While the first may be spiritually richer than the second, the second is likely to come out on top in the World Bank's ranking of countries by GNP per capita. A society that bestows its greatest honors on those that excel in the study of scripture provides an additional incentive to choose that occupation compared to a society that bestows these honors on those who control large business enterprises. The consequence is, of course, a greater concentration of the ablest people studying scripture in the first society than in the second, with the consequent production differences.

2.2 An instrumental approach to incorporating social arrangements

Many social scientists have pointed to the importance of the phenomena discussed above. Adam Smith said "It is not wealth that men desire, but the consideration and good opinion that wait upon riches;"[2] and Max Weber is known for his work linking the structure of Protestantism and individuals' incentives to engage in commercial enterprise. Despite the recognition of the importance of social arrangements, they are largely neglected in modern economics.[3] This neglect is to a large extent an outgrowth of the current methodological foundations of economics. Economics has been remarkably successful compared to other social sciences, and much of the success results from the parsimonious, general models that are the basis of modern economics. Most economists have a great deal of sympathy with the idea that social arrangements may be relevant, but worry that incorporating them into our models is too "soft". The concern is that if we incorporate so-

[2] Adam Smith, *The Theory of Social Sentiments*.
[3] There are exceptions of course; see, e.g., Akerlof (1997) and Bernheim (1994).

cial arrangements by putting things like status into the utility function, we can explain anything. It's certainly true that one can imagine thousands of meaningless regressions being run; nevertheless, it is possible to incorporate social arrangements into economic models without losing the discipline that standard economic modelling provides.

The first hurdle in integrating agents' concern for others' opinions into economic models is determining the form that that concern might take. Consider the problem of dressing for a dinner party. Is the aim to dress similarly to others or to be the most elegantly dressed? If it is to be the best dressed does it matter *how much* better dressed I am than others, or is it only the ordinal rank that matters? If the magnitude matters, is it the difference between me and the next best dressed, the worst dressed or the average? There are probably plausible arguments for any set of answers to these questions. In the absence of any compelling modelling assumptions about what form a concern for others' opinion takes, there is little that can be said; different assumptions lead to different conclusions. This is the lack of discipline that economists fear necessarily accompanies an extension of economic models to include social arrangements.

There is a modelling strategy that ameliorates the difficulties that stem from our ignorance of the fine detail of an individual's concern about how his or her choices compare to others. The discussion above outlined the case for people caring about the opinion of others. The discussion proceeded as though others' opinion of you was an argument of your Lancastrian utility function. Starting from the position that the opinion of others is an argument of the Lancastrian utility function leads to the difficulty in determining the form of the concern outlined above. There is an alternative that captures the idea that people care about others' opinions without those opinions being direct arguments of the Lancastrian utility function, however. Individuals can care about the opinion of others because of indirect effects of those opinions; that is, the concern for others' opinions may be instrumental. I care about the opinions that an editor of a journal has about my work, but his or her opinion is not an argument of my utility function. What *does* affect my utility is whether my paper gets published or not, and the editor's opinion affects *that*. The day the person stops being an editor, I am indifferent to his or her view of my work.

This suggests a modelling strategy that eliminates the need to put non-standard variables into the utility function, or to put it another way, eliminates the freedom to do so. In this modelling approach, one begins with a standard utility function over goods and services, but includes the social arrangements in the model. The extent to which the social arrangements affect peoples' decisions, and the way that they do so, is determined within the model by the interrelationship between the social arrangements and economic consequences.

Goldin (1992) provides a nice example of this approach. Women in the U.S. in the 1950's who went to college typically worked a relatively short time before leaving the workforce for a number of years while they raised children. Because of this absence from the workforce, the return to a woman's college investment was lower than it was for men. Goldin estimates that because of the time out of the workforce, the rate of return was only about half what it was for men at this time, about 4-6%, as compared to 10% for men. But this standard estimate ignores the social arrangements in place at this time. Going to college had a substantial impact on the kind of man a woman ended up marrying. Goldin estimates that attending college at this time increased the income of the man a woman would marry by about 40%. There are many assumptions that one can make about what share of this increase is captured by women, about how long marriages last, etc., but Goldin estimates that this indirect return to college in the form of higher spousal income approximately doubles the rate of return on investment in college, putting it in the same range as that for men.

Our interest is not so much whether that particular estimate is absolutely correct or not. Rather, it's that the paper illustrates how economic decisions are made in the context of a set of social arrangements - how matching between men and women takes place. Taking account of the effect of the social arrangements on economic decisions changes our estimate of the rate of return on investment in education by women by a factor of 100%.

This example illustrates a central point - that we must consider and include social arrangements if we want to understand behavior in many economic situations. The example, however, doesn't shed insight into the question of how different social arrangements in different countries or groups might lead to different behavior. We describe in the next section formal models that address this question.

3 Formal models of social arrangements

This section provides examples of formal models that demonstrate how the absence of complete markets can allow for a role for differing social arrangements that affect economic behavior. The aim here is not to provide a survey of work in this area, but rather, to illustrate how social organization can be formally modelled and included in economic analyses.

