

COST-BENEFIT ANALYSIS OF FUEL-ECONOMY IMPROVEMENTS

Mark A. Delucchi

Institute of Transportation Studies

University of California, Davis

Discussion paper

August 2007

Introduction

Fuel economy improvements, such as provided by hybrid electric vehicles, increase the cost of making a motor vehicle, but reduce expenditures on fuel over the life of the vehicle. Generally, if the present value of the benefit of reduced fuel expenditures exceeds the extra cost of making the fuel economy improvements, then the improvements are economically worthwhile. This note discusses the proper way to perform a cost-benefit analysis of fuel economy improvements, using hybrid vehicles as a nominal example. It delineates the difference between the doing the analysis from the perspective of an individual consumer and doing the analysis from the perspective of society. It also shows that the high implicit discount rate that consumers appear to apply to fuel-economy-purchase decisions is best understood *not* as an explicit expectation of a very high rate of return on the investment foregone by spending money on fuel economy, but rather as the implicit equivalent of a series of “conservative” assumptions about fuel prices, fuel economy improvement, resale value, and so on, combined with an expectation of a normal rate of return on foregone investments.

Overview of the method

The objective is to compare the present value of the costs of fuel-economy improvements with the present value of the benefits of the fuel-economy improvements. The present cost of fuel-economy improvements is equal to:

- 1) the initial extra retail cost of the fuel economy measures, less
- 2) the present salvage value of the measures, plus
- 3) the present value of any periodic costs of maintaining the *extra* fuel economy (i.e., not just maintaining the vehicle generally, but maintaining the extra fuel economy per se) over the life of the vehicle.

The benefit of fuel-economy improvement is more complicated. Generally, the private benefit is equal to the present value of reduced expenditure on fuel, which can be calculated on the basis of:

- 1) the miles driven,
- 2) the consumer price of fuel,
- 3) the fuel economy improvement,
- 4) vehicle resale and the preferences of subsequent owners, and
- 5) the discount rate.

This benefit potentially is realized over the life of the vehicle. Because the factors in the benefit calculation will change over time, the benefits ideally should be calculated at regular intervals (say, annually) over the life of the vehicle. The benefits calculated at each interval (year) will have to be discounted at a pertinent consumer discount rate to obtain a present benefit value properly compared with the extra initial (present) cost.

The estimation of the annual miles driven, the price of fuel, the fuel economy improvement, the life of the vehicle, and the discount rate over the life of the vehicle depend on whether one is taking the perspective of the initial purchaser of the vehicle or the perspective of society (as represented by the calculations of a “neutral” analyst).

For example, the initial purchaser is likely to be “conservative” in the sense of being more averse to mistakenly overestimating the benefits and incurring unanticipated losses than to mistakenly underestimating benefits and foregoing unanticipated gains. (This is “loss aversion,” in which people want more to avoid a loss of \$X than to forego a gain of \$X. Put another way, people probably want to make sure that fuel-economy improvements are worthwhile under worst-likely-case scenarios.) If so, then consumers will tend to estimate the future price of fuel and the fuel economy gain more conservatively than will supposedly “neutral” analysts. Such factors will prove to be significant.

Details of the calculation: costs

1. The extra initial cost. The extra initial cost is the difference between the retail cost of the improved-fuel-economy vehicle (a hybrid electric vehicle, in our case) and the retail cost of the non-improved (non-hybrid) counterpart.¹ Note that from both the perspective of society and the perspective of the individual consumer, the relevant cost is the *retail* cost, not the manufacturing cost. However, in the consumer cost-benefit analysis all taxes should be included, whereas in the social cost-benefit analysis at least some taxes (such as the general sales tax) should be excluded because they don’t represent real resource costs to society associated with making motor vehicles.

2. Salvage value. In general, the first purchaser can expect to get a cost refund equal to the present value of the future salvage value of the fuel economy improvements at the physical end of life of the vehicle. However, motor vehicles have almost no positive salvage value at scrappage, and even if they did, it would occur so far in the future that it would be discounted to near zero anyway. For example, even given a \$5,000 initial

¹ The difference in retail cost should be based not just on costs related to engines, transmissions, and fuel-economy-improvement measures, but on the cost of the entire vehicle, because changes in fuel economy can have cost effects that are not immediately obvious. For example, there is some evidence that vehicles with greater fuel economy have less costly emission-control systems, because higher fuel economy is associated with lower g/mi engine-out emissions and hence a reduced need for tailpipe control to achieve a given g/mi tailpipe emissions standard (DeLuchi et al., 1992). These sorts of costs differences should be included in the analysis of retail cost differences.

