


CHICAGO

JOHN M. OLIN LAW & ECONOMICS WORKING PAPER NO. 342
(2D SERIES)



What Does Happiness Research Tell Us about Taxation?

David A. Weisbach

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

June 2007

This paper can be downloaded without charge at:
The Chicago Working Paper Series Index: <http://www.law.uchicago.edu/Lawecon/index.html>
and at the Social Science Research Network Electronic Paper Collection:
http://ssrn.com/abstract_id=995319

What Does Happiness Research Tell Us About Taxation?

David A. Weisbach

The University of Chicago Law School

June 2007

Abstract

This paper analyzes the consequences of the findings from research into self-reported well being or happiness for taxation. It primarily considers two findings: that happiness depends on status as well as income, and that individuals may adapt to disability, exhibiting relatively small losses in happiness from disabilities. In each case, it examines how adding these concerns to standard tax models changes the results and then compares the empirical findings of the happiness literature to see whether they provide the type of data needed to parameterize the models. In both cases, the theoretical models ask for different types of data than the happiness studies emphasize. The paper also looks at Robert Frank's arguments for a progressive consumption tax based on the findings of the happiness research. It finds that these claims are not supported by the current findings.

Preliminary Draft. Please do not cite without permission.
Send comments to: d-weisbach@uchicago.edu

What Does Happiness Research Tell Us About Taxation

*David A. Weisbach**

June 2007

Happiness research has the *potential* to change our views about taxation, possibly significantly. Standard tax models assume a very simple utility function which is uniform across all individuals, increases in consumption, and decreases in work effort. Individuals vary only in their ability to earn income. Happiness research show that reported levels of happiness (which I will take here to mean utility), is a complex function of many different variables. Utilities may be interdependent because of status concerns. Individuals may adapt to income levels or to disabilities. Very low work effort such as unemployment may reduce happiness, even holding consumption constant. A more complex view of utility functions will lead to different tax conclusions. This has long been noticed, and a number of tax models incorporate some of these ideas, particularly those related to status.

The question this paper will address is whether the findings of happiness research to date, taking them as valid, tell us what we need to know to be confident making tax policy recommendations. The answer, I think, is not yet. It is clear that incorporating ideas about happiness into standard tax models changes the results. The empirical research on happiness, however, does not yet tell us in sufficient detail how to model utility to get concrete tax results. For example, the most common method of incorporating the findings of happiness research into tax models is to include status competition. There are a variety of ways of modeling status, however, each with different implications. Within each modeling strategy, varying the parameters dramatically changes the results. Happiness research is not yet at the point that it can determine which modeling strategies are best and how to set the parameters.

*Walter J. Blum Professor, The University of Chicago Law School. I thank Sam Bagenstos, Omri Ben Shahr, Louis Kaplow, Brian Leiter, Adam Samaha, Mike Stein, Cass Sunstein, Adrian Vermuele, participants at the University of Virginia Law and Economics workshop, and participants Tel Aviv University Law School law and economics workshop for comments, and Anne King for excellent research assistance.

Two caveats are in order. First, I am taking happiness research as valid and as measuring utility. There are significant issues in both regards, issues that have been widely discussed in the literature. It is not my comparative advantage to evaluate these issues. If the research does not validly measure utility, its implications for taxation will be more limited. To explore the question of its implications for taxation, I will assume that the research is at least potentially relevant.

Second, I am asking whether there is empirical support for including the findings of the happiness research in tax calculations. Most tax theory, however, is based on limited empirical support and simplifying assumptions are often made to make the problem mathematically tractable rather than because they are supported by the data. It is no more valid to assume a standard utility function (i.e., only consumption and labor effort effect utility) than it is to assume a more complex utility function that reflects the tentative findings of happiness research. That is, asking for empirical support for including happiness information but not for what are currently more standard models unfairly stacks the deck. Indeed, our priors on the utility function should probably resemble those suggested by the happiness research rather than the usual simple functions.

A standard response is to use Occam's Razor to argue for a higher standard of proof on models that add complexity to the utility function, but this is not my reason for subjecting the happiness research to scrutiny. Instead, my goal here is to be constructive. By showing how the happiness research and the theoretical tax policy literature do not yet line up, my hope is that future research might fill the gap. Models should reflect known facts, and empirical research, to the extent the goal is to have policy implications, should ask questions that are relevant to normative theories. I do not mean to excuse unsupported modeling using standard utility functions – instead, my goal is examine what we would need to know for the happiness research to have tax policy relevance.

Part I of this paper will provide some very brief background on tax theory, background which is necessary for considering how happiness research might change the standard learning. I will then consider two aspects of the tax structure that might be informed by happiness research. Part II will consider how status concerns might affect marginal tax rates on labor income. There are a number of papers modeling tax rates given status concerns. Part II will try to show what drives these models and what information is needed to specify the utility function. I will then briefly discuss what the literature on happiness tells us about these aspects of utility functions. Part II will also examine claims made by Robert Frank about how status concerns affect the optimal tax structure, claims which to some extent are distinct from those related to optimal tax rates.

Part III will consider the extent to which the literature on adaptation to disability affects how the tax/transfer system should take disability into account. A simple claim might be that the literature showing adaptation to disability means that the tax/transfer system need not do anything in particular regarding disability, at least beyond the initial period in which adaptation takes place. If the utility of a disabled person and a non-disabled person are roughly the same (after some period of adaptation), there is no need to give any special dispensations to the disabled through taxes or transfers. Policymaking regarding disability, however, needs to know much more about the effect of disability on utility than we have learned so far from the literature on adaptation.

Part IV concludes by offering a grab bag of other possible implications of happiness research for tax policy. The list is merely suggestive and I do not examine any of the possibilities in any detail.

I. A Very Brief Primer on Tax Policy¹

The now standard approach to taxation (including negative taxation or transfers), due to Mirrlees (1971), is to view taxation as a problem of information. Individuals are assumed to vary by their ability to earn income, which is assumed to be private information. The government wants to redistribute from those with high ability to those with low ability either because the social welfare function is concave or because marginal utility is declining in income. With complete information, the government could assign higher unavoidable taxes to those with higher ability and, thereby redistribute from those with high ability to those with low ability. The Second Welfare Theorem would apply, and we could achieve a first best result.

If ability is private information, however, the government cannot simply assign high taxes to high ability individuals because it cannot tell who has high ability. Instead, it must rely on observable signals of ability. In standard models, the only signal of ability is labor income, and, therefore, the government has no choice but to use this for the tax base. The central problem is that if marginal rates on income are set too high, individuals will work less. That is, high ability individuals can masquerade as low ability individuals, defeating redistributive goals.

The optimization problem is to maximize a function of utilities subject to a resource constraint (total consumption has to equal total income), and an incentive constraint (to prevent masquerading). The incentive constraint is that an individual of a given ability cannot prefer to earn the income of someone with lower ability by working less. Solving this maximization is mathematically complex and the first order conditions, in their general form, are hard to interpret. It is common to use numerical simulations to get a sense of the resulting tax schedule.

¹For extensive reviews of the literature, see Kaplow (forthcoming), Stiglitz (1987), Tuomala (1990).

A simplifying assumption often used to understand the first order conditions is that utility is quasi-linear in consumption, so that utility can be expressed as $u = c + v(l)$ where c is consumption and l is labor effort. The major effect of this assumption is to remove income effects. Using this assumption, Diamond (1998) expressed the first order conditions for the marginal rate at a given ability level n as

$$\frac{t'_n}{(1 - t'_n)} = ABC$$

where A is a measure of labor supply elasticity at that ability level, B is a measure of the social cost of taking a dollar away from everyone above that ability level, and C is a measure of relative the size of the populations at that ability level and above it. Saez (2001), following Stiglitz (1987), suggested the following interpretation. Suppose that tax rates are set optimally and we consider a small increase of the tax rates at some income level, such as between \$50,000 and \$50,001. Individuals at that level will face a higher marginal tax rate, which will distort their work effect, as measured by A , a measure of labor supply elasticity. Individuals above that level will not face a change in marginal rates, so their work effort remains the same (there is no income effect). They have less income, however because of the higher rate they pay at the \$50,000 level, and B measures the social cost of lowering their income. C gives the ratio of the size of the two effects.

The first order conditions become considerably more difficult to interpret in general cases because labor supply elasticity can vary with ability and because of income effects. The quasi-linear case, however, is restrictive because it eliminates declining marginal utility of consumption (utility is linear in consumption), which is an important motivation for redistribution.

The focus of the discussion below is on how happiness-informed utility functions change the conclusions about taxation. It is worth, therefore, making several notes about the utility functions and assumptions

used in the standard analysis. First, some unemployment may be optimal in this model. Utility is assumed to be a nondecreasing function of consumption and leisure. Nothing else matters. High rates on low income individuals may be desirable because the cost of their lost work effort is low and the high rates are inframarginal for everyone else. These high rates may cause them to have zero labor effort, but as long as there are sufficient transfers, they are better off not working.

Second, utilities are independent of one another, which rules out status concerns as well as altruism. In particular, individuals care only about their own consumption and leisure. Third, individuals are assumed to be identical to one another other than with respect to innate ability to produce income (which is assumed to vary by a known distribution). This means that there is no heterogeneity other than with respect to income.

