Envy and efficiency

Joseph Heath[1]

The Pareto principle states that if a proposed change in the condition of society makes at least one person better off, and does not make anyone else worse off, then that change should be regarded as an improvement. This principle forms the conceptual core of modern welfare economics, and exercises enormous influence in contemporary discussions of justice and equality. It does, however, have an Achilles' heel. When an individual experiences envy, it means that improvements in the condition of others may worsen the condition of that individual. As a result, envy has the potential to block a vast range of changes that we might intuitively be inclined to regard as Pareto improvements. (Or more precisely, envy results in *too many* states getting classified as Pareto-optimal, not because, intuitively, they cannot be improved upon, but because no one's condition can be improved upon without making someone else envious.) For example, a market exchange between two people might not wind up being classified as a Pareto improvement if the benefits produced for the two parties generated envy in some otherwise uninvolved third.

Some theorists take this as sufficient reason to dismiss the Pareto standard entirely. These views are not my concern here. The more common response, among both economists and philosophers, has been to treat preferences that involve envy as illegitimate. This is not to deny that agents have such preferences. The claim is simply that these preferences should carry no weight, from the normative point of view. The most memorable formulation of this idea was provided by Robert Goodin (1986), with his image of "laundering preferences." The claim, roughly, is that a theory of justice will take preferences as "input," then apply some set of principles, or a social welfare function, in order to generate a normative ranking of possible states as "output." Such a theory may contain both an input and an output filter. The former will filter out illegitimate (e.g.

[1] University of Toronto – joseph.heath@utoronto.ca

"sadistic or meddlesome") preferences, so that only certain preferences will count when it comes to determining social welfare. The latter will filter out recommendations that are trumped by other normative considerations (e.g., certain improvements in welfare might get ruled out, because achieving them would require unacceptable violations of individual rights).

In principle, theorists who find envy objectionable might choose to filter it out on either the input or the output side. For example, if envy truly is one of the seven deadly sins, then there might be a deontic prohibition against it. This prohibition could be sufficiently strict that it trumped the social welfare gain associated with a particular transformation, and so filtered out envy on the output side. However, no contemporary theorist seems to have advanced such a view. The general consensus is that envy is to be filtered out on the input side. Among philosophers, two proposals have dominated the discussion. John Rawls stipulates that rational agents behind the veil of ignorance "do not take an interest in one another's interests" [Rawls (1971), p. 13], while David Gauthier excludes any "tuistic" preferences from consideration in his bargaining theory [Gauthier (1986), p. 87].[2]

I would like to register some concerns about these two filtering strategies. First, I believe that they are far too broad. The way that "envious" preferences get filtered out has the effect of completely invalidating any concern people may have about their *relative standing* in the economy or in any other domain of social interaction. This has meant that theorists like Rawls and Gauthier have had almost nothing useful to say about *status*, along with other concerns which, for better or for worse, occupy a very central role in people's value systems. Furthermore, there are many cases in which people might have a legitimate interest in adopting institutional remedies for collective action problems that involve envy. Laundering out envy from their preferences has the perverse

[2]For an overview of discussions outside philosophy, see McAdams (1992).

consequence of rendering the structure of these collective action problems invisible, from the normative point of view.

More generally, the broad exclusion of envy can be faulted for its implicit reliance upon a substantive theory about what sort of preferences people have. The general assumption is that envy-based preferences will be far more the exception than the rule. They are treated as essentially an aberration (like cruelty), which can safely be set aside. By and large, this presupposes a world in which people are primarily concerned with their own private affairs, and only secondarily with the affairs of others. Unfortunately, this sort of privatism is unrealistic. Agents often do not have a conception of their own "business" that is sufficiently independent of that of others [Frank (2005)]. Thus the interests of one cannot be cleanly detached from those of others. Furthermore, in many cases it is not the nature of people's desires as such that creates interdependence of preferences, but rather the nature of the goods that they desire. Thus we need to find a way to take the social embeddedness of individuals – and thus the relative character of many of their goals and projects – more seriously, without thereby undermining our ability to identify important Pareto improvements. This will require a more nuanced treatment of envy, one that filters out genuinely anti-social preferences, but which recognizes the legitimate concerns individuals may have over questions of relative standing.

1. Why not envy?

It is widely accepted that a theory of justice must contain some mechanism to balance the often conflicting interests of individuals. However, because people generally disagree in their conceptions of the good, it is not possible for the political philosopher simply to specify what these interests are. So-called "objective list" accounts of individual interests, or conceptions of the good, do not have a strong track record when it comes to achieving consensus.[3] What Rawls calls the "fact of reasonable pluralism," [Rawls (1991), p. 36], viz. the existence of reasonable disagreement over fundamental questions of value, does appear to be a permanent feature of liberal democratic societies. Thus most theorists of a broadly welfarist persuasion consider it necessary to assign considerable weight to individuals' own judgments, when it comes to specifying their interests. If a person says that she wants to climb Mt. Everest, then we should take this as *prima facie* grounds for taking her to have an interest in doing so (even if we can see no particular reason why anyone should want to climb a mountain). If she is for some reason prevented from doing so, we should regard her interests as having been frustrated (and thus, all things being equal, the world in which she is unable to do so as normatively inferior to one in which she is).

The need to assign individuals authority in the specification of their own interests can, however, be taken too far. The quickest way to accord individuals such authority is simply to take their preferences, as represented by their utility functions, as an articulation of their interests. Many welfare economists have adopted this strategy, using some version of the doctrine of "consumer sovereignty" in order to justify the decision.[4] The problem is that the standard decision-theoretic definition of preferences contains no substantive restrictions on the content of these preferences. Thus treating agents' utility functions as a representation of their interests leads to agents having "interests" at stake in all sorts of interactions that, from an intuitive perspective, we would regard as

[3] This why so-called "perfectionists" are so studiously vague when it comes to actually stating the nature of the good. Anyone who goes beyond stating platitudes to providing a list of concrete goods inevitably also invites controversy. Consider, for example, the list of "essential human goods" that Martha Nussbaum [(1992), p. 222] provides. Is it really essential that we all be "able to laugh, to play, and to enjoy recreational activities"? Perhaps some adults are just serious by nature.

[4] For a recent example, see Kaplow and Shavell (2002).

being absolutely none of their business. In some versions of decision theory, anything that the agent is capable of forming a belief about is also something that she can have a preference over. Thus, for example, the way that a complete stranger chooses to arrange the books on his shelves is something that I could have an 'interest' in. Obviously this is far too broad. No theory of justice can take such 'interests' seriously. In particular, as Amartya Sen (1972) argued long ago, any form of welfarism this broad has the potential to conflict in significant ways with the classical negative liberties.

A concern over how a stranger shelves her books is an example of what are sometimes called "intrusive preferences." They are preferences that concern states of the world that are, in some sense, unrelated to the agent, other than through the preference. (In Sen's example [(1972), p. 155], it is the fact that two individuals have preferences over what book the *other one* reads that generates the conflict between welfare and individual liberty.) If such preferences were given weight when it came to the design of institutional constraints, then the resulting rules would generate intolerable interferences in individual liberty. (Imagine that this person, while shelving her books, was *obliged* to consider the impact that her action had upon a perfect stranger who, for some idiosyncratic reason, had strong feelings about her shelving system.) However, while we all have a very strong intuitive sense of which preferences are intrusive and which are not, providing an explicit specification and defense of this intuition has proven to be extremely difficult. Probably the first concerted effort to deal with this problem was made by John Stuart Mill, in *On Liberty*, where he tried to determine what should count as a "harm" to another person. The issue remains extremely controversial to this day [Feinberg (1987)].

