Some Implications of Preference-Shifting for Optimal Tax Theory

Theodore P. Seto

This paper is part of a larger project: to explore the extent to which the claims of optimal tax theory are sensitive to the assumptions that underlie them. The paper focuses on two canonical assertions of the standard model: (1) that taxes produce deadweight loss (Harberger 1964), and (2) that 100 percent of all taxes are borne by human beings, the only question being which. The assumption it relaxes is the standard welfarist assumption that preferences are fixed and exogenous and reflect welfare. Although this assumption is not widely accepted in other social sciences, economists generally treat situations in which it does not hold (situations involving "internalities") as limited exceptions, and therefore of limited interest.

The explanatory power of the standard model would be enhanced if the possibility that preferences do not reflect welfare could be more generally incorporated. This paper outlines one possible approach to doing so in a context in which disparities between preferences and welfare are arguably widespread. Specifically, it postulates the existence of "preference-shifting": where businesses shift the preferences of potential consumers in ways that are not commensurately welfare-enhancing and then satisfy the resulting shifted preferences. If preference-shifting exists, the paper demonstrates (1) that a tax less than or equal to the amount of any such preference-shift will not result in any tax deadweight loss and will reduce preference-shifting deadweight loss; as a result, such a tax will increase aggregate welfare, and (2) that a portion of the welfare impact of the tax up to the amount of the preference-shifting dead-weight loss will be recovered through elimination of that loss and will not be borne by any human being.

Part I: Preference-Shifting

Economists have yet to offer a persuasive general model of the effects of advertising on welfare. Advertising may, of course, convey information. As a result, observed choices will change even if consumers' underlying preference sets do not. The additional information allows consumers to conform their behaviors more closely to their preexisting preference sets.

Not all effects of advertising, however, can be thus explained. Many of the most effective advertising campaigns contain remarkably little informational content. "Zoom. Zoom." "The ultimate driving machine." "The real thing." "Got milk?" "Super Size Me!" "You are now free to move about the country." "It's Miller time." "When you're having more than one." "Bullish on America." "Does she … or doesn't she?"

"They'rrrre GR-R-REAT!" "Snap! Crackle! Pop!" "Breakfast of Champions." "Bet you can't eat just one." "Good to the last drop." "The computer for the rest of us." "We bring good things to life." "Red Bull gives you wings." "Finger-lickin' good." "Have it your way." "I'm lovin' it." "Where's the beef?" "We do chicken right." "You're in good hands." "Like a good neighbor." "The Rock." "A diamond is forever." "I can't believe I ate the whole thing."

Such campaigns change observed choices. Businesses spend large sums to effect such changes. Advertisers, businesses, and consumers generally interpret such changes as reflecting changed preferences. This interpretation, however, poses a serious challenge to welfare economics. If consumers' pre-advertising preferences correctly reflect consumer welfare, post-advertising preferences arguably cannot.

Gary Becker has offered the best-known response (Becker 1996). Becker hypothesizes that non-informational advertising creates intellectual capital, which in turn generates imputed income. This imputed income reduces the shadow price of the good itself. For example, if a given car is worth $300 a month to a given consumer pre-advertising, the advertising is assumed to produce a stock of intellectual capital that will throw off $100 per month of imputed income in the consumer's mind if he buys the car. When he is observed to pay $400 per month for the car, the theory implies, this change in observed behavior does not reflect a change in preferences; the car is still only worth $300 a month. $100 of that $400 is being paid for the $100 of imputed income – not for the car. This reconstruction allows Becker to continue to assert that preferences do not change – and therefore allows him to continue to assume that those fixed and exogenous preferences reflect welfare.

Becker's theory is nonfalsifiable. Both the hypothetical stock of intellectual capital and the hypothetical flow of imputed income are inherently unobservable. The theory also violates Occam 's razor; indeed, it is Ptolemaic in its convolution. What the rest of the world calls increased desire Becker calls "imputed income" – different in kind, not merely in magnitude, from the base level of desire originally observed. Because this imputed income is not a "preference," Becker can assert that the consumer's "preferences" remain unchanged: fixed, exogenous, and welfare-reflecting as they have always been. It is unlikely that any non-economist finds this plausible.

Becker's theory also leads to the rather odd conclusion that a country whose residents have been conditioned to buy things that they otherwise would not possesses more capital – and is therefore is wealthier – than an otherwise identical country whose residents have not been so conditioned. Even economists, however, do not include Becker's hypothesized intellectual capital in their national wealth accounts.

Behavioral economics, by contrast, acknowledges the possibility of irrationality – that is, differences between preferences and welfare – but

limits its focus to distortions triggered by specific frames and heuristics. The consumer psychology literature that services the advertising and marketing industries takes a less theoretical, more practical approach, largely indifferent to the specific heuristics or irrationalities used to trigger the desired behaviors (Haugtvedt 2008). A large part of that literature describes and tests practical techniques for inducing consumers to buy more than they otherwise would, at higher prices than they otherwise would pay. It is possible that at least some such techniques have the effect of shifting the preferences of consumers in ways that are not welfare-enhancing.

