

Envy and efficiency

Professor Joseph Heath
Department of Philosophy
University of Toronto

The Pareto principle states that if a proposed change in the condition of society makes at least one person happier, and does not make anyone else unhappy, 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 the happiness of others itself becomes a source of unhappiness. As a result, envy has the potential to block any and all Pareto improvements. Making one person better off will automatically make someone else worse off, so there will be no point talking about efficiency gains.

Some people take this to be sufficient reason for dismissing 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 they should carry no weight, from the normative point of view. The most memorable formulation of this idea was provided by Robert Goodin, with his image of "laundering preferences." The claim, roughly, is that a theory of justice will take preferences as "input," then generate a social welfare function (or a normative ranking of possible states), as "output." Such a theory will contain both an input and an output filter. The former will filter out illegitimate (e.g. "sadistic or meddlesome") preferences. The latter will filter out recommendations that are trumped by other normative considerations (e.g. respect for individual rights.)

In principle, theorists who find envy objectionable might choose to filter it out either at the front or the back end. 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 at the “back end.” 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. Thus John Rawls assumes that rational agents behind the veil of ignorance “do not take an interest in one another’s interests.” Similarly, David Gauthier excludes any “tuistic” preferences from consideration in his bargaining theory.

I would like to register some concerns about this filtering strategy. First, I believe that envy has been defined far too broadly. In many cases, the way that envy is defined 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. My second claim is that even if we adopt a more restrictive definition of envy, there are still 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 consequence of rendering the structure of the collective action problem invisible, from the normative point of view.

More generally, the 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. This treats them as

essentially an aberration (like cruelty), which can safely be ignored. 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. (After all, Rawls' formula is just a nice way of stating the injunction that agents should, by and large, mind their own business.) 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. Thus the interest of one often cannot be cleanly detached from that of others. 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. Rather than excluding envy from our judgments of social welfare, we should think more seriously about how to make efficiency judgments among agents who are susceptible to serious envy, or at least strongly relative preferences.

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, are subject to intractable controversy.¹ Thus it is necessary to assign at least some 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

¹ Insert some ref. here.

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. The problem is that the standard decision-theoretic definition of preferences contains no substantive restrictions on the content of these preferences. Thus treating the agent’s utility function as a representation of his interests leads to agents having “interests” at stake in all sorts of interactions that, from an intuitive perspective, we would regard as 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 will 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. As Amartya Sen showed long ago, any form of welfarism this broad essentially makes it impossible for individuals to enjoy any of the classical negative liberties.²

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). If such preferences were given weight when it comes to the design of institutional constraints, then the resulting constraints 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

² Sen, Amartya., “The Impossibility of a Paretian Liberal,” *Journal of Political Economy* 78 (1972): 152-157.

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 defence 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.³

However, if there is one point on which there is extremely widespread agreement, it is that the dissatisfaction generated by feelings of envy should not count as a harm (or, conversely, that the fact that the outcome of someone’s actions may inspire envy in me should not be taken to give me an legitimate interest the outcome of that person’s choice).⁴ To see how this exclusion works, consider the standard argument for the market. 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.

Imagine in a situation in which my neighbour 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 neighbour and his family. Does this undermine the win-win character of the transaction between the neighbour 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

³ See Joel Feinberg.

⁴ For especially clear formulation, see Robert Nozick, ASU, p. 274.

neighbour's new acquisition should not count as a consideration that speaks against the transaction.

On the other hand, the air-conditioner might also make a lot of noise, which keeps me awake at night. Then we might not want to regard the purchase as purely a private matter between the neighbour and the merchant. I become an unwilling participant, and my loss of welfare, it seems, should count for something. Economists would say that in this case the transaction creates a "negative externality." Thus when we talk about markets, the "laundering" of preferences normally occurs in the decision that we make about which external effects of a transaction to treat as "externalities" – and thus as part of the "social cost" of an action.

