Schmeduling

Jeffrey B. Liebman and Richard J. Zeckhauser*
Harvard University and NBER

July 2003

We thank David Weisbach, Joshua Meltzer, and participants at the USC Behavioral Public Finance conference for comments on an earlier draft, and Alberto Abadie and Erzo Luttmer for helpful conversations. We are grateful to Jesse Shapiro for his assistance with the analysis of the data from the food stamp cash out experiment.

* Contact information: jeffrey_liebman@harvard.edu and richard_zeckhauser@harvard.edu

Abstract

Complicated pricing schedules can make it very difficult for consumers to know what price they are paying. Such schedules are in widespread use in important economic domains such as taxation, assistance to the poor, and utility pricing. When people have limited understanding of the actual schedules that they face, we argue, they are likely to perceive them in a crude or smeared fashion. We define the term "schmedule" to be a smeared – or inaccurately perceived – schedule. Our focus is on two schmeduling practices: ironing and spotlighting. Ironing arises when an individual facing a multipart schedule perceives and responds to the average price to the point where he consumes. Spotlighting occurs when consumers identify and respond to immediate or local prices, and ignore the full schedule, even though future prices will be affected by current consumption.

We analyze the welfare implications of schmeduling in three settings: a profit-maximizing monopolist, a Ramsey-pricing regulator, and a social-welfare maximizing tax authority. We show that with convex schedules, outcomes that are Pareto superior to the rational responders' outcome are available in all three contexts, though the schmedule setter will not necessarily choose such outcomes.

We provide empirical tests of ironing using the 1998 introduction of the child tax credit, and of spotlighting using data from a food stamp cash out experiment. In both cases, the data suggest that schmeduling behavior is occurring.

"If line 11 is equal to or more than line 12, enter the amount from line 8 on line 14 and go to line 15. If line 11 is less than line 12, divide line 11 by line 12. Enter the result as a decimal (rounded to at least three places)."

Internal Revenue Service (2002)

"Beginning with your November bill and continuing through April 2001 your gas adjustment factor will be $0.68530 per therm. The local distribution adjustment factor will be $0.00820. . . .For an average customer on Rate R-3 this will amount to a $33.83 increase in your bill."

Keyspan Energy Delivery (2000)

"Roaming rates apply to calls placed and received outside this area. Long distance charges for calls received while roaming are calculated from your home area code to the location where you received the call. Due to delayed reporting between carriers, usage may be billed in a subsequent month and will be charged as if used in the month billed.. . . . Other charges, surcharges, assessments, universal connectivity charge, and federal, state and local taxes apply."

AT&T Wireless (2002)

Medigap Monthly Premiums for Plan C in Denver Colorado Zip Code

	Age65	Age70	Age75	Age80
AARP	129	129	129	129
Equitable Life	96	113	123	134
5 Star Life	52	78	58	88
Union Banker	204	230	271	326

www.centuraseniors.org

The demand curve is a bedrock concept in economics. It tells how much of something a person will buy at each price. The efficiency of the market equilibrium requires that the demand curve accurately reflect the willingness to pay. Yet, in a vast array of circumstances people have little or no idea what price they are paying. For example, we conjecture that few consumers have any idea of how much it would cost to run their dishwasher twice a day rather than once a day or to keep the thermostat in their home set one degree higher during the winter. Similarly, we suspect that few people know with any precision how close they are to running out of their monthly allotment of zero-cost cellular phone minutes. Moreover, there is ample evidence (discussed below) that taxpayers and welfare benefit recipients often have little understanding of their marginal wages net of taxes and transfers. In some cases, consumers may simply ignore pricing schedules. Medigap is a product that is standardized by the federal government. The Denver Medigap example atop this paper demonstrates, however, that a wide range of schedules are available, and some persist in the market despite being clearly dominated.[1] In all of these cases, and in many other ones, it is likely that individuals are making suboptimal choices. Interestingly, in important cases these suboptimal choices reduce deadweight loss, and thus increase collective welfare.

In this paper we undertake four tasks. First, we develop a theory that describes the circumstances under which people are likely to fail to perceive the true prices that they face. In particular, we argue that misperception of prices is most likely to occur when pricing schedules are complex, when the connection between consumption and payoffs is remote, and when other features of the economic environment make it difficult to learn from past experience. We illustrate this theory with examples from five areas of economic behavior.

Second, drawing upon experimental results in psychology as well as evidence on how

---

[1] AARP's premiums do not vary with age; the others do. Moreover the rates charged by 5 Star Life are well below those of the other insurers, while those charged by Union Banker are well above. These insurers were those that listed rates for the most plans. The $58 per month premium at age 75 is the premium listed on the web site. It is not a typographical error by us.

people perceive the incentives created by existing tax, transfer, and regulatory systems, we posit several behavioral rules for how people actually perceive and respond to schedules. We argue that when people have limited understanding of the actual schedules that they face, they are likely to perceive them in a crude or smeared fashion. We define the term "schmedule" to be a smeared – or inaccurately perceived – schedule. Thus, even when confronted with well-defined schedules, people often behave as if they were facing a schmedule. We call this practice schmeduling, and those who do it schmedulers.

Our focus is on two schmeduling practices: ironing and spotlighting. Ironing in real life is intended to make something flat. The schmeduling variant of ironing arises when an individual facing a multipart schedule perceives only the average price to the point where he consumes. Thus, an individual earning $80,000, and therefore in the 30 percent marginal tax bracket, observes that his taxes are $16,005. He irons (flattens) out his perceived tax schedule and operates as if his rate were a constant 20 percent. Spotlighting occurs when consumers respond to immediate or local prices and ignore the full schedule that they face. In particular, spotlighting frequently occurs when individual make choices in response to current prices, but do not take into account the effect of current choices on future prices. Thus, a food stamp recipient may consume more calories in the early days of the month, when food appears to have a much lower cost.[2]

---

[2] Though we address schmeduling and spotlighting in their pure forms, we recognize that most schmedulers mix in some element of rational response to schedules, and that some individuals are closer to full rationality than others are. Thus, our representative food stamp recipient understands that he will run out of stamps, so implicitly attaches a positive shadow price to current food stamp consumption. However, his shadow price is likely too low given his own preferences. Similarly, our wage earner may not iron perfectly flat. He may misperceive his 30 percent marginal tax rate as 25 percent, not as 20 percent.

Third, we analyze the welfare implications of pure ironing and spotlighting behavior. We study ironing in three settings: a profit-maximizing monopolist, a Ramsey-pricing regulator, and a social-welfare-maximizing tax authority. We refer to the behavior of a sophisticated schedule setter who takes account of the schmeduling behavior of consumers in setting the schedules, as "schmedule setting." We show that ironing behavior eliminates some of the deadweight loss from high marginal prices. When the optimal schedule with rational consumers is convex, ironing improves the outcomes available to the schmedule setter. Indeed, with convex schedules, outcomes that are Pareto superior to the rational responders' outcome are available in all three contexts, though the schmedule setter will not necessarily choose such outcomes. We also show, using calculations from the 1998 IRS public use sample of tax returns, that the welfare implications of the ironing variant of schmeduling are potentially very large in the tax example. We defer our welfare analysis of spotlighting to a subsequent draft.

Fourth, we provide findings from two empirical tests of rational versus schmeduling behavior. The first uses data from before and after the 1998 introduction of the child tax credit to test for ironing. The second uses data from the San Diego food stamp cash out experiment to test for spotlighting. We recognize that in real life some individuals are serious schmedulers, whereas others perceive schedules accurately and respond appropriately. Thus, if data reveals evidence of schmeduling, it is despite the presence of some people who have behaved rationally. We turn now to the conditions that give rise to schmeduling, recognizing the caution that some people are more prone to such behavior than others.