One consequence of assuming no heterogeneity is that the government should only tax labor income (or equivalently, consumption). Individuals vary only in this dimension and, therefore, taxing other attributes can only serve as a bad proxy for directly taxing labor income. Thus, for example, taxing capital income as a proxy for taxing individuals with high ability (as is commonly suggested by those supporting a conventional income tax) does not make sense absent heterogeneity. Replacing such a tax with a direct tax on labor income leaves the distributive effects and effects on labor the same while reducing the distortions in savings patterns.² Introducing heterogeneity significantly complicates the picture because it may be desirable to tax items other than labor income. Models with more than one dimension of difference among individuals, however, are in their infancy and, because of mathematical difficulties, it is not clear when general propositions will be available.

²An exception would be if taxing capital income can reduce the distortions caused by labor income taxation because capital income is a complement to leisure. This possibility is generally regarded as unlikely and similar considerations may lead to a subsidy rather than a tax on capital income.

Heterogeneity will be relevant below in the discussion of adaptation to disability.

The assumption that the government cannot observe anything other than income means that the government will not, in the standard analysis, attempt to tax ability directly. As discussed below, most models of status and taxation continue this tradition by assuming that the government cannot directly tax status seeking activities and instead, must adjust the income tax as a proxy. Frank (1985, 2000) makes a slightly modified assumption, which is that particularly expensive items, such as watches or large houses, indicate status consumption and that the tax system can be designed to weigh more heavily on purchases of these items.

II. Taxation and Status

The observation that status matters to individuals goes back to at least *The Theory of Moral Sentiments*. Smith (1759). Status has been the subject of numerous of studies since then, with such figures as Mill, Pigou, Friedman and Savage, and Becker contributing to the literature. Proposals for taxing status consumption have been traced back to John Rae (1834). The intuition is that improving one's status imposes a cost on others, at least if status comes in fixed supply – we cannot all have the highest income, be the best looking, the smartest, or whatever it is that confers status. Status seeking creates negative externalities because gains in one's status reduces someone else's. Taxing status seeking activities, therefore, might be welfare enhancing much like any other Pigouvian tax. Frank (2000; 152-153) compares the problem to an individual standing up in a stadium to see better. Each individual separately gains by standing, but if everyone stands, everyone is worse off – they cannot see any better and they have to stand instead of sit (and in Frank's story, they stand on the seats, straining their heads until someone falls off). Status competition might make everyone work too hard or consume goods they otherwise would prefer not to consume, but at the end of the day, there is no overall gain in status.

This simple intuition does not tell us anything about the likely effects of status on the tax rate schedule. For example, increasing progressivity would move everyone closer together. This might decrease status competition, because the gains from competition are smaller – it would be harder to separate yourself from the group. On the other hand, it might increase status competition. If you are closer to beating someone in a status race, you might try harder.³ Thus, we can imagine status considerations leading to either a more progressive tax system or a less progressive tax system. More careful modeling is necessary. The question for this section is how status has been (or can be) incorporated into the optimal tax models and whether the empirical research on status effects gives us the relevant information.

Although numerous papers discuss taxing status consumption, only a handful embed the analysis in an optimal tax framework. Boskin and Sheshinski (1978), Oswald (1983), Tuomala (1990); Ireland (1998, 2001) and Allgood (2006).⁴ Each of these authors models status differently. To keep the discussion manageable, I will focus on two papers, Oswald (1983) and Ireland (2001).

A. *The Models*

Oswald takes the standard utility function (utility is a function of consumption and leisure) and adds a concern for the consumption of others measured by a function α . Thus, utility of an individual of type n is equal to $u(c, l, \alpha)$ where c is a vector of consumption goods, and l is leisure. The comparison function α is the sum of the consumption of all other individuals in society, weighted by $\omega(n)$. If all individuals are weighted equally, the comparison function measures average

³See, for example, Hopkins and Kornienko (2004), who find this type of result in a game theoretic model of status competition.

⁴Outside of the optimal tax framework, major papers on taxation and status include the many papers by Frank, and Hopkins and Kornienko (2004). Abel (2005) considers the effect of relative consumption concerns on taxation in a growth model.

consumption. Alternatively, the consumption of the rich or the poor could be weighted more heavily. Depending on the sign of α as consumption changes, it can represent either be altruism (utility goes up as others' consumption increases) or envy (utility goes down as others' consumption increases).

A key to this model (and all other models of status) is that individuals compare themselves to the consumption of others, not utility. Moreover, there are no particular status goods, which means that overall income is what confers status. We might call this model of status as the big fish/small pond theory of status. Holding income constant but moving to a place where comparison income is lower improves welfare. Moving, however, is not easy within the model: all individuals subject to taxation are required by the model to have the same comparison group. If the tax is a national tax, the comparison group is necessarily the nation. That is, contrary to Frank (1985), individuals within Oswald's world cannot choose their pond.

One implication of using this model of status is that redistribution does not necessarily improve (or change at all) how status affects utility. For example, suppose that status is based on average consumption, and we perfectly equalized all incomes. Those whose status goes up would have increased utility and those whose status goes down would have decreased utility. Depending on the distribution of individuals and exactly how the comparison function enters utility, overall status effects may be higher, lower, or unchanged. Also, if we, say, destroyed \$1 billion of Bill Gates's fortune, everyone's (but Gates's) utility would go up because average income would go down. Depending on how strongly status matters and how social weights are computed, overall social welfare might even go up because of this destruction of wealth. Giving the same \$1 billion to the poor, however, would leave average income unchanged but move the poor closer to the middle. The effect on social welfare is indeterminate and would depend on the factors just mentioned.

This status-based utility function is then run through the optimal tax analysis. The status term in the utility function leads to an extra constraint in the optimization.⁵ Like in the more general case, income effects make interpretation difficult, although it is clear that the additional term alters the general formula. To get some intuitions, Oswald makes two simplifying assumptions. First, he considers only the case where the comparison is average income. Second, he considers the case where envy (or altruism) has no effect on consumption decisions or labor effort. Instead, it merely reduces utility. (Technically, it is additively separable in status.)

With these assumptions, Oswald is able to derive a very simple term for the effect of envy or altruism on tax rates. In particular, suppose that average income increases by a dollar and consider the effect on utility of each individual (assuming the individual's income, however, is constant). The sum of these effects across the whole population determines the cost of lowering taxes by a margin.⁶ For example, if individuals are generally envious, their utility goes down as average income goes up (holding their income constant). Lowering taxes increases average income and, therefore, lowers utility, all else equal. Thus, if individuals are generally envious, marginal tax rates should be higher than otherwise.

Note the information that the model needs. We need to know the shape of $u_{\alpha,n}$, which is the marginal utility from status for each type of individual n . For example, we need to know how having status affects

⁵In particular, Oswald gets $t' = \phi p \omega(n) / \lambda + \text{normal term for } t'$, where p is the vector of producer prices and λ is the multiplier for the revenue constraint. Because p and λ are both positive, the sign of the additional term depends on the sign of ϕ and ω . To determine these, he needs the additional assumptions.

⁶In particular, the shadow price of envy (or altruism) is equal to (minus) the sum of the effects on marginal utility of changes in average income, or $\phi = - \int u_{\alpha} f(n) dn$ where u_{α} is marginal effect of utility from status.

those with twice average income compared to how it affects those with three times average or one-half average income. If everyone compares themselves to the average, we would want to know, for example, whether status benefits decline with distance from the average or increase with distance.

Finally, note that the model abandons the stadium theory of status competition because the effect of status on work effort is eliminated (through the simplifying assumption of a separable utility function). Because status has no effect on work effort, the change in the tax function due to status is also unrelated to labor effort. Status, in this model, acts like a taste for redistribution. This should be controversial. It is one thing to ask everyone in the stadium to sit down. It is another to give weight to a preference that others be worse off.⁷

Ireland (2001) (and also 1998) models status concerns differently. He starts with a standard utility function where individuals maximize a function of consumption and leisure. He then assumes that individuals care about how others perceive them. Others, however, can only observe particular types of visible status consumption, s , such as large automobiles. Others observe this status consumption and use this signal to infer utility. The individual giving the signal knows this and gets utility from their inferences. Status consumption is nonvaluable other than as a signal to others. (It may be closely related to valuable consumption – it may involve a car for example, but it is the portion beyond the optimal amount which is consumed solely for signaling purposes.) Ireland allows utility to vary by setting overall utility equal to a weighted average of own consumption and the benefits of status consumption. Thus, an individual of type n weighs own consumption by $1-\alpha$ and status by α , maximizing:

$$z(n) = (1-\alpha)U(c, h) + \alpha U(v(s)^*, h^*)$$

⁷Tuomala (1990) is able to generalize the paper to allow envy to have behavioral effects, but the generalized form of his conditions do not allow easy interpretation.

where U is the standard utility function, c is consumption, h is leisure, s is status consumption, and $v(s)$ is the total consumption others assume he has when they observe s . The asterisks symbolize the amounts others impute when they observe s . Ireland considers only the case where the signal separates types, and then shows that signaling and utility increase with ability types (higher ability individuals signal more and have higher utility). A key fact to note that status in this model is not zero sum. Status is merely perceived utility, which can increase for everyone.