Unfortunately, measuring the size of any resulting economy-wide wedge between preferences and welfare would first require that we define "welfare" – a task that eluded both philosophers and scientists for all of recorded history. This paper will not attempt any such task; nor will it attempt to estimate the size of any such wedge. Instead, it will ask and answer a simpler question: whether such a wedge would require modification of the two standard assertions identified at the outset: that taxes produce deadweight loss and that 100 percent of such taxes are borne by human beings. If the answer is affirmative, the task of measuring the size of any such wedge becomes more urgent.

This paper will, however, offer two examples of possible preference-shifting to support its premise that the problem is worth considering. The first involves changed demand for an intermediate good – consumer loans; the second, changed demand for a product as to which the objective welfare effects have been rigorously measured – a prescription drug.

Marianne Bertrand et al (Bertrand 2010) found that including a small picture of a woman in a consumer loan solicitation increased loan demand by about as much as a 25% reduction in the interest rate charged, compared with an identical solicitation without the picture. In other words, the picture permitted a 33% market equilibrium price increase. Since consumer loans are intermediate goods, it is hard to see how respondents who accepted the solicitation at the higher price could have received any commensurate welfare increase. Inclusion of the picture seems to have shifted respondents' utility and demand curves without any commensurate increase in the welfare they derived from the loans.

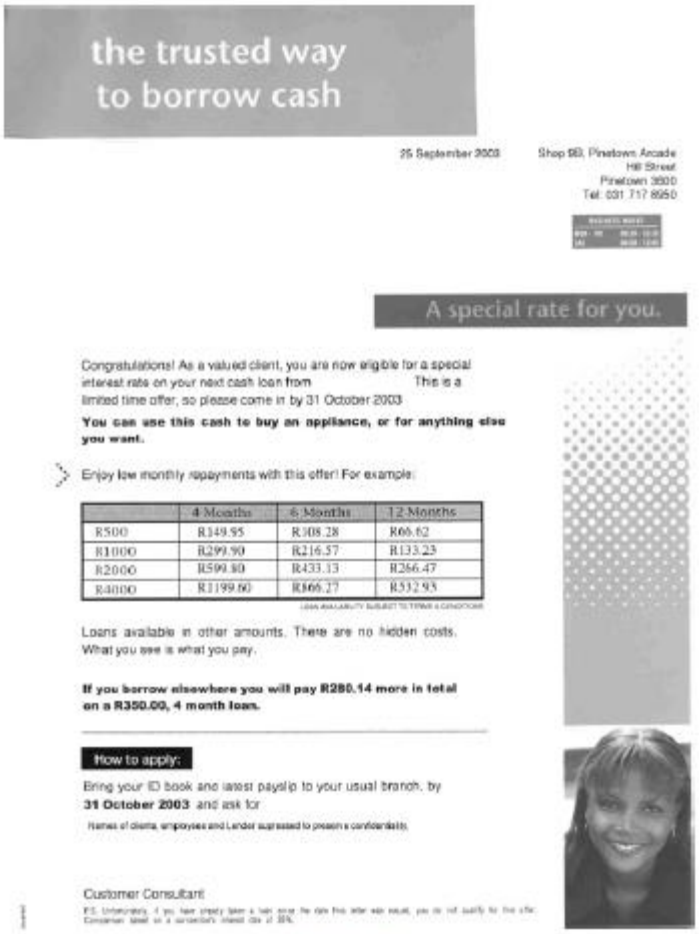
the trusted way
to borrow cash

25 September 2003

Shop 99, Pinetown Arcade
Hill Street
Pinetown 3600
Tel: 031 717 8950

A special rate for you.

Congratulations! As a valued client, you are now eligible for a special interest rate on your next cash loan from This is a limited time offer, so please come in by 31 October 2003

You can use this cash to buy an appliance, or for anything else you want.

Enjoy low monthly repayments with this offer! For example:

	4 Months	6 Months	12 Months
R500	R149.95	R108.28	R66.62
R1000	R299.90	R216.57	R133.23
R2000	R599.80	R433.13	R266.47
R4000	R1199.60	R866.27	R532.93

Loans available in other amounts. There are no hidden costs. What you see is what you pay.

If you borrow elsewhere you will pay R280.14 more in total on a R350.00, 4 month loan.

How to apply:

Bring your ID book and latest payslip to your usual branch, by **31 October 2003** and ask for

Customer Consultant

FIGURE II
Example Letter 2

A second example: omeprazole, a drug sold under the brand name "Prilosec," is used to decrease the production of gastric acid. AstraZeneca's patent on omeprazole was due to expire in 2001. AstraZeneca therefore developed and patented an approximate enantiomer (mirror-image version) of omeprazole, esomeprazole, which it marketed as "Nexium." Chemically, the two drugs operate the same way when ingested. Extensive studies have not shown any statistically significant differences in efficacy at equivalent doses (Drug Class Reviews 2009).