A main theme of this paper is that alternative social arrangements can have a direct effect on economic performance. Cole, Mailath and Postlewaite (1992) (hereafter CMP) specifically addresses this point. That paper analyzes a model in which there is matching between men and women, and shows that in some circumstances there may be different social arrangement that govern the matching that are stable in the sense that for all agents, matching according to the prevailing social arrangements yields higher utility than not. Furthermore, when there are multiple stable social arrangements, the growth rate can vary substantially depending on which social arrangements govern behavior. We turn now to a more detailed description of that paper.

3.1 Modelling a concern for relative rank[4]

The model in Cole, Mailath and Postlewaite (1992) focuses on a simple multi-generational society in which, at each period, there is a continuum of men and a continuum of women who live for one period. The agents are matched into pairs, with each pair producing two offspring, one male and one female. Besides the matching decision, agents make standard economic decisions: what to consume and what to invest, with the proceeds of the investment bequeathed to the next generation. The agents in this model differ from neoclassical agents only in that the matching affects the consumption possibilities of agents and their descendants.

The model is asymmetric with respect to men and women in two respects. First, women are endowed with a nontraded and nonstorable good. Women's endowment of this good is exogenous and may differ from woman to woman. Second, only the welfare of male offspring enters the couple's utility function.[5]

In addition to this female good, there is a durable good that can be consumed or invested. Both the female good and the male good are jointly

[4] The material in this section draws heavily on Cole, Mailath and Postlewaite (1992).

[5] The asymmetry between males and females is for analytic tractability, although it might be defended as historically plausible for some societies. A similar model that treats males and females symmetrically is analyzed in Mailath and Postlewaite (2000), which is discussed below.

consumed by the couple. (That is, all consumption is public consumption within matched pairs.)[6] The problem facing a matched couple is how much of their endowment (inherited from the male's parents) should be consumed and how much should be bequeathed to their son. Bequests to the son benefit him in two ways: he has greater opportunity for consumption, and his greater wealth may attract a wealthier mate.

CMP introduces a notion of *status* which is used to rank the men. This ranking is used in the prescription of how matching between men and women is to take place. Specifically, however status is assigned to men, the prescribed matching is that the wealthier a woman is, the higher the status of the man she is to match with. It is not assumed, however, that prescriptions for behavior, either matching or consumption, will necessarily be followed. Agents act in a self-interested way and follow norms and prescribed behavior only when it is in their interest to do so. A *norm equilibrium* is then a description of men's status is determined and a prescription of matching and savings behavior such that no man and woman can deviate from the prescription and increase their utility.

3.1.1 Wealth as status

There is one straightforward norm that assigns status in a way that is easily seen to be consistent with agents' incentives. A *wealth-is-status* norm assigns status to the men in each generation strictly according to their wealth. Given this norm, the wealthiest men will be highest status, and accordingly are to match with the wealthiest women. Clearly, no man or woman has any incentive to deviate from the prescribed behavior. A man who is at the k^{th} percentile in the wealth distribution will be at the k^{th} percentile in the status distribution and is to match with the k^{th} percentile woman in the female wealth distribution. He would clearly prefer matching with a wealthier woman, but any wealthier woman would prefer her prescribed match, who by assumption, will be wealthier. Similarly, women would prefer to marry wealthier men, but those men would not want to do so.

It is not surprising that the wealth-is-status norm is an equilibrium, since it prescribes behavior that is precisely what agents would want to do in the

[6] The assumption that consumption is joint within matched pairs is made primarily for ease of exposition. There must be *some* jointness in consumption to induce a preference for wealthier mates, but the qualitative properties of the equilibria of the model are not driven by consumption within couples being entirely of a public good nature.

absence of any prescriptions. Nonetheless, the equilibrium has interesting properties. First, in any generation, the matching aspect of the model induces couples to consume less (and, consequently, invest more) than they would in the absence of matching. As mentioned above, when a couple leaves a bequest to their son, it benefits him in two ways. First, larger bequests increase his consumption possibilities, and second, larger bequests improve the quality of his match. Since status is assigned on the basis of wealth, large bequests translate into higher status, and, as a consequence, into wealthier mates. Since consumption is joint, a larger bequest results in greater consumption of the female good that accompanies higher status. This component of the son's benefit arises only because matching is incorporated into the model.

While each couple saves more because of the concern for their son and the effect of their savings on his matching prospects, the increased savings by other families nullifies their efforts. It is shown in CMP that the equilibrium savings function in any generation is monotonic, so that the percentile rank of any man after all families increase their savings because of concern for ranking is the same as it would be if all families ignored ranking concerns. It is essentially a prisoners' dilemma game in which all couples in a given generation would be better off if everyone committed to not "oversaving", but any single couple would have an incentive to unilaterally deviate and save a bit more in order to increase their son's prospects in the matching in the next period. The wealth-is-status norm induces what might be called a rat-race of the rich in which there are no winners.