I. Conditions that Give Rise to Schmeduling

We begin by listing general conditions that are conducive to schmeduling. We then describe five significant examples in which we believe schmeduling is prevalent, and show how these illustrate the listed conditions. In subsequent sections, we discuss more thoroughly what behavior is likely to result when people do not perceive their true marginal prices, and derive the welfare implications of this behavior.

We conjecture that there are nine conditions that are often present in circumstances in which people have difficulty perceiving incentives -- e.g., prices or taxes – operating at the margin. While not all of these conditions need to be present for schmeduling to arise, we expect that it will be rare to observe significant schmeduling unless several of these conditions are present, and that schmeduling will arise more often and in more extreme forms when a great number of the conditions occur.

The nine conditions fit into three broad categories:

Category A: Complexity

The first three conditions involve characteristics of schedules that make it non-trivial to determine the marginal price and make it costly to calculate a person's exact location on the schedule.

1. ***Nonlinear pricing***. Schmeduling is more common when there is the potential to confuse average and marginal prices.

2. ***Schedule complexity***. Schmeduling is more common when there are more rates in the schedule or if the consumer is operating on two or more schedules simultaneously.

3. ***Frequent revisions of schedules***. Schmeduling is more common if the pricing schedule is revised frequently, implying that rates may not be known, or that groping toward the

optimum is less likely to be successful.

Category B: Remote Connection Between Consumption and Payoff

The next two conditions make it difficult to perceive prices from one's own market transactions. Both conditions are found, say, with household consumption of electricity and water.

4. ***Delayed payoffs***. Schmeduling is more common when the payoff from a decision is separated in time from the consumption choice.

5. ***Bundled consumption***. Schmeduling is more common when the payoff from each choice is bundled with many other choices. These other choices can either be different types of choices or they can be similar choices at different points in time.[3]

Category C: Environment is not conducive to learning

The remaining four conditions are ones that make it difficult for a person to learn the marginal price he faces from personal experience or the experience of acquaintances.

6. ***Nonstationary economic environment***. Schmeduling is more common if the environments in which people are making choices are changing so that people are operating at different points on the schedule each time they make a choice.

7. ***Heterogeneity in offered schedules.*** Schmeduling is more common when one's acquaintances face different schedules or are operating at different points on the schedule

[3] As we discuss later, this criteria is closely related to Herrnstein and Prelec's (1991) concept of distributed choice.

than you are. When population members face different rates, it is more difficult to learn one's true price by asking a friend or by comparing one's payoff to the payoff received by a friend who made a different consumption choice.

8. ***Obscure pricing units***. Schmeduling is more common if the units for which people are charged are different from the units in which people make consumption choices. Indeed, given this condition, some forms of schmeduling are likely to arise even if prices are constant.

9. ***False signals***. Schmeduling is more common in circumstances in which other information is presented to the consumer that could be misinterpreted as the marginal price. This includes circumstances when the consumers are presented with average prices along with or instead of marginal prices. It also includes circumstances in which there are multiple payoffs received per accounting period, but the payoffs in the early period do not reflect the marginal payoff conditional on expected future behavior.

We now present five examples of areas of economic behavior in which we expect to observe schmeduling. As we discuss each one, we will refer to Table I, which identifies the conditions above that apply to each.

1. Tax Systems

A substantial body of research indicates that people do not understand what their tax schedules look like. Interviews with taxpayers in the UK (Brown, 1968; Lewis, 1978), Italy (Bises, 1990), and Sweden (Brannas and Karlsson, 1996) and with EITC recipients in the U.S.

(Liebman, 1996; Olson and Davis, 1994; Romich and Weisner, 2002) all suggest substantial confusion about marginal rates.[4] Fujii and Hawley (1988) compare responses to a survey question about marginal tax rates to calculated marginal tax rates using Survey of Consumer Finances data; they conclude that there are significant differences.[5] De Bartolome (1995) shows that people confuse average and marginal tax rates when asked to make calculations using a tax table similar to those published by the Internal Revenue Service with the 1040 tax form. Individuals' actual choices often reveal traces of schmeduling. For example, the evidence that taxpayers generally do not bunch at kink points (Heckman, 1983; Liebman, 1998; Saez 2002) and that people locate at places on the budget constraint where theory says that they should not reside (Macurdy et al 1991) is usually interpreted as suggesting that taxable income elasticities are small (Saez 2002) or that the specification of preferences is wrong (Heim and Meyer 2002).[6] Schmeduling offers a different explanation. Lack of bunching at concave kink points and the presence of people at convex kink points could also result from people not knowing or

[4] Rosen's (1976) evidence suggests that people do not ignore taxes altogether. Break (1957) finds that solicitors and accountants in the UK are aware of their marginal rates (but that taxes have little impact on their work hours).

[5] The Fujii and Hawley study is open to alternative interpretations. The authors do not observe itemized deductions in their data set. Hence, measurement error could contribute to the discrepancies that they present. Moreover, the papers presents average marginal tax rates using both the survey and the calculated approach but do not show the distribution of individual level discrepancies. Therefore, their study is not as informative as it could be for the question we are asking.

[6] Saez (2002) acknowledges the possibility of a behavioral explanation for the lack of bunching.

misperceiving the details of their tax schedule.[7]

More generally, the complexity of the tax code makes it unlikely that most taxpayers calculate their marginal tax rates accurately. In addition to the seven statutory marginal tax rates (in 2002 these are 0, 10, 15, 27, 30, 35, and 38.6), there are 22 provisions (not including the alternative minimum tax) that "give rise to deviations between effective marginal tax rates and statutory marginal tax rates," and the Joint Tax Committee estimates that 1 in 8 filers faces an effective marginal tax rate that differs by more than 10 percent from the taxpayer's statutory marginal rate (Barthold, et al 1998).[8] These calculations do not take into account other tax and transfer programs that can alter incentives (state income taxes, food stamps, TANF, student loans, housing assistance, etc.). Neither do these calculations factor in 1) complicated dynamic issues such as the lifetime marginal tax rate if some of a person's earnings are going to be converted to future consumption via saving; 2) tax rates on future returns to human capital investment or on-the-job experience; or 3) the relationship between current payroll taxes and future OASDI benefits (Auerbach and Kotlikoff 1985; Kotlikoff, 1996).

The implication of these features of the tax system is that there is little chance that most people accurately figure or somehow know their marginal tax rates. Therefore, the important

---

[7]The simulations in Saez (2002) suggest that the uncertainty about what annual income will turn out to be is not large enough to explain the lack of bunching at kink points if elasticities are at least moderately large. Preliminary simulations of our own suggest a similar result for schmeduling. In most cases, it would take very large amounts of uncertainty for us to be unable to distinguish a rational consumer from a schmeduler.

[8] In some of the cases in which tax credits are phased out, it is very difficult for a taxpayer to figure out the rate at which the credit is phased out. The first quotation atop the paper is from the calculations that determine the phaseout of education tax credits. Even with the tax form in hand, it would be very challenging for taxpayers to figure out what phaseout rate applies (in the unlikely event that they wished to do so).

question is what alternative methods taxpayers could or do employ, and how well these alternative methods allow them to approximate their true marginal tax rates.

People may come to perceive their true net wages by observing how their after-tax income changes from year to year in response to changes in effort, even without doing calculations using tax tables or trying hypotheticals in TurboTax. We suspect, however, that few people undertake calculations.[9] Even if they do attempt them, their economic environments are not sufficiently stationary to make this sort of calculation useful. Even if earnings stay constant, changes in marital status, family composition, housing consumption, life-cycle earnings patterns, and tax laws mean that people will often be on different segments of the budget constraint. Thus, we think it worth considering ways people might go astray in making these decisions.