However, if there is one point on which there is very widespread agreement, it is that the dissatisfaction generated by the feeling of envy should not count as a harm to the person who suffers it (or, conversely, that the fact that the outcome of someone's actions may inspire envy in someone else should not be taken to give that second person a legitimate interest the outcome of the first's choice).[5] To see how this exclusion works, consider the standard argument for market exchange. The practice of trade is often defended on the grounds that when two individuals exchange goods, each is made better off from his or her own perspective. The reason that they agree to the exchange is precisely that, in the initial allocation of goods, each prefers what the other has to what he or she has. Thus the voluntariness of the exchange reflects the fact that it improves the welfare of everyone involved. It is a win-win transformation, a Pareto improvement.

Yet there are many cases in which market exchange generates so-called "externalities" – welfare effects on third parties uninvolved in the transaction. Imagine a situation in which my neighbor acquires an air conditioner. This is a purely private transaction between himself and the merchant. Unfortunately, as a consequence of this purchase, I may find myself, on sweltering days, glaring enviously across the yard, resenting the comfort enjoyed by my neighbor and his family. Does this undermine the win-win character of the transaction between the neighbor and the merchant? There is a very strong moral intuition which suggests that, in this sort of case, the loss of welfare that I experience from my neighbor's new acquisition should not count as a consideration that speaks against the transaction (and thus that my loss of welfare should not really "count" as an externality).

On the other hand, the air-conditioner might also make noise, which keeps me awake at night. Then we might not want to regard the purchase as purely a private matter between the neighbor and the merchant. I become an unwilling participant, and my loss of welfare, it seems, should count for something. Here we could probably all agree that the transaction creates a negative externality. Thus when economists and philosophers talk about markets, the "laundering" of preferences normally occurs in the (usually implicit) decisions that are made about which external

[5] For an especially clear formulation, see Nozick [(1974), pp. 245-46].

effects of a transaction are to be treated as "externalities" – and thus as part of the "social cost" of an action – and which are to be ignored.

Yet despite the intuitive attractiveness of these distinctions, it is difficult to formulate a precise articulation of the underlying logic. In both cases, my loss of utility is due to the character of my preferences. If I wasn't susceptible to envy, then I would not be upset by my neighbor's new-found comfort. But if I wasn't so sensitive to noise, or if I liked humming noises, then the fan would not bother me either. Thus the fact that we choose to treat noise as an externality, and not envy, cannot simply be due to the fact that it is *my* preference that is causing the loss of my own welfare (and thus that it is my responsibility to deal with it). There must be some deeper rationale. Furthermore, we would be ill-advised simply to rest the distinction upon our intuitions about which preferences are "intrusive" and which are not. After all, my sensitivity to noise may be so extreme that I wind up bothering my neighbor with all sorts of complaints about his routine activities. In this case, we would be inclined to handle the "intrusiveness" of my preferences by balancing my sensitivities against his liberty interests. Yet in the case of envy, we generally do not try to balance interests, we simply rule out one set of them *ab initio*. So what is it about envy-based preferences that makes them completely illegitimate?

2. Attempts to define envy

Let us take a very concrete case. In a study conducted by Daniel Zizzo and Andrew Oswald (2001), groups of four experimental subjects were given sums of money which they then wagered on a series of gambles. The gambles were arranged so that they would generate, at the end of the experiment, two winners and two losers. The losers were then given the option of using some of the money they still held in order to "buy" a reduction in the winnings of the others. They were offered rates between $0.02 and $0.25 to reduce one of the winners' take by $1. Disconcertingly, 62 per cent of the experimental subjects choose to "burn" the winners' money even at the high price of $0.25 on the dollar. Thus a majority clearly derived strong satisfaction from seeing the holdings of others reduced – enough that they were willing to pay for the pleasure.

Many people quite rightly see this as a deeply disturbing experimental finding. Without getting into its broader social implications for theories of fairness, I think it is sufficient to note that the type of preferences that this experiment brings to the fore are precisely the ones that are being targeted when we try to filter out envy on the input side of a social welfare function.[6] The pleasure that people take in seeing the wealth of others destroyed should not count as a "plus," from the normative point of view. Yet while there may be agreement that this is the phenomenon we are concerned with when we try to launder out envy from a social welfare function, there is no agreement over how we should specify the underlying preferences, or determine what makes them worthy of exclusion.

The common-sense approach is to treat envy as a specific type of emotion, a feeling that we get under certain circumstances. Such emotions are often regarded as anti-social, and so would fall into the same category as cruelty. We might regard such states as being intrinsically wrong, and thus

---

[6] These sorts of examples do, however, suggest that it is misleading to use the term "envy-freeness" as a way of talking about fairness [see, e.g., Varian (1974), Dworkin (2001)]. The concept of envy-freeness is really a way of talking about a set of compossible preferences with regard to a proposed distribution of goods. Thus "envy" in this sense is just preference, and has nothing to do with envy in the colloquial sense of the term, or in the sense that it is used in this paper. Use of the term envy-freeness, however, appears to lend credence to the old accusation that partisans of equality are really just motivated by envy [e.g. Nozick (1974), p. 240]. It can also lead to a highly unwelcome conflation, such as one finds in Sussangkarn and Goldman (1983), of the distinction between essentially anti-social feelings of envy and a principled concern for social justice.

unworthy of consideration in the theory of justice because of their very nature. Thus we could say, for example, that experimental subjects in the example above acted out of spite, and that motives of this type do not deserve to be recognized and respected.

The problem with this view is that it depends upon precisely the sort of ethical judgment that the move toward first-person authority (or subjectivism) in the definition of interests was intended to avoid. Who are we to say what emotions other people should and should not be experiencing? It is highly problematic to suggest that there is something intrinsic in the nature of certain feelings that disqualifies them from consideration, when it comes to calculating social welfare. As a result, philosophers have generally tried to develop more formal definitions of envy, in order to determine the structural characteristics of the preference that make it normatively problematic. The goal is to exclude envy on more technical grounds, so that we are not simply left saying "it's *wrong* to feel that way."

From this perspective, the problem with envy is often said to arise from the way that these preferences are connected up with *other people's* preferences and interests. According to this view, it is the fact that my neighbor is able to improve his condition through the purchase of an air conditioner that serves as the source of distress in the envious person. There have been two major attempts to formalize this intuition. The narrower version excludes envy on the grounds that it is a type of *interdependent preference*. This is David Gauthier's view (1986).[7] A somewhat broader version excludes envy on the grounds that it is a *relative preference*. Rawls's formulation (1971) lends itself to this interpretation. The distinction is subtle, but quite important.

1. *Interdependent preferences*. It is relatively easy to specify what counts as an interdependent preference. In standard utility theory, an agent's preferences can be defined in terms of a ranking of sets of

[7] It is also the definition adopted by Zizzo and Oswald (2001).

possible worlds [Jeffrey (1983)]. The preferences of one individual depend upon those of another if the individual is *not* indifferent between two worlds that are identical except for the way that they are ranked by some other agent. The easiest way to ensure this is simply to define possible worlds in a way that excludes the intentional states of all other agents. This way we guarantee that each agent's ranking will be completely independent of everyone else's.[8] (This is, in fact, the standard practice in decision theory.)