Even though the increased saving that the wealth-is-status norm induces does not change the matching that occurs in any generation, there are important economic consequences of the norm. Specifically, savings and investment in every period are greater than they would be in the absence of matching considerations. Hence, if an economy is governed by a social norm similar to the wealth-is-status norm, estimates of the rate of growth of capital based on the primitives of the economy, but ignoring matching considerations, underestimates the rate of growth of capital.

More importantly for our purpose in this paper, the actual rate of growth of the economy at a particular time depends on the distribution of wealth in an interesting way. When a couple considers the consequences of leaving a slightly smaller bequest to their child, they are comparing the increase in their own utility from doing so with the decrease in their son's utility that results. But part of the decrease in the son's utility is a consequence of his being lower in the ranking among men of his generation. But how *much* lower he will be depends on the particular distribution of wealth in that generation. A very concentrated distribution of wealth might result in a very substantial drop in the son's rank, while a very dispersed distribution would result in a comparatively much smaller drop in rank. Thus, the oversaving that results from the competition among families concerned about their sons' rank in the wealth distribution is greater when the wealth distribution is more concentrated than when it is dispersed, with the result that the part of savings that is due to rank concerns will be greater when wealth is more evenly distributed. This provides an avenue by which the wealth distribution affects investment that is generally overlooked.

The discussion illustrates the economic consequences of particular social arrangements that may play an important role when markets are incomplete. It is the absence of a market for mates that induces the tournament-like competition among families to save more than they would otherwise. As argued in the previous paragraph, ignoring this particular social arrangement leaves out a potentially important incentive couples face when choosing between consumption and saving. We next show that there may be other social arrangements when the markets for mates are absent that qualitatively differ from the wealth-is-status norm.

3.1.2 Aristocratic assignment of status

Besides the wealth-is-status norm described above, CMP considers an alternative norm characterized by an assignment of status not based on wealth. Suppose that in the first period, men are assigned status in some way, not necessarily correlated with their wealth. As before, the prescription for matching will be that the wealthiest women match with the highest status men. More specifically, the woman who is in the k^{th} percentile female wealth distribution is to match with the man who is in the k^{th} percentile of the status distribution. If couples match according to this prescription, their son inherits his father's status; if a couple deviates from the prescription, their son have the lowest status.

As we indicated above, we are interested in norms that are equilibrium, that is, norms for which agents do not blindly follow the prescriptions, but follow the prescriptions only when it is in their interest to do so. It was straightforward to see that the wealth-is-status norm was an equilibrium, since it prescribed behavior that the agents would choose in the absence of norms. The aristocratic norm described above is different, however. Here, the prescription is for the wealthiest woman to match with the highest status man, independent of his wealth. If he is not the wealthiest man, this woman would prefer to match with the wealthiest man instead. Further, the wealthiest man has an interest in matching with the wealthiest woman. How can it be, then, that the aristocratic norm could be an equilibrium? The answer is that parents care about their son's welfare, and when the wealthiest woman matches with the highest status man, her son inherits that status. The consequence of his inheriting the highest status is that he matches with the wealthiest woman in the next period. Given the aristocratic norm, when the wealthiest woman contemplates deviating and marrying the richest man rather than the highest status man, she must balance the immediate increase in wealth that accompanies the deviation with the future cost to her male descendants. The future cost of deviating may be sufficient to support the prescriptions of the aristocratic norm as an equilibrium.[7]

The possibility that there may exist different norms - wealth-is-status norm and aristocratic norm - for a given economy illustrates the points made earlier. The fundamentals of the economy do *not* necessarily determine its economic performance. If the social arrangements that govern matching between men and women are as in wealth-is-status, savings will be higher (and, *a fortiori,* growth will be higher) than if social arrangements are governed by the aristocratic norm.

The model described above illustrates the potential importance of social arrangements in understanding economic behavior, and suggests a number of important issues that it is not, unfortunately, particularly well suited to address. For example, casual observation suggests that the social arrangements governing behavior change over time, sometimes rapidly. While CMP shows that there may at a given time be multiple possible social arrangements, it does not address the question of how or why a change from one to another takes place. In the next section, we present a simple model that allows such an analysis.

3.2 Endogenizing status

We present a simple model that is analyzed in Mailath and Postlewaite (2000) (refereed to as MP hereafter). The model is a simple dynastic model similar

[7] See Cole, Mailath and Postlewaite (1992) and Cole, Mailath and Postlewaite (1995) for conditions under which the aristocratic norm can be an equilibrium.

to that described above. There are continua of men and women in each generation, each exogenously endowed with a stochastic income. Unlike CMP, there is no saving: men and women match, consume jointly their combined income, have children and die. Individuals have a strictly concave utility for current (joint) consumption, and in addition, they care about their childrens' utility. Because the children have a strictly concave utility function over current consumption, the randomness of income reduces the utility from what it would be in the absence of the randomness; we assume that the income risk is not insurable.