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 neighbour's new-found comfort. But if I wasn't sensitive to noise, 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 welfare (and thus my responsibility to deal with it). There must be some deeper rationale. Furthermore, we would be ill-advised to simply 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 trying to dictate to my neighbour what volume he should keep his television at within the privacy of his own home. In this case, we would be inclined to handle the "intrusiveness" of my preferences by balancing my interests against his liberty interests, and would probably conclude that his interests in this matter outweighed mine. 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 so illegitimate?

2. Defining envy

Let us take a very concrete case. In a study conducted by Daniel Zizzo and Andrew Oswald, groups of four experimental subjects were given sums of money which they then wagered on a series of gambles.⁵ 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. Thus a majority clearly derived strong satisfaction from seeing the holdings of others reduced – enough that they were willing to pay for the privilege.

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 preference that this experiment brings to the fore are precisely the ones that are being targeted when we try to filter out envy from a social welfare function. 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 that we are concerned with, 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 normally regarded as anti-social, and so

⁵ “Are People Willing to Pay to Reduce Others’ Incomes,” Daniel John Zizzo and Andrew Oswald (forthcoming).

would fall into the same category as cruelty. We might regard such states as being intrinsically wrong, and thus 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 in the definition of interests was intended to avoid. Who are we to say what emotions other people should and should not be experiencing? 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.”

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. In the case of the air conditioner, it is the fact that my neighbour is able to satisfy one of his desires (i.e. derives satisfaction from it) 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. A somewhat broader version excludes envy on the grounds that it is a *relative preference*. Rawls’s formulation 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 possible worlds. 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. (This is, in fact, the standard practice in decision theory.)

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 ranking of possible world moves down in one agent's ranking in response to its moving up in some other's. If the state of affairs came to pass, 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). *Schadenfreude* (or simply malevolence) can be defined as the opposite of envy – when the ranking of 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 problem with this definition is that it is extremely “welfarist.” It looks only at agent's utility functions. Thus envy is said to arise only when the *happiness* of others is a source of unhappiness for others. However, in many concrete cases, what provokes envy is not so much the happiness of others as it is their possessions. This is what we may refer to, somewhat more precisely, as *covetousness*. I may envy my neighbour'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. 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 much too narrow.

2. *Relative preferences.* In order to expand the definition of envy to preclude covetousness, it is necessary to consider not only the structure of agents' utility functions, but also the economic 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 stipulates that, under each category of primary good, agents prefer an allocation that gives them more to one that gives them less, independent of what others receive in the distribution scheme. 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," because their preferences over any distribution are determined entirely by the particular allocation that they receive, and not by any concern about what others receive.

It should be noted that Rawls' definition, unlike Gauthier's, has the effect of eliminating all distributional preferences. This is in accordance with our intuition that my desire to see Tom receive more than Bill cannot possibly count as an argument in favour of Tom receiving more than Bill. Nevertheless, such a distributional preference may not depend, in any crucial way, upon either Bill or Tom's preferences. Thus it must be excluded by specifying that, where g_x is the quantity of primary goods that individual x receives under a particular allocation scheme $U = \{g_1, g_2, g_3, \dots, g_n\}$, the Rawlsian constraint amounts to saying that, 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, a preference for equality cannot count as an argument for

an egalitarian distribution. Many people, unfortunately, especially economists, make this argument.⁶)

The two strategies for identifying the envy-based preferences differ considerably, yet are 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 salient features of certain social interactions.

3. Problems with interdependence

To see the problem with laundering out interdependent preference, 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 every point on the two axes that intersect s . Speaking loosely, these can be described as the set of win-win outcomes (speaking more 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.

⁶ See, for example, Louis Kaplow and Steven Shavell, *Fairness versus Welfare* (Cambridge, MA: Harvard University Press, 2002).