After patenting Nexium, AstraZeneca undertook a direct-to-consumer marketing campaign, built around the phrase "The Purple Pill™." By 2009 it had succeeding in shifting consumer preferences to the point that it could charge an average U.S. price of $190 per prescription for Nexium. Omeprazole, with medically equivalent effects at equivalent doses, commanded a price of only $31 per prescription. The price premium for

Nexium, therefore, exceeded 500 percent. By 2009, Nexium was the second-largest-selling prescription drug in the United States, with over $5 billion per year in sales.

Nexium's $159 per prescription premium was arguably evidence of preference-shifting: an upward shift in observed utility and demand curves without any corresponding increase in consumer welfare. The advertising theme used to effect this shift – "The Purple Pill™" – was largely devoid of informational content. With respect to Nexium, at least, Becker's accumulation-of-intellectual-capital explanation has little explanatory power. One might reasonably doubt that consumers received $159 per prescription of additional welfare by reason of their decisions to purchase Nexium rather than omeprazole.

Neither this nor any other study or anecdote establishes the existence of a general wedge between preferences and welfare. For present purposes, it is sufficient that economically significant preference-shifting is plausible. If preference-shifting can be shown to change standard results materially, the need for tools that will allow us to estimate the size of the wedge between preferences and welfare becomes more pressing. Does "Super Size Me!" shift demand upward without a commensurate increase in objective consumer welfare? What about "The Ultimate Driving Machine"? "Got Milk?" Such questions are empirical and largely beyond the capacity of currently available econometric tools. This paper asks instead whether the development of such tools might be worth the effort – specifically, whether differences between observed choices and welfare, if they do exist, require adjustments to important parts of the optimal tax canon.

Part II: Modeling Preference-Shifting

It would be useful to begin by making explicit the assumptions that underlie the standard model of consumer and producer surplus. The standard model assumes that the height of the demand curve at any quantity equals the marginal social benefit of consumption of that quantity and that the height of the supply curve equals the marginal social cost of production of that quantity, all costs and benefits having been internalized. Under these assumptions, total surplus provides a measure of the net increase in social welfare by reason of market exchanges at the equilibrium price and quantity. If so, then in the absence of market failure, supply and demand adjust automatically to maximize social welfare. Critically, this conclusion depends on the assumption that the choices embodied in observed demand reflect welfare.

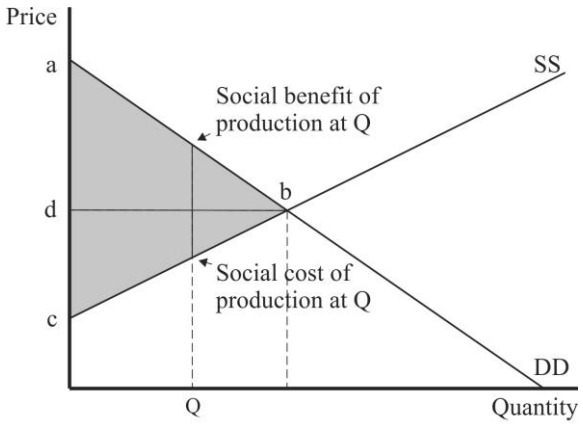

Figure 1

In Figure 1, production and exchange at quantity Q increase total welfare, since the height of the demand curve exceeds the height of the supply curve at that quantity. Triangle abc represents the aggregate increase in welfare resulting from exchanges of the good in an efficient market – in standard nomenclature, "total surplus." Triangle abd represents the welfare increase to consumers – "consumer surplus" – and triangle bcd the welfare increase to producers – "producer surplus."

The effects of preference-shifting can be modeled by postulating two demand curves: a non-preference-shifted, welfare-reflecting curve and a preference-shifted curve. In the absence of intervention, supply and demand will equilibrate at the intersection of supply and preference-shifted demand.

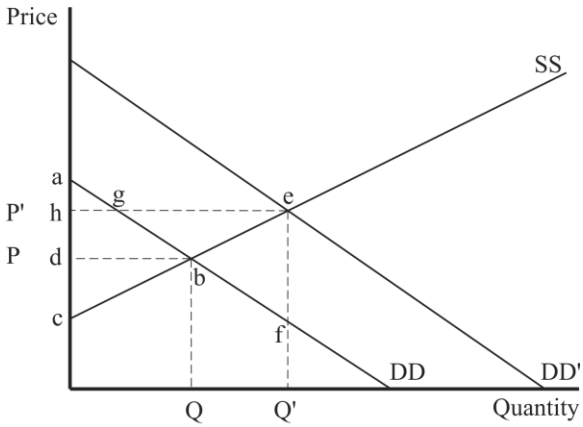

Figure 2

In Figure 2, DD is the non-preference-shifted, welfare-reflecting demand curve and DD' the preference-shifted curve. In a non-preference-shifted world, the market would equilibrate at price P and quantity Q, with

the same welfare results as in Figure 1. Because preferences have been shifted, however, the market equilibrates at price P' and quantity Q'.