We add to this quite standard model an attribute that is heritable, that is, an attribute that children can acquire if and only if at least one parent possesses the attribute. Children with two parents possessing the attribute acquire it for sure, children with neither parent possessing it never acquire it, and children with one parent possessing it acquire it with some probability. The attribute is independent of income and does not enter any agent's utility; hence, it is irrelevant from an economic point of view. We show, however, that if agents are sufficiently risk averse, there are stable social arrangements in which the attribute matters. The particular stable social arrangement we identify is one in which high income agents with the attribute match with each other, and poor agents without the attribute match. But high income agents *without* the attribute match with low income agents *with* the attribute. This last matching provides each agent in the pair with a benefit: the low income, high attribute agent gets higher current consumption than if he or she matched with a similar type, while the high income, low attribute type gets the chance to pass on to his children the high attribute of the low income, high attribute mate. If he or she matched with another high income, low attribute agent, children would be certain to not have the attribute. The benefit of (possibly) passing on the attribute to the children is that — under these social arrangements — high attribute agents consume more than low attribute agents. Hence, the attribute is providing partial insurance against the randomness in income that is assumed to be otherwise uninsurable.

We are interested both in the case in which the interpretation of the transmission of the attribute from parents to children is genetic and that in which it is cultural. In the case that the attribute is transmitted genetically, it is natural to assume that when only one parent possesses the attribute, a child will have the attribute with probability 1/2. On the other hand, when it is culturally transmitted (think of teaching a child to play the piano), the probability may be much higher than 1/2. It may also be lower if there are "economies of scale" in the transmission. By economies of scale we mean that the probability that a child acquires the attribute more than doubles when we move from a single parent possessing the attribute to both parents possessing it (think here of acquiring good conversational skills).

If the social arrangements that use the attribute in matching are stable, we show that if the probability of passing the attribute on to children is below 1/2, but not too much lower, the social arrangement that utilizes the attribute will be stable but that the attribute asymptotically vanishes from the society. When the attribute is passed on with probability greater than 1/2, we show that the social arrangement that utilizes the attribute cannot be stable, that is, such an attribute cannot affect matching in a stable social arrangement.

3.2.1 Model

There is an infinite number of generations each consisting of a continuum of men and a continuum of women. There is a single consumption good which is not storable. Each individual in each generation is endowed with some of this good. In each generation, men and women match and consume their combined wealth (that is, the good is a public good within couples). In addition, each couple has two offspring. The common consumption utility function for all individuals is $U : R \to R$. Individuals care about their descendants' welfare; the utility to any matched couple is their utility from consumption plus the (discounted) average utility of their children. This means, of course, that their utility depends on the consumption of all future generations.

Within each generation, half the men and half the women have high income (H) and half have low income (L). We assume that the probability that any individual has high wealth is independent of past history. Additionally, there is an attribute that is independent of income (for example, height). Each individual is equally likely to be endowed with one of two possible levels of this attribute, high h and low, l; the attribute is independent of income. For now, we assume that this attribute is heritable. By heritable, we mean that if both parents have attribute h or both have attribute l, their children inherit the parents' attribute with probability 1. If the parents have different attribute levels, each child is assumed (independently) to have each level of attribute with probability 1/2. In summary, each generation contains four equal measure sets of agents: (H, h), (H, l), (L, h), and (L, l), and each of these sets contain equal measures of men and women.

Our interest is in the matching of men and women. We assume that matching is voluntary, that is, no unmatched pair can increase their utilities by matching (taking into account, of course, the consequences to their descendants). We call a matching that satisfies this *stable;* given a stable matching, we call the match between any man and woman of the population stable. Since the nonincome attribute is independent of income, there trivially is a stable matching of men and women that ignores this attribute. For example, a matching that has each man of type $(\cdot,\cdot)$ matching with a woman of the same type is stable. An (H,h) man can do no better than match with an (H,h) woman, since this gives each consumption this period of $2H$, the highest possible, and given this matching rule, the current period's match has no effect on their descendants' consumption since income is independent across generations. Similarly, a match between an (H,l) man and an (H,l) woman cannot be improved upon. Since no individual of type $(H,\cdot)$ matches with an individual of type $(L,\cdot)$, matching of men and women of identical types $(L,\cdot)$ is stable.

We next compute the utilities of high and low income people in this society. A high income person H matches with another high income person. The result of the match is consumption of $2H$ and two offspring, each of which is equally likely H or L. Hence the utility for a person of high income (including the utility coming from his children) is (the superscript I denotes that the values are for the case in which the status or ranking is on the basis of income alone):

$$V_H^I = U(2H) + \beta[1/2V_H^I + 1/2V_L^I]$$

Similarly, an L matches with and L and receive utility

$$V_L^I = U(2L) + \beta[1/2V_H^I + 1/2V_L^I]$$

We normalize the utility function u so that $U(2L) = 0$ and $U(2H) = 1$. Solving, we get

$$\begin{aligned} V_H^I &= \frac{2-\beta}{2(1-\beta)}, \text{ and} \\ V_L^I &= \frac{\beta}{2(1-\beta)} \end{aligned}$$

This matching of rich with rich is natural since the only other variable — the nonincome attribute — is unrelated to the sole argument of individuals' direct utility functions, income. There may be, however, other stable matchings. We can think of the matching described above as "income-is-status"; that is, agents are ranked by income and the highest ranked men match with the highest ranked women. Our aim is first to show that there is another ranking in which not only income, but the nonincome attribute affects status, and hence, matching.