Table I shows that the tax system features most of the conditions we predict should give rise to schmeduling. The tax system creates nonlinearities in the price of leisure. Tax schedules are complex. They are revised frequently. The payoff from a decision this January may not be realized until April of the following year (or August or October in the case of a taxpayer who requests extensions). Often very different decisions (labor effort of two people, sale of capital assets, degree of tax avoidance undertaken) together determine a single annual payoff and individuals can be on a very different point on the schedule than their friends or neighbors. Finally, taxpayers may be misled by their monthly, weekly, or biweekly payroll statements to think that they face a constant tax rate equal to the quotient of their net pay divided by their gross

---

[9] We have been told by several colleagues that the summary statistics automatically produced by TurboTax when the taxpayer has finished filling out the tax return include the taxpayer’s average tax rate, but not the marginal tax rate. Thus, it takes deliberate effort (redoing the tax return with an alternative income level) for a taxpayer to learn his marginal tax rate from this software.

pay. In simple cases, this would lead taxpayers to confuse average and marginal tax rates. Moreover, the withholding schedule ignores non-wage income and the gap between gross and net pay that the taxpayer observes on his paystub may involve other payroll deductions for things like life insurance, dependent care accounts, medical savings accounts, parking, and the like. This makes it possible that taxpayers are reacting to a number that is neither their marginal nor their average tax rate.

2. Welfare Programs

Income transfer systems create, to our knowledge, the most complex schedules faced by ordinary citizens. Many recipients receive benefits from multiple programs, each with its own schedule of how benefits fall with earnings. Thus, even when the benefit-reduction schedule from a single program is linear, the combined schedule is highly nonlinear. Moreover, each program has complicated rules about amounts of income that are disregarded before applying the benefit-reduction schedule. For example, the food stamp program disregards the first $134 dollars of income (in 2000) plus 20 percent of earnings (among a long list of other deductions) and then assesses its benefit reduction schedule on remaining earnings. This implies that the way in which earnings are allocated across months affects benefit payments. The Earned Income Tax Credit (EITC) initially rises with earnings, is constant at its maximum value for a range of earnings, and then is phased out as earnings rise even further. Payment is usually made as part of the worker's annual tax refund check. Therefore, workers often do not know how much of their refund was due to the EITC and how much was due to excess withholding of income tax during the year. More importantly, when making decisions during the year, they do not understand how the decisions will affect their EITC refund.

Even economists preparing professional papers have a hard time computing effective marginal tax rates for recipients. The complexity of the income disregards largely explains the wide range of estimates of effective tax rates for AFDC/TANF recipients in the empirical literature – ranging from the work of Dickert, Houser, and Scholz (1994) who find cumulative rates of 15 to 40 percent to the work of Giannarelli and Steuerle (1995) who find rates of 75 percent or more.

The economic environments of welfare recipients are often nonstationary, making it hard for recipients to know where they are on the schedule in the current period. In part, this is due to program rules. For example, benefit reduction rates often vary based upon the length of time a person has been receiving benefits.[10] Earnings variability also contributes. Recipients often experience large discrete jumps in earnings levels, implying that knowing the marginal rate on one additional dollar of income may be a very bad estimate of the payoff to the change (say from part-time to full-time work) that the person is actually contemplating.

There are features of welfare programs, however, that make their incentives easier to perceive than those of the tax system. Their accounting period is usually one month, whereas tax systems generally use one year. Thus, a person who increases his or her earnings this month will, within a month or two, see the effect on his or her welfare benefits. Kling, Liebman, and Katz

---

[10] Time limits make the welfare recipient's decision problem into a complex dynamic programming problem of how to consume out of a fixed potential-benefit stream. See Grogger and Michalopoulos (1999) on welfare time limits and Pollack and Zeckhauser (1996) on the more general problem of how to consume out of a fixed budget over multiple periods. Each paper finds that complex, nonintuitive strategies are optimal.

(2001) report that a large share of recipients of housing assistance know that their rent will go up by exactly 30 percent of any increase in their earnings. They speculate that the difference between respective monthly and annual accounting periods explains why people receiving housing assistance know their marginal benefit reduction rate whereas, for example, EITC recipients generally have no concept of the EITC phaseout. A second possible explanation is that the 30 percent housing tax rate has been around unchanged for many years; moreover, it applies nationwide to everyone in public housing. Public housing is thus unlike more complicated tax or transfer programs, which place individuals at different points and slopes on the schedule, and rates vary across locales. Your neighbor in public housing or your cousin in another city can accurately tell you what your effective tax rate is from housing assistance.

3. Utility Pricing

Utility pricing schedules – though not as complex – bear similarities to tax schedules in that they have multi-tiered, nonlinear pricing. However, utility schedules have four additional features that make it difficult for consumers to perceive the marginal price of consuming additional water, electricity, or heating fuel. First, the pricing schedules are sometimes not published on the monthly bill and therefore take effort to discover. Second, consumers often locate at very different points on the schedule in different seasons, since their demand for natural gas in the winter and for electricity in the summer can be large multiples of demand in other seasons. Third, the pricing schedules often vary from season to season, with higher prices in the winter for natural gas, and higher prices for water during the summer drought/lawn watering season. Fourth, and most importantly, the link between the decisions made by consumers (how long to stay in the shower, whether to run the dishwasher when it is only half full, where to set the thermostat) and the consumption of water, gas, oil, or electricity is hard to observe. How many gallons of water are used per shower and how much does it cost to heat the hot water that is used? How many additional therms of natural gas are used up if you keep your thermostat at 71 rather than 68 if the outside temperature is 20^{o}F? Bills are presented in consumption units that are not directly observable (and are often incomprehensible) to the consumer and monthly payments aggregate hundreds of disparate individual consumption decisions (turning on the light, leaving the computer on over night, running the dishwasher, buying a new refrigerator).[11] These last three factors – which correspond to a nonstationary economic environment, delayed payoff, and bundled consumption on our list of conditions – combine to make it almost impossible to determine one's marginal price by observing how bills vary with behavior.[12] Thus, as in the case of tax and transfer systems, the question remains of what people actually do when making their decisions.

4. *Nonlinear Penalties, Fines, and Insurance Contracts*

---

[11] Even simple values are not known. How many of us would wager at even odds that they could guess their cost of operating a 100 watt bulb for an hour if given a range of X to 2X, where they pick the X?

[12] Interestingly, 100 years ago people were much more likely to know the marginal price of a bath or shower since they would pay for showers and baths at public facilities. Shifrin writes in *Victorian Turkish Baths* (2002): "In effect the bathers were divided by admission charge into two classes. From 10:00 am until 4.00 pm cost 2s 6d; from 4:00 till 9:00 the charge was reduced to 1s 6d. Typically, women bathers were disadvantaged, being admitted on Tuesday and Fridays only, from 10:00 am till 1:00 pm, and at the higher charge of 2s 6d."

In 1959, one of us (RJZ) brought his car to Cambridge at the start of his sophomore year. Early in the fall semester, he let his roommate borrow his car. His roommate returned the car, and mentioned that he had received a ticket, but that it was free. Indeed, the schedule was something like two free tickets, then $5 for the 3rd, $10 for the fourth, and $20 for any ticket thereafter. The roommate asserted he owed nothing. RJZ, believing that he would probably get four or five tickets in the year, suggested that $15 might be the expected marginal cost due to the ticket, and that they could always settle up when the cost became known at the end of the year.[13]

This sort of penalty structure is common. For example, automobile insurance rates often start rising after a person has received more than 3 points on his license from moving violations, or has made a certain number of claims under his comprehensive insurance policy. Criminal sentencing guidelines often impose higher prison sentences on convicts who have previous convictions; recent state "three-strikes" laws are a particularly salient example. Medical flexible spending accounts similarly have the feature that the out-of-pocket costs of initial units of consumption are low (zero), but the true marginal cost of additional consumption early in the time period conditional on expected consumption later in the time period is high (positive). Insurance policies often have a deductible, cover some fraction of expenses over a particular range of losses, and then pay nothing once a high claim limit is reached. We conjecture that confusion similar to that of the roommate arises commonly in these cases. A key feature for each is that the within-accounting-period payoffs present false signals of the ultimate marginal price.[14]

---

[13] The roommate, who was not studying economics, did not counter with the argument that the owner would be excessively careless, not internalizing costs incurred by the roommate.