Note that this definition of interdependent preference uses the term in a somewhat stricter sense than some economists have been inclined to do. Robert Pollak, for instance, refers to interdependent preferences as "preferences which depend on other people's consumption" [Pollak (1976), p. 309). Similarly, Zizzo (2005) defines interdependent preferences informally as ones in which "*A*'s utility function depends on those of other agents," but then goes on to present a formal definition in which *A*'s utility is in part a function of *B*'s "material consumption," rather than *B*'s utility. This elides an important philosophical distinction. Thus it is worth emphasizing that Gauthier uses the term literally, to refer to preferences that depend upon the preferences – not the consumption – of others. (Speaking loosely, one could say that he is concerned with "I want what you want," rather than "I want what you have" preference structures.)

This notion of interdependent preference can be used to rule out three problematic types of preference. We can stipulate that a preference involves *envy* if the valuation of a possible world moves down in one agent's ranking in response to its moving up in some other's. If that possible world became actual, we would say that the first agent's loss of utility is *due to* the other's gain (with the implicit counterfactual, that the first agent *would not have* lost utility had the other not gained).

[8] Thus Gauthier's concept of "mutual unconcern," in *Morals by Agreement*, amounts to the constraint that the agent's "utility function, measuring his preferences, is strictly independent of the utility functions of those whom he affects" [(1986), p. 311].

Schadenfreude (or simply malevolence) can be defined as the opposite of envy – it occurs when the valuation of a possible world moves up in one agent's ranking in response to its moving down in some other's. Finally, sympathetic preferences are those that move some state up in one agent's ranking in response to its moving up in some other's, or move it down in that agent's ranking in response to its moving down in some others.

The limitation of this definition is that it is extremely welfarist, precisely because it looks only at agents' utility functions. Thus envy is said to arise only when the *happiness* of others (in the preference-satisfaction sense of the term) is a source of unhappiness for ourselves. However, in many concrete cases, what provokes envy is not so much the happiness of others as it is their possessions. I may envy my neighbor's car even after he has become indifferent to its charms. It is the fact that he owns it, and not me, that is the source of my distress. If he doesn't even like it, that may exacerbate the situation. Of course, in the more usual run of cases people simply have no idea how other people feel about various situations, whether they are really happy or not. Thus a definition of envy in terms of interdependent preferences runs the risk of being too abstract to be of any use.

2. Relative preferences. In order to expand the definition of envy to deal with these concerns, it is necessary to consider not only the structure of agents' utility functions, but also the material environment in which they are situated. Rawls, for example, assumes not only agents with preferences, but also a system of distribution under which each individual receives an allocation of primary goods. He then launders out envy by stipulating that, under each category of primary good, agents prefer an allocation that gives them more over one that gives them less, independent of what anyone else receives in the distribution scheme. He refers to this as "mutual disinterest" [Rawls

(1971), p. 13].[9] Agents' evaluation of distribution schemes behind the veil of ignorance is governed by the same constraint. Thus they "do not take an interest in one another's interests," and they do not "try to gain relative to each other" [Rawls (1971), p. 144], because their preferences over any distribution are determined entirely by the individual allocation that they receive, and not by any concern about what others receive. (In the spirit of Rawls's work, one can think of these preferences as the ones the individuals would hold over allocations if all information about what everyone else received was hidden behind a veil of ignorance.[10])

It should be noted that Rawls' definition, unlike Gauthier's, has the effect of laundering out any preferences that people may have over the distribution of wealth (such as a preference for an equal rather than an unequal distribution). This is in accordance with our intuition that my desire to see Tom receive more than Bill cannot possibly count as an argument in favor of Tom receiving more than Bill. Nevertheless, such a distributional preference may not depend, in any crucial way,

[9] Later, when he considers problems of implementation, he provides the following definition of envy: "We may think of envy as the propensity to view with hostility the greater good of others even though their being more fortunate than we are does not detract from our advantages" [Rawls (1971), p. 532].

[10]Serge-Christophe Kolm (1995) has proposed an interesting variant of this strategy, distinguishing what he calls "intrinsic preferences" from those that are in some sense relative to the standing of others. Rather than restricting the normatively salient set of preferences to those that the agent would have over allocations were he to be ignorant of what anyone else receives, Kolm restricts the set of preferences to those that the individual would have over allocations under the assumption that everyone else would receive the same. This has the effect of eliminating any role that invidious comparisons might play in determining the agent's preferences. As will be seen below, this proposal has the same problems as Rawls's when it comes to dealing with positional goods.

upon either Bill or Tom's preferences, and so it may not exhibit any objectionable sort of interdependency. The Rawlsian constraint excludes it, however, by specifying that, where g_x is the quantity of primary goods that individual x receives under a particular allocation scheme $U = (g_1, g_2, g_3, \ldots, g_n)$, for any other allocation scheme U', if $g'_x \succ_x g_x$ then $U' \succ_x U$. Envy is essentially a limit case of a distributional preference. (This is why, incidentally, it is unwise to subsume concerns over fairness or equality under the concern for efficiency, by treating them as merely a "taste for equality," and thus as one more preference that needs to be satisfied. Many people, unfortunately, especially economists, make this argument.[11])

These two general strategies for identifying envy-based preferences differ considerably, yet are often not adequately distinguished in the literature. The former (Gauthier's) is much narrower, and precludes only a rather specific, essentially misanthropic, preference structure. The second (Rawls's) is much broader, and is much closer in spirit to the idea that agents should "mind their own business." Nevertheless, there are problems with both approaches. My concern is not with the question of whether these definitions adequately capture what we consider objectionable about envy-based preferences. My concern is that both definitions can have the effect of laundering out normatively significant features of certain social interactions.

3. Problems with interdependence

To see the problem with laundering out interdependent preferences, it is helpful to recall the basic structure of the Pareto standard. This principle, it may be recalled, does not provide a complete ordering of possible social states, as Figure 1 shows. If point s is the status quo, then any point to the north-east of that point (quadrant 1), represents a Pareto-improvement. This includes

[11] For an especially clear version, see Kaplow and Shavell [(2002), p. 444], also Sussangkarn and Goldman (1983).

every point on the two axes that intersect *s*. Speaking loosely, these can be described as the set of win-win outcomes (speaking precisely, they are the set of no-lose outcomes). Every point to the south-west (quadrant 3) is Pareto-inferior. These are the lose-lose, or no-win outcomes. Points in quadrant 2 and 4 are Pareto-noncomparable to *s*. They represent win-lose outcomes.

The elimination of envy-based preferences is designed to correct situations like those shown in the move from *s* to point *a*. Player 2's move from y_0 to y_1 is private and self-regarding, yet it provokes envy in player 1, whose utility is thereby reduced from x_0 to x_1. Thus taking player 1's preference into account, from a social welfare perspective, would allow him to block player 2's effort to improve her own welfare. Disregarding player 1's preference would make the move to *a* normatively equivalent to the move to *a'*. So while in reality the move away from the status quo represents a Pareto-noncomparable transformation, the laundering of preferences will represent it as a Pareto-improvement, from the normative point of view.