3.2.2 A mixed attribute status ranking

We consider a matching that is not positively assortative on income and show that for some initial data of the problem, it is stable. The ranking is one in which individuals of types (H, h) and (L, l) match only with individuals of their own type. Individuals of type (H, l) match with (L, h). We next determine when this might be a stable matching.

This ranking induces a preference for mates with attribute h, all other things equal, since this increases the chance that their children have attribute h, and an individual with attribute h matches with a high income (H) individual regardless of his or her own income. Agents with attribute l, on the other hand, match with low income (L) agents. Hence, agents of type (H, h) prefer to match with individuals of the same type, and any proposed matching that does otherwise cannot be stable. Since both high income and attribute h are desirable, those agents with neither, (L, l) must be matched with the same type agent in any stable matching. What remains to be determined is whether both agents in a matched pair whose types are (H, l) and (L, h) prefer to match with each other or with agents who are of their own type. As before, we compute the value function for each type, given the proposed matching.

Agents of type (H, h) consume $2H$ and have two children, each of whom has attribute h, and income either H or L with equal probability. Hence, the utility of an agent of type (H, h) (including future generations) is

$$V(H, h) = U(2H) + \beta[1/2V(H, h) + 1/2V(L, h)]$$

Similarly, an agent of type (L, l), who matches with an agent of the same type, will consume $2L$ and have two children with attribute l, and have income either H or L with equal probability. Their utility is

$$V(L,l) = U(2L) + \beta[1/2V(H,l) + 1/2V(L,l)]$$

Agents of types (H,l) and (L,h) will have the same utility, since they match with each other and jointly consume $H+L$ and have children that are equally likely to have attribute h or l. Their utility will be

$$V(H,l) = V(L,h) = U(H+L) + \beta[\frac{1}{4}V(H,h) + \frac{1}{4}V(L,h) + \frac{1}{4}V(H,l) + \frac{1}{4}V(L,l)]$$

We continue as before to consider the utility function normalized so that $U(2L) = 0$ and $U(2H) = 1$. We denote the utility of the third possible consumptions level, $H+L$, by u; u is in [1/2,1) if u is (weakly) concave, with u higher the more concave the utility function is.

We denote $V(H,h)$ and $V(L,l)$ by V_h^M and V_l^M respectively (the superscript M denotes that the values are for the case in which we have a mixed status ranking), and denote by V_m^M, $V(H,l)$ and $V(L,h)$. We can rewrite the value function equations then as

$$\begin{aligned} V_h^M(1-1/2\beta) &= 1 + 1/2\beta V_m^M \\ V_m^M(1-1/2\beta) &= u + 1/4\beta[V_h^M + V_l^M] \\ V_l^M(1-1/2\beta) &= 1/2\beta V_m^M \end{aligned}$$

Hence,

$$
\begin{aligned}
V_m^M(1-\frac{1}{2}\beta) &= u + \frac{\frac{\beta}{4}(1+\beta V_m^M)}{(1-\frac{1}{2}\beta)} \\
&= \frac{(1-\frac{1}{2}\beta)u + \frac{\beta}{4}(1+\beta V_m^M)}{(1-\frac{1}{2}\beta)} \\
&\implies V_m^M = \frac{(1-\frac{1}{2}\beta)u + \frac{\beta}{4}}{1-\beta} \\
V_l^M &= \frac{\frac{\beta}{2}(2-\beta)u + \frac{\beta^2}{4}}{(1-\beta)(2-\beta)} \\
V_h^M &= \frac{1}{1-\frac{1}{2}\beta} + V_l \\
&= \frac{2(1-\beta) + \frac{\beta}{2}(2-\beta)u + \frac{\beta^2}{4}}{(1-\beta)(2-\beta)}
\end{aligned}
$$

As mentioned above, for the proposed matching to be stable, both agents in a matched pair whose types are (H, l) and (L, h) must prefer to match with each other rather than with agents who are of the same type they are. For a type (H, l) to prefer matching with an (L, h) rather than with another (H, l), it must be true that

$$
\begin{aligned}
u + \beta\left[1/4V_h^M + 1/2V_m^M + 1/4V_l^M\right] &\geq 1 + \frac{1}{2}\beta(V_m^M + V_l^M) \\
&\implies 4u + \beta V_h^M \geq 4 + \beta V_l^M \\
&\text{solving, we get} \\
\frac{\beta}{2-\beta} &\geq 2(1-u) \\
&\textit{or} \\
u &\geq 1 - \frac{\beta}{2(2-\beta)}
\end{aligned}
$$

We next consider the inequality that must be satisfied for the proposed matching to be stable from the point of view of a type (L, h). An (L, h) who matches with an (H, l) will consume $H + L$ and have children who are equally likely to be h or l. Thus, the inequality that must be satisfied for stability is

$$\begin{aligned} u + \beta[1/4V_h^M + 1/2V_m^M + 1/4V_l^M] &\geq \frac{1}{2}\beta[V_h^M + V_m^M] \\ &\implies 4u + \beta V_l^M \geq \beta V_h^M. \end{aligned}$$