cost differential (fairly high), a 1% salvage value (again high), a 4% discount rate (low), and a 12-year life (low), the present salvage value is only \$30. With more likely parameters, the present salvage value is less than \$5. Nonetheless, it is straightforward to represent this formally, and we do so here. To maintain generality, we multiply four factors, for each year: the extra initial cost of fuel economy improvements, the salvage value percentage (Delucchi [2000a] uses 0.3%), the present value factor (based on the discount rate, discussed below), and the overall probability that the vehicle and the individual survive to the end of life of the vehicle (discussed below in the section on benefits).

3. Maintenance costs. This third factor comprises any (net) extra operating and maintenance costs for the hybrid, and the administrative cost portion of any additional insurance premium specifically for the extra value of the hybrid components. The extra operating and maintenance cost probably is small, on the order of \$5. The insurance administrative-cost premium can be shown to be small, too: assuming that the administrative cost is 20% of the premium, that the hybrid vehicle premium increases by 10%, and that the base premium is \$50/month, the result is a cost of only \$1/month. (The portion of the extra premium other than that for administrative costs covers reimbursements paid by the insurance companies to the insured, and so in effect is recuperated by the vehicle owners, on average.)

Details of the calculation: benefits.

As mentioned above, the calculation of benefits depends on miles driven, the price of fuel, the fuel-economy improvement, vehicle resale and the valuations of subsequent owners, and the discount rate. Because all of the parameters in the benefit calculation vary over time, the calculation should be done for each year over the entire life of the vehicle.

1. Miles driven. The benefits of improved fuel economy must be calculated on the basis of the actual miles driven in each year. In this regards, there are a number of analytical subtleties.

First, we are interested in the vehicle miles of travel (VMT) that would have been driven in the lower fuel economy (in our case, non-hybridized) counterpart vehicle, which is not necessarily the same as the VMT in the hybrid vehicle. The reason for the potential difference is that the lower fuel-cost-per mile of the hybrid *may* induce additional VMT (see the discussion below). However, reduced fuel expenditures must be calculated relative to what would have been spent driving the non-hybrid counterpart, and *that* in turn is based on the VMT that would have been driven in the non-hybridized counterpart. (The potential benefit of the any extra VMT due to the lower fuel cost-per-mile in the hybrid is discussed in the section “other considerations”.)

Second, the annual VMT schedule is not easy to estimate from the data typically available. What we wish to know is how the annual VMT of a current model-year vehicle will change over time, whereas what the available survey data typically tell us is the miles driven this year by vehicles of different ages. In order to estimate the annual VMT of a given model year, one needs to have several years of survey data and other information. Table 5 of Delucchi (2000a) presents the results of such an analysis, which we use here.

Third, we must account for several risks that there will be no VMT, and hence no benefit, at all. Every year, there is some probability that accidents, vehicle scrappage due to old age, death of the vehicle owner, or theft of the vehicle will cause zero VMT for the vehicle owner. If there is zero VMT then (with some qualifications discussed momentarily) there is zero benefit; hence, we need to know what I will call the “realized” VMT schedule. The realized VMT in any year is equal to the VMT given no accidents, vehicle scrappage due to old age, death, or theft multiplied by the probability in that year of no accidents, no vehicle scrappage due to old age, no death of the vehicle owner, and no vehicle theft. Yearly survival probabilities reflecting loss due to accidents and old age are available (Davis and Diegel, 2007). This survival probability is extremely high for the first few years of a vehicle's life. Theft and death probabilities must be estimated separately; here, I assume that in the individual's survival probability

is 99.6% in the first year, and declines by 0.4% per year (relative terms) thereafter, and that the non-theft probability is 98.5% in the first year and increases by 0.15% per year thereafter (assuming that thieves are more likely to steal newer vehicles). (The assumptions regarding vehicle theft result in a fleet-average theft probability of 0.5% per year, which is the same as the total number of motor-vehicle thefts in the U. S. in a year [www.fbi.gov/ucr/05cius/offenses/property_crime/motor_vehicle_theft.html] divided by the number of registered vehicles [www.fhwa.dot.gov/policy/ohim/hs05/htm/mv1.htm].)