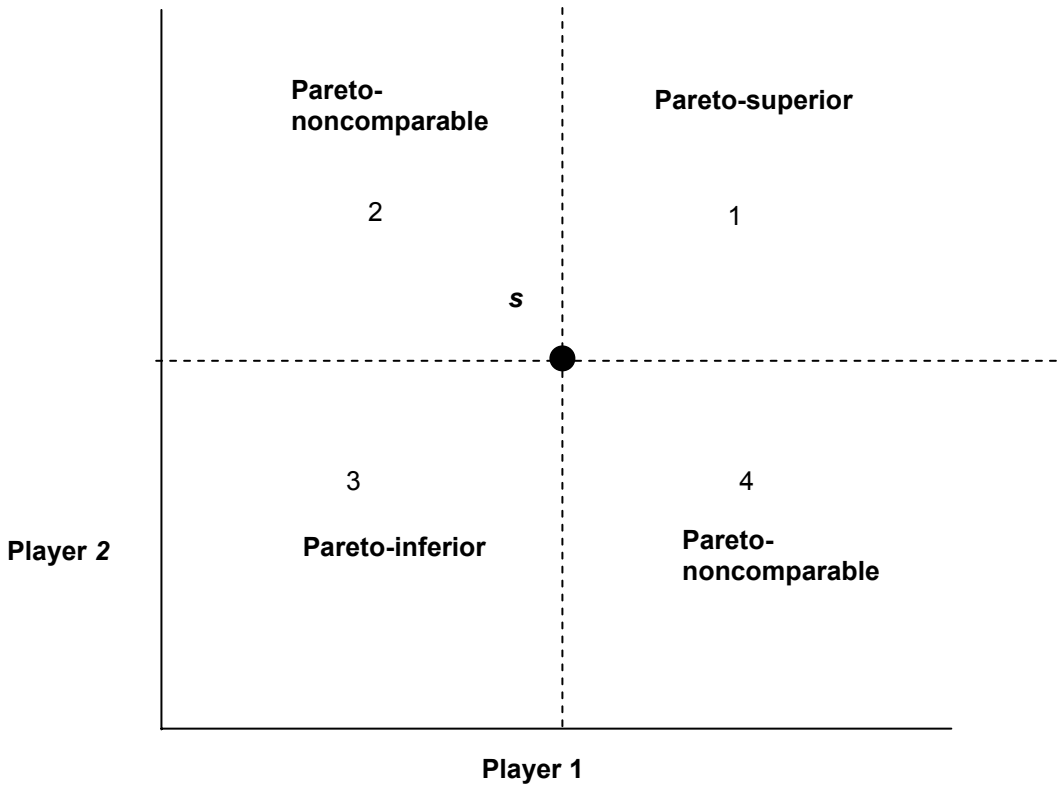


Figure 1. The Pareto principle

The elimination of envy-based preference is designed to correct situations like those show in point a , in Figure 2. Here, when player 2 moves from y_0 to y_1 , it provokes envy in player 1, whose utility is 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 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.