The marginal social benefit of consumption of quantity Q' is given by the height of the non-preference-shifted demand curve DD at Q', the marginal social cost by the height of the supply curve at Q'. For all quantities between Q and Q', marginal social cost exceeds marginal social benefit. Total social loss by reason of preference-shifting is therefore represented by triangle bef in Figure 3 below, the "preference-shifting deadweight loss." (Costs incurred by producers to effect the preference shift also constitute social loss; this paper ignores such costs.)

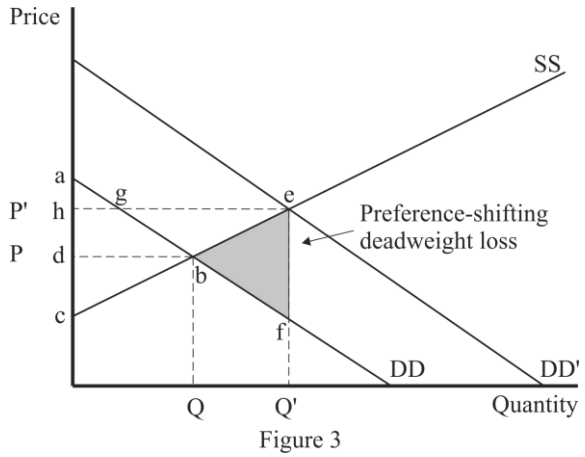

Figure 3

More generally:

Preference-shifting deadweight loss = $\int_{Q}^{Q'} d(q) - s(q)$,

Where:

Q = quantity exchanged at the welfare-maximizing equilibrium
Q' = quantity exchanged at the observed equilibrium
d(q) = marginal social benefit as a function of quantity consumed, and
s(q) = marginal social cost as a function of quantity produced.

If Q' > Q and $\frac{ds}{dq} > \frac{dd}{dq}$ between Q and Q', preference-shifting will produce net deadweight loss.

In figure 4 below, producer surplus is represented by quadrilateral cbgh. In addition, an amount represented by triangle beg is redistributed from consumers to producers. This is not "surplus," because it does not represent an increase in social welfare by reason of market exchanges; it is pure redistribution. Total gain to producers from exchanges in a

preference-shifted market, represented by triangle ceh, is the sum of producer surplus and this preference-shifting redistribution. Or:

Producer gain from market exchanges = $\int_0^{Q'} P' - s(q)$

Producer preference-shifting profits = Q * (P' – P) + $\int_Q^{Q'} P' - s(q)$

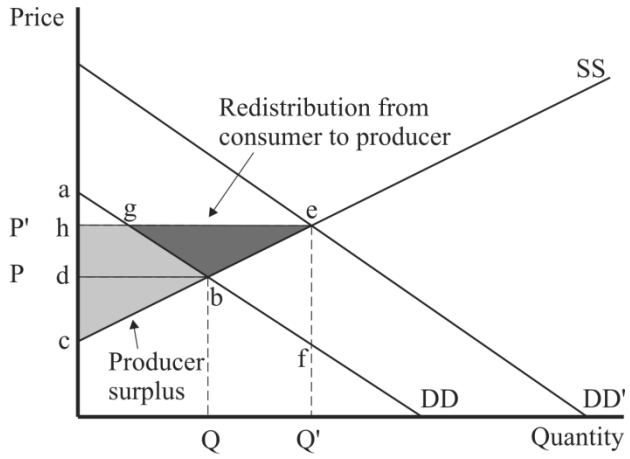

Figure 4

The computation of consumer gain or loss in a preference-shifted market begins with consumer surplus, represented by triangle agh. From this, however, must be subtracted both the preference-shifting redistribution from consumers to producers, represented by triangle beg, and the preference-shifting deadweight loss, represented by triangle bef. Or:

Consumer gain from market exchanges = $\int_0^{Q'} d(q) - P'$

Note that, as in Figure 5, the result may be negative. If so, market exchanges of the preference-shifted good reduce consumer welfare.

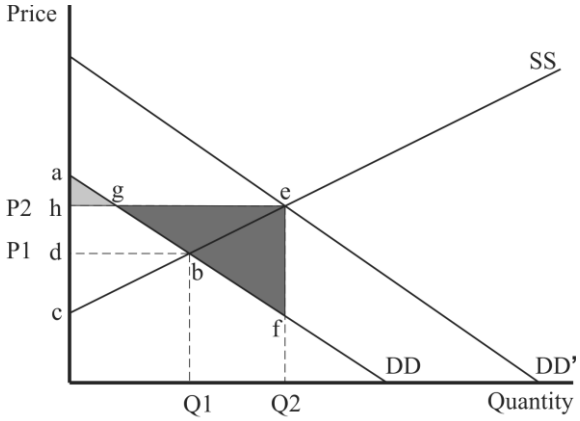

Figure 5

Finally, in Figure 6 net social benefit or loss by reason of market exchanges of the good is represented by total surplus, triangle abc, less preference-shifting deadweight loss, triangle bef. More generally:

Social gain from market exchanges = $\int_0^{Q'} d(q) - s(q)$

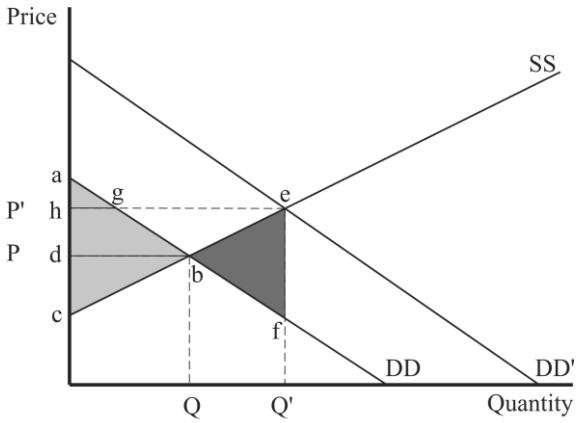

Figure 6

There is no inherent upper limit to preference-shifting deadweight loss. If preference-shifting deadweight loss is sufficiently large, as in Figure 7, the net social welfare consequences of exchanges in a preference-shifted market can be negative. (Recall that the 2009 price premium for Nexium was over 500 percent.)

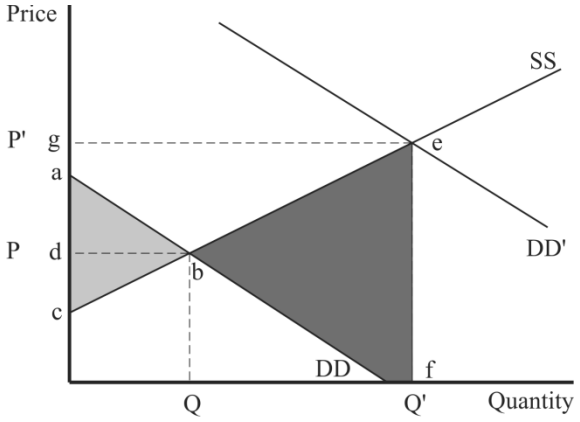

Figure 7

The effects of preference-shifting can therefore be summarized as follows.

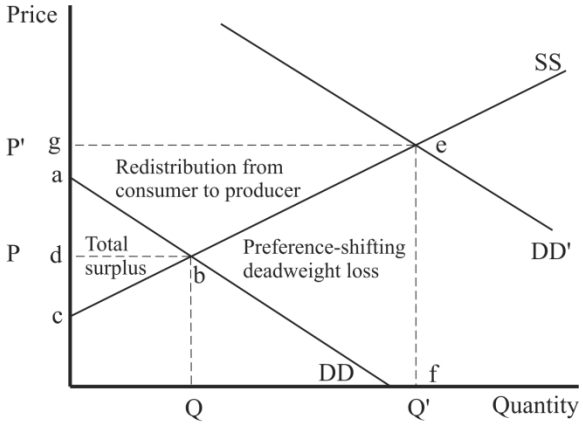

Figure 8

Total surplus is represented by the area between the supply curve and the welfare-reflecting demand curve to the left of quantity Q.

Total surplus = $\int_0^{\mathrm{Q}} d(q) - s(q)$

Preference-shifting deadweight loss is represented by the area between the supply curve and the welfare-reflecting demand curve between quantities Q and Q'. No additional surplus is generated by exchanges of quantities above Q; such exchanges instead produce social loss.

Preference-shifting deadweight loss = $\int_{\mathrm{Q}}^{\mathrm{Q}'} d(q) - s(q)$

Finally, preference-shifting redistribution from consumers to producers is represented by the area above the supply curve and the welfare-reflecting demand curve bounded by price P'.

$$\text{Preference-shifting redistribution} = \int_0^{\mathrm{Q}} P' - d(q) + \int_{\mathrm{Q}}^{\mathrm{Q}'} P' - s(q)$$

The foregoing results obtain whenever preferences do not correctly reflect objective welfare, regardless of whether intentional preference-shifting has occurred, so long as observed demand exceeds welfare-reflecting demand. These results obtain even if the activities that shift consumer preferences also enhance welfare (for example, by conveying useful information or shifting shadow prices) and thereby shift the welfare-reflecting demand curve upward, so long as observed demand is increased more than welfare is increased (that is, so long as the objective welfare increase is not commensurate).

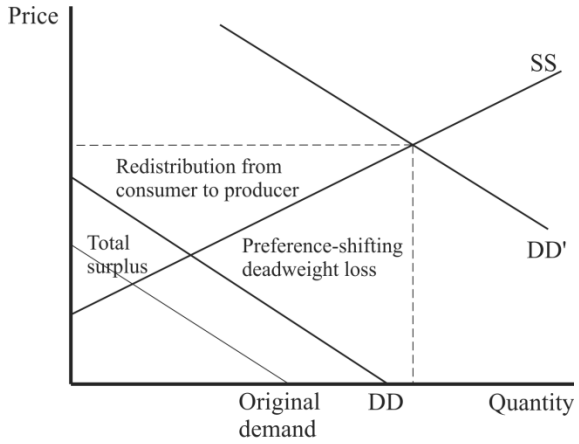

Figure 9

In Figure 9, producers' preference-shifting activities shift demand in ways that increase consumer welfare, but not commensurately with the increase in observed demand. The new welfare-reflecting demand curve is DD, the preference-shifted demand curve DD'. Total surplus, preference-shifting deadweight loss, and redistribution from consumers to producers are then the same as in Figure 8.