There are however complications, some of which bear on the difference between a social perspective and an individual consumer's perspective. In some cases of theft or accident, an insurance company will reimburse the complete value of the extra fuel economy improvement measures. This matters to the individual consumer. However, the consumer also will pay a higher premium in order to ensure against the loss of the extra value of vehicle. Over time over a large population, the higher insurance premiums will be at least as great as insurance-company payouts for reimbursement. Hence, if we expand the analysis to include the higher insurance premiums, as we should, then any insurance company reimbursement for loss simply cancels the higher insurance premiums, and the consumer still is left with the loss of the stream of benefits that the higher fuel economy was supposed to provide. And the retirement of the vehicle due to old age, and the death of the vehicle owner himself, definitely end the stream of benefits from fuel economy improvements. Hence, from the standpoint of the first purchaser, we may assume that any VMT not realized due to accidents, theft, death, or old-age scrappage terminates the benefit stream and should *not* be counted as realized VMT in the benefit calculation.

However, the matter may be somewhat different from the perspective of society. Theft, for example, may be viewed as an involuntary and therefore economically non-optimal transfer of the fuel economy benefit from one person to another. Whereas in the case of theft the cost to the victim is the entire fuel economy benefit, the cost to society is that resulting from the non-optimal pattern of vehicle production and use.

This pattern is non-optimal because the thief's willingness-to-pay WTP for the stolen vehicle presumably is much less than either the cost of the vehicle or the WTP of the original purchaser. Thus, on practical and perhaps ethical grounds, society might decide that theft at least reduces the benefit stream. However, this is difficult to quantify, and so for simplicity I will assume that from the perspective of society, theft does *not* affect the benefit stream.

Society's perspective also may differ from the individual's in the matter of death. Whereas the first purchaser may consider whether he is going to be around long enough to enjoy the benefits of his investment in fuel economy, society may not care who gets the benefits and may implicitly assume that the benefits simply are costlessly transferred to someone else in the household. (Of course, an individual consumer may consider this, too, and so give less weight to his own enjoyment of the benefits. However, I will assume that this is *not* the case.)

I summarize this as follows. From the perspective of the consumer or vehicle owner, the “raw” or unadjusted annual VMT should be multiplied by the vehicle survival probability (which accounts for accidents and scrappage due to old age), the individual's survival probability, and the non-theft probability, for each year. From the perspective of society, the raw annual VMT should be multiplied by the vehicle survival probability only.

2. The retail price of fuel. Because the future price of fuel is uncertain, and the conservative assumptions of car buyers probably will be different from those of “neutral” analysts, one will have to estimate three sets of values here:

- i) the best "neutral" analytical projections, for the social-cost analysis;
- ii) the first-purchaser's perspective, and
- iii) the first purchaser's assessment of the subsequent purchaser's perspective.

For the social cost perspective, one can use the projections of the Energy Information Administration (EIA), available in their *Annual Energy Outlook* (www.eia.doe.gov/oiaf/aeo/index.html) (we assume that prices start at \$2.70/gallon, and increase at 0.6% per year). The second and third probably are best represented as fractions of the “neutral” analytical projection. Thus, one might multiply the EIA projections by, say, 0.85 to represent the conservatism of the first purchasers, and then again by 0.90 to represent the first purchaser's further conservative assessment of subsequent buyers. Thus, if the neutral analyst estimates a price of \$2.70/gallon, the conservative first purchaser assumes a price of \$2.30/gallon for himself and a price of \$2.07/gallon for anyone he has to sell the car to.

If the analysis is to be done in nominal dollars, which I suggest (because it is more psychologically natural to use a nominal interest rate, and the basis of the interest rate must be the same as the basis of the fuel price), then the projections of fuel price should be in current-year dollars.

In the consumer cost-benefit analysis, the price of fuel should be the complete retail price at the pump, including all taxes. In the case of the social cost-benefit analysis, one can argue that retail excise taxes, sales taxes, and producer surplus (PS) ought to be deducted from the full retail price because they are not real resource costs to society associated with making the fuel. (Taxes are a transfer from consumers to the government, and PS is a transfer from consumer to producers.)