Ireland runs this utility function through an optimal tax analysis based on Diamond (1998) and, making the same quasi-linear in consumption simplification, gets the same expression for marginal tax rates given above: $t'/(1-t') = ABC$. The only difference is in the B term. Recall that B was the social cost of taking a dollar away from everyone above income level n when we are trying to determine the marginal rate at that level. In Ireland's model, only the "normal" or "own" utility matters in this calculation.

For example, suppose that we are computing the tax rate for individuals at some level n . If μ^n is the average of marginal social weights on people of type higher than n , the Diamond expression for B would be $1 - \mu^n$. When we add status signaling, the expression is $1 - (1 - \alpha)\mu^n$, where α is the status weight and $1 - \alpha$ is "own" weight in the utility function. We only count the own utility cost of taking a dollar away from higher-type individuals. Thus, the more high income individuals are concerned about status, the less we weigh their welfare and the more we are willing to impose high *average* rates on them.⁸

⁸Ireland's equations imply higher average rates on high types because the status enters through the B term. The B term, recall, measures the cost of taking a dollar away from everyone of higher type by raising marginal rates at type n . This might result in higher marginal rates at high income levels, but it depends on the distribution of types and labor supply elasticities. Ireland (2001) offers simulations to illustrate.

To illustrate, assume the government is utilitarian (along with Ireland's assumption that preferences are quasi-linear in consumption). Without status, the optimal marginal tax rate would be zero. No redistribution would be desirable because marginal utility does not decline with income (quasi-linear preferences) and the government does not otherwise care about inequality (utilitarianism). Mathematically, B would be zero. With status, however, B is a function of α . If caring about status is constant (so that α and, therefore, B , is constant) and labor supply elasticity is constant (so that A is constant), tax rates would depend on the distribution of skill types.⁹ We get positive marginal tax rates in the status case but not in the normal case because taxes can reduce signaling costs: it is cheaper to signal type, thereby reducing wasteful status consumption.

Note that what is driving the model is not that status seeking is zero sum or that status consumption does not increase overall utility. Consuming a status good in this model is very much like consuming any other good in that it increases utility and also that there are no particular external effects. The difference between status consumption and other consumption is that status consumption increases utility indirectly by signaling to others and it is others' esteem that increases utility. We can tax the signal and not reduce its benefits. Indeed, by taxing the labor income of low types, we can make status signaling cheaper for high types.¹⁰ Note that unlike in the Oswald model, there can be labor supply effects of status seeking, which are reduced through taxation. Thus, we might think of the Ireland model as the half-stadium model. There are labor supply effects, but status is not zero sum.

The key information about status in the Ireland model is the distribution of $\alpha(n)$: how preference for status varies across the

⁹Ireland uses a Pareto distribution as an example to illustrate relative effects of status. With this distribution, marginal rates would be constant as well.

¹⁰A similar effect can be seen in Hopkins and Hornienko (2004). These authors consider only a corrective or Pigouvian tax rather than a complete optimal income tax, but their corrective tax has the similar feature of reducing signaling costs by high earners.

population. Ireland illustrates this in several examples. In one example, he compares a society with five types in a Pareto distribution, and a constant elasticity of labor supply with and without a constant status parameter (α) equal to 25 percent. Taxes are uniformly higher when status matters, but marginal rates increase more slowly – when status matters, tax rates are higher but the schedule is flatter. The reason the schedule is flatter relates to the comment above, that higher marginal tax rates on low types reduces signaling costs for high types. On the other hand, if α increases with type, the tax schedule is both higher and steeper than in the case where status does not matter. Thus, the distribution of α is critical.

As noted, there are a number of other models of status and taxation, but at this point, we can ask whether the empirical happiness research can provide the information needed by these models (and whether the models are supported by the information we have). In both cases, the key information is how status matters for different individuals in the population.

B. *The Evidence*

The empirical evidence for status starts with the Easterlin paradox.¹¹ Easterlin (1974, 1995). The paradox is that there is a positive relationship between income and happiness within a country for different individuals at a given period of time, but, once a country has reached some minimum level of wealth, a very small (or zero) relationship between overall wealth in a country and happiness (both across countries and within a country over time). One way to reconcile the data is assume that happiness is relative: increasing income within a society improves status and, therefore, happiness, but changing the overall wealth of society does nothing for status rankings within the country and, therefore, does not affect happiness. Although the Easterlin paradox suggests status matters, it does not provide any direct evidence. Moreover, the Easterlin paradox

¹¹For a survey of the literature, see Clark, Frijters and Shields (2006).

is too broad-based and crude, to give us the information needed to solve the optimal tax problem. Instead, we must look at more direct evidence.

There are a large number of studies on this topic, and I cannot do justice to all of them. Clark, Frijters, and Shields (2006) provides a good survey. I consider here the recent studies by Ferrer-i-Carbonell (2004), Luttmer (2005), and Clark and Oswald (1996).

Luttmer (2005) examines how reported well-being correlates with neighborhood income. He uses the National Survey of Families and Households from 1987-88 and 1992-94 which included a question on well-being. He is able to construct panel data for about 10,000 individuals living in more than 550 separate areas. He matches this data with information about local earnings. To get local earnings, he estimates information from the Public Use Microdata Areas (PUMAs) with CPS data on national earnings by industry, occupation, and year.¹² In simple regressions on this data, he finds a coefficient of 0.20 on log own household income and a coefficient of -.17 on average log of predicted household income in one's locality. Thus, relative income matters approximately as much as own income: a dollar of increased income increases happiness about the same amount as a dollar reduction in average income in a neighborhood. The finding is robust to a variety of controls and highly significant.¹³ In a specification that hints at some of the information required by the optimal tax models, he examines how local earnings affect happiness for households above and below the local median. He finds almost identical effects: wealthier and poorer families respond identically to a change in predicted local earnings.

¹²PUMAs range in size from about 127,000 to 144,000. [Using predicted local earnings opens up the possibility that individuals whose income falls relative to predictions are not doing with in their careers and that utility drops for this reason rather than merely because they compare themselves to others. That is, predicted earnings might act as information.]

¹³Might be just about definition of happiness. Hard to interpret questionnaire. Tests this by looking at other measures of well-being.

Luttmer imposes PUMA's as the comparison group. It would be nice to know how this is effecting the results and whether the neighborhood is the right comparison. In unreported regressions, he says that he finds that within a neighborhood, individuals compare themselves to smaller subgroups, in his specification, college educated or not. Further examination of this issue would likely be helpful because neighborhoods as comparison groups raise the "right pond" issue. If neighbors are the comparison group, comparison groups are endogenous because you can choose where to live. If individuals know about the comparison income affect, however, we might expect sorting to take advantage of this. Thus, a wealthy person might live in a poor neighborhood to increase his subjective well-being. The effect of such sorting on overall happiness, if it were to happen, would be unclear, but it is also contrary to the casual observation that individuals sort into neighborhoods by wealth, not against wealth. Luttmer cites a paper by Loewenstein, O'Donoghue, and Rabin (2003) for the claim that individuals make forecasting errors when choosing neighborhoods, although these authors only casually suggest this possibility and do not provide evidence for it. An alternative explanation is that reference groups are not endogenous – they are the type of individuals you compare yourself to and would be even if you did not choose to live near them. Therefore, there is no cost to sorting into neighborhoods by wealth. The neighborhood effect is simply picking up the fact that comparison groups and neighborhoods coincide.

Ferrer-i-Carbonell (2005), uses panel data from the German Socio-Economic Panel to estimate an equation in which subjective well-being is a function of own income, a comparison, and a set of controls. The sample includes about 16,000 individuals from the former East and West Germany during the years from 1992 to 1997. He tests three comparisons: average income of the reference group, the difference between the log of own income and log of the average reference group income, and an asymmetric measure in which being below the average of reference group income matters more than being above. The reference group is one of 50 different groups, categorized by education level (five categories), age

bracket (10 year windows), and whether the individual lives in the former East or West Germany.

Like Luttmer (and other studies), he finds that reference group income matters. Moreover, like Luttmer, he finds roughly similar coefficients for own income and reference group income (of opposite signs). For example, if own income and reference group income go up by the same amount, subjective well being stays roughly constant. The asymmetric test is interesting for the optimal tax analysis because it might help us understand how status effects vary across the population. Unfortunately, the data seem inconclusive: He finds some evidence for asymmetry in the West German sample but not for the East German sample.

A final paper that attempts to measure relative preferences directly is Clark and Oswald (1996) They use 1991 British Household Panel survey questions on job satisfaction. Individuals rated job satisfaction from 1 to 7 for seven job-items (pay, prospects, etc.) and then answered a question about overall job satisfaction. The authors use this last question as a measure of subjective well-being or utility.¹⁴ The reference group is an estimate of typical income of someone with the individual's observable characteristics.¹⁵ With this reference group, the authors find a negative relationship between reference group income and job satisfaction. They then confirm this result with information from outside the data set about expected earnings.

¹⁴Note that it is not clear that this makes sense. So, for example, C&O find that job satisfaction does not correlate with income – get a U-shaped line with the lowest wage earners the most satisfied. Log income gets a negative coefficient. This seems odd if job satisfaction is utility but may not be if it is just a part of utility. We can imagine high earners sacrificing job satisfaction for, say, sending kids to college. Overall, might be happy with choice of job but not when asked about job alone.