[14] We suspect that the strategies basketball coaches employ in taking out players in foul trouble may sometimes represent schmeduling errors. Often players taken out early in the game

In Table 1, we see that these examples match the conditions of nonlinear pricing, delayed payoff, and bundled consumption as well.

5. Nonlinear Pricing of Consumer Goods

For most consumer goods, such as milk or clothing, consumers are told the price at the time of contemplated purchase, and the per-unit price does not vary with quantity. However, even with ordinary goods, consumers are sometimes offered quantity discounts. Such pricing may lead consumers to schmedule..

In the typical case, the schmeduler is hurt by failing to rationally optimize. Say he uses ironing. If the schedule is convex, he consumes units whose marginal cost exceed marginal benefits. If the schedule is concave, he fails to consume units whose marginal benefits exceed marginal cost. This assumes that the schedule setter does not change the schedule in response to the schmeduler's behavior. If he does respond, the schmeduler may be better off than he would be in a rational world. As we will see in the theory section below, the setter of the schedule may be benefitted or hurt, depending on the goal, the shape of the schedule, and the particular schmeduling behavior.

In some cases, merchants hurt themselves by presenting their pricing in a confusing fashion that produces schmeduling behavior. Shutterfly.com describes its holiday cards as costing 82 cents per card if 100 are purchased and 69 cents per card if 200 are purchased. Thus the marginal cost of the second 100 cards is only 56 cents per card. We suspect that most

---

end up not fouling out, implying that they could have played more minutes. There is also a misperception about time (points) late in the game being more important, which requires an unusual model. In this discussion, we are ignoring issues of fatigue.

consumers deciding between 100 and 200 cards perceive the marginal price as 69 cents and that Shutterfly could increase both its sales and its customers' consumer surplus by describing its pricing as a two-part schedule, so as to get some additional people to respond to the schedule rather than a schmedule.

In other cases, it is likely that businesses capitalize on the confusion that pricing schedules create, causing some customers to pay more or purchase more than they otherwise would for the service. Cell phone packages which have prices that rise steeply if a customer uses more than his or her allocated amount of monthly minutes presumably fall in this category. It is next to impossible for a customer to know how many minutes he or she has used up so far in the month. Moreover, as the quotation at the top of the paper shows, there is no necessary connection between when the customer makes the calls and which months the calls are assigned to for billing purposes.[15]

[15] It is not clear exactly what strategy cell phone companies are following here. One possibility is that they are counting on people consuming more than they plan and then paying high marginal rates on extra minutes. Another possibility is that they are trying to get people to purchase more than their expected number of minutes so as to protect themselves against going over the limit. This is a pricing policy similar to the option to pay for a full tank of gas ahead of time when using a rental car rather than paying an exorbitant amount if you return the car without a full tank of gas. The existence of firms offering "one-rate" plans suggests that there may be some uncertainty over whether it raises or lowers profits to have confusing plans – or it may simply indicate that there is some segment of the market that is willing to pay a premium to eliminate the uncertainties of nonlinear pricing. This process is severely complicated because the companies are in competition with one another. They offer plans that differ significantly in structure. Normally, the companies would suffer because consumers would adversely select against them (high users taking unlimited plans, people who call each other frequently taking family plans, etc.) Schmeduling reduces such behavior.

The Salad Bar Problem

Schmeduling may even arise where there is a simple per-unit cost, if the unit is hard to measure. Many salad bars present soggy ingredients, soggier than customers would choose for themselves quite apart from cost. Presumably this enables them to post more modest per ounce prices. Presumably also, a customer would prefer to pay $1.00/ounce for appropriately presented ingredients rather than $0.80 for soggy ingredients that weigh 25% more. Salad bar entrepreneurs are assuming that customers respond to price/ounce rather than price/unit of valued ingredient, or to how much the whole salad costs. In other words, soggy producers are assuming no discount for preferred presentation and no costs of straining ingredients as they are selected.

II. How People Respond to Schmedules

Some people are rational or face very simple pricing schemes and therefore know exactly where they are on their various schedules. In addition, some affluent people may hire people to do the calculating and optimizing for them. There are also cases in which a rational heuristic may be almost as good as being rational. For example, some individuals may engage in first-differencing as a means to estimate marginal prices.[16]

We have argued, however, that there are many circumstances in which people are unlikely to understand the true marginal prices or incentives that they face. An important question, therefore, is how people behave in these situations. Our proposition is that people

---

[16] In other words, consumers may infer marginal prices by calculating the change in payoff divided by the change in quantity consumed over subsequent accounting periods. Alternatively, they may infer marginal prices by comparing their own situation to that of a similar person who made a slightly different choice.

facing pricing schedules often engage in two prominent variants of schmeduling: ironing and spotlighting.

With ***ironing***, people smooth over the entire range of the schedule, perceiving the average price rather than the marginal price. Thus one decides whether or not to raise the temperature in one's home by noting that $300 per month represents an average price of 60 cents per therm, rather than 89 cents for the last (and next) therm of natural gas.[17] With ***spotlighting***, people act as if the instantaneous payoff in the current sub-period applies to the entire accounting period. Thus users of medical flexible spending accounts act as if consumption in January is free, ignoring the fact that by the end of the year they will be paying the full cost of marginal care.

Although our initial impetus for studying these two forms of schmeduling came from a combination of intuition and casual empiricism, it turns out that these same "irrational" behaviors are well documented in the experimental psychology literature for a wide range of species including pigeons, rats, monkeys, and humans. In particular, schmeduling is closely related to Richard Herrnstein's theories of melioration, and distributed choice.

Herrnstein (1961) demonstrated what he called the matching law: that hungry pigeons, choosing which of two response keys to peck, peck on each lever in proportion to the amount of reinforcement (food) obtained by pecking on that lever. In this experiment, food became available from each key after different intervals of delay and the intervals were independent of the behavior of the pigeons. For example, a peck on key A would produce food only if one

[17] In fact, in the home heating example we doubt that people ever do the conversion to price per therm. Nonetheless, in thinking about marginal consumption decisions, we believe people make those decisions by thinking of increments to the $300 monthly bill which (since it includes the price of inframarginal consumption) will result in behavior that is responsive to average rather than marginal prices.

minute had passed since the previous time food had been delivered in response to a peck on key A, whereas food was delivered with a two minute delay on key B. In this case, two-thirds of the reinforcement would be obtained from pecking on key A and the pigeons would peck on key A approximately two-thirds of the time. By 1976, similar results had been obtained from rats, monkeys, and humans (De Villiers and Herrnstein, 1976), and the matching law had been shown to apply to variations in the quantity of reinforcement obtained (as opposed to the frequency of reinforcement) and to the decision of how often to engage in a single activity (as opposed to the choice between two different options).[18]

Herrnstein and Vaughn (1980) and Herrnstein (1982) proposed a theory, which they called melioration, to explain the phenomenon of matching. Melioration states that subjects act to equalize the average reinforcement (utility) across choices. Under this model, a pigeon in the example above might start dividing his pecks equally across the two alternatives. But doing so would result in a higher average return per peck on the A key, causing the pigeon to reallocate pecks to the A key until the average returns were equalized. This would occur when two-thirds of the pecks occurred on the A key. Herrnstein (1982) argues that the calculations necessary for melioration are much simpler than those required for maximization: "[Melioration requires] the organism to respond only to the difference between local reinforcement rates from individual behaviors. Maximization, in contrast, requires the selection of the biggest aggregation of reinforcement across behaviors."