Excise taxes and sales taxes on motor fuel are easy to estimate, but PS is not. Delucchi (2004) analyzes original modeling done by others, and estimates that PS in the crude oil industry is about 40% of pre-tax price-times-quantity payments. We start with that figure here. Assuming that roughly half of crude oil payments are to foreign producers, whose welfare doesn't count in a U. S. social-cost analysis (and hence whose PS gain is a real net resource loss to U. S. consumers), then about 20% of the pre-tax price-times-quantity payment for gasoline is a transfer from U. S. consumers to U. S. crude oil producers. Allowing for some PS in the refining industry, we assume that a total of 25%

of the pre-tax price-times-quantity payment for gasoline is a transfer from U. S. consumers to U. S. producers. This amount is deducted in the social welfare analysis.²

3. The fuel economy improvement. The discussion here follows that for the retail price of fuel. First, we assume that the fuel economy of the baseline or non-hybrid vehicle is 20 mpg, and that the neutral analyst estimates that a hybrid gets 40% better fuel economy. Then, we assume that the individual's estimates of fuel economy improvements are more conservative than those of the objective analyst, and represent this with adjustment factors analogous to those applied to the fuel price, above. (That is, as mentioned above, the initial purchaser is likely to be "conservative" in the sense of being more averse to mistakenly overestimating the benefits and incurring unanticipated losses than to mistakenly underestimating benefits and foregoing unanticipated gains.) In the absence of studies of the value of these conservatism adjustment factors, I assume that the first purchaser estimates the fuel economy improvement (i.e., the percentage increase in fuel economy over the baseline) to be 66% of the neutrally analytically estimated percentage increase, and that the first purchaser assumes that subsequent purchasers would estimate the fuel economy improvement to be 85% of the first purchaser's estimate. Thus, if a neutral analyst estimates that a hybrid vehicle gets 40% better fuel economy – in our case, 28 mpg vs. 20 mpg for the non-hybrid version – then the first purchaser assumes an improvement of only 26.4% (25.3 mpg) for himself and 22.4% (24.5 mpg) for anyone he has to sell the car to.

4. Vehicle resale and subsequent owners. From the perspective of society, the fuel savings benefit should be calculated over the life of the vehicle from initial purchase to scrappage. However, some analysts argue that from the perspective of the initial

² This is the amount to be deducted in a social-cost analysis done from a U. S. perspective. In this U. S. perspective, we also count the so-called pecuniary externality of oil use (which is higher payments to foreign oil producers for non-transportation uses of oil), but we estimate GHG-emissions damages for the U. S. only, rather than for the whole world. By contrast, if we take a global perspective in our social-cost analysis, then all producer surplus – even that accruing to foreign producers – is a transfer, and we use global rather than U. S. GHG-emissions damages, but we ignore the pecuniary externality because it is a transfer within the global system.

purchaser, the fuel savings benefit should be calculated for only those years that the initial purchaser owns the vehicle. These analysts assume that the initial purchaser does not expect to be paid anything extra for the vehicle's fuel-economy improvements at the time of resale. I believe that it is more accurate psychologically to say that initial purchasers have some degree of *skepticism* regarding the value of the fuel economy improvements to second buyers. Fortunately, it is straightforward to represent this skepticism formally with full generality.

In principle, the maximum amount that the second purchaser of the vehicle will pay for the fuel-economy improvement features is the present value (at the time of resale, or second purchase) of the stream of benefits that the fuel economy improvement will provide to the second purchaser. This stream of benefits is calculated in the same way as is the stream of benefits for the initial purchaser of the vehicle. Hence, the rational (and even skeptical) initial purchaser will estimate the benefit stream all the way to the actual end of life of the vehicle, because those benefits will be capitalized into the resale value that the initial purchaser receives.

However, at this point we depart from conventional analyses to introduce psychological realism (skepticism), but in a way that still maintains complete generality. The key point is that we must estimate the post-resale benefit stream *not* using neutral analytically derived “social” parameter values, and *not* even using the parameter values that we assume the first purchaser uses himself during his initial ownership, but rather using parameter values that the first purchaser is likely to ascribe to *subsequent* purchasers. This is because the first purchaser's initial decision is based in part on what he expects to get at resale, which in turn he (the initial purchaser) believes is what the subsequent purchasers are willing to pay for the fuel economy benefits as *they* – the second purchasers – perceive them. In going through this process, the initial purchaser is likely once again to be conservative, or skeptical, and to assume that subsequent purchasers, whether out of ignorance or even greater conservatism, value the key benefit parameters *even less than he (the initial purchaser) does*.