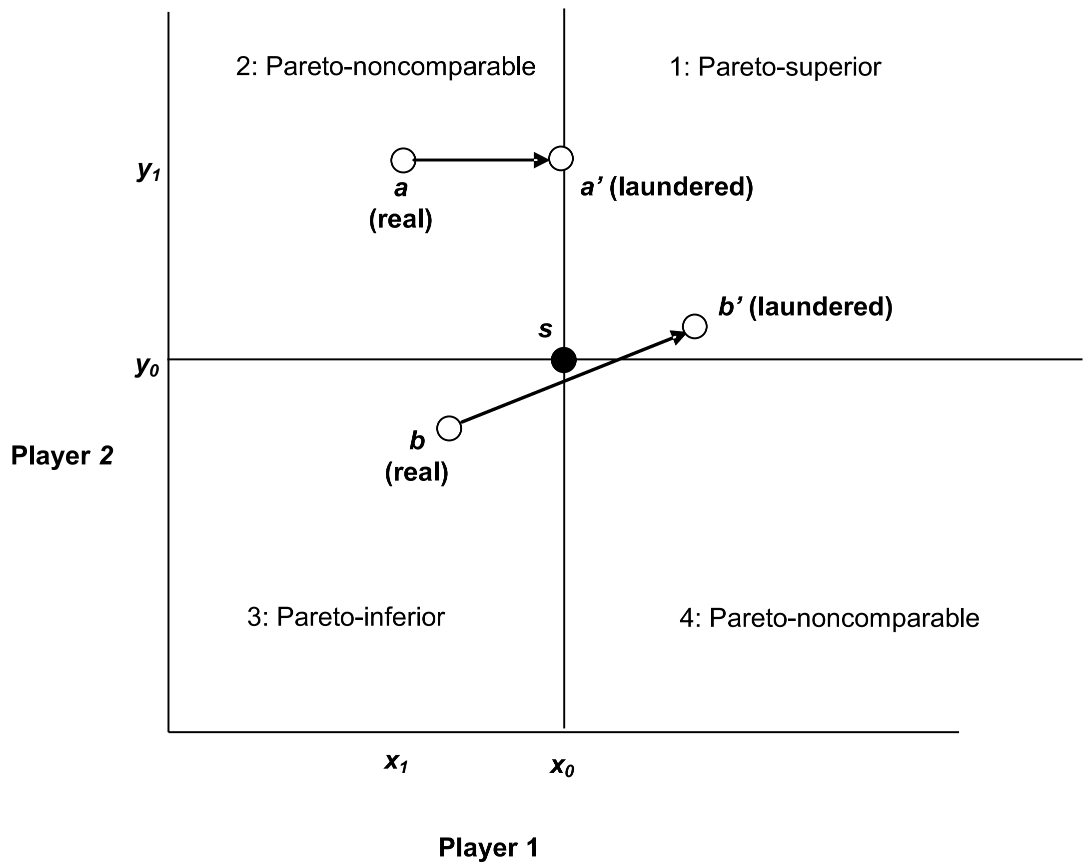

Figure 1. Laundering preferences

This seems perfectly plausible, as long as the laundering procedure simply turns Pareto-noncomparable states into Pareto-superior ones. It seems unreasonable that someone should be able to block a project that makes someone else happy, simply because he resents that person's gains. (Think, for example, of the normative intuitions underlying the "Lockean proviso" [Nozick (1974), pp. 174-182] in various versions of social contract theory.) Thus the laundering procedure prevents one person from unilaterally setting the terms of interaction.

Our intuitions become much murkier, however, when the proposed "real" transformation is not to a Pareto-noncomparable outcome, but rather to a Pareto-inferior one. Imagine, for example, two lovers who, when they quarrel, like to hurt one another by kissing someone that each knows the other hates. They both do it in order to provoke the other's envy. While each derives some slight enjoyment from kissing the other's rival, this is entirely outweighed by the displeasure caused by the look of triumphant self-satisfaction on the face of his or her own rival. Thus the interaction has the structure of a collective action problem. We might easily imagine this couple, in a cooler moment, swearing that they will each desist from such behavior in the future. They might even arrange things so that the person who breaks the agreement is punished in some way for doing so. However, if we look at the agreement from a social-welfare perspective, the underlying rationale for it will unfortunately be laundered out. Because the unhappiness that the couple seeks to avoid is the result of envy, the suboptimality of the interaction is invisible. Points *b* and *b'* in Figure 1 show the effects of such laundering. The outcome in which the lovers exacerbate their conflict shows up as a Pareto-improvement, from the normative point of view, whereas in reality it is suboptimal.

Thus the initial motive for laundering out envy – to prevent individuals from being obstructionist and blocking changes that stand to benefit others (in cases like *a*) – does not exist here. In the case of mutual envy, there is no third party to the interaction whose gains would be blocked. Everyone involved has an incentive to preserve the status quo. Yet perversely, laundering their preferences may make a transformation in which they all lose look like one in which they all gain. Agents have good reason to enter into an agreement that allows them to avoid such outcomes. And in principle there is no reason why they should not be able to use formal institutional arrangements to enforce such agreements.

4. Problems with relative preference

The example of the quarreling lovers is somewhat unusual, in that the individuals are assumed to know what the other is feeling, and act with the specific intention of provoking envy in the other. In the more normal run of cases, people do not really know or care what others are feeling, they are concerned with external states. They may want to earn *more than* their siblings, or drive a car that is *faster than* their coworkers', or have a house that is *larger than* their neighbors'. These are all relative preferences. In order for one person to satisfy such preferences, it is necessary that someone else's be frustrated. When one person buys a faster car (or, as we say, a car with "better pickup"), it makes everyone else's slower by comparison. Thus each gain in welfare for one is typically accompanied by a loss of welfare for others (assuming that people assign at least some value to the same set of properties in such goods).

However, the welfare loss suffered by individuals with comparative preferences is normally laundered out as a type of envy-based preference. According to the standard Rawlsian view, it is reasonable for people to care about the size of their own houses, or salaries, but it is none of their business how large their neighbors' are. When one thinks about it, however, this prohibition is somewhat puzzling. After all, the goal of laundering out envy was to eliminate intrusive preferences, so that individuals are not able to interfere arbitrarily in the affairs of others. But from the fact that I have no legitimate stake in the question of how much Tom earns, relative to Bill, it does not follow that I have no legitimate stake in how much Tom earns, relative to *me*. While it would be pretty easy to show that the former is none of my business, it is much harder to do so in the latter case, simply because the relative size of my salary has important consequences in terms of the quality of life that I may aspire to (especially if I find myself in a bidding war with Tom for a house that I would like to buy).

To illustrate, consider the little quiz that economist Robert Frank asks his readers. Imagine two worlds: In World *A*, you would earn $110,000 a year, while everyone else would earn $200,000 a year. In World *B*, you would earn only $100,000, $10,000 less than in World *A*. But everyone else in World *B* would earn only $90,000. In short, your absolute income would be a little lower in World *B*, but your relative income would be much higher. (To add a science-fiction touch, assume that prices are the same in the two worlds, so the income difference represents a real difference in living standards [Frank (1999), p. 129]).