Part III: Effects of Tax on Deadweight Loss

In the standard account, a commodity tax can be modeled either as a downward shift in demand or an upward shift in supply, depending on the party on which it is nominally imposed, in each case by the amount of the tax. The incidence of the tax will depend on the relative elasticities of supply and demand, not on the government's choice of nominal taxpayer.

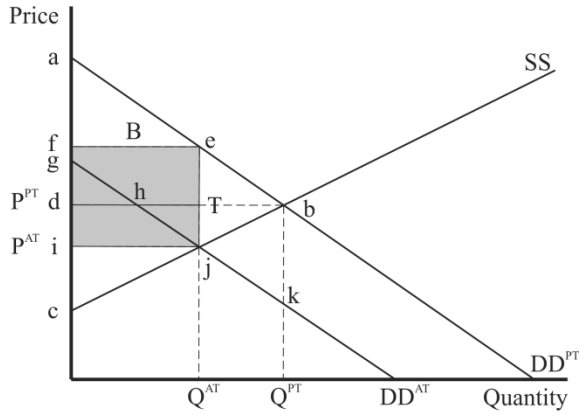

Figure 10

In Figure 10, buyer must pay tax at rate T in addition to the market price. As a result, the demand curve faced by producers is shifted downward by amount T from pre-tax demand DD^{PT} to after-tax demand DD^{AT}. The market equilibrates at after-tax price P^{AT} and quantity Q^{AT} rather than at pre-tax price P^{PT} and quantity Q^{PT}.

Prior to imposition of the tax, total surplus is represented by triangle abc, consumer surplus by triangle abd, and producer surplus by triangle bcd.

After imposition of the tax, consumer surplus is represented by triangle aef and producer surplus by triangle cij. Rectangle efij represents tax revenue collected by the government, tax rate T times base B. So long as the revenues thus collected are expended on items producing equivalent welfare, this appropriation is welfare-neutral. Triangle bej, however, represents the additional amount of social welfare that would have been generated through market exchanges at price P^{PT} and quantity Q^{PT} in the absence of the tax. This foregone social welfare is the "tax deadweight loss."

Tax deadweight loss = $\int_{QAT}^{QPT} s(q) - d(q)$

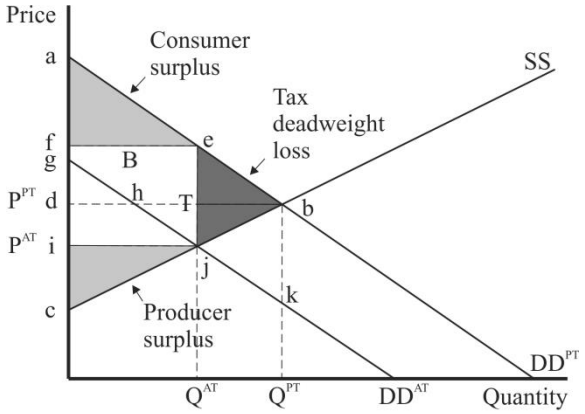

Figure 11

We can now model imposition of a tax on a good for which demand is preference-shifted. Prior to imposition of the tax,

Total surplus = $\int_0^{Q} d(q) - s(q)$, and
Preference-shifting deadweight loss = $\int_{Q}^{Q'} d(q) - s(q)$,

Where:

Q = quantity exchanged at the welfare-maximizing equilibrium
Q' = quantity exchanged at the observed pre-tax equilibrium
d(q) = marginal social benefit as a function of quantity consumed, and
s(q) = marginal social cost as a function of quantity produced.

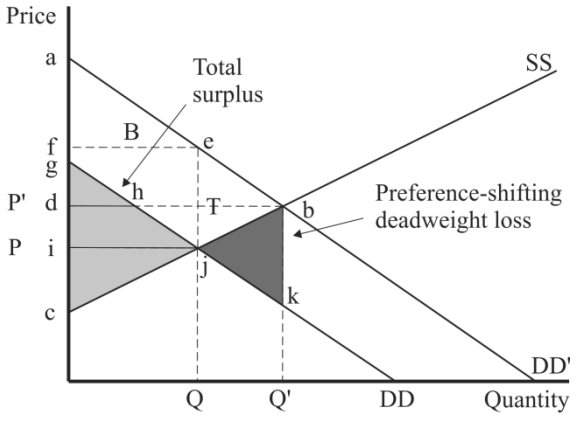

Figure 12

Consider tax at rate T in an amount equal to the preference-shift. The tax forces the market to reequilibrate at price P and quantity Q – the same equilibrium price and quantity the market would have reached in the absence of preference-shifting. Buyers pay price P plus tax at rate T, and at that effective price only demand quantity Q.