¹⁵There might be real problems with using this as a reference group because the comparison between own income and this measure might merely show that the individual is underperforming relative to expectations. Job dissatisfaction might arise because of negative signals from bosses or peers instead of from relative preferences.

There are a number of other studies as well as problems inherent to all of the studies.¹⁶ We can, at this point however, ask what we get out of the empirical literature. Almost every paper, including all three reviewed here, finds that relative income matters. Moreover, relative income seems to be close to a zero sum game. Thus, if both an individual's income and the reference group's income goes up, subjective well-being seems to stay constant. This means that there is little support for the Ireland (1998, 2001) formulation of status, which was not zero sum.

¹⁶Blanchflower and Oswald (2004) look at U.S. General Social Survey data by state and over time. Their paper is not focused on relative income – it is a general study of the determinants of happiness. They find, for example, that the overall trend has been negative for the U.S. and that work and marital status have large and well-defined effects. They test the relative income hypothesis by examining how the ratio of an individual's income to state per capita income affects happiness. They find a positive coefficient.

McBride (2001) uses GSS data (only 324 observations) to estimate well being as function of log income + log (past standard of living) + log(cohort income) The reference group individuals within 10 year age group. Income above \$75,000 doesn't count. Say that if they find relative income effects at low end, that is enough.

He finds negative effects of cohort income on happiness. He finds some interesting results: absolute income seems to matter more for low income individuals and relative income more for high income individuals. These finds are suggestive of data required by the optimal tax models. Unfortunately, his data is crude – SWB is in only three categories – his sample size very small, and he does not have data for income above \$75,000, making the study suggestive but in need of further confirmation.

Clark, Frijters, and Shields (2006) list some endemic problems with the empirical estimates of relative preferences. For example, most studies impute a reference group rather than allowing the individuals to make this selection. In addition, they do not take into account that reference groups might be endogenous and chosen to maximize long run utility (for example, an individual might choose a high reference group, making him unhappy today, but with the benefit of inducing harder work and happiness in the long run – would you rather go to a school with a bunch of smart people who will inspire you to learn more or with a bunch of mediocre people that will give you immediate status benefits?). In addition the problems listed by Clark et al, reference group earnings might be information about performance rather than creating envy.

The empirical work has used different formulations to measure relative income, in almost all cases, imposed by the researcher. The only paper to attempt to measure different formulations is Ferrer-i-Carbonell (2004). Although he imposes a reference group, he looks at different possible comparisons, finding support for a comparison to average income of the reference group and some limited support for an asymmetric response, with relative income mattering more for the poor.¹⁷ Similarly, there is no clear notion of how to measure the reference group. In all cases, it is just imposed by the researcher.

Most importantly, none of the studies tell us what we need to know for the existing optimal tax models. We need to know something about the distribution of relative preferences across incomes. That is, we need to know whether individuals who are poor, middle class, or rich are more or less envious and who their reference groups are. More fundamentally, Frank's stadium analogy suggests that what we need to know is what happens if reference group income declines but leisure increases because the goal of a tax based on envy is in part to reduce the work externality. No study gives this information. They measure a change in subjective well being when reference group income changes, but we do not know why this is occurring. If reference group income goes up (but own income does not), this could be because the reference group has increased hours or increased hourly wages. Without knowing which, we cannot sort out the effect of a tax. That is, an increased labor tax will make people work less and earn less. Showing that subjective well-being goes up when the reference group earns less might not show that it would go up if the reference group also works less.

Before turning to Robert Frank's proposals on taxation and status, there are two final comments on optimal taxation. First, a prediction of those arguing for status taxation is that higher labor income tax rates

¹⁷McBride (2001) finds the opposite, that relative income matters more for the wealthy, but McBride's paper has a number of problems that make it less convincing than Ferrer-i-Carbonell (2004).

should, up to a point, increase overall happiness. I do not know of any study that has attempted to measure this. Although one can imagine many complications, we should be able to compare data such as that used by Easterlin to labor tax rates in those countries to see whether higher taxes do indeed increase happiness.

Second, to understand the effect of taxing status, we need a better understanding of why status concerns arise. They might just be an evolutionary detritus akin to a peacock's feathers. They might, however, continue to serve a useful function, such as sorting or providing incentives. For example, there is a literature arguing that rank order tournaments are, in certain circumstances, efficient. Bolton and Dewatripoint (2005). Rank order status concerns might provide similar incentives. Before concluding that taxing status is desirable, we should know more about why it arises.

C. *Frank's status taxation proposals*

Robert Frank has been very prominent in thinking about the link between status seeking and taxation. His proposal, however, is distinct from the optimal labor income tax discussed above. Although never outlined in detail, Frank has argued in several papers and books for a progressive consumption tax because of status concerns. Frank (1985, 1997, 1999, 2000, 2005). The question for this section is how Frank's arguments relate to the optimal taxation arguments given above.

The optimal taxation literature models a tax on labor income while Frank argues for a tax on consumption. The two, however, are closely related and, in basic cases, identical. The reason is that leaving aside bequests and gifts (neither of which is an apparent concern here), labor income and consumption have the same present value: you can only spend what you earn. This means that taxes on labor income and consumption also tend to have the same present value and, therefore, impose the same burden.

To illustrate, suppose that an individual earns labor income y in period 0 and can consume it in either of two periods, period 0 and period 1. If the individual waits until period 1, he invests it at rate r . If C_i is consumption in period i , we know that

$$y = C_0 + C_1/(1+r)$$

If we impose a tax on consumption at rate t , labor income must equal the present value of consumption inclusive of the tax. Therefore,

$$y = C_0(1+t) + C_1(1+t)/(1+r)$$

If we divide by $(1+t)$, we see that a tax on consumption is equivalent to reducing labor income by $1/(1+t)$, which is equivalent to a tax rate of $t/(1+t)$ on labor income. This identity used a flat rate tax t but the same holds for progressive taxes on labor income. In particular, if an individual faces a some tax rate on labor income, this tax rate is equivalent to tax at flat rate on his consumption, even if other individuals face other labor income tax rates. (As will be discussed below, the reverse is not quite true: a progressive tax on consumption need not translate into a particular labor income tax rate.)

There are two relevant differences between Frank's arguments and the optimal income tax arguments.¹⁸ First, Frank argues that holding labor income constant, individuals in status races will consume too early and save too little. The idea is that spending is observable while savings is not, so individuals concerned with status will spend too much. On this

¹⁸There are a number of other differences between a progressive, individual-level consumption tax and a labor income tax. For example, as Summers (1981) points out, the timing of government revenue flows is different in the two systems. There are also administrative differences. With a labor income tax, employers could withhold taxes while a withholding system would be difficult to incorporate into a progressive consumption tax. Although labor taxes are used throughout the world, no country currently uses a progressive consumption tax of the sort Frank proposes. None of these is directly relevant to the discussion.

basis, he argues that we should shift from the current income tax, which burdens savings as well as labor, to a consumption tax.

The optimal income tax models take for granted that we should tax labor income and not the return to savings. They are models of labor income taxes, not conventional capital income taxes. In addition, there are good reasons, independent of status concerns, for taxing only consumption or labor and not savings.¹⁹ Nevertheless, the optimal income tax models generally do not have savings (or even time), so they do not consider whether status concerns reduce savings and, if so, whether the tax system should be modified as a result. Frank argues for lower taxes on savings, so, if we otherwise believe the tax on savings should be zero, perhaps we should have a savings subsidy.

A difficulty with this argument is that it treats savings like any other good. Suppose that there are two goods, apples and oranges and consumption of apples is an observable signal of status and the consumption of oranges is not. We might expect individuals to consume more apples than otherwise. In Frank's world, savings is like the orange – a good the consumption of which is unobservable and, therefore, underutilized. Savings, however, is just future consumption: it is like two apples in the future, not like the orange. Consuming more today to win a status race means consuming less in the future and losing the status race in the future. Frank has to be arguing that status competition distorts discount rates because individuals would have to be willing to give up future status for current status. This could be true, but as far as I know, there is no empirical support for this. None of the studies reviewed above say anything about discount rates. It does not seem implausible and I am not aware of a study that rejects the claim (or that even addresses it), but it is not yet supported by evidence. We do not know, therefore, whether status considerations argue for a lower tax or even a subsidy on savings.

¹⁹See Bankman and Weisbach (2006).

The second relevant difference between a progressive consumption tax and a labor income tax is that if the rate on consumption is progressive within a given period, lumpy consumption is taxed at a higher rate than level consumption. To illustrate, compare a person who spends \$100 each period to a person who spends \$200 every other period with the same overall total. Suppose that we impose a progressive consumption tax with a zero tax rate on the first \$100 consumed each period and a 20 percent tax rate on everything above \$100. The first individual would owe no tax. The second individual, with the same overall consumption but with a lumpier pattern, would pay \$20 of tax every other period. Lumpy consumption, by pushing individuals into higher tax brackets, is taxed at a higher rate than level consumption. It is for this very reason that Vickrey (1947) proposed an averaging scheme to complement progressive taxation: income averaging prevents this effect.