---

[18]The introductory essays by Rachlin and Laibson (1997) are indispensable in understanding this literature.

As Herrnstein (1990) notes: “It should soon be evident that the fundamental difference between matching and utility maximization is that matching is based on average returns (in utility or reinforcement) over some extended period of activity, while maximization requires a sensitivity to marginal returns at each moment. Where the marginal and average returns to response alternatives are equal . . . we would expect to find no large difference in the predictions of theories relying on one or the other of them.” Herrnstein and Heyman (1979), Mazur (1981), and Vaughn (1981) conducted experiments with pigeons in which melioration and maximization predict different behaviors. In all three of these studies, behavior followed the predictions of melioration.[19]

Pigeons Ironing

In the Mazur (1981) experiment, pecks on either a red or green key would occasionally result in a three-second period of darkness. During a random fraction of these periods of darkness, food was delivered. The percentage of dark periods yielding food was different for the red key and the green key. The periods of darkness occurred randomly and were assigned equally across the two keys. Moreover, once a period of darkness had been assigned to a key, nothing happened until the pigeon pecked on that key. Thus the maximizing strategy for the pigeon was to alternate pecks in equal proportion across the two keys (since pecking on the currently scheduled key was the only way to get the schedule to advance to the next dark period). Although the maximizing strategy was simple and the pigeons were all initially conditioned to respond at a ratio between that predicted by melioration and maximization, all of the pigeons

[19] Our accounts of these studies are based upon Herrnstein (1982).

moved nearly all of the way to the pecking ratio that would be predicted by melioration (proportional to the relative ratio of dark periods providing reinforcement) and far from the ratio that would be predicted by maximization. Thus instead of equalizing marginal returns they equalized average returns.

Pigeons Spotlighting

The Vaughn (1981) study is particularly relevant to our schmeduling discussion because pigeons faced a complicated payoff schedule similar to the sorts of nonlinear pricing schedules that are our focus. In this study, pigeons faced the overall reinforcement schedule shown in Figure 1. A pigeon who consistently spent between 12.5 percent and 25 percent of the time pecking on the right key would receive the highest payoff, whereas a pigeon who spent between 75 percent and 87.5 percent pecking on the right key would receive the lowest payoff. Every four minutes, the fraction of pecks on each key in the previous four minutes were tallied and the payoff delivered during the next four minutes was based upon the fraction of pecks on the right key in the previous four minutes. However, if the relative time spent on the right during the previous period was less than 75 percent, the payoff during the current period was disproportionately delivered after pecks to the right key. The three pigeons in this experiment all ended up spending between 75 and 80 percent of their time on the right key – responding to the local reinforcement instead of the global payoff schedule and thereby achieving the minimum possible payoff.[20] In other words, they act exactly like the health care consumer who interprets

[20] In the other condition in the experiment, the global payoff was held (approximately) constant while the local reinforcement was set to be higher for left if the fraction of right was 25 percent or more. In this case, the pigeons selected right 20 percent of the time.

the current month’s copayment rate as the price of medical care, ignoring expected consumption later in the year.

Humans Spotlighting

In section I, we argued that schmeduling is particularly likely to occur when many different choices are aggregated into a single pricing schedule. Herrnstein and Prelec (1991) similarly argue that when a “choice is an aggregate of many smaller decisions, distributed over a period of time” the choice “may be reliably and predictably suboptimal, in terms of the person’s own preferences.” The basis for this claim is a series of experiments on humans (see Herrnstein 1991) in which subjects had to choose between pressing the right or left arrow on a computer keyboard in exchange for monetary rewards. The subjects observed their monetary reward accumulate via a computer screen which showed pennies falling. A penny fell each time a key was pressed, but the key could not be pressed again until the preceding penny had finished falling. Pushing the right key always caused a penny to fall 2 seconds faster than did pushing the left key. However, the greater the fraction of right key presses in the past 10 choices, the slower the coins fell regardless of the choice. The exact parameters made it optimal to exclusively choose the left option. However, almost all of the subjects exclusively picked the right option. Thus rather than optimizing over the entire schedule, the subjects responded to the immediate reinforcement rate.

Distributed Choices and Accounting Periods

Herrnstein (1982) notes that melioration with long accounting periods approaches maximization. He writes that for an organism that has "the capacity to detect correlations between behavior and its correlations between behavior and its consequences over increasingly long time spans, melioration approaches maximization." However, he is dubious that this phenomenon – which is equivalent to the first-differencing strategy for learning one's marginal prices that we described above – is widespread: "A meliorating organism is a maximizing organism if it has an infinite capacity to redefine response categories to suit prevailing contingencies of reinforcement, for then the optimal distribution of responses in any situation would be treated as a single response category it its own right and would be chosen exclusively as a result of melioration . . .However, no evidence has been provided for infinite response plasticity in any species." (Herrnstein 1987). Herrnstein and Prelec (1991) doubt that humans can first-difference when faced with distributed choices: "It is possible in principle to attempt to compare the average returns associated with different choice distributions. But to conduct this introspective exercise, one would have to mentally consume, say, a 40-60 mix of the two meals, and compare it to a 60-40 mix, and an 80-20 mix, and so on. Can a person discriminate among satisfaction levels produced by meal series that differ in the relative frequencies of items?"

Herrnstein's theory of distributed choices also suggests a prediction for when we will expect to observe ironing and when we will expect to observe spotlighting. Ironing will tend to occur when there is a single payoff for all of the bundled choices. Spotlighting will occur when there are multiple within-accounting-period payoffs.

In sum, we believe that the evidence from experimental psychology – mostly with pigeons but some with humans – establishes ironing and spotlighting as plausible models of how people behave when faced with complicated schedules. Whether these theories can in fact explain people's behavior in the applications that are our focus remains to be seen. Before turning to empirical tests of these theories in section IV, we first discuss the potential welfare implications of such behavior.

III. Welfare Implications of Schmeduling

We argued in the prior section that both studies in psychology and intuition about decision making suggest that schmeduling behavior is to be expected. Indeed, given how often people do not know the location of schedules, it is almost inevitable. In the next section we provide empirical evidence of such behavior. In this section, we study the welfare implications of ironing, leaving the welfare implications of spotlighting for a later draft.

We posit a model with responders who fall into two types, depending on where they would land on the schedule. We call them respectively, HI and LO, for those who buy (earn) a large amount and those who buy (earn) a small amount. We begin by studying the behavior of a profit-maximizing monopolist. We then turn to the case of a Ramsey-pricing public utility and of a social-welfare-maximizing tax authority. For simplicity, for the monopolist and Ramsey-pricing cases, we shall assume that goods are produced at constant marginal cost, and that there are no economies of scale on the consumption side (e.g., in delivery).

We consider schedules with two linear segments, and shall focus on situations where optimal schedules are convex. Such convexity could arise for efficiency reasons, e.g., because

HI's elasticity of demand is lower, or to meet distributional concerns, e.g., in a progressive optimal income tax. There are, of course, situations where schedules are concave, as they are when quantity discounts are offered. For such situations, not considered here, ironing usually produces opposite results.

Profit-Maximizing Monopolist

We assume that the monopolist is limited to setting a price schedule in which the first K units cost p_1 dollars each and all subsequent units cost p_2 dollars each. While more elaborate schedules might be optimal under some models, schedules in the real world tend to have only a limited number of brackets, perhaps because people dislike complexity.[21] Moreover, most of the qualitative results we derive would apply for more elaborate schedules with monotonically increasing or decreasing prices.