Solving, we get the condition $u \geq \frac{\beta}{4-2\beta}$. It is easy to verify that this inequality is satisfied whenever the inequality associated with the incentive constraint for the type (H, l) is satisfied. To sum up, this matching is incentive compatible whenever

$$u \geq 1 - \frac{\beta}{2(2-\beta)} = \frac{4-3\beta}{2(2-\beta)}.$$

For $\beta \in (0, 1)$, $1 - \frac{\beta}{2(2-\beta)} > 1/2$. Thus, as long as the utility function for consumption is sufficiently concave, the proposed matching is stable. Note that the necessary amount of concavity becomes arbitrarily small as β approaches 1.

3.2.3 Welfare comparison of wealth and mixed attribute status rankings

We demonstrated above that there is always a status ranking based only on wealth. We also showed that if the utility function is sufficiently concave (relative to the discount factor β), there is an additional ranking that depends nontrivially on the "irrelevant" characteristic. We next compare the utilities for agents in these two equilibria.

It's clear that for the mixed attribute status ranking case, $V_h^M > V_m^M$. We show that when parents care sufficiently about their children's welfare, V_m is larger than V_H, the utility of the high income people in the wealth is status case.

Recall

$$\begin{aligned} V_m^M &= \frac{(1 - \frac{1}{2}\beta)u + \frac{\beta}{4}}{1-\beta} \text{ and} \\ V_H^I &= \frac{2-\beta}{2(1-\beta)}. \end{aligned}$$

Hence,

$$
\begin{aligned}
V_m^M &> V_H^I \iff \\
\frac{(1-\frac{1}{2}\beta)u+\frac{\beta}{4}}{1-\beta} &> \frac{2-\beta}{2(1-\beta)} \iff \\
2(1-\frac{1}{2}\beta)u+\frac{\beta}{2} &> 2-\beta \iff \\
2u-\beta u &> 2-\frac{3}{2}\beta \iff \\
(2-\beta)u &> 2-\frac{3}{2}\beta \iff \\
u &> \frac{4-3\beta}{2(2-\beta)}
\end{aligned}
$$

In summary then, if $u > \frac{4-3\beta}{2(2+\beta)}$, individuals of type (H,l) and (L,h) have higher utility in the mixed status regime than when rank is determined by income alone. It is obvious that $V_h^M > V_m^M$, and consequently individuals of type (H,h) have higher utility in the mixed regime than in the income only ranking case. Hence, if $u > \frac{4-3\beta}{2(2+\beta)}$, only individuals of type (L,l) might be worse off in the mixed status regime. We compute next the conditions under which this type is also better off in the mixed regime.

$$
\begin{aligned}
V_l^M &> V_L^I \iff \\
\frac{\frac{\beta}{2}(2-\beta)u+\frac{\beta^2}{4}}{(1-\beta)(2-\beta)} &> \frac{\beta}{2(1-\beta)} \iff \\
\beta(2-\beta)u+\frac{\beta^2}{2} &> \beta(2-\beta) \iff \\
\beta(2-\beta)u &> 2\beta-\frac{3}{2}\beta^2 \iff \\
u &> \frac{4-3\beta}{2(2-\beta)}
\end{aligned}
$$

This is the same inequality that guarantees that $V_m^M > V_H^I$; hence if u satisfies this inequality, all individuals are better off in the mixed ranking regime than in the ranking by income alone.

We first note that the right hand side of the inequality 1 when $\beta = 0$ and $1/2$ when $\beta = 1$, and is decreasing as β increases. In other words, for any

$\beta \in (0,1)$, there is $\hat{u} \in (1/2, 1)$ such that for $u \in (\hat{u}, 1)$, the mixed ranking regime Pareto dominates the income only ranking regime. The interpretation is straightforward: higher u corresponds to more a concave utility function, that is, higher risk aversion. Higher risk aversion translates into higher value for the insurance that is provided by the mixed ranking regime, and for a fixed β, if agents are sufficiently risk averse, all agents have higher utility in the mixed regime than in the income only ranking regime. We could have altered the roles of β and u in the interpretation of course. That is, for any $u \in (1/2, 1)$, there is a $\hat{\beta}$ such that for all $\beta \in (\hat{\beta}, 1)$ such that for the problem with u and β, the mixed regime is preferred by all individuals to the income only ranking regime. In other words, for any risk aversion, if agents' discount at sufficiently low rates, the mixed regime is Pareto superior.

We note that for any $\beta < 1$, $\beta < 4 - 3\beta$. Hence the inequality for the stability of the mixed attribute matching, $u \geq \frac{\beta}{4-2\beta}$, is automatically satisfied whenever the inequality for the Pareto superiority of the mixed attribute ranking is satisfied, $u > \frac{4-3\beta}{2(2-\beta)}$. In other words, whenever the mixed attribute ranking is Pareto superior, it is stable.