Is it really so unreasonable to select World *B*? More pointedly, is a preference of this sort, informed as it is by a concern over relative standing, so unreasonable that we should want to completely ignore it entirely when making judgments of social welfare? On the contrary, the person who selects *A* would seem to be either someone with slightly idiosyncratic preferences, or else someone who misunderstands how society really works. All of the major components of the good life have a deeply competitive structure. The person with the above-average income can buy the house of his choosing, can walk into any restaurant without checking the prices, has access to exclusive clubs, has access to a wider range of sexual and marriage partners, and so on. The important point is that being able to pay *more than* others does not just mean that one is able to purchase proportionately more, it often means that one is able to acquire a disproportionate share of goods that are qualitatively superior, because one is able to win a greater number of competitions.

Thus the person who chooses World *B* can hardly be characterized as irrational. In fact, Frank reports that a solid majority of people, after some deliberation, opt for *B*. Furthermore, there is some empirical evidence to suggest that *B* is the wiser choice. Various attempts to measure subjective well-being – feelings of happiness – have shown that there is essentially no correlation between increases in average wealth and average happiness in wealthy industrialized nations [Frey and Stutzer (2001)]. People living in very poor countries, who lack the basic necessities of life, are on average much unhappier than those who live in wealthy countries. But once annual income reaches an average of about US$10,000 per year, the correlation between economic growth and increased happiness disappears completely. The one thing that does remain constant, however, is the correlation between happiness and relative wealth. Beyond a certain point, absolute wealth counts for very little, but relative wealth continues to be important [Easterlin (1974); Frank (1999), pp. 111-115; Frank (1985), pp. 26-32; Scitovsky (1992) ,pp. 134-136] .

In other words, while money can't buy happiness, having more money than others does increase your chances. What could explain this phenomenon? There are no doubt a number of different forces at work.[12] One factor that seems especially important, however, is that in very poor countries, the increase in satisfaction due to economic growth stems largely from the fact that the sorts of goods produced at the advent of industrialization are valued primarily for their intrinsic properties. Access to food, clothing, shelter, clean water, plumbing, electricity, and so forth, brings about an absolute increase in well-being. However, as people become richer, they tend to spend a

[12] Two major explanations that should be noted, but that I will not discuss further, are: 1. the "rising aspirations" account, which argues that the social and economic environment in a given society serves as a benchmark for the development of individual preferences, and thus economic growth simply ratchets up expectations, leaving individual welfare largely unchanged [Easterlin (1995)]; and 2. that human psychology equips us only to respond to *changes* in consumption, not the absolute level (in the same way that when traveling in a car we can only feel acceleration, not velocity), and thus every increase in the baseline generates only a temporary gain in satisfaction [Scitovsky (1992)]. These theories are similar enough that they are often lumped together under the heading "adaptation of preference." My personal view is that this emphasis on adaptation pays too much attention to the individual psychological dimensions of consumption, and not enough upon its social aspects (for example, the way that individuals use their consumption choices as a way of communicating with others, or establishing a social identity). This is why, in the discussion that follows, I privilege the "neo-Veblenian" account, with its emphasis upon the social. The important thing to note, however, is that none of these explanations are mutually exclusive; not only does each capture an important dimension of the truth, but there is also substantial overlap between them.

greater proportion of their income on goods that are valued more for their relative qualities. Food becomes increasingly valued for its taste, housing for its charm, real estate for its location. As a result, the positive-sum character of the economy diminishes [Hirsch (1976), p. 67]. This would explain the fact that, as society becomes richer, relative income remains an important determinant of happiness, while absolute income declines in significance.

It is not difficult to observe this effect in the world around us. If one looks at the major categories of household expenditure, one can see that almost all members of wealthy industrialized societies have access to pretty much the full range of goods. There are very few people who cannot either directly afford, or else secure access to, housing, food, clothing, transportation, etc. The major distinction between social classes lies not in what they possess, but rather in the status associated with the goods that they do possess (the brand, the style, the quality, the location, etc.). In other words, differences in consumption are now largely demarcated by differences of taste. Material scarcity is no longer the issue, "social" scarcity is what consumers are largely striving for. Elites in our society typically just have nicer versions of the same stuff that the lower classes have. Their cars offer a smoother ride, their houses are more spacious, their clothing is more tasteful. These qualities are all relative, because any judgment made with respect to them is implicitly benchmarked to other goods in the same category. A "good" bottle of wine, for example, is simply a bottle that is better than what most people drink. Since Thorstein Veblen's seminal work [(1967), p. 169], this process has been a well-recognized feature of our consumption habits.

But if envy is defined in terms of a concern over relative standing, then a desire to drink only good wine, or live in a spacious home, or work in a prestigious occupation, necessarily commits the agent to a degree of envy in his preferences. Thus the insistence that individuals "not take an interest in one another's interests" seems unreasonable restrictive. Furthermore, if the relative structure of many of our preferences is responsible for much of the failure of economic growth to produce increases in happiness, then preference laundering has the effect of robbing us of any insight into what is in fact one of the major features of our economy. After all, we make enormous sacrifices as a society in order to achieve a high rate of economic growth. We accept unemployment and social inequality, serious degradation of the natural environment, commercialization of our culture, and so on. But if this growth in absolute living standards is not leading to any increase in real (i.e. unlaundered) welfare, what is the point? We might be better off sacrificing some of this growth in order to achieve other objectives. For example, improvements in air quality could be achieved at the expense of slower growth. This might be welfare-enhancing, since air quality is a good that has intrinsic value (or at least no obvious relative dimension). However, it is impossible to make the argument for air quality *versus* private consumer goods if we launder out the preferences that give many markets for consumer goods a zero-sum structure at the level of individual welfare.

5. Legitimating envy

Consider the following, which I refer to as Hirsch's Law: "As the level of average consumption rises... the satisfaction that individuals derive from goods and services depends in increasing measure not only on their own consumption but on consumption by others as well" [Hirsch (1976), p. 2]. If this is correct, then it creates something of a dilemma. There are clearly cases in which we do not want to grant any authority to preferences that are based *merely* on a concern about what others receive. Such preferences are often intrusive. And yet we cannot ignore the fact that people on many occasions have very legitimate concerns about their relative standing, especially as society becomes more affluent. How do we decide where to draw the line?

One option is simply to say that "unilateral envy" will be ignored, but that "mutual envy" will provide a legitimate basis for collective action. Thus the onus would be upon those who suffer from envy to show that the interaction with those whom they envy has the structure of a collective action problem. If everyone could agree to a particular measure, then that should carry the argument, regardless of the structure of the underlying preferences. This would certainly work if our goal was to prevent certain forms of conspicuous consumption that have the characteristics of a race to the bottom. The fact that everyone benefits from the countermeasures would outweigh the somewhat dubious basis of the constraints.

This proposal, however, generates a number of subtle difficulties. The basic problem is that it gives people who suffer from "unilateral envy" (and thus find their preferences being laundered out of the social welfare function) an incentive to transform their interactions into Pareto suboptimal ones. For instance, if my envy toward my neighbor's new air conditioner is being disregarded, then I might choose to invest in some expensive landscaping, in order to make his front yard look shabby by comparison. I could then offer to stop if he agrees to stop buying things that I want. These are clearly not attractive incentive effects.