The upshot is that one must maintain at least three sets of parameter values for every year: those that represent society's valuations, those that represent the initial purchaser's valuations, and those that represent the initial purchaser's estimation of subsequent purchasers' valuations. In order to maintain maximum generality, we will want to weight the initial and subsequent purchaser values each by the appropriate probabilities, which can be derived from the year-by-year likelihood of first resale. Thus, in every year, the weight on the first-purchaser's parameter values is equal to one minus the resale probability in that year, and the weight on the subsequent-purchaser's parameter values is equal to the resale probability in that year. The resale probability for year X is simply equal to the percentage of vehicles that are resold for the first time in year X of their life. I have specified this so that by year 4, about 50% of vehicles will have been sold.³

Finally, to be fully psychologically realistic, we should recognize that the first purchaser will assume that in reality he will not capture at resale the full present value of the subsequent purchaser's benefit stream, however conservatively estimated, but will have to give up a little "benefit" to the subsequent purchaser as an inducement, so that the subsequent purchaser does not merely break even. This concession can be represented as a fixed negative dollar amount, which will be multiplied in every year by the probability of resale. This fixed amount also can be understood to include any psychological or actual "transaction" costs associated with marketing the fuel economy improvements per se. I speculate that \$100 is a good fixed amount to account for both inducement and transaction costs specific to the fuel economy improvements. I assume that this is the cost in the first year, and then inflate this by 2% per year.

5. The discount rate. In most cost-benefit analyses of fuel-economy improvements, the discount rate is the most contentious parameter. This, however, is because most analysts try to account in the discount rate for all of the analytical and

³ Stolyarov (2002) states that the probability of resale is very low for very new vehicles, but then rises rapidly and peaks at around 3 to 5 years of age. He states that data from the 1995 Nationwide Personal Transportation Survey indicate that the average resale rate is 5.8% for one-year old cars, and 14.4% for two-year old cars. His graphs indicate that for three-to-five-year old cars, the rate is 10% to 20%.

“conservatism” (or risk-aversion) factors discussed above (probability of vehicle scrappage, probability of vehicle theft, probability of death, conservatism regarding fuel prices, conservatism regarding fuel economy improvements, skepticism regarding the valuations of subsequent buyers, and transaction costs). However, all of these factors are better treated explicitly separately, as above. When this is done, estimation of the discount rate becomes relatively straightforward.

In an analysis that accounts explicitly for all of the analytical and psychological factors discussed above, the discount rate itself has only one component: the consumer opportunity cost of money. The pertinent opportunity cost of money to consumer is the rate of earnings on safe investments that are alternative to expenditures on transportation. In *nominal* terms, this probably averages between 4% and 7% (based on interest rate data from the U. S. Federal Reserve, [www.federalreserve.gov/releases/H15/]), with 5.5% (before taxes) being a decent best estimate. To obtain a real (inflation-free rate), one can to a first approximation simply subtract about 2.5% (e.g., www.bls.gov/cpi/).

However, in a social-cost benefit analysis, a strong case can be made for applying a lower discount rate to some streams, such as the valuation of future reductions in GHG emissions (Sherwood, 2007). I use a nominal social discount rate of 4% per year.

Technically, the social-cost analysis probably should be done with before-tax rather than after-tax interest rates, because taxes can be viewed as obligatory purchase of government services that should be excluded, so that we are left with the real (before-tax) market yield. However, if one wishes to do an after-tax analysis, then one not only should subtract from the interest earnings the portion that is paid in taxes (perhaps 30% at the margin), but also should adjust any of the other cost or benefit parameters (such as vehicle costs, which in some cases are deductible) to an after-tax basis (a point which is sometimes overlooked).

In sum, in a formulation with explicitly separate accounting of the full range of psychologically relevant consumer factors, one does not add a "risk" premium to the consumer interest rate because all of the conceivable risks – of dying, of lower-than-expected gasoline prices, of having the vehicle stolen, of risk-aversion itself, and so on – already have been explicitly, separately represented – by, for example assuming that the consumer's estimates of fuel economy gains and future oil prices are "conservative" (risk-averse) with respect to "neutral" analytical estimates.

Other considerations

Performance, safety, and emissions. Fuel economy improvements may affect the performance, safety, and emissions of a vehicle. Changes in these attributes have implicit values, which in principle can be estimated and incorporated into a complete social cost-benefit analysis. Moreover, changes in performance and safety certainly will affect the marketability of the vehicles, and hence are pertinent in a consumer cost-benefit analysis as well. However, it probably is reasonable to assume that the safety and performance of hybrids vehicles are the same as the safety and performance of their non-hybridized counterparts. Thus, a cost-benefit analysis of hybrids, whether conducted from the perspective of the consumer or society, probably can ignore performance and safety. (However, this conclusion is not generalizable to all kinds of fuel economy improvements.)