Though we assume convexity, there are some arenas where schedules are concave. Economies of scale in producing for a particular customer, say due to delivery costs, would produce them (we continue to assume constant costs). Concave schedules will also arise where large buyers have the more elastic demand. This might occur for technological reasons; the larger buyer may have lower transactions costs in switching to an alternative supplier. It could

[21] One of us (RJZ) hired an architect several years ago. The architect usually charged $100 an hour for his services or a flat fee per job. Knowing the lessons of contract theory, RJZ proposed a contract that was a combination: $2000 plus $60 per hour, the architect's presumed opportunity cost. The architect accepted this contract, and it worked splendidly, with the architect making just the appropriate effort on such matters as design specificity, and RJZ imposing the right number of change orders. A couple of years later RJZ saw the architect and asked him what he had thought about the unusual contract. The architect said it was terrific, but that he had not used such a contract in any other of his projects because it was too complicated. RJZ hired him several years later. He paid him a flat $5000 for the new project.

also simply be that tastes are such that people with the greater taste for the good also happen to have higher elasticities.[22]

Our convexity assumption requires that $p_2 > p_1$. This rising-price case is most salient in the tax policy applications we turn to later in the paper, and therefore allows for the most direct comparisons across the three models. Such schedules are rarely observed in segments of the real world where big users could easily break up their use to curb their costs. However, they are quite common with utility pricing or taxes, where it would be hard, or illegal, to break one's use or earnings into little pieces.

Figures 2a and 2b illustrate the profit maximizer's problem. Figure 2a shows the budget constraint for consumers who must allocate income between consumption of the good produced by the monopolist and consumption of all other goods, with both goods measured in dollars. The first K units of the monopolist good are available price at price p_1 and any remaining units can be purchased at a higher price p_2. We assume that there are two consumers who differ in their taste for the good.[23] Consumer HI receives higher total and marginal utility from the monopolist

[22] Insurance offers an interesting case in which it is often not possible to buy half as much from two sources since there are prohibitions against insuring the same thing twice (or at least against collecting if you do). Hence, people buy all their insurance for a home through a single insurer who in turn can charge more the greater the percentage you insure your house for. The same holds true for mortgages. An 80 percent mortgage costs more than twice a 40 percent mortgage, but you can't buy two separate 40 percent mortgages that will behave like a single 80 percent mortgage. The Rothschild-Stiglitz insurance results, and their push for nonlinear pricing, all hinge around these issues. With life insurance, in contrast, you can buy two smaller policies that have the same impact as one larger policy, and, largely due to negative correlations between income and mortality risk, prices tend to fall with the amount insured

[23] For the purposes of drawing this figure we assume that they have the same income because this allows us to depict them as facing the same budget constraint, but our results do not require them to have the same income.

good than does consumer LO. Thus, if g indexes the quantity of monopolist good consumed, then $U_H(g) > U_L(g)$ and $U_H'(g) > U_L'(g)$ for all g.

As depicted in Figure 2a, the monopolist selects p_1 and a kink point K such that the low type consumes at the kink (point A) and then a second slope p_2 such that the HI type is tangent to the second segment of the budget constraint. In maximizing its profits, the monopolist faces several constraints that apply both when the consumers are rational and when they are schmedulers. First, he is restricted by assumption to a pricing schedule that starts at zero and that rises with quantity consumed. This prevents the monopolist from charging each consumer a lump sum and then pricing at marginal cost. Second, the monopolist must offer a segmented linear schedule, rather than two points. Third, points A and B must be such that both consumers prefer those points to zero consumption. Fourth, there is a no-envy condition. Consumer HI must prefer point B to point A. Finally, observe that in optimizing against type A, the monopolist has two different policy tools, the price and the length of the pricing segment. However, it turns out that it is never optimal to prevent LO from consuming as much as he wants at p_1, as we explain in conjunction with Figure 2b.

Figure 2b shows the solution to the monopolist's profit maximization problem both when confronted with rational consumers and with schmedulers. The vertical axis measures net revenue, i.e., marginal cost is subtracted out. We first consider the solution when consumers are rational. Look at the outcome for LO, and the Feasibility Rational LO curve. This shows how profits to the monopolist from sales to LO vary with p_1. At the right-most end, p_1 is low, quantity demanded is high, but revenues just cover costs. As we move left on the curve, p_1 is rising.

LO will always consume on this feasibility frontier. In other words, the kink point will always occur at the point that LO would choose if offered the opportunity to consume an unlimited quantity at a price of p_1 per unit.[24]

Point S indicates the point where net revenue is maximized, taking into account only the sales to LO. p_{1s} is the slope of the curve from the origin through S. Posit that the optimal outcome is to have LO consume at R and HI at E. The monopolist will set $p_1 > p_{1s}$ (along Feasibility Rational LO) i.e., at a level that is higher than would be optimal if he were optimizing against the low type in isolation. Doing so will allow him to lower p_2, thereby increasing the quantity consumed and profits from the HI type. Raising the price on inframarginal units and lowering them on marginal ones is like imposing a fixed cost on HI and then a lower price beyond that allowing higher profits from HI. But this comes at the cost of lower profits from LO. Point R represents the optimal balancing of profits from HI and profits from LO, and therefore will always be to the left of S, the maximum on the Feasibility Rational LO curve. Note also that the higher p_1 results in lower utility for type LO than if the monopolist were optimizing against LO in isolation.

Now consider the curve labeled "Feasibility Rational HI." This curve, added on to point R, shows how profits to the monopolist from sales to the high-type consumer vary with p_2. Convexity requires that $p_2 > p_1$. When p_2 becomes sufficiently high, i.e., steeper than the

[24]To see this, consider interior point T as a possible kink. The monopolist would secure more from LO by offering the alternative p_1 that runs to R, the point vertically above T on the frontier. This is also the point that LO would choose at this new p_1. The monopolist also gets more from HI. Say that the best the monopolist can do against HI given a kink at T is V. With a kink instead at R, he could receive greater net revenues from HI, e.g., at E.

"Feasibility Rational HI" curve at point R, this consumer prefers the kink point to any point on the p_2 portion of the schedule and consumes at point R. Point E indicates the HI type consumer's consumption at the level of p_{2*} , the value that maximizes the monopolist's profits.

Now we turn to the profit maximizer's problem when consumers are ironing. We first observe that a naive schedule setter who merely varies his two-part price schedule through trial and error will reach the same outcome as a monopolist who carefully calculates his consumers' ironing behavior.[25] The ironing solution has HI perceive only the average price. Hence, his feasibility curve assumes that a price line pivots starting at the origin. His feasibility frontier lies strictly above that for rational HI to the northwest of D. That is because he perceives a lower price at the margin for any amount of revenue raised. Thus, for example, if the HI ironer were offered the price schedule represented by R, with p_{2*} beyond, he could consume at G, a point beyond E. He perceives the price as lower than would a rational HI.

Given that HI is responding as an ironer, the location of the first segment of the schedule does not matter. Thus, it is optimal to move the first segment to S, with a caveat about envy, discussed below. HI will be offered the schedule running from S through F. He consumes at F where his indifference curve is tangent to the price line from the origin through F. The envy caveat applies if HI prefers S to F. The point S_e on the Feasibility Ironing HI curve shows the point where HI is indifferent to S. In this case, F is preferred to S. If it weren't, it would be optimal to move S to the left and F to the right until envy was just eliminated.

---

[25]In contrast, we will see that in the Ramsey pricing and optimal income tax models, the schedule setter cares about social welfare and will therefore need to know the preferences underlying the behavior he observes. He may draw erroneous inferences about preferences if he does not realize that consumers are ironing.