To support Frank's proposal separately from the findings of the optimal income tax literature, it would have to be the case that lumpy consumption is particularly related to status and that it is desirable to tax lumpy consumption at higher rates than level consumption.²⁰ As far as I know, however, there is no data supporting a claim that status concerns lead to lumpy consumption. None of the studies are able to pinpoint how comparisons are made and whether big splurges create more status than constant, everyday spending. Perhaps big splurges are more visible and, therefore, create more status, but this is merely an assertion, not something yet supported by the literature. Frank often cites the Easterlin data for the claim that preferences are relative, but this data says nothing about lumpy as compared to smooth consumption.

²⁰The statement in the text is relative to the optimal income tax models. Frank also wants to change the current income tax, which imposes a burden on savings, to a consumption tax. This change is taken for granted in the optimal income tax models because they are models of a labor income tax. There are, however, very good independent reasons for shifting to a consumption tax.

Moreover, whether something counts as lumpy depends on the accounting system. Durable goods, such as houses, cars, and watches, all frequent targets of Frank, are often counted as lumpy because their purchase is made all at once. Because they are durable, however, they actually offer consumption over a period of time. For example, one can alternatively buy a durable good or rent it. A progressive consumption tax would likely treat buying it as lumpy and tax it at a high rate but treat renting it as smooth, taxing it a low rate. The lumpiness of these goods, however, is simply an arbitrary construct of accounting rather than anything fundamental. Moreover, it is hard to see how owning a fancy car compared to leasing it changes its status enhancing properties.

To summarize, there are two claims that Frank makes that are distinctive from the optimal tax literature reviewed above. The first, is that status concerns change discount rates, causing individuals to care more about status today than in the future. Although possible, there is not yet any evidence to support this claim. The second is that status concerns lead to lumpy consumption, as conventionally measured by tax systems. There is also, to my knowledge, no evidence to support this claim, particularly because tax systems' measurements of lumpiness is arbitrary. Without these distinctive elements of Frank's proposals, the analysis of taxation and status reverts to the discussion of optimal taxation considered above.

IV. Adaptation to Disability and Taxation

A second important issue raised by the happiness literature for the tax and transfer system is the idea that individuals may adapt to circumstances. The most striking version of this claim is that individuals with severe disabilities such as tetraplegia, are just as happy as healthy individuals. Although the data do not support this precise claim, they do show significant adaptation to disability.

Most countries, including the United States, spend substantial (sometimes huge) resources on individuals with disabilities. These

include direct transfers through disability insurance programs, tort liability, and safety regulations (to prevent disabilities in the first place). In a very rough estimate, Weisbach (2007) concluded that the major disability programs in the United State cost around \$275 billion per year in direct costs and that these expenditures by no means exhaust the spending on the disabled. Total annual spending in the United States could easily be over half a trillion dollars a year. If the adaptation literature is correct, much of this may be wasted.

There is already some writing on how adaptation affects tort damages and safety regulation. Bagenstos and Schlanger (2007), Sunstein (2007). I will focus here on tax/transfer systems using the same approach used above. I will first look at the theory to see what the potential effects of disability are and what we would need to know to determine how to set the tax parameters. I will then ask whether the evidence on subjective well-being to date tells us what we need to know. I will give the same answer, which is that the empirical literature does not yet tell us what we need to know but it might in the future.

A. Optimal taxation and disability

Since Diamond and Mirrlees (1978), it has been standard in the economics literature to model disability as a wage rate of zero but not otherwise affecting an individual. The planner's problem is to provide insurance against the risk of having a zero wage rate at some time in the future while reducing the moral hazard problems created by offering the insurance. This model of disability is not very helpful in the present context because I want to examine the distinctive element disability has in happiness. Disability has to be thought of as affecting individuals more generally: We want to compare an individual with a disability to an individual with the the same wage rate without a disability and determine, as between these individuals how they should be taxed. That is, we have to expand the optimal tax theory to allow individuals to vary in two dimensions: wage rate and disability.

This problem, of differences among individuals other than wages, was noted by Mirrlees (1976), but in the thirty years since then, modeling has remained its infancy. The key problem is that if individuals vary in more than one, unobservable dimension, incentive constraints may no longer bind in a single direction, making the model impossible to solve in any general form. As illustrated above, even in the single dimension case, simplifying assumptions are often used to help interpret the results. Adding a second dimension compounds the difficulty significantly.

The most straightforward way to simplify the problem is to assume that disability is observable. This will be true of some but not all disabilities. (Even with a given medical diagnoses, however, individuals may vary dramatically, so observing many disabilities may be difficult.) If disability is observable, the standard optimal tax results hold because we can divide the population into categories. Within each category, individuals would be the same except with respect to their earning ability, and we can apply the optimal tax results to each category. The overall tax schedule across categories would be linked by a common shadow price of revenue. Kaplow (2006), Boadway and Pestieau (2003).²¹

This simplification means that we can determine the treatment of the disabled by examining how disability affects the parameters in the optimal tax formula. Recall, that the optimal rate at some income level n is set according to the equation

$$t'/(1-t') = ABC$$

where A was a measure of labor supply elasticity, B was a measure of the social value of taking money from individuals of higher ability than n , and C was a measure of the distribution of individuals. C should be easily observed and is not affected by the happiness literature. It is simply a

²¹Kaplow (2006) briefly discusses the case where the second dimension of difference (here disability) is not observable. He characterizes the solution as imposing commodity taxes with effects that roughly mimic the case where disability is observable.

matter of counting. The A term relates to labor supply elasticity. Although important, this term is also not particularly of interest to the happiness literature. (There is an important policy dimension embedded in this term, however, because the labor supply of the disabled may depend on policy toward the disabled, such as whether buildings are accessible.)

All of the action is in the B term. To see how this works, we have to expand the term. The term weighs the cost of taking a dollar from individuals above some income level, n . We need to know two factors to determine this: the social welfare weights on individuals and their change in utility when they lose a dollar. That is, we need to know $W'(u)u_c$, where $W(u)$ is the social weighting of an individual and u_c is their marginal utility from consumption. We need to add these terms up for all individuals with income above n .²²

Determining $W'(u)$ depends on philosophical theories and not empirical facts about happiness, so there is little to say about it here. We also need to know the level of utility of a disabled and nondisabled individuals of a given income and the marginal utility of those same individuals. That is, to determine the B term, we need to know how disability affects individuals. The happiness research may have a lot to say about this.

B. Evidence on the subjective well-being of the disabled

There is a cottage industry examining the effect of disability on subjective well being, originating with the well-known study by Brickman

²²In the continuous case, for a tax rate on a type- n individual, we get

$$B = \frac{\int_n^\infty \left(1 - \frac{W'(u)u_c}{\lambda}\right) dF[n]}{1 - F[n]}$$

. The division of $W'(u)u_c$ by λ converts the social costs to dollars. It is the λ term that links the tax schedules of the different types of individuals.

et al (1978). The overwhelming majority of these report only changes in subjective well being, not controlled for any differences among individuals, and, in particular, not controlled for income. A typical example is Dijkers (1997), which is a meta-analysis of 22 studies but does not provide any significant controls.

There are a handful of studies that begin to give us the data we need. Smith et al (2005) compares changes in subjective well being due to a disability for individuals with higher than and lower than median wealth (within the sample). They find that the wealthier half of their sample experienced a smaller decline in subjective well-being than the poorer half. They conclude that wealth has a larger effect on subjective well-being at the onset of disability than it does more generally. This would mean that disability increases the marginal utility of wealth.

Oswald and Powdthavee (2006) look at the British Household Panel Survey from 1996 to 2002 and examine how complete disability (meaning inability to work) affects well-being, measured on a scale from 1 to 7. A significant and new aspect of the study is that their data allows them to do a panel study, while prior work tends to be cross-sectional. Moreover, their study uses a much larger data set than prior studies. The raw data show the disabled to be less happy: the nondisabled have an average of well-being score of 5.3 while disabled have a score of 4.3. The authors then run a time series analysis, which shows some but not full recovery. Controls for gender, age, and education do not significantly change the results. They then control for income and find similar effects. Although they do not run the regression, their data should allow them to compute the interaction of disability and income and, therefore, the marginal effect on reported well-being of income for the disabled and the nondisabled.²³

²³Note that there is a separate issue of whether data on marginal happiness scores tells us anything about marginal utility. The problem is that there might be some nonlinear translation of happiness reports to actual happiness.

A problem with the Oswald study is that at their measured onset of disability, individuals already have a significantly lower happiness score than individuals who never experience a disability. At the onset of the disability, their well-being goes down from this lower starting point and then recovers somewhat, at most up to the lower starting point. For example, an individual who never experiences a disability reports a happiness score of 5.3. An individual who experiences a disability starts at 4.2 prior to the onset, goes down to 3.9 immediately after the onset, and then recovers to almost 4.1. As compared to the pre-disability number, the recovery is significant. The lower starting point, however, suggests that the disability began prior to the reported time, or that for some reason, those that are likely to become disabled are unhappy for some other reason. If we compare the fully recovered well-being to the well-being of the nondisabled (on the assumption that if we could correctly measure the onset of disability, the individuals would start at the nondisabled level), recovery is not even close to complete, even for the less seriously disabled. This problem, however, may not matter for the tax analysis. The tax analysis is not particularly focused on adaptation. Instead, we need to know the level of utility and marginal utility for disabilities. Even if the study cannot measure the onset of disability, it can provide information about the long-term welfare of individuals with disability.