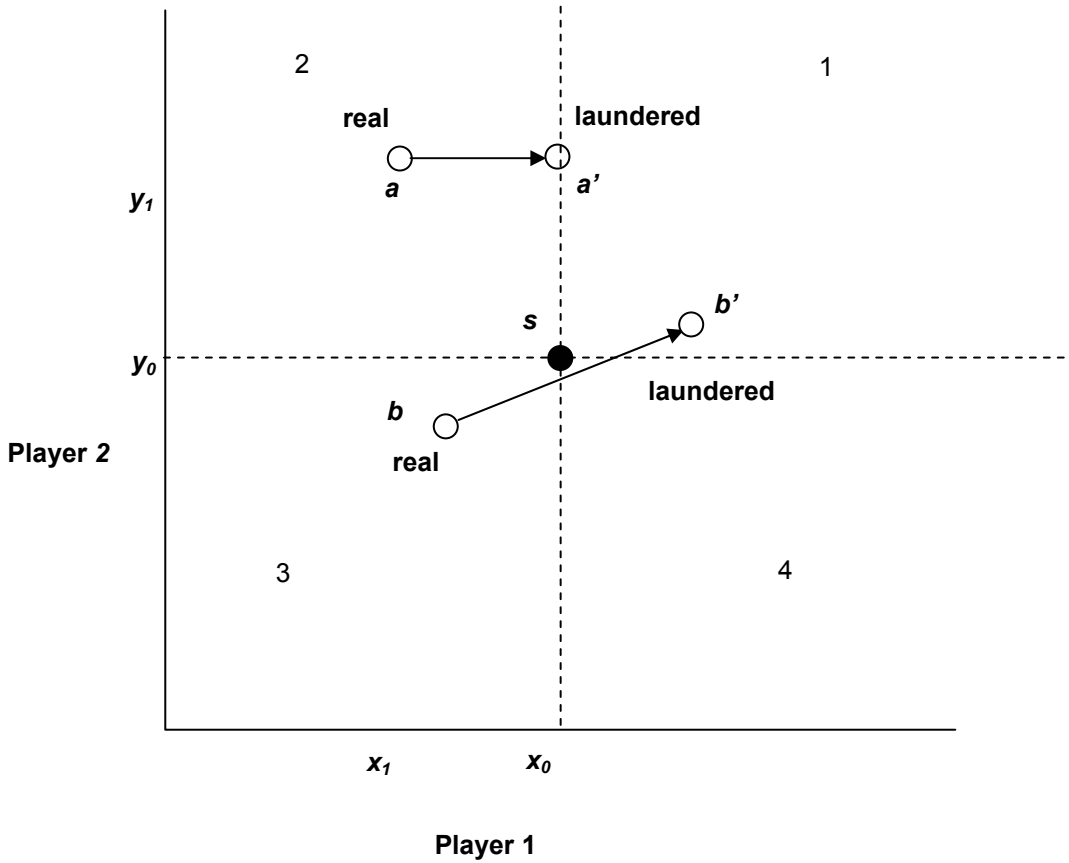


Figure 2. Laundering preferences

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. Thus the laundering procedure prevents one person from unilaterally setting the terms of interaction.

Our intuitions become much murkier, however, when the issue is not a Pareto-noncomparable outcome, but rather a Pareto-inferior one. Imagine, for example, two lovers who, when they quarrel, like to drive the other crazy by kissing someone that each knows the other hates. They both do it in order to provoke the other's jealousy, but as a result, wind up simply exacerbating their quarrel. We might easily imagine this couple, in a cooler moment, swearing that they each each desist from such behaviour 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 2 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.

The goal of laundering out envy was to prevent individuals from blocking the Pareto-improvements of others. But 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 to power of the state to enforce such agreements.

4. Problems with relative preference

The example of the quarrelling lovers is somewhat unusual, in that the individuals actually know what the other is feeling, and act with the specific intention of provoking jealousy or 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 neighbours'. 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, it makes everyone else's slower by comparison. Thus each gain in welfare for one is accompanied by a loss of welfare for others.

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 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 neighbours' 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 would be much harder to do so in the latter case, simply because the relative size of my salary has important consequences in terms of the level of consumption that I may aspire to.

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 \$85,000. In short, your absolute income would be a little lower in World B, but your relative income would be much higher. Prices are the same in the two worlds, so the income difference represents real differences in living standards.”⁷

Is it really so unreasonable to select B? More pointedly, is this preference so unreasonable that we should want to completely eliminate it from any judgements 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 – and often of goods that are qualitatively superior – because one is able to win a greater number of competitions.

Thus the person who chooses 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 average wealth and average happiness in wealthy industrialized nations.⁸ People living

⁷ Luxury Fever.

⁸ See Bruno Frey and Alois Stutzer, *Happiness and Economics* (Princeton: Princeton University Press, 2002).

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 very important.

In other words, while money can't buy happiness, having more money than others does increase your chances. What is the most plausible explanation for this phenomenon? There are no doubt a number of different factors at work. One hypothesis that seems especially plausible 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, bring about an absolute increase in well-being. However, as people become richer, they tend to spend a greater proportion of their income on goods that are valued more for their relative qualities. As a result, the positive-sum character of the economy diminishes. This would explain the fact that, as society becomes richer, relative income remains important while absolute income declines in significance.