Two consequences ensue. First, because no exchanges occur in excess of quantity Q, preference-shifting deadweight loss is eliminated. Second, market exchanges at price P' and quantity Q' would not have generated any additional social welfare; the forced reequilibration at price P and quantity Q therefore does not result in any tax deadweight loss. As a result, imposition of tax at a rate less than or equal to the preference-shift increases aggregate social welfare.

Part IV: Incidence of the Tax

If the tax rate equals the preference-shift, the resulting increase in aggregate social welfare will equal the preference-shifting deadweight loss that would have been incurred in the absence of the tax.

Total welfare gain from tax = $-\int_{\mathrm{Q}}^{\mathrm{Q}'} d(q) - s(q)$

This welfare gain is allocated among consumers, producers, and the government as follows.

The government receives the amount of the tax.

Government receipts = Q * T

Producer welfare is reduced by the amount of producers' pre-tax preference-shifting profits, represented by quadrilateral bdij in Figure 13 below. This consists of (1) the portion of consumer surplus appropriated by producers by reason of the preference-shift, represented by quadrilateral dhij, plus (2) the preference-shifting redistribution from consumers to producers, represented by triangle bhj. As a result, producer welfare gains by reason of market exchanges at the after-tax equilibrium exactly equal the producer surplus that would have been generated by market exchanges in the absence of both tax and preference-shifting.

Producer incidence = – Q * (P' – P) – $\int_{\mathrm{Q}}^{\mathrm{Q}'} P' - s(q)$

If preference-shifting is a form of market failure, from producers' perspective the tax exactly corrects the market failure.

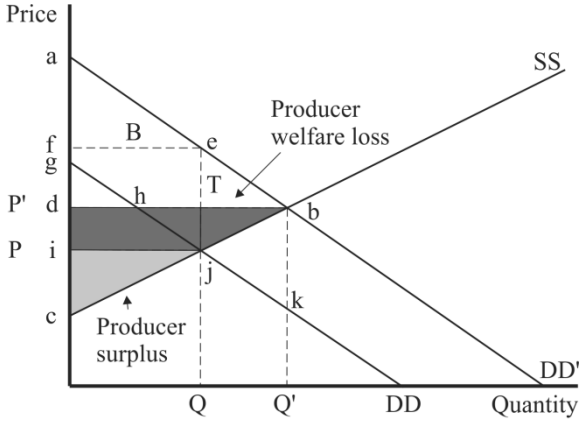

Figure 13

Net consumer welfare gain or loss by reason of the tax, depicted in Figure 14 below, equals (1) the preference-shifting deadweight loss plus the preference-shifting redistribution from consumers to producers in excess of quantity Q, reduced by (2) the portion of the tax in excess of the change in price (P' minus P).

Consumer incidence = $[\int_{Q}^{Q'} P' - d(q)]$ – [Q * (P + T - P')]

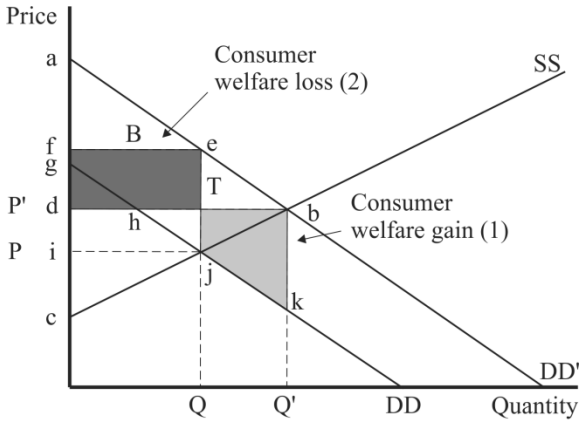

Figure 14

If (1) is greater than (2),

$\int_{Q}^{Q'} P' - d(q)$ > Q * (P + T - P'),

the tax will result in net welfare gain to consumers in the amount of the difference. In such event, producers will bear the entire burden of the tax,

but only in the amount of their preference-shifting profits. If (2) is greater than (1),

$$\int_{Q}^{Q'} P' - d(q) < Q * (P + T - P'),$$

consumers will bear a portion of the tax equal to such excess. A portion of the tax equal to the preference-shifting deadweight loss will be recovered through elimination of that loss and will not be borne by any human being.

Part V: Conclusions and Further Work

This paper has shown that, in the case of a commodity tax, imposition of tax at a rate less than or equal to the preference-shift increases aggregate social welfare. Imposition of tax in an amount equal to the preference-shift will return producer surplus to the amount that would have accrued to producers in the absence of preference-shifting at the welfare-maximizing equilibrium. Producers will bear the burden of such a tax to the extent of their pre-tax preference-shifting profits. Consumers will bear the burden of such a tax to the extent such burden is not borne by producers, but will also benefit from the elimination of preference-shifting deadweight loss. A portion of such tax equal to the preference-shifting deadweight loss will be recovered through elimination of that loss and will not be borne by any human being.