3.2.4 Dynamic social arrangements

As mentioned above, norms in a society change over time. A common lament is that "people just don't care about the things that used to be important." The analysis above shows that while there is always a ranking based on income alone, there may be other rankings that depend on payoff irrelevant characteristics that are stable for some values of the parameters v and β. One could simply assert that the change in values is captured by a switch from one equilibrium ranking system to another. There are two objections to this approach.

The first objection is that we would like the change in norms in a society to be endogenous, that is, we would like the change to arise from the underlying characteristics of the society. If we simply assume that a society switches from one equilibrium to another, that provides no understanding of *why* the change took place.

The second objection is more serious. For each of the two ranking systems to be stable, there was an incentive constraint that no unmatched pair of individuals would prefer to match together rather than follow the suggested matching under the proposed ranking regime. The verification of the incentive constraints was based on the assumption that the ranking system was permanent. If agents understand that the ranking system may change in the future, this should be incorporated into the incentive constraints.

More concretely, in the mixed ranking system we analyzed above, individuals of type (H, l) prefer to match with a type (L, h). The type (H, l) is trading off the present period utility cost of not matching with a type (H, l) and getting higher consumption with the benefit of matching with a (L, h) that results in positive probability of having offspring with the desirable attribute, h, which assures those offspring higher consumption. The higher expected consumption of offspring that compensates for the immediate lower consumption is less valuable if there is a chance that future generations won't "honor" the claim to higher consumption expected for individuals with attribute h.

The approach in MP is to construct an equilibrium in which the ranking is stochastic, with the change in the ranking system arising from changes in the environment. The basic idea is that, as we showed above, the possibility that a mixed ranking system is stable depends on the relationship between u and β. The discount factor β is fixed, but we introduce income growth into the basic model. A high income individual who matches with a low income individual has lower utility from consumption than if he had matched with another high income individual. The utility difference, however, generally depends on the two income levels. If there is rising income, the "risk premium" an individual will pay to ameliorate the riskiness in future generations' consumption may decrease. If this risk premium *does* decrease, it may be that the mixed ranking status equilibrium may no longer be stable. We next analyze how this may happen in equilibrium.

We extend our basic model to include stochastic income growth. We maintain the two income distribution that we have analyzed above, but allow the possibility that there is a one-time income increase that occurs at a random time. (We consider the case of perpetually increasing incomes below.) As above, initially there are two income levels, $\{H, L\}$. There is a pair of higher potential incomes, $\{\alpha H, \alpha L\}$, $\alpha > 1$; in each period there is probability p that the income increases from $\{H, L\}$ to $\{\alpha H, \alpha L\}$. Once incomes increase, they stay at that level permanently.

This particular income growth process preserves the relative incomes; only the level changes. Thus if the utility function u exhibits constant relative risk aversion, the incentive constraint that determines whether the mixed ranking system is stable is satisfied at the initial income level if and only if it is satisfied at the higher level. That is, the introduction of stochastic income growth doesn't change the possibility of equilibria with rankings other than the income alone ranking.

Suppose, however, that the utility function u exhibits decreasing relative risk aversion. In this case the risk premium associated with the random consumption of future generations is smaller after the income increase than before, and the incentive constraint requiring a type (H, l) to prefer matching with a type (L, h) to matching with another (H, l) may not be satisfied after the income increase. If this is the case, only the ranking that depends only on income is stable after the income increase.

Can it be the case that prior to the income increase the mixed ranking system can be stable? As before, we want to consider "rational expectations" equilibria, that is, we are interested in whether prior to the income increase, the mixed ranking can be stable when the agents *know* that there is a chance that the norm breaks down in any period, and hence, that it *must* break down eventually. We designate by $V^M(t|p)$, $t \in \{(H,h),(H,l),(L,h),(L,l)\}$ the value functions for each type of individual for the mixed ranking system in the case that there is probability p that in any period incomes do not increase, and accordingly, we change to the income-only ranking system with probability $(1-p)$. It must then be the case that

$$\begin{aligned}
V^M((H,h)|p) &= u + \beta p \left[V^M(\frac{1}{2}(H,h)|p) + \frac{1}{2}(H,l)|p) \right] \\
&\quad + \beta(1-p) \left[\frac{1}{2} V^I_H(\alpha(H,L)) + \frac{1}{2} V^I_H(\alpha(H,L)) \right] \\
V^M((H,l)|p) &= u + \beta p \left[V^M(\frac{1}{4}(H,h)|p) + \frac{1}{2}(H,l)|p) + \frac{1}{4}(L,l)|p)) \right] \\
&\quad + \beta(1-p) \left[\frac{1}{2} V^I_H(\alpha(H,L)) + \frac{1}{2} V^I_H(\alpha(H,L)) \right] \\
V^M((L,l)|p) &= u + \beta p \left[V^M(\frac{1}{2}(H,l)|p) + \frac{1}{2}(L,l)|p) \right] \\
&\quad + \beta(1-p) \left[\frac{1}{2} V^I_H(\alpha(H,L)) + \frac{1}{2} V^I_H(\alpha(H,L)) \right]
\end{aligned}$$

That is, the value functions are a convex combination of the value functions for the mixed ranking case when income is unchanged, and the value functions for the income only ranking case with higher incomes, with the weights given by the probabilities of the two regimes in the next period. But as $p \to 1$, the value functions $V^M((\cdot,\cdot)|p)$ must converge to the value functions for the case in which $p = 1$, that is the value functions calculated in the previous section. Hence, if the incentive constraint for the case in which income is unchanging is satisfied with strict inequality, then for sufficiently high probability p, it be satisfied for the case in which incomes increase with probability p.