If this example seems a bit fanciful or problematic, it is presumably because it focuses upon (potentially capricious) subjective preferences. A more promising strategy for picking out "mutual envy" would be to assume an economic environment (rather than just looking at utility functions), then examine more closely the nature of the goods that serve as a focus for concern over relative standing. The first thing that springs to mind in this context is the distinction, introduced by Fred Hirsch, between "positional" and "material" goods. Material goods are those that, through increased production, we are able to create more of (and as good). Positional goods, by contrast, are "either (1) scarce in some absolute or socially imposed sense or (2) subject to congestion or crowding through more extensive use" [Hirsch (1976), p. 27]. This means that we are unable simply to expand the supply of positional goods, and so access to them will be determined by relative ability to pay. Thus consumers with preferences for positional goods must, by virtue of that very fact, have relative preferences. Furthermore, because these preferences are imposed by the nature of the good, one might safely assume that the envy among individuals is mutual in that regard.[13]

Consider, for example, the market for real estate. In a mature municipal region, while it is possible to increase the density of dwellings, it is impossible to create more good locations. As a result, the demand for real estate in a desirable location necessarily has a positional structure, since access to it will be determined by relative ability to pay.[14] The realtor's saying "location, location,

---

[13] Because of this mutual envy, markets for positional goods tend to have a zero-sum structure from the standpoint of consumers. On the supply side, however, they often have a winner-take-all structure, since "the demand for top rank can be satisfied by only a limited number of products in any given category… [and] people are often willing to pay substantial premiums for top-ranked products" [Frank and Cook (1995), p. 41].

[14] Recently, Elizabeth Warren and Amelia Warren Tyagi (2003) have taken critics of consumerism to task for failing to understand the true causes of the dire financial straits that many Americans find themselves in. They observe that while consumers may "waste" a certain amount of money on branded consumer goods, this constitutes an insignificant component of the average household budget. Furthermore, it is a component that can easily be cut out in the event of loss of employment, or any other financial disaster. What really drives Americans into bankruptcy is housing and cars – the two major categories of consumer expenditure. This is certainly a point that merits emphasis. However, the authors fail to observe that housing is the single most important category of competitive consumption (exacerbated, in the United States, by the fact that the quality of public education varies so greatly from neighborhood to neighborhood). Furthermore, a significant aspect of car ownership is the instrumental value of these vehicles in the competition for other positional goods, such as desirable suburban real estate locations. Thus their analysis, while serving as an important corrective to the old "consumerism as brainwashing" theory, is entirely

location," reflects the fact that it is this positional character of the good that makes up the preponderance of its value. Sometimes this is due to intrinsic physical scarcity, such as the existence of a waterfront, a ravine, or a neighboring park. Sometimes the scarcity is socially imposed. Downtown real estate, for example, will always be scarce, simply because downtown is generally where the people who want to live where everyone else is living live. Thus people who want to live downtown necessarily have an interest in the sort of consumption choices made by others, because the object that they desire is constituted by the choices of others. People who want to live in the wilderness are in the same situation, except that their desire is to live where no one else wants to live – that's just what "the wilderness" is.

Does it follow that a person who buys a house downtown, or has one built in the wilderness, creates a normatively salient "consumption externality" for other consumers, by virtue of squeezing them out of the market for some intrinsically scarce resource? In the case of goods that are merely subject to severe inelasticity of supply, the answer is "no." The externality produced is a pecuniary externality, which is in a sense not a real externality, since it does not result in any Pareto-inefficiency. (There are various ways of articulating the reasoning underlying this claim. One is to say that there is no externality because the social cost of consumption is fully internalized, in the form of the very high price paid for the positional good.) However, in the case of goods that are positional because they are subject to crowding or congestion, the case is quite different. Here there is a negative network externality present – the value of the good is a decreasing function of the number of others who consume it. This is akin to a technological externality [Scitovsky (1954); Holcombe and Sobel (2000)], which is very much a "real" externality. An example would be a good that is subject to what Harvey Leibenstein refers to as a "snob effect" (1950), such as wanting to live in a

---

consistent with the traditional Veblenian critique of consumerism.

"good" neighborhood. Here, exclusivity in consumption is one of the desired characteristics, and so any expansion of supply results in a degradation of quality (i.e. crowding).

The most important examples of goods that are subject to crowding effects are those that are consumed in order to confer status or prestige. Status is necessarily hierarchical – in order for someone to be on top, someone else must be below. Since position is relative, when one person moves up it necessarily pushes someone else down. Thus goods that are consumed for the status they confer are necessarily positional. Their value is based on their exclusivity. As a result, if too many of them are produced (e.g. "knockoffs"), then the status that is associated with their consumption will be displaced onto some new good that is still in limited supply. The very term "exclusivity" refers to the socially imposed scarcity of these goods – the point of owning them is that others are *excluded* from doing so. This is a case of what Hirsch calls "pure" social scarcity, where it is the scarcity itself that is implicitly valued [Hirsch (1976), p. 28). It is worth noting as well that not only luxury goods have this character. The central status hierarchy in contemporary urban consumption is "cool," which represents a system of classification that reflects an almost obsessive concern with exclusivity [Heath and Potter (2005)].

Social scarcity is a lot more common than we imagine, simply because we often don't realize that many of the aesthetic standards used to determine the value of goods have an implicit social benchmark. What counts as a "spacious" home, for example, is entirely determined by how large homes in the relevant comparison class are. As Frank points out, an apartment that seems gigantic in Manhattan would appear impossibly cramped by the standards of Palo Alto. The poorest 25 per cent of Americans enjoy an average living space of 112 m^2, compared to the overall European average of only 93 m^2; yet presumably the former feel more cramped than the latter. Pierre Bourdieu (1979) has argued, more generally, that all aesthetic judgments rely implicitly upon the social *distinction* that they establish, between those with good taste and those with bad taste. Thus beauty has an inherently comparative structure, and so goods that are valued for their aesthetic properties are all subject to quality degradation through crowding effects. As a result, we all have what amount to relative preferences, even though we often don't recognize them as such.

One can certainly see this thesis illustrated quite clearly in the area of physical beauty. The emphasis that feminists and other critical theorists have placed upon the role of "archetypes" in defining norms of physical beauty is highly misleading, in that it distracts attention away from the intrinsically relative structure of beauty [Heath (2001)]. There are very tangible advantages to be derived from beauty, yet these advantages are obtained, not by people who are beautiful in some absolute sense, but rather by people who are *more beautiful than* their classmates, their coworkers, the other singles in the bar, etc. In other words, beauty is relative. One can see this more clearly in the case of face-lifts. How old a person looks depends entirely upon what the average person that age looks like – there is no fixed criterion. As a result, the combination of cosmetics and plastic surgery simply makes people look *younger than* other members of their cohort. This is clearly a zero-sum game – making yourself look younger just makes others your age look older. Someone must be bumped down in order for anyone else to move up. Thus cosmetic surgery of this type generates a negative network externality, making any increase in absolute spending collectively self-defeating.

Once one begins to look at the economy in this way, it becomes clear that a concern with distinction is ubiquitous (so ubiquitous that the term "snob effect" is quite misleading; there is often nothing like snobbery involved). Even Rawls includes "the social bases of self-respect" among his list of primary goods. Yet this basis is determined primarily by what Veblen referred to as "the pecuniary standard of decency," which is an intrinsically comparative measure. Earning an income in the bottom quintile deprives individuals of one of the most important bases of self-respect, yet 20 per cent of the population must necessarily fall into this category. Or to pick an example from Hirsch, being in a position of authority on the job is an important source of self-respect, yet in order for there to be leaders there must always be followers [(1976), p. 41]. Thus "fulfilling" work is intrinsically scarce. As a result, we cannot assume that people will always prefer more of some primary good to less, independent of what others are receiving. In some cases, the nature of the good itself requires that they take an interest in how much others are receiving.[15] Wanting the good *amounts to* taking an interest in the interests of others. Thus it might be tempting to suggest that relative preferences should be regarded as a legitimate concern, when it comes to determining the distribution of such goods. In this way, the emphasis is displaced from the preferences of the individual to the nature of the goods.