It is not difficult to imagine why this might be so. If one looks at the major categories of household expenditure, one can see that almost all members of our society 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, a car, a telephone, 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. Elites in our society simply 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 finely tailored. But of course the threshold of expenditure continues to recede as more wealth becomes available. When machine-made clothing becomes available to the general public, hand-stitching becomes the mark of distinction. Since Thorstein Veblen’s seminal work, this process has been a well-recognized feature of our consumption habits.

But if the relative structure of our preferences is indeed what is responsible for 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 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 is not leading to any increases in welfare, then 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 (and no obvious relative dimension). However, it is impossible to make the argument for air quality vs. consumer goods if we launder out the preferences that give most markets for consumer goods a zero-sum structure at the level of individual welfare.

5. Legitimizing envy

These observations create something of a dilemma. There are clearly cases in which we do not want to grant any authority to preferences based on envy. And yet we cannot ignore the fact that people on many occasions have very legitimate concerns about their relative position. How do we decide when to draw the line?

Two possibilities present themselves. The first would be strictly formal. We could simply 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 were to prevent certain forms of conspicuous consumption, which have the structure of a race to the bottom. The fact that everyone benefits from the countermeasures would outweigh the somewhat dubious source of these benefits.

This proposal 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 neighbour’s new air conditioner is being disregarded, then I might choose to start making lots of noise at night, in order to keep him awake. I could then offer to stop if he agrees to stop buying things that I want. Of course, we would not want individuals to have the power to “self-launder” their own envy-based preferences in this way. Even if we specify that the suboptimality must stem from envy-based preferences on both sides, such an arrangement gives me an incentive to do whatever I can to provoke my neighbour’s

jealousy, in order to legitimate my own preferences. This seems wrong as well. If envy is to count as a legitimate type of harm in a suboptimal interaction, then it should also count as legitimate in a unilateral interaction. Otherwise we give individuals the ability to legitimate harms that they suffer by imposing losses upon others.

The more promising option is to assume an economic environment, 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 is the distinction, introduced by Fred Hirsch, between “positional” and “material” goods. Material goods are those which, through increased production, we are able to create more of. Positional goods are, for either intrinsic or social reasons, necessarily in fixed supply. Thus in a typical municipal region, it is possible to build more houses, or to increase the density of dwellings, but it is impossible to create more land. We can create more beautiful buildings, but we cannot create more beachfront real estate. As a result, demand for positional goods necessarily has a competitive structure, since access to them will be determined by relative ability to pay. Thus particular markets necessarily have a zero-sum structure, even if the economy as a whole is positive sum.

With real estate, for instance, the realtor’s saying “location, location, location,” reflects the fact that it is the positional character of these goods that makes up the majority of their value. Sometimes this is due to intrinsic physical scarcity, sometimes it is due to social scarcity. 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 to want to live downtown necessarily have interdependent preferences, because the object that they desire is constituted by the preferences 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.

The most obvious example of a system that creates "social scarcity" is status or prestige. Status is necessarily hierarchical – in order for someone to be one top, someone else must be on the bottom. Since position is relative, when one person moves up it necessarily pushes someone else down. Thus goods that are consumed for the status that they confer are necessarily positional. Their value is based on their exclusivity. As a result, if too many of them are produced ("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. It is worth nothing as well that it is not only luxury goods that 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.

Social scarcity is a lot more common than we imagine, simply we often don't realize that many of the aesthetic standards we use to determine the value of goods have an intrinsically social referent. 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. Pierre Bourdieu has argued, more generally, that all judgements of taste 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. As a result, we all have what amount to relative preferences, even though we 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. 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 they 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.

Thus it seems to be the case that certain types of goods – what Hirsch calls positional goods – have an intrinsically zero-sum structure. People who want these goods automatically commit themselves to having relative preferences. Yet these goods are also ubiquitous. Even Rawls includes “the social bases of self-respect” among his list of primary goods. Yet this basis is determined primarily by what Veblen refers to as “the pecuniary standard of decency,” which is an intrinsically comparative measure. Earning an income in the bottom decile deprives individuals of one of the most important bases of self-respect, yet 10 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. Thus “fulfilling” work is intrinsically scarce. As a result, we cannot assume that people will always prefer more of some primary good to less. In some cases,

the nature of the good itself requires that they take an interest in how much others are receiving. Wanting the good *amounts to* taking an interest in the interests of others.