Hybrids will have lower emissions of criteria pollutants and greenhouse gases than will their non-hybridized counterparts. For the "average" consumer in a real market, the implicit monetary value of these lower emissions probably is considerably less than the full value to society (as estimated by the standard damage-function approach), but greater than zero. It would be interesting to determine *private* willingness-to-pay (WTP) for lower emissions, in the real world, and incorporate it into the consumer cost-benefit analysis. Of course, a social cost-benefit analysis would include the full value of the reduced emissions, estimated by the standard damage-function approach (Delucchi, 2000b).

The effect of fuel economy on VMT. Finally, there is the intriguing and sometimes incompletely analyzed question of how changes in the costs of transportation affect decisions about how much transportation to “consume” -- that is, in this case, how much to drive. In the case of improved fuel economy, some analysts note that the higher fuel economy reduces the cost-per-mile of fuel, and then assume that the lower fuel-cost-per mile induces greater consumption of miles, in the form of increased driving (the so-called “rebound” effect). Any increase in driving has a private benefit that ought to be counted in a consumer cost-benefit analysis: the consumer surplus associated with the extra miles. (The consumer surplus is the difference between the full consumer value of the extra driving and the full consumer cost, including time cost, of the extra driving.) Note that the benefit here is *not* the reduced gasoline cost associated with the extra driving, because these miles would not have been driven in a non-hybridized vehicle. Rather, the cost-benefit analysis estimates the benefit of reduced fuel costs on the basis of the miles driven in the non-hybridized vehicle (as discussed above), and then estimates the benefit of the extra driving as the associated consumer surplus.

Any increased VMT may have pollution and accident costs that the consumer does not account for and hence which are external to the consumer cost-benefit analysis. These external costs ought to be counted in a full social cost-benefit analysis.

Note, though, that this discussion of the cost-benefit analysis of extra VMT is based on the *assumption* that the fuel-cost-per mile is the only cost argument in the consumer's driving calculus that is affected by the improvement in fuel economy. As I will discuss next, this assumption may not be correct.

Now, it is true that the cost of fuel is the only cost factor that is a continuous direct function of miles of driven *and* is affected by improvements in fuel economy. If this were the end of the story, then it would be true that an improvement in fuel economy would reduce the cost of a mile of travel, with no other relevant effect, and thereby induce consumption of more miles. However, if the consumer has a fixed monthly

amount to spend on transportation (i.e., is up against a transportation budget constraint), then *any* cost factor affected by the improvement in fuel economy becomes relevant. For example, an improvement in fuel economy affects the amortized initial monthly cost of the vehicle, as well as the fuel cost-per-mile. If the consumer faces a transportation budget constraint, this amortized initial cost will be an argument in the consumer's driving calculus. The consumer will find that his monthly gasoline expenditures are lower, but he also will find that his monthly income is lower, because he had to take more money out of savings to buy the hybrid. If the consumer has a fixed transportation budget, then the amount of extra driving that he can afford to buy is equal to the reduction in fuel expenditures on the original miles *less* the loss of monthly income due to the higher initial cost of the hybrid. If the lower income roughly cancels the lower gasoline expenditures (which to a first approximation appears to be the case), then in the face of a budget constraint there is no room for the consumer to spend more on additional driving.

It is important to note that the fact that econometric analyses have found a relationship between fuel economy or cost per mile and VMT does *not* indicate that there is no budget constraint operating. This is because the cost of the econometrically estimated extra VMT always is less than the saved fuel cost on the original VMT, so that if there is a budget constraint, it is not reached.

What is the upshot of this discussion of consideration of the potential effect of extra driving? Given the difficulties of estimating the value of the consumer surplus of any extra VMT on the one hand, the possibility that a fixed budget constrains consumers' ability to buy more VMT on the other, and considering that recent estimates of the rebound effect already are quite low (e.g., Small and Van Dender, 2005), one perhaps might suggest that the rebound effect be ignored.