The problem with this view is that most goods are complex, and are valued for a variety of different properties, some of which will tend to be positional, but others not. Diamonds, for example, are obviously valued primarily for their exclusivity[16], but there are many people who claim

[15] This problem with Rawls's proposal is shared by Kolm's as well (see note 8 above). Both have the effect of laundering out preferences for positional goods. In Kolm's case, the agent's intrinsic preferences are limited to those that an individual would have under the assumption that everyone else was to receive the identical allocation. This precludes many things, including a concern for distinction. Would I be quite as enthusiastic about the tie that I am wearing, if every man that I knew owned the same one? Would anyone want to buy a diamond, or original artwork, if everyone else owned the same item? These are not decisive considerations – what they show is simply that Kolm's proposal does not merely excise the sort of envy that we find objectionable, but results in a very substantive rearrangement of the sort of desires and values that individuals are entitled to present as legitimate interests.

[16] Pigou noted, many years ago (1903), that is not merely the scarcity of diamonds that determines their value. Even more important is the question of *who* owns them. If the quantity of diamonds were miraculously doubled, but this merely resulted in those who already owned diamonds having

to value them solely for their "brilliance." Expensive automobiles, although valued primarily for the status they confer, tend also to have interesting mechanical properties. Homes in prestigious neighborhoods are often built with better materials. Furthermore, people who consume such goods tend to defend their purchases, when criticized, through reference to these "material" qualities, and there is no way of verifying the true basis of their preference. Thus they could object to any institutional arrangement that treated their consumption of these goods as grounds for legitimating the relative preferences (i.e. envy) of others, on the grounds that the latter was "unilateral envy." "Why should my consumption be an object of public concern," they might argue, "just because *other people* happen to value my possessions for their positional qualities?" In the absence of any clear and decisive way of distinguishing "material" from "positional" goods, it is difficult to see how this objection could be countered.

6. Why is this important?

Faced with these difficulties, there may be a temptation to go back to laundering out (i.e. ignoring) all relative preferences. This is certainly an option. The problem, as I see it, is that the welfare losses associated with competitive consumption, driven by relative preference, have become so massive that they cannot in good conscience be ignored. It is perhaps only a slight exaggeration to say that they stand poised to absorb all future gains in the output of our economy. Thus it might be better to use a "rough-and-ready" means of identifying those positional goods that are associated

---

twice as many, no one's demand curve for diamonds would change substantially. The situation would be quite different if the increase in quantity resulted in diffusion of ownership into the "so-called lower class" [Pigou (1903), p. 60]. For an interesting formal development of this "class diffusion" model, see Swann (2001). Swann's model also allows one to specify quite nicely one class of negative network externalities generated by crowding.

with significant crowding effects, even though it is potentially unfair to some people, than it would be to go on ignoring relative preferences. An enormous quantity of resources is being channeled into forms of consumption – the quest for status, beauty, or style – that have the same basic structure as an arms race. Furthermore, Coasean solutions are unlikely to arise among private parties affected by the relevant consumption externalities, not only because of the scale of the collective action problems, but for a variety of other reasons, including the fact that people are generally loathe to see themselves as purchasing status, or as receiving compensation in return for low status. Yet we refrain from adopting any public measures to prevent this waste, on the grounds that doing so confers legitimacy upon a set of supposedly misanthropic preferences. Efforts to curb competitive consumption are described as a capitulation to "the politics of envy" [Frank (1985), 222-223]. As a result, a wide range of opportunities to improve social welfare wind up being ignored.

Take for example taxation policy. The general problem with taxes is that they distort economic incentives, and thus generate deadweight losses. When the tax revenue is used to provide public goods, these losses may be outweighed, so that the overall outcome is a Pareto-improvement. Nevertheless, this improvement is not as great as it could be, and the size of the deadweight loss places constraints on the range of market failures that the state can reasonably aspire to correct or offset. These deadweight losses are not an inevitable feature of taxation systems, however. In cases where particular goods generate (Pareto-relevant) negative externalities, it is possible to raise revenue by imposing a tax on such goods, while at the same time improving efficiency. The tax changes economic incentives, but does so in the right direction. (By increasing the price of these goods, these taxes bring consumption levels closer to what they would be if all costs were internalized.) Thus the incentive effect of the tax itself generates a *prima facie* efficiency gain. This gain can then be amplified when the tax revenue is used to provide goods that are undersupplied by the market because of positive externalities. The classic example of this would be a gasoline tax that is used to subsidize public transportation.

These sorts of "pigovian" taxes represent a significant opportunity for the state to improve the welfare of its citizens. Whenever there is a significant collective action problem caused by negative externalities, there is the potential for the beneficial imposition of such taxes. Traditionally, however, our state policies have put most emphasis on collective action problems in the domain of health, safety and the environment. Collective action problems that are the result of relative preference structures have largely been ignored. This is not ideologically neutral [Frank (1985), p. 231]. If we begin to keep track of such problems, and to regard them as a serious source of welfare loss, we can see opportunities for new forms of pigovian taxes (or new justifications for old ones, see [Kolm (1971)]). Indeed, luxury consumption has long been identified as a potential source of non-distortionary taxation. As John Stuart Mill observed a long time ago, in his *Principles of Political Economy,* "the consequence of cheapening articles of vanity, is not that less is expended on such things, but that the buyers substitute for the cheapened article some other which is more costly, or a more elaborate quality of the same thing; and as the inferior quality answered the purpose of vanity equally well when it was equally expensive, a tax on the article is really paid by nobody: it is a creation of public revenue by which nobody loses" [Mill (1965), p. 869; see also Mason (1998), pp. 26-27]. Unfortunately, "luxury taxes" are often criticized, and are being phased out in countries such as the United States, on the grounds that they are simply punitive towards those with expensive tastes. Yet if we treat them as pigovian taxes, it should become clear that they can be imposed in ways that generate no harm and do not violate neutrality.

In the same way that an arms control agreement need not undermine the security of any nation (insofar as it leaves relative positions unchanged), a tax on goods that are valued for their exclusivity need not change anyone's access to them (again, insofar as it leaves relative positions unchanged).[17] Real estate taxes in certain markets could be thought of in this way, especially given the competitive premium on desirable locations (buyers are compensated for the tax by lower house prices). Not only can taxes of the general sort reduce the level of wasteful "intermediate" consumption that people undertake in the struggle to obtain positional goods, the revenue raised can also be used to provide public goods, which are less likely to be subject to competitive consumption. Interesting proposals have also been made concerning the possibility of non-distortionary taxes focused on the "brand premium" that certain products enjoy [Bagwell and Berheim (1992)]. Of course, the design and implementation of taxes targeting consumption externalities is difficult, and would generate its own deadweight losses [McAdams (1992), pp. 70-76]. But if Hirsch's Law is correct, then the size of these losses can be expected to diminish with time (while the case for such taxes will become more urgent).