The problem with this view is that there are very few goods that are purely positional – and thus in absolutely fixed supply. It is true that in order for one person to live in a penthouse apartment, many others must be willing to live in the floors below (and thus penthouses are socially scarce). Yet it is always possible to build new buildings, and so expand the supply of penthouses. Similarly, the preponderance of the value of diamonds is due to their scarcity, but we are constantly uncovering more of them. Even in the case of social status, despite its hierarchical structure, we could imagine more of it being produced (e.g. if a broader segment of the population was admitted into a higher social class). Most of what get called positional goods are actually just goods that are subject, for one reason or another, to extreme supply inelasticity. As a result, there is no clean and simple way to displace the emphasis from the preferences of the individual to the goods that are desired.

One way of handling to problem would be to conceive of goods with a strong positional dimension as the source of “consumption externalities.” With pure positional goods, we could then take the magnitude of the externality to be equal to the total cost of the good. The externality could of course be distributed out across many consumers, but the total would still be zero-sum. With goods that are simply inelastic, the size of the externality could then be treated as a fraction of the value of the good, reflecting the fact that there is a slight net gain achieved through the consumption.

6. Why is this important?

It should be clear, however, that there is no clean and simple way of determining which goods generate relative preferences. These difficulties might lead one to question the value of the entire enterprise. Rather than drawing such complex distinctions between categories of goods, it may be preferable simply to go back to ignoring envy completely.

This is 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 be in good conscience 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. An enormous quantity of resources are being channeled into form of consumption – the quest for status, beauty, or style – that have an essentially zero-sum structure. Furthermore, we refrain from taking any measures to prevent this waste, on the grounds that doing so confers legitimacy upon a set of misanthropic preferences. Efforts to curb competitive consumption are described as a capitulation to “the politics of envy.” And opportunities to invest in areas where we can still achieve across-the-board welfare gains are passed up.

Take for example taxation policy. The general problem with taxes is that they distort economic incentives, and thus generate “deadweight losses” (i.e. Pareto-suboptimalities). 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 losses 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 negative externalities, it is possible to raise revenue by imposing a tax on this good. These so-

called “pigovian” taxes distort economic incentives, but do so in the right direction. By increasing the price of the good that generates the externalities, 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. This discourages congestion both by reducing automobile use and by encouraging alternate forms of transportation. This benefits everyone, including those who pay the tax.

Thus pigovian taxes represent a significant opportunity for the state to improve the welfare of its citizens. Whenever there is a major negative externality, there is the potential for the beneficial imposition of such taxes. Unfortunately, our policies have put most emphasis on environmental and health externalities. Externalities produced because of relative preferences structures (i.e. envy) have largely been ignored. However if we begin to keep track of the “consumption externalities,” and regard them as a serious source of welfare loss, we begin to see opportunities for new forms of pigovian taxes (and new justifications for old ones). For example, “luxury taxes” are often criticized, and are increasingly being phased out, on the grounds that they are simply punitive towards those with expensive tastes. Yet if we treat these as pigovian taxes, it should become clear that they harm no one. Real estate taxes could also be thought of in this way, especially given the competitive premium on desirable locations.

Of course, the details of the design and implementation of such taxes 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” or

optimal taxation schemes. This taboo is unjustifiable, and leads us to overlook vast regions of the economy where there is room for significant welfare gains.

7. Conclusion

The reasons for wanting to launder out envy from our social welfare judgements are for the most part sound. It is very 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. The problem is that our preferences cannot be separated cleanly from one another, 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. The goals and projects that we develop are deeply informed by what we see around us, what we see the typical person achieving. As a result, the satisfaction of these desires necessarily generates externalities for others.

If all of our desires were of this type, and everything were relative, then there would be no problem. 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 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 introduce a category of “consumption externalities,” and begin to take seriously

the way in which the satisfaction of one person's preferences can impede the satisfaction of some other's.