In the standard account, once it has been shown that commodity taxation produces deadweight loss, the result is generally extended to income and other forms of taxation without further demonstration. If preference-shifting were to operate randomly with respect to income, it would be necessary to customize commodity taxes to reflect such random preference-shifting. It is plausible, however, that preference-shifting is more likely with respect to non-essentials and therefore increases with income. If so, progressive income taxation may reduce income-correlated preference-shifting deadweight loss and thereby increase aggregate welfare.

Standard accounts of the welfare effect of labor tax rates (*e.g.*, Mirrlees 1971) assume that preferences reflect welfare. Such accounts typically begin at the individual utility curve level. Demand curves, of course, are simply aggregations of individual utility curves. The model offered in this paper assumes two sets of utility curves – a set of non-preference-shifted, welfare-reflecting utility curves and a set of preference-shifted utility curves. If a wedge exists between these two sets, Mirrlees' analysis and its progeny may require modification.

A more immediate and perhaps less theoretically contestable implementation of the commodity tax results might be to limit retail sales and value added tax exemptions to non-advertised goods. This would

merely require that we assume that, in general, preference-shifting by reason of advertising equals or exceeds the retails sale or value added tax rate. Since such rates are normally small relative to price, such an assumption is plausible. The Atkinson-Steiglitz theorem, which suggests that governments abstain from differentiated commodity taxation if non-linear income taxation is an option, does not apply in this context because it too is based on the assumption that preferences are fixed and exogenous and reflect welfare (Atkinson-Steiglitz 1976).

In any event, the fact that preference-shifting implies results very different from those of the standard model makes more urgent the need for tools to measure the wedge, if any, between observed preferences and welfare. This may, in turn, require further definition of "welfare" – a task that economics has heretofore largely avoided. The standard account asserts that whatever "welfare" means, if observed preferences reflect welfare, taxes result in tax dead-weight loss. This paper has shown that if observed preferences do not reflect welfare, taxes do not necessarily result in dead-weight loss – indeed, that taxes may enhance "welfare," again regardless of what that term means. To move from theory to measurement, however, may require further definition.

One possibility, suggested by the happiness literature, would be to use steady-state happiness as a measure of objective welfare without actually defining the term itself – in other words, to assume that the body knows what it needs *ex post*, even if the mind does not *ex ante*. Under such an approach, what is sometimes taken to reflect declining marginal utility in the happiness literature might instead be interpreted as reflecting the possibility that preference-shifting increases once income is adequate to purchase necessities.

"Capitalism," said Andrew Carnegie, "is about turning luxuries into necessities." Preference-shifting may be necessary to capitalism. Nothing in this paper is intended to suggest that either is undesirable. Its thesis is purely technical: that if economically significant preference-shifting exists, the optimal tax canon requires adjustment. If important consequences of the standard account depend on an assumption – that preferences reflect welfare – that may be false in economically significant ways, it would be inappropriate to continue to assert the results of the standard model without appropriate caveat.

References

Atkinson, A. B.; Stieglitz, J. E., The Design of Tax Structure: Direct versus Indirect Taxation, Journal of Public Economics 6 (1-2): 55–75 (1976).

Becker, Gary, Accounting for Tastes (Harvard Univ. Press 1996).

Bertrand, Marianne, et al, What's Advertising Content Worth? Evidence From a Consumer Credit Marketing Field Experiment, Quarterly Journal of Economics (2010) 125 (1): 263-306.

Bowles, Samuel, Endogenous Preferences: The Cultural Consequences of Markets and other Economic Institutions, 36 J. Econ. Lit. 75-111 (1998).

Cohen, Michael D., and Axelrod, Robert, Coping with Complexity: The Adaptive Value of Changing Utility, 74 Am. Econ. Rev. 30-43 (1984).

Day, Richard H., On Endogenous Preferences and Adaptive Economizing, in Day, Richard H., and Eliasson, Gunnar (eds.), The Dynamics of Market Economies 153-170 (North-Holland, Amsterdam 1986).

Dixit, Avinash K., and Norman, V., Advertising and Welfare, 9(x) J. Econ. 1-17 (1978).

Drug Class Reviews, Proton Pump Inhibitors, Final Report (Update 5, May 2009).

Gintis, Herbert, Welfare Criteria with Endogenous Preferences: The Economics of Education, 15 Int'l Econ. Rev. 415-430 (1974).

Harberger, Arnold C., The Measurement of Waste, 54 Am. Econ. Rev. 58 (1964).

Haugtvedt, Curtis P., et al, eds., Handbook of Consumer Psychology (Psychology Press 2008)

Lee, Frederic S., and Keen, Steve, The Incoherent Emperor: A Heterodox Critique of Neoclassical Microeconomic Theory, 62 Rev. Soc. Econ. 169-199 (2004).

Mirrlees, J. A., An Exploration in the Theory of Optimum Income Taxation, 38 Rev. Econ. Stud. 175-208 (1971).

Pollak, Robert A., Endogenous Tastes in Demand and Welfare Analysis, 68 Am. Econ. Rev. 374-379 (1978).

Sen, Amartya, Internal Consistency of Choice, 61(3) Econometrica 495-521 (1993).