It is readily seen that the monopolist is better off with ironing behavior. He could always offer the optimal rational schedule. Under that schedule, an ironing HI would operate at point G, which offers more net revenue than E, the point produced by the rational HI. Since the monopolist also selects S rather than R as the kink point, LO is definitely better off with ironing. HI, however, is likely to be worse off as an ironer.

Ramsey Pricing Utility

We deal next with the Ramsey pricing model because it is the closest to the profit-maximizing case. The essential difference between the Ramsey pricer and the profit-maximizing monopolist is that the Ramsey pricer minimizes deadweight loss, subject to the constraint that profits cover fixed costs, whereas the monopolist maximizes profits.

We continue to study the convex schedule case where the higher volume user pays a higher per unit charge on marginal units.[26] We assume that the cost of new capacity is above that of old capacity. Therefore, in order to cover fixed costs, we have a rising marginal cost curve for consumers. The rising marginal cost may also help achieve distributional objectives, but that is not our concern here.

---

[26] Duke Power in North Carolina offers such a rate schedule (http://www.dukepower.com/misc/rates/ncrates/NCScheduleRS.pdf). The reverse case, where there are quantity discounts, is equivalent to a public utility that must get outside support, and that is below capacity, hence has low marginal costs. To encourage utilization, as we used to do with electric utilities, we charge $2 for the first 100 units and $1 for everything thereafter. If the 200 unit users perceive that their marginal price is $1.50 per unit, efficiency is lost. Ceteris paribus, their utility will be lower, and fewer of the fixed costs of the utility can be covered.

Figure 3 shows how the Ramsey pricer's solution changes with ironing, assuming that the envy constraint is not binding. The feasibility constraints in the rational case for both consumers are identical to those in the profit maximizing example since the only thing that has changed is the producer's objective. Let points A and B represent the optima, assuming rationality, for the LO and HI type consumers respectively. In other words, these points reflect the inverse elasticity rule. Note that point A lies to the right of the revenue maximizing point because the Ramsey pricer is trying to maximize social welfare, not revenues, and therefore takes LO's utility into account.

It is easy to see that the Ramsey pricer can do better if HI is an ironer. He simply offers a schedule with the second segment going from A to C, where C lies on the line from the origin through B. HI will consume at C. HI is better off, since points strictly better than B – those in the triangle formed by extending horizontal and vertical lines from B to the line connecting A and C – were available to him. In this solution, there is no envy, since $U_H(C) > U_H(B) > U_H(A)$. Moreover, revenues are higher, implying that some could be given back to LO and/or HI. This merely shows how to beat the rational outcome. By adjusting the locations of A and C, the schedule setter with ironing can further improve the outcome, while making sure not to adjust so far that HI prefers A to C.

Optimal Income Tax

We assume that the tax schedule is convex: i.e., that marginal tax rates increase somewhere and decrease nowhere. The presumed justification is that the marginal utility of money is decreasing. We are dealing with a situation where both taxpayers pay positive amounts, though allowing for net negative taxes would merely involve rescaling the axes. Figure 4 shows our analysis with pre-tax income on the horizontal axis and after-tax income on the vertical axis. The scales on the two axes are drawn so that post-tax income equals pre-tax income (the usual 45 degree line) along the steep dotted line ending at F. This makes the diagram easier to read. Assume that the tax schedule depicted with the solid lines is the optimal schedule if the two taxpayers are rational. Thus the L type taxpayer chooses point A and the H type taxpayer chooses point B.[27]

We will now show that there is a superior outcome available if HI is a schmeduler. We draw a straight line from the origin through point B. There is a point, C, on this line that is on the schmeduler's feasibility curve (in other words, the schmeduler's indifference curve is tangent to the average tax rate line at this point) and provides higher utility than at B. In particular, the schmedule setter can get type H to choose this point, by offering a tax schedule with the same tax rate through point A as in the rational case and then setting the second tax rate so that the tax schedule beginning at A goes through point C (the lightly dotted line). With this new schedule, the schmeduler not only has higher utility, but also generates more tax revenue (he has higher pre-tax income and is paying the same average tax rate as at B). Thus, we see that the tax scheme with schmeduling is Pareto superior to the one without.[28]

---

[27] Note that the marginal tax rate for the second, higher, bracket will be below the revenue maximizing one, since H's welfare counts somewhat and the loss in revenue from moving away from the revenue-maximizing point is initially zero.

[28] Though we have shown that ironing behavior allows for a Pareto-superior outcome, the optimal outcome given ironing may not be Pareto superior. For example, if ironing gets rid of most of the deadweight loss associated with taxing HI, the optimal scheme may cut his welfare while substantially raising welfare for LO.

What of the progressivity of the optimal tax schemes under the two scenarios? Can we be confident that schmeduling will lead to a more progressive regime? Part of the difficulty in answering that question is knowing how to measure progressivity. We believe that many answers are possible depending on the progressivity measure employed. We are confident of one result. In comparison with the optimal tax scheme with rational responses, there exists a Pareto-superior schmeduling scheme that simultaneously collects more taxes from HI, has a higher average tax rate imposed on HI, and leaves HI better off. This is achieved at a point slightly below C on the schmeduler's feasibility frontier. He is still better off than he was at B, but pays more in taxes and has a higher average tax rate.

The arguments regarding no-envy conditions and naive schedule setters in the optimal income tax case follow directly from those in the previous two models and we therefore omit discussion of them here.

The results from our geometric presentations are straightforward. Ironing behavior basically eliminates some of the deadweight loss from high marginal prices, or taxes. This implies that when the optimal rational schedule is convex, superior outcomes are available for the monopolist, for the Ramsey pricer, and for the setter of an optimal income tax. In all three contexts, although they may not be chosen, outcomes that are Pareto superior to the rational-responders' outcome are available.

Tax Policy Implication of Schmeduling: Deadweight Loss Effects

With a convex tax schedule, ironers will perceive a tax rate that is lower than the true marginal tax rate. Therefore, they will supply more income (work harder) and the tax system will impose a lesser deadweight loss. To assess the quantitative importance of this effect, we conduct simulations using the 1998 IRS public use sample of tax returns and NBER's internet Taxsim model. We "age" the sample to reflect 2000 income levels and use tax schedules for that year.[29] We follow Feldstein (1999), and calculate deadweight loss using the Harberger-Browning approximation

$$DWL = \frac{1}{2} TY \varepsilon \frac{\tau^2}{1-\tau},$$

where TY is taxable income, ϵ is the elasticity of taxable income with respect to the after-tax share, and τ is the tax rate. We use a value of 0.4 for ϵ, based on the estimates of Gruber and Saez (2002). Since deadweight loss is linear in ϵ, readers who prefer alternative values can easily use them.

We initially assume that taxpayers are of the ironing variant of schmedulers and mistake their average tax rates for their marginal tax rates. Then we ask what would happen if we informed these taxpayers of their true marginal rates. Table II presents our results. In the data, taxpayers have taxable income of $4.233 trillion and pay income tax of $974.7 billion. The Harberger-Browning formula using the perceived tax rates yields deadweight loss of $56.7 billion or 5.8 percent of revenue raised.[30]

---

[29] 2000 is currently the most recent year covered by Taxsim. We dropped a couple dozen observations from the sample for whom Taxsim calculated marginal tax rates below -40 percent or above 50 percent.

[30] For comparability with the main results in Feldstein (1999) these results ignore the payroll tax. Treating the personal income tax as an increment on top of the payroll tax would produce larger deadweight loss estimates.