To summarize, if at the initial income levels, the mixed ranking system is stable with a strict inequality in the incentive constraint, and if at the higher income level the incentive constraint is not satisfied, there is an equilibrium in which matching is based on the mixed ranking system until incomes increase, at which point the ranking system must change to the income only ranking.

We can generalize this observation to the case of perpetually (stochastically) increasing incomes. In each period there are two income levels. In the first period, the incomes are $H_1 = H$ and $L_1 = L$. In period t, the incomes will be $\{\alpha H, \alpha L\}$, $\alpha_t \geq 1$. As before, the relative wealth levels stay the same but the incomes grow over time. The income factors α_t are stochastic with $\alpha_t = \alpha_{t-1}$with probability $p \in (0,1)$, and $\alpha_t = \alpha_{t-1} + \gamma$, $\gamma > 0$, with probability $1 - p$.[8] Suppose that the utility function u exhibits decreasing relative risk aversion. Then the value of the insurance to a type (H,l) individual that results from a match with a high attribute partner is decreasing. If at some point, it is not sufficient to offset the immediate utility loss from consumption that results from a match with a type (L,h), the mixed ranking status regime is not stable, and matching will be based on income alone. However, if the initial income levels are such that the incentive constraint on the mixed matching are satisfied with strict inequality, then for sufficiently small p, there is an equilibrium in which the norm governing matching is the mixed ranking as long as that incentive constraint is satisfied with strict inequality, and matching governed by the income only ranking after that. We note that if $R(x) = -x\frac{u"(x)}{u'(x)} \longrightarrow 0$ as $x \longrightarrow 0$, and the incentive constraint for the mixed ranking case is eventually violated with probability one.

This result can be interpreted as the sure (eventual) demise of ranking systems that depend on non-payoff relevant criteria if there is asymptotically vanishing relative risk aversion. It is interesting to note that at the point at

[8] We assume that the increases in income, γ, do not depend on the period or the current income level for expositional ease only. We could allow the size of the increases to depend on these without changing any of the analysis as long as the increases are bounded above. Similarly, the probability that incomes may rise at any time may depend on the period and the current income level; the constraint will be on the maximum probability of an income change in any period.

which the matching regime changes, there may be only a small change in the income distribution, but a large change in the distribution of consumption. Under the income only ranking system, all high income individuals match with other high income people, while in the mixed regime, half the high income individuals match with low income individuals. The collapse of the mixed ranking system is accompanied by a large increase the "assortativeness" of matching.

3.3 Discussion

The last model abstracts from many interesting and potentially important aspects of social arrangements. The results are not meant to be definitive statements about particular features that must be exhibited by social arrangements, but are meant to illustrate how, in the presence of incomplete markets, payoff irrelevant aspects of social environments can have economically important consequences under plausible social arrangements. The equilibrium discussed in section 3.2.4 also demonstrates how particular social arrangements in a society can be inherently unstable, necessarily disintegrating as the economy in which they operate grows and changes. The simple model analyzed above focussed on two different social arrangements, neither of which involved productive attributes, that is, attributes that affected output. Mailath and Postlewaite (2000) examine in more detail social arrangements in which attributes such as education can be both productive, and simultaneously, play a role in social arrangements. This allows the possibility that disintegration of particular social arrangements as illustrated in the last section can have important and immediate effects on aggregate economic variables.

4 Bibliography

Akerlof, G. (1997) "Social Distance and Social Decisions," *Econometrica*, pp 1005-1024.

Bernheim, D. (1994) "A Theory of Conformity," *Journal of Political Economy*, pp 841-877.

Cole, H., G. Mailath, and A. Postlewaite (1992) "Social Norms, Savings Behavior, and Growth," *Journal of Political Economy,* 100, pp. 1092-1126.

Cole, H., G. Mailath, and A. Postlewaite (1995)Erratum: "Response to 'Aristocratic Equilibria'," *Journal of Political Economy*, 103 pp 439-443.

Goldin, C. (1992) "The Meaning of College in the Lives of American Women: The Past One-Hundred Years," paper written for the Conference on Women's Human Capital and Development, May, 1992, Bellagio, Italy.

Lancaster, K. (1971) *Consumer Demand: A New Approach*, Columbia University Press, New York.

Mailath, G., and A. Postlewaite (2000) "Social Assets," in progress.

Parente, S. and E. Prescott (1999) *The Barriers to Riches*, manuscript.