These studies remain preliminary. Moreover, we have a wealth of data on the income, consumption, and employment of the disabled, and it would be nice to be able to reconcile this data with the well-being data. For example, we know that disability makes it much less likely that one is employed and that unemployment leads to a significant decline in subjective well-being. If disability does not lead to a significant decline in well-being, however, the reports seem inconsistent.

V. Additional Tax Issues and Conclusion

There are a number of other ways that the happiness literature might change basic tax results. This final section speculates on some of the possibilities.

1. The literature examined so far makes the assumption that we cannot directly tax status seeking activities. If status seeking simply involves earning more labor income, then the labor income tax directly taxes the activity, but if status seeking involves consumption of particular goods, adjusting the labor income tax only indirectly addresses the problem. Adjusting the labor income tax is justified on the grounds that what goods act as status goods is highly contingent. If we try to tax them, we make them more expensive and the goods that are used to signal status will shift.²⁴

Ireland (1994) suggests taxing status goods but does not address what these goods might be or whether they could feasibly be taxed. In a recent paper, Tomer Blumkin and Efraim Sadka (2007) suggest that charitable donations are status goods. Although charitable donations have a positive externality (they help the recipient as well as provide utility to the donor) and, therefore, might be subsidized, if they are status goods, we might want to tax them. Whether the net result is a tax or a subsidy depends on the parameters.

Frank (2000) suggests large homes, fancy automobiles, and mechanical watches as status goods. Observers of social behavior note many others. (Ireland (1994) recounts a story about a brand of basketball shoes in high demand among the urban poor.) We do not, however, understand how these arise and the extent to which they can shift if taxed.

2. Standard models of optimal taxation are not particularly concerned about unemployment. Unemployment is simply labor supply of zero. Holding income constant (say, through transfers), utility goes up

²⁴Ng (1987) suggests that some goods, so-called diamond goods, are valued for their value. If taxed, individuals will consume a lower quantity of the good but spend the same total amount. He uses the example of a diamond. He suggests that an individual who wants to spend \$1,000 on a diamond will spend the same \$1,000 whether this buys one carat or half a carat. As Ng carefully points out, diamond goods are not the same as status goods because diamond goods can be consumed privately and status goods may not have the “valued for value” feature of diamond goods.

with labor supply goes down, even to zero. Consider a high marginal rate on a low ability person. If the person reduces work effort, say to zero, there is little lost productivity. The high rate at the low end, however, is inframarginal to a large number of individuals. Therefore, high marginal rates on low incomes may be desirable.²⁵

The happiness literature generally finds a large negative effect on subjective well being from (involuntary) unemployment, controlling for income. For example, Clark and Oswald (1994) find that “joblessness depressed well-being more than any other single characteristic, including important negative ones such as divorce and separation.” This suggests that the standard utility function used in the optimal tax literature might not accurately reflect utility when labor supply is low. I do not know of any models of taxation that attempt to include these effects in utility functions, but if the data is correct about the effect of unemployment, incorporating these effects may change the results significantly.

3. Similar to unemployment, standard tax models do not attribute any special benefits to marriage. At most we see economies of scale or benefits from mutual altruism. Kaplow (1998), Bittker (1975). The happiness literature, however, shows a large effect on subjective well being.

Tax systems with increasing marginal rates generally must choose between marriage penalties or marriage bonuses. A possible implication of the happiness literature is that the tax system should lean toward marriage bonuses. More evidence, however, is needed. For example, we can imagine a marriage bonus as simply paying a couple an amount each year if they stay married. It is not yet clear from the happiness literature whether this would be effective or whether only marriage that would occur absent special subsidies increase subjective well being.

²⁵This is to be distinguished from average rates. Average rates could be low or negative because the poor may receive transfers.

4. There have been some suggestions that taxes should be age-based. Kremer (2005). The argument is that by separating individuals by age, we can take better advantage of the parameters in the optimal tax equation. In particular, labor supply of the young might be very elastic while the labor supply of the middle-aged might be very inelastic. (This is the A term in the optimal tax equation given above.) Moreover, distribution of wage rates is likely to be different for the young than for the middle-aged or old. If the young have a higher percentage of individuals with low wages, high rates on low wages distorts more for the young than for others. (This is reflected in the C term of the optimal tax equation given above.) Kremer (2005) suggests that, all else equal, taxes should be lower for the young than for the middle aged.

The happiness literature, however, suggests that individuals are least happy in middle age. If this is true, the distributional consequences of imposing higher taxes in middle age may be negative. We would be making unhappy people more unhappy. The B term in the optimal tax equation would vary by age and would offset the effect of age-based A and C terms. Theorizing about age-based taxation is in its infancy, so we do not yet understand the full effects. Nevertheless, the happiness literature should be relevant to future learning in this area.

It is clear that the findings of happiness research have the potential to change tax policy. The research is interesting by itself, but if one of the goals of the research is to have policy implications – to find out how to make individuals happier – the key message of this paper is that the research must line up better with the normative models.

References

Abel, Andrew B. "Optimal Taxation when Consumers Have Endogenous Benchmark Levels of Consumption." *Review of Economic Studies* 72 (2005): 21-42.

Allgood, Sam. "The Marginal Costs and Benefits of Redistributing Income and the Willingness to Pay for Status." *Journal of Public Economic Theory* 8 (2006): 357-377.

Bagenstos, Samuel R. and Margo Schlanger, *Hedonic Damages, Hedonic Adaptation, and Disability*, working paper (2007)

Bankman, Joseph, and David A. Weisbach. "The Superiority of an Ideal Consumption Tax over an Ideal Income Tax." *Stanford Law Review* 58 (2006): 1413.

Bittker, Boris I. "Federal Income Taxation and the Family." *Stanford Law Review* 27 (1975): 1389-1463

Blanchflower, David G., and Andrew J. Oswald. "Well-being over time in Britain and the USA." *Journal of Public Economics* 88 (2004): 1359-1386.

Blumkin, Tomer and Efraim Sadka, *On the Desirability of Taxing Charitable Contributions*, CESifo Working Paper No. 1900 (2007)

Boadway, Robin, and Pierre Pestieau. "Indirect Taxation and Redistribution: The Scope of the Atkinson-Stiglitz Theorem." in *Economics for an Imperfect World* (2003). (Richard Arnott, Bruce Greenwald, Ravi Kanbur and Barry Nalebuff eds.)

Bolton, Patrick, and Mathias Dewatripont. *Contract Theory*. Cambridge, MA: The MIT Press, 2005.

Boskin, Michael J., and Eytan Sheshinski. "Optimal Redistributive Taxation when Individual Welfare Depends Upon Relative Income." *The Quarterly Journal of Economics* (Nov. 1978): 589-601.

Brickman, Philip, Dan Coate, and Ronnie Janoff-Bulman. "Lottery Winners and Accident Victims: Is Happiness Relative." *Journal of Personality and Social Psychology* 36.8 (1978): 917-927.

Clark, Andrew E., Paul Frijters, and Michael A. Shields. *Income and happiness: Evidence, explanations and economic implications*, Paris-Jourdan Sciences Economiques, working paper No. 2006-24 (2006)

Clark, Andrew E., and Andrew J. Oswald. "Unhappiness and Unemployment." *Economic Journal* 104 (1994): 648-59.

Clark, Andrew E., and Andrew J. Oswald. "Satisfaction and comparison income." *Journal of Public Economics* 61 (1996): 359-381.

Diamond, Peter A. "Optimal Income: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *The American Economic Review* 88.1 (Mar. 1998): 83-95.

Dijkers, Marcel. "Quality of life after spinal cord injury: a meta analysis of the effects of disablement components." *Spinal Cord* 35 (1997): 829-840.

Easterlin, Richard. "Does economic growth improve the human lot?" in *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz* (1974).

Easterlin, Richard A. "Will raising the incomes of all increase the happiness of all." *Journal of Economic Behavior & Organization* 27 (1995): 35-47.

Ferrer-i-Carbonell, Ada. "Income and well-being: an empirical analysis of the comparison income effect." *Journal of Public Economics* 89 (2005): 997-1019.

Frank, Robert H. *Choosing the Right Pond*, NY: Oxford University Press, 1985.

- - -. *Luxury Fever*. New York, NY: Free Press, 2000.

- - -. "The Demand for Unobservable and Other Nonpositional Goods." *The American Economic Review* 75.1 (Mar. 1985): 101-116.

- - -. "The Frame of Reference as a Public Good." *The Economic Journal* (Nov. 1997): 1832-1847.

- - -. "Positional Externalities Cause Large and Preventable Welfare Losses." *AEA Papers and Proceedings* 95.2 (May 2005): 137-141.

- - -. "Progressive Consumption Taxation as a Remedy for the U.S. Savings Shortfall." *The Economists' Voice* 2.3 (2005): 1-10.

Ireland, Norman J. "Status-seeking, income taxation and efficiency." *Journal of Public Economics* 70 (1998): 99-113.

- - -. "On limiting the market for status signals." *Journal of Public Economics* 53 (1994): 91-110.

- - -. "Optimal income tax in the presence of status effects." *Journal of Public Economics* 81 (2001): 193-212.