Some results of the analysis

I have constructed a simple spreadsheet model that performs a consumer and a social cost benefit analysis of fuel economy improvements with hybrid vehicles. This spreadsheet can be used to determine:

- i) whether the benefits of fuel economy improvement are at least as great as the initial cost, from the perspective of the first purchaser or from the perspective of society, for a given set of assumptions (cost-benefit analysis);
- ii) the price of gasoline, or the interest rate, or the initial cost differential, that is required to make the benefits of fuel-economy improvement equal to the cost, from the perspective of the consumer or of society (breakeven analysis);
- iii) the difference between an analysis that formally represents consumer risks and conservatism and one that doesn't (methodological/scenario analysis).

The third type of analysis yields an interesting result. It turns out that the effect of representing consumer risk and "conservatism" parameters separately, as discussed above, with a nominal 5.5% interest rate, is the same as folding all of those parameters into the discount rate and making the interest rate 19%. That is, when there is no conservative underestimation of future gasoline prices or fuel economy improvement benefits, no resale transaction cost, and no probability of zero VMT in any year due to death or theft, then a consumer interest rate of 19% is required to produce the same net private benefit as in the baseline case with explicit representation of those parameters and a 5.5% interest rate. A discount rate of 19% is consistent with cost-benefit studies that find values on the order of 20% for consumer discount rates that implicitly incorporate risk and conservatism and are used to value fuel-economy improvements (e.g., see Greene [1983], who reviews work done before 1983 and finds "implicit" real discount rates of 4 to 40% depending on income).

Thus, the high implicit discount rate that consumers appear to apply to fuel-economy-purchase decisions is best understood *not* as an explicit expectation of a very high rate of return on the investment foregone by spending money on fuel economy, but rather

as the implicit equivalent of a series of conservative assumptions about fuel prices, fuel economy improvement, resale value, and so on, combined with an expectation of a normal rate of return on foregone investments. This suggests that we can explain, finally, the difference between the high “implicit” discount rates found in many studies and the much lower actual rates of interest in the economy.

Acknowledgments

This paper was inspired by conversations with David Greene, Paul Leiby, and others, in 2005. I drafted the paper in September 2005 and circulated it for comment but did not publish it. On the basis of recent conversations with Greene and Leiby, and of a recent unpublished paper by David Greene on risk aversion and fuel economy improvements, I have in this version highlighted my discussions of conservatism and risk aversion.

References

S. C. Davis and S. W. Diegel, *Transportation Energy Data Book*, ORNL-6978 (edition 26 or ORNL-5198), Center for Transportation Analysis, Oak Ridge National Laboratory, Tennessee (2007). <http://cta.ornl.gov/data/index.shtml>.

M. A. Delucchi, *Motor-Vehicle Goods and Services Priced in the Private Sector*, UCD-ITS-RR-96-3(5), rev. 1., Institute of Transportation Studies, University of California, Davis (2004). Available at www.its.ucdavis.edu/people/faculty/delucchi.

M. A. Delucchi, *Electric and Gasoline Vehicle Lifecycle Cost and Energy-Use Model*, UCD-ITS-RR-99-4, report to the California Air Resources Board, revised final report, Institute of Transportation Studies, University of California, Davis, April (2000a). www.its.ucdavis.edu/people/faculty/delucchi.

M. A. Delucchi, “Environmental Externalities of Motor-Vehicle Use in the U. S.,” *Journal of Transport Economics and Policy* **34**: 135-168, May (2000b).

M. A. DeLuchi, Q. Wang, and D. L. Greene, *Motor Vehicle Fuel Economy, The Forgotten HC Control Strategy?*, ORNL-6715, Oak Ridge National Laboratory, Oak Ridge, Tennessee, June (1992).

D. L. Greene, "A Note on Implicit Consumer Discounting of Automobile Fuel Economy: Reviewing the Available Evidence," *Transportation Research B* **17B**: 491-499 (1983).

S. Sherwood, "Discounting and Uncertainty: A Non-Economist's View," *Climatic Change* **80**: 205-212 (2007).

K. A. Small and K. Van Dender, *The Effect of Improved Fuel Economy on Vehicle Miles Traveled: Estimating the Rebound Effect Using U. S. State Data, 1966-2001*, UCEI Energy Policy and Economics Paper 014, University of California Energy Institute, Berkeley, California, September (2005). www.ucei.berkeley.edu/PDF/EPE_014.pdf.

D. Stolyarov, "Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?," *Journal of Political Economy* **110**: 1390-1413 (2002).