Such taxes would be subject to objections from those who claim that their consumption of these goods is non-competitive, although the credibility of such claims would have to be taken with a grain of salt. One way of avoiding the problem would be through a general "progressive consumption tax," along the lines suggested by Frank [(1996), p. 211]. The argument for such a tax stems from the intuition that, as individuals become richer, their marginal propensity to consume positional goods (or goods that are valued for their distinction) increases. If we stop laundering out the welfare consequences of this increasing shift toward consumption of positional goods, then the traditional efficiency arguments against progressivity in the tax system begin to look more dubious.

Of course, the details and the analysis of these tax proposals are enormously complicated. The philosophical point, by contrast, is quite simple. To date there has been a strong taboo against treating envy as a source of externalities in the design of "second-best" taxation schemes, or even when arguing for progressivity in existing schemes. Unfortunately, this taboo has led us to overlook

[17] There is a nice example of this involving sports cars in Frank [(1999), p. 221].

vast regions of the economy where there is room for significant welfare gains. It must therefore be rethought.

7. Conclusion

The desire to launder out envy from our social welfare judgments is based upon a plausible intuition, since it is important that we be able to identify win-win transformations in social outcomes, without being held back by people who get upset at the mere fact that somebody else is winning. Misanthropic preferences, such as those revealed by Zizzo and Oswald (2001), should not be allowed to guide social policy. The problem is that what counts as an individual's "interests" cannot be separated cleanly from the interests of others, simply because our judgments – the very concepts that we use to articulate our needs and desires – have a deeply relative character. Human beings have a natural bias that leads us to notice *changes* rather than *states*, and to judge in *relative* rather than *absolute* terms. Thus we have a persistent tendency to value goods, not for their intrinsic properties, but for those qualities that make them *better than* (or even just *distinct from*) what other people have.

Of course, if all of our preferences were of this type, and every judgment was purely relative, then there would be very little to debate. The human race would have been locked into a state of hedonic homeostasis since its inception. The problem is that the relativity of our desires admits of degrees. As a result, it is possible to achieve Pareto-improvements by shifting resources out of areas that have the *de facto* structure of a zero-sum game, and into areas where improvements in absolute welfare level are still possible. One need only glance at the area of environmental policy to see that we have not exhausted these gains. So ironically, our desire to have a free hand to effect win-win transformations is leading us to ignore a vast set of potential Pareto-improvements. The way to correct this is to begin to think of various forms of consumption on the analogy of an arms race, and to begin to take seriously the way in which the satisfaction of one person's preferences can impede the satisfaction of someone else's. Insofar as we are able to identifies categories of consumption where social scarcity and crowding effects determine the predominant qualities of the goods that are exchanged, there may be a legitimate grounds for collective action aimed at shifting resources out of these sectors, and into more productive employments.

Bibliography

Bagwell, L. S. and B. D. Berheim., 1992, "Conspicuous Consumption, Pure Profits and the Luxury Tax," NBER Working Paper #4163,

Bourdieu, P., 1979, *La distinction*, Minuit.

Dworkin, R., 2000, *Sovereign Virtue*, Harvard University Press.

Easterlin, R., 1974, "Does economic growth improve the human lot? Some empirical evidence," in P. David and M. Reder, eds. *Nations and Households in Economic Growth,* Academic Press.

Easterlin, R., 1995, "Will Raising the Incomes of All Increase the Happiness of All?" *Journal of Economic Behavior and Organization*, vol 27.

Feinberg, J., 1987, *Harm to Others,* Oxford University Press.

Frank, R. H., 1985, *Choosing the Right Pond,* Oxford University Press.

Frank, R. H., 1999, *Luxury Fever,* Free Press.

Frank, R. H., 2005, "Positional Externalities Cause Large and Preventable Welfare Losses," *American Economic Review,* vol 95.

Frank, R. H., and P. Cook, 1995, *The Winner-Take All Society,* Penguin.

Frey, B. S., and Stutzer, A., 2001, *Happiness and Economics,* Princeton University Press.

Gauthier, D., 1986, *Morals by Agreement,* Oxford University Press.

Goodin, R., 1986, "Laundering Preferences" in J Elster and A Hylland, ed. *Foundations of Social Choice Theory,* Cambridge University Press.

Heath, J., 2001, "Problems in the Theory of Ideology," in J. Bohman and W. Rehg, eds. *Pragmatism and Critical Theory,* MIT Press.

Heath, J. and A. Potter., 2005, *Révolte consommée,* Éditions Naïve.

Hirsch, F., 1976, *The Social Limits to Growth,* Harvard University Press.

Holcombe, R. G. and R. S. Sobel., 2000, "Consumption Externalities and Economic Welfare," *Eastern Economic Journal*, vol 26.

Jeffrey, R., 1983, *The Logic of Decision* , 2nd edn. University of Chicago Press.

Kaplow, L. and Shavell, S., 2002, *Fairness versus Welfare,* Harvard University Press.

Kolm, S.-C., 1971, "La taxation de la consommation ostentatoire," *Revue d'economie politique*, vol 1.

Kolm, S.-C., 1995, "The Economics of Social Sentiments: The Case of Envy," *Japanese Economic Review* vol 46.

Leibenstein, H., 1950, "Bandwagon, Snob and Veblen Effects in the Theory of Consumers' Demand," *Quarterly Journal of Economics* vol 64.

Mason, R., 1998, *The Economics of Competitive Consumption,* Edward Elgar.

McAdams, R. H., 1992, "Relative Preference," *Yale Law Journal*, vol 102.

Mill, J. S., 1965 [1848], *Principles of Political Economy,* J. M Robson (ed.) University of Toronto Press.

Nozick, R., 1974, *Anarchy, State and Utopia,* Basic Books.

Nussbaum, M., 1992, "Human functioning and social justice: In defense of Aristotelian essentialism," *Political Theory*, vol 20.

Pigou, A. C., 1903, "Some Considerations on Utility," *Economic Journal*, vol 13.

Pollak, R., 1976, "Interdependent Preferences," *American Economic Review*, vol 66.

Rawls, J., 1971, *A Theory of Justice,* Belknap.

Rawls, J., 1991, *Political Liberalism,* Columbia University Press.

Scitovsky, T., 1954, "Two concepts of external economies," *Journal of Political Economy* vol 62.

Scitovsky, T., 1992, *The Joyless Economy,* 2nd edn. New York: Oxford University Press.

Sen, Amartya. 1972. "The Impossibility of a Paretian Liberal," *Journal of Political Economy* 78: 152-157.

Sussangkarn, Chal and Steven M. Goldman. 1983. "Dealing with Envy," *Journal of Public Economics*, vol 22.

Swann, G. M. P., 2001, "The Demand for Distinction and the Evolution of the Prestige Car," in U. Witt (ed.) *Escaping Satiation,* Springer.

Varian, H., 1974, "Equity, Envy and Efficiency," *Journal of Economic Theory*, vol 9.

Veblen, T., 1967 [1899], *The Theory of the Leisure Class,* Penguin.

Warren, E. and Warren Tyagi, A., 2003, *The Two Income-Trap,* Perseus.

Zizzo, D. J. and Oswald, A., 2001, "Are People Willing to Pay to Reduce Others' Incomes?"*Annales d'économie et de statistique*, vol 63-64.

Zizzo, D. J., 2005, "Interdependent Preferences," SSRN: http://ssrn.com/abstract=823685