Kaplow, Louis. "Optimal Distribution and the Family" *Scandinavian Journal of Economics* 98 (1996): 75-92.

- - -. Optimal Policy with Heterogeneous Preferences, working paper (2006)

- - -. Taxation and Redistribution. (2006)

Kremer, Michael, Should Taxes be Independent of Age, working paper (2005)

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. "Projection Bias in Predicting Future Utility." *The Quarterly Journal of Economics* (Nov. 2003): 1209-1248.

Luttmer, Erzo F. "Neighbors as Negatives: Relative Earnings and Well-Being." *The Quarterly Journal of Economics* 120 (2005): 963-1001.

McBride, Michael. "Relative-income effects on subjective well-being in the cross-section." *Journal of Economics Behavior & Organization* 45 (2001): 251-278.

Mirrlees, James A. "An exploration in the theory of optimal income taxes." *Review of Economics Studies* 38 (1971): 175-208.

Mirrlees, James A. "Optimal Tax Theory." *Journal of Public Economics* 6 (1976): 327-358.

Ng, Yew-Kwang. "Diamonds are a Government's Best Friend: Burden-Free Taxes on Goods Valued for the Values." *American Economic Review* 77.1 (Mar. 1987): 186-191.

Oswald, Andrew J. "Altruism, Jealousy and the Theory of Optimal Non-Linear Taxation." *Journal of Public Economics* 20 (1983): 77-87.

Oswald, Andrew J., and Nattavudh Powdthavee. "Does Happiness Adapt? A Longitudinal Study of Disability with Implications for Economists and Judges." IZA Discussion Paper Series No. 2208 (July 2006). 3 May 2007 <<http://ssrn.com/=921040>>.

Rae, John, *The New Principles of Political Economy* (1834)

Saez, Emmanuel. "Using Elasticities to derive optimal income tax rates." *Review of Economic Studies* 68.1 (2001): 205-230..

Smith, Adam. *The Theory of Moral Sentiments*. Cambridge: Cambridge University Press, 1759.

Smith, Dylan, et al. "Health, Wealth, and Happiness." *Psychological Science* 16.9 (2005): 663-666.

Stiglitz, Joseph. "Pareto Efficient and Optimal Taxation and the New New Welfare Economics." in *Handbook of Public Economics* 2 (1987): 991-1042. (Alan J. Auerbach and Martin Feldstein eds.)

Summers, Larry H., "Taxation and capital accumulation in a life cycle growth model." *American Economic Review* 71 (1981): 533-554.

Sunstein, Cass. *Illusory Losses*, working paper (2007)

Tuomala, Matti. *Optimal Income Tax and Redistribution*. Oxford: Clarendon Press, 1990.

Vickrey, William. *Agenda for Progressive Taxation*. New York: The Ronald Press Company, 1947.

Weisbach, David. *A Welfarist Approach to Disabilities*, manuscript (2007)

Readers with comments should address them to:

Professor David A. Weisbach
University of Chicago Law School
1111 East 60th Street
Chicago, IL 60637
d-weisbach@uchicago.edu

Chicago Working Papers in Law and Economics
(Second Series)

For a listing of papers 1–200 please go to Working Papers at <http://www.law.uchicago.edu/Lawecon/index.html>

201. Douglas G. Baird and Robert K. Rasmussen, Chapter 11 at Twilight (October 2003)
202. David A. Weisbach, Corporate Tax Avoidance (January 2004)
203. David A. Weisbach, The (Non)Taxation of Risk (January 2004)
204. Richard A. Epstein, Liberty versus Property? Cracks in the Foundations of Copyright Law (April 2004)
205. Lior Jacob Strahilevitz, The Right to Destroy (January 2004)
206. Eric A. Posner and John C. Yoo, A Theory of International Adjudication (February 2004)
207. Cass R. Sunstein, Are Poor People Worth Less Than Rich People? Disaggregating the Value of Statistical Lives (February 2004)
208. Richard A. Epstein, Disparities and Discrimination in Health Care Coverage; A Critique of the Institute of Medicine Study (March 2004)
209. Richard A. Epstein and Bruce N. Kuhlik, Navigating the Anticommons for Pharmaceutical Patents: Steady the Course on Hatch-Waxman (March 2004)
210. Richard A. Epstein, The Optimal Complexity of Legal Rules (April 2004)
211. Eric A. Posner and Alan O. Sykes, Optimal War and *Jus Ad Bellum* (April 2004)
212. Alan O. Sykes, The Persistent Puzzles of Safeguards: Lessons from the Steel Dispute (May 2004)
213. Luis Garicano and Thomas N. Hubbard, Specialization, Firms, and Markets: The Division of Labor within and between Law Firms (April 2004)
214. Luis Garicano and Thomas N. Hubbard, Hierarchies, Specialization, and the Utilization of Knowledge: Theory and Evidence from the Legal Services Industry (April 2004)
215. James C. Spindler, Conflict or Credibility: Analyst Conflicts of Interest and the Market for Underwriting Business (July 2004)
216. Alan O. Sykes, The Economics of Public International Law (July 2004)
217. Douglas Lichtman and Eric Posner, Holding Internet Service Providers Accountable (July 2004)
218. Shlomo Benartzi, Richard H. Thaler, Stephen P. Utkus, and Cass R. Sunstein, Company Stock, Market Rationality, and Legal Reform (July 2004)
219. Cass R. Sunstein, Group Judgments: Deliberation, Statistical Means, and Information Markets (August 2004, revised October 2004)
220. Cass R. Sunstein, Precautions against What? The Availability Heuristic and Cross-Cultural Risk Perceptions (August 2004)
221. M. Todd Henderson and James C. Spindler, Corporate Heroin: A Defense of Perks (August 2004)
222. Eric A. Posner and Cass R. Sunstein, Dollars and Death (August 2004)
223. Randal C. Picker, Cyber Security: Of Heterogeneity and Autarky (August 2004)
224. Randal C. Picker, Unbundling Scope-of-Permission Goods: When Should We Invest in Reducing Entry Barriers? (September 2004)
225. Christine Jolls and Cass R. Sunstein, Debiasing through Law (September 2004)
226. Richard A. Posner, An Economic Analysis of the Use of Citations in the Law (2000)
227. Cass R. Sunstein, Cost-Benefit Analysis and the Environment (October 2004)
228. Kenneth W. Dam, Cordell Hull, the Reciprocal Trade Agreement Act, and the WTO (October 2004)
229. Richard A. Posner, The Law and Economics of Contract Interpretation (November 2004)
230. Lior Jacob Strahilevitz, A Social Networks Theory of Privacy (December 2004)
231. Cass R. Sunstein, Minimalism at War (December 2004)
232. Douglas Lichtman, How the Law Responds to Self-Help (December 2004)
233. Eric A. Posner, The Decline of the International Court of Justice (December 2004)
234. Eric A. Posner, Is the International Court of Justice Biased? (December 2004)
235. Alan O. Sykes, Public vs. Private Enforcement of International Economic Law: Of Standing and Remedy (February 2005)
236. Douglas G. Baird and Edward R. Morrison, Serial Entrepreneurs and Small Business Bankruptcies (March 2005)
237. Eric A. Posner, There Are No Penalty Default Rules in Contract Law (March 2005)
238. Randal C. Picker, Copyright and the DMCA: Market Locks and Technological Contracts (March 2005)
239. Cass R. Sunstein and Adrian Vermeule, Is Capital Punishment Morally Required? The Relevance of Life-Life Tradeoffs (March 2005)
240. Alan O. Sykes, Trade Remedy Laws (March 2005)
241. Randal C. Picker, Rewinding *Sony*: The Evolving Product, Phoning Home, and the Duty of Ongoing Design (March 2005)
242. Cass R. Sunstein, Irreversible and Catastrophic (April 2005)
243. James C. Spindler, IPO Liability and Entrepreneurial Response (May 2005)

244. Douglas Lichtman, Substitutes for the Doctrine of Equivalents: A Response to Meurer and Nard (May 2005)
245. Cass R. Sunstein, A New Progressivism (May 2005)
246. Douglas G. Baird, Property, Natural Monopoly, and the Uneasy Legacy of *INS v. AP* (May 2005)
247. Douglas G. Baird and Robert K. Rasmussen, Private Debt and the Missing Lever of Corporate Governance (May 2005)
248. Cass R. Sunstein, Administrative Law Goes to War (May 2005)
249. Cass R. Sunstein, Chevron Step Zero (May 2005)
250. Lior Jacob Strahilevitz, Exclusionary Amenities in Residential Communities (July 2005)
251. Joseph Bankman and David A. Weisbach, The Superiority of an Ideal Consumption Tax over an Ideal Income Tax (July 2005)
252. Cass R. Sunstein and Arden Rowell, On Discounting Regulatory Benefits: Risk, Money, and Intergenerational Equity (July 2005)
253. Cass R. Sunstein, Boundedly Rational Borrowing: A Consumer's Guide (July 2005)
254. Cass R. Sunstein, Ranking Law Schools: A Market Test? (July 2005)
255. David A. Weisbach, Paretian Intergenerational Discounting (August 2005)
256. Eric A. Posner, International Law: A Welfarist Approach (September 2005)
257. Adrian Vermeule, Absolute Voting Rules (August 2005)
258. Eric Posner and Adrian Vermeule, Emergencies and Democratic Failure (August 2005)
259. Douglas G. Baird and Donald S. Bernstein, Absolute Priority, Valuation Uncertainty, and the Reorganization Bargain (September 2005)
260. Adrian Vermeule, Reparations as Rough Justice (September 2005)
261. Arthur J. Jacobson and John P. McCormick, The Business of Business Is Democracy (September 2005)
262. Adrian Vermeule, Political Constraints on Supreme Court Reform (October 2005)
263. Cass R. Sunstein, The Availability Heuristic, Intuitive Cost-Benefit Analysis, and Climate Change (November 2005)
264. Lior Jacob Strahilevitz, Information Asymmetries and the Rights to Exclude (November 2005)
265. Cass R. Sunstein, Fast, Frugal, and (Sometimes) Wrong (November 2005)
266. Robert Cooter and Ariel Porat, Total Liability for Excessive Harm (November 2005)
267. Cass R. Sunstein, Justice Breyer's Democratic Pragmatism (November 2005)
268. Cass R. Sunstein, Beyond Marbury: The Executive's Power to Say What the Law Is (November 2005, revised January 2006)
269. Andrew V. Papachristos, Tracey L. Meares, and Jeffrey Fagan, Attention Felons: Evaluating Project Safe Neighborhoods in Chicago (November 2005)
270. Lucian A. Bebchuk and Richard A. Posner, One-Sided Contracts in Competitive Consumer Markets (December 2005)
271. Kenneth W. Dam, Institutions, History, and Economics Development (January 2006, revised October 2006)
272. Kenneth W. Dam, Land, Law and Economic Development (January 2006, revised October 2006)
273. Cass R. Sunstein, Burkean Minimalism (January 2006)
274. Cass R. Sunstein, Misfearing: A Reply (January 2006)
275. Kenneth W. Dam, China as a Test Case: Is the Rule of Law Essential for Economic Growth (January 2006, revised October 2006)
276. Cass R. Sunstein, Problems with Minimalism (January 2006, revised August 2006)
277. Bernard E. Harcourt, Should We Aggregate Mental Hospitalization and Prison Population Rates in Empirical Research on the Relationship between Incarceration and Crime, Unemployment, Poverty, and Other Social Indicators? On the Continuity of Spatial Exclusion and Confinement in Twentieth Century United States (January 2006)
278. Elizabeth Garrett and Adrian Vermeule, Transparency in the Budget Process (January 2006)
279. Eric A. Posner and Alan O. Sykes, An Economic Analysis of State and Individual Responsibility under International Law (February 2006)
280. Kenneth W. Dam, Equity Markets, The Corporation and Economic Development (February 2006, revised October 2006)
281. Kenneth W. Dam, Credit Markets, Creditors' Rights and Economic Development (February 2006)
282. Douglas G. Lichtman, Defusing DRM (February 2006)
283. Jeff Leslie and Cass R. Sunstein, Animal Rights without Controversy (March 2006)
284. Adrian Vermeule, The Delegation Lottery (March 2006)
285. Shahar J. Dilbary, Famous Trademarks and the Rational Basis for Protecting "Irrational Beliefs" (March 2006)
286. Adrian Vermeule, Self-Defeating Proposals: Ackerman on Emergency Powers (March 2006)
287. Kenneth W. Dam, The Judiciary and Economic Development (March 2006, revised October 2006)
288. Bernard E. Harcourt: Muslim Profiles Post 9/11: Is Racial Profiling an Effective Counterterrorist Measure and Does It Violate the Right to Be Free from Discrimination? (March 2006)
289. Christine Jolls and Cass R. Sunstein, The Law of Implicit Bias (April 2006)

290. Lior J. Strahilevitz, "How's My Driving?" for Everyone (and Everything?) (April 2006)
291. Randal C. Picker, Mistrust-Based Digital Rights Management (April 2006)
292. Douglas Lichtman, Patent Holdouts and the Standard-Setting Process (May 2006)
293. Jacob E. Gersen and Adrian Vermeule, *Chevron* as a Voting Rule (June 2006)
294. Thomas J. Miles and Cass R. Sunstein, Do Judges Make Regulatory Policy? An Empirical Investigation of *Chevron* (June 2006)
295. Cass R. Sunstein, On the Divergent American Reactions to Terrorism and Climate Change (June 2006)
296. Jacob E. Gersen, Temporary Legislation (June 2006)
297. David A. Weisbach, Implementing Income and Consumption Taxes: An Essay in Honor of David Bradford (June 2006)
298. David Schkade, Cass R. Sunstein, and Reid Hastie, What Happened on Deliberation Day? (June 2006)
299. David A. Weisbach, Tax Expenditures, Principle Agent Problems, and Redundancy (June 2006)
300. Adam B. Cox, The Temporal Dimension of Voting Rights (July 2006)
301. Adam B. Cox, Designing Redistricting Institutions (July 2006)
302. Cass R. Sunstein, Montreal vs. Kyoto: A Tale of Two Protocols (August 2006)
303. Kenneth W. Dam, Legal Institutions, Legal Origins, and Governance (August 2006)
304. Anup Malani and Eric A. Posner, The Case for For-Profit Charities (September 2006)
305. Douglas Lichtman, Irreparable Benefits (September 2006)
306. M. Todd Henderson, Paying CEOs in Bankruptcy: Executive Compensation when Agency Costs Are Low (September 2006)
307. Michael Abramowicz and M. Todd Henderson, Prediction Markets for Corporate Governance (September 2006)
308. Randal C. Picker, Who Should Regulate Entry into IPTV and Municipal Wireless? (September 2006)
309. Eric A. Posner and Adrian Vermeule, The Credible Executive (September 2006)
310. David Gilo and Ariel Porat, The Unconventional Uses of Transaction Costs (October 2006)
311. Randal C. Picker, Review of Hovenkamp, *The Antitrust Enterprise: Principle and Execution* (October 2006)
312. Dennis W. Carlton and Randal C. Picker, Antitrust and Regulation (October 2006)
313. Robert Cooter and Ariel Porat, Liability Externalities and Mandatory Choices: Should Doctors Pay Less? (November 2006)
314. Adam B. Cox and Eric A. Posner, The Second-Order Structure of Immigration Law (November 2006)
315. Lior J. Strahilevitz, Wealth without Markets? (November 2006)
316. Ariel Porat, Offsetting Risks (November 2006)
317. Bernard E. Harcourt and Jens Ludwig, Reefer Madness: Broken Windows Policing and Misdemeanor Marijuana Arrests in New York City, 1989–2000 (December 2006)
318. Bernard E. Harcourt, Embracing Chance: Post-Modern Meditations on Punishment (December 2006)
319. Cass R. Sunstein, Second-Order Perfectionism (December 2006)
320. William M. Landes and Richard A. Posner, The Economics of Presidential Pardons and Commutations (January 2007)
321. Cass R. Sunstein, Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas) (January 2007)
322. Cass R. Sunstein, Completely Theorized Agreements in Constitutional Law (January 2007)
323. Albert H. Choi and Eric A. Posner, A Critique of the Odious Debt Doctrine (January 2007)
324. Wayne Hsiung and Cass R. Sunstein, Climate Change and Animals (January 2007)
325. Cass R. Sunstein, Cost-Benefit Analysis without Analyzing Costs or Benefits: Reasonable Accommodation, Balancing and Stigmatic Harms (January 2007)
326. Cass R. Sunstein, Willingness to Pay versus Welfare (January 2007)
327. David A. Weisbach, The Irreducible Complexity of Firm-Level Income Taxes: Theory and Doctrine in the Corporate Tax (January 2007)
328. Randal C. Picker, Of Pirates and Puffy Shirts: A Comments on "The Piracy Paradox: Innovation and Intellectual Property in Fashion Design" (January 2007)
329. Eric A. Posner, Climate Change and International Human Rights Litigation: A Critical Appraisal (January 2007)
330. Randal C. Picker, Pulling a Rabbi Out of His Hat: The Bankruptcy Magic of Dick Posner (February 2007)
331. Bernard E. Harcourt, Judge Richard Posner on Civil Liberties: Pragmatic (Libertarian) Authoritarian (February 2007)
332. Cass R. Sunstein, If People Would Be Outraged by Their Rulings, Should Judges Care? (February 2007)
333. Eugene Kontorovich, What Standing Is Good For (March 2007)
334. Eugene Kontorovich, Inefficient Customs in International Law (March 2007)
335. Bernard E. Harcourt, From the Asylum to the Prison: Rethinking the Incarceration Revolution. Part II: State Level Analysis (March 2007)
336. Cass R. Sunstein, Due Process Traditionalism (March 2007)
337. Adam B. Cox and Thomas J. Miles, Judging the Voting Rights Act (March 2007)

338. M. Todd Henderson, Deconstructing *Duff & Phelps* (March 2007)
339. Douglas G. Baird and Robert K. Rasmussen, The Prime Directive (April 2007)
340. Cass R. Sunstein, Illusory Losses (May 2007)
341. Anup Malani, Valuing Laws as Local Amenities (June 2007)
342. David A. Weisbach, What Does Happiness Research Tell Us about Taxation? (June 2007)