We estimate that if taxpayers were informed of their true marginal tax rates, taxable income would fall by about 5 percent to $4.020 trillion and revenue would fall by about 6 percent. Deadweight loss would rise to $109 billion or 11.9 percent of revenue raised.[31]

A similar calculation can be done for the marginal excess burden of taxation. Consider a 10 percent increase in all marginal tax rates (to illustrate, a 20 percent marginal tax rate would become 22 percent). Under the schmeduling model, revenue increases by $82.5 billion and deadweight loss increases by $13.9 billion for a marginal excess burden of 17 cents per $1 of additional revenue. In the rational model, revenue increases by $68.9 billion and deadweight loss increases by $27.2 billion for a marginal excess burden of 39 cents.[32]

Tax Policy Implications of Schmeduling: Interpreting Existing Natural Experiment Evidence

It is worth emphasizing that existing estimates of the elasticity of taxable income come from studies of behavioral responses to tax changes. These elasticities are calculated by dividing the change in behavior by the change in after-tax share. The changes in after-tax shares in these calculations are based on marginal tax rates. Changes in after-tax shares calculated based on perceived tax rates (i.e., average tax rates) would be smaller, resulting in larger elasticities

---

[31] Our estimates for the rational case are quite similar to Feldstein's (1999) estimates. Feldstein, using an elasticity of 1.04, estimates that DWL from the personal income tax in 1994 was 32.2 percent of revenue. Multiplying our DWL estimate by (1.04/0.4) produces an estimate of 30.9 percent. Interestingly, our estimate that DWL under schmeduling is 48 percent lower than under the rational case is very similar to that of de Bartholeme (1995) who does an illustrative calculation for a representative worker with mean earnings using parameter estimates from Hausman (1981) and finds that DWL falls by 43 percent when taxpayers substitute average rates for marginal rates.

[32] With a taxable income elasticity of 1.04, the marginal excess burden is $1.89 per dollar of revenue raised under the rational model and 52 cents under the schmeduling model.

relative to the perceived change in after-tax shares. Therefore, it might be appropriate to use larger elasticities in the calculations above. This would produce higher estimates for the deadweight loss in the current U.S. tax system. However, it would not alter the estimates of the relative amount of deadweight loss under schmeduling and the rational model (since we would simply be using higher elasticities in both calculations).

There is one piece of natural experiment evidence that is potentially inconsistent with the predictions of our ironing model. Feldstein (1995), Eissa (1995), and Auten and Carroll (1997) provide evidence that high-income taxpayers increased their incomes substantially in response to the reduction in marginal tax rates from the Tax Reform Act of 1986 (TRA86). Since TRA86 was designed to be distributionally neutral, it affected average tax rates only slightly at most income levels. Thus, our ironing model would predict little behavioral response to this tax reform.

This evidence does not lead us to abandon our schmeduling model. First, we have argued that while many individuals schmedule, some individuals are rational. The very high-income taxpayers studied in the TRA86 literature are likely to be among the most rational of all taxpayers, both as players for themselves, and because they hire expert advisors. Thus they are the ones whom we would least expect to observe engaging in schmeduling. Second, there is a large literature by Slemrod (1990), Goolsbee (2000) and others suggesting that the TRA86 evidence is a product of widening income inequality, and the shifting of income between the corporate and individual income tax bases, not of behavioral responses to taxation. If we are able to accumulate evidence demonstrating that taxpayers often engage in ironing, this will, in our view, increase the probability that these alternative interpretations of TRA86 are correct.

IV. Empirical Tests

This section conducts two empirical tests of schmeduling. The first uses data from before and after the 1998 introduction of the child tax credit to test for ironing. The second uses data from the San Diego food stamp cash out experiment to test for spotlighting. In addition to providing empirical evidence on whether schmeduling is occurring in these two instances, these examples illustrate the kinds of conflicting predictions that could allow one to distinguish between the rational and schmeduling models more generally.

A. 1998 Introduction of the Child Credit

Beginning in 1998, U.S. taxpayers with children could claim a $500 per child tax credit. In most cases, this credit was not refundable. Thus a taxpayer with $500 or less of tax liability could not take advantage of the full credit.[33] Figure 5 illustrates the impact of the introduction of the child credit on marginal and average tax rates at different income levels for a taxpayer in 1998. For the purpose of this figure, the taxpayer is assumed to be married with two qualifying children,[34] claim the standard deduction, and have only wage income.

Before 1998, taxpayers with incomes between about $18,000 and $25,000 owed income tax, and therefore faced a 15 percent marginal tax rate. But beginning in 1998, the child credit eliminated the entire tax liability for these taxpayers and reduced their marginal tax rate from the federal personal income tax to zero. Thus, their marginal tax rate fell by 15 percentage points. All taxpayers with income above $18,000 experienced a reduction in tax liability and therefore a

[33] The credit was partially refundable for some taxpayers with three or more children.

[34] Only children age 17 and under can qualify a taxpayer for the child credit.

reduction in average tax rates. The reduction in average tax rates grows with income from $18,000 until the point at which a taxpayer can use the entire $1000 (2 children X $500) credit. After that point the reduction in average tax rates falls gradually as the reduction in tax liability remains $1000, but the denominator in the average tax rate calculation, the person's income, rises.

The rational model would predict that the reduction in marginal tax rates would induce people with income between $18,000 and $25,000 to increase their earnings. We would also expect to see some bunching at $25,000, the point at which the marginal tax rate jumps from zero to 15 percent after the reform. Because income effects are generally thought to be close to zero, we would expect to see little effect on the earnings of people with incomes above $25,000, and any effect would be a reduction in earnings due to the income effect.

In contrast, the schmeduling model predicts increased work by anyone whose average tax rate fell – everyone with income above $18,000. In particular, we would expect to see increased work by people with incomes above $25,000 and no bunching at that point – two predictions that depart from those of the rational model.

To test these predictions, we use data from the 1997 and 1999 IRS public use Statistics of Income tax files. These files are based around random samples of individual tax returns, but are blurred in various ways to protect taxpayer confidentiality. Our basic approach is to examine whether the change in the distribution of taxpayers by income between 1997 and 1999 looks more like what would be predicted by the rational model or by the ironing model.

In order to be able to predict how individual behavior will change in response to the change in budget constraints, we need to model people's preferences. In particular, given our interest in the bunching of taxpayers at kink points, we cannot simply predict the change by multiplying the percentage change in the after-tax share times an elasticity. We follow Diamond (1998) and Saez (2002) in assuming that preferences take the quasilinear form:

(1) $$U = C - \frac{L^{1+k}}{1+k} \ ,$$

where C is consumption and L is labor effort. Under this specification, there is a single preference parameter, *k*, which is equal to $1/\epsilon$, where ϵ is the labor supply elasticity. There is no income effect in this model. We view this model as the simplest structural analog to elasticity calculations with a constant elasticity.

The Rational Model

With rational taxpayers, the first order condition from this model is

(2) $(L^*)^k = w(1-t^*)$

where t* is the tax rate on the segment of the budget constraint where the taxpayer's optimum lies and w is the taxpayer's wage. By multiplying both sides by w^k and rearranging, it is possible to express w as a function of observable quantities (We observe pre-tax income, wL*, in our data set; given wL* we know t* since we know the tax schedule that the taxpayer faces.):

(3) $$w = \left[\frac{(wl^*)^k}{1-t} \right]^{\frac{1}{k+1}}$$

Thus given an elasticity, ϵ, and a distribution of income under a known tax schedule, we can derive the wage distribution and simulate the distribution of income under any other budget constraint.

This approach encounters two complications. First, if a taxpayer locates exactly at the kink point between segments with tax rates of t_a and $t_{b,}$ we do not know his exact wage, only that it lies between the two values that would occur from substituting t_a and t_b into the equation above. In practice, only a couple of people in our data set locate exactly at a kink and we randomly assign those people to a wage between the two implied by t_a and $t_{b,}$ The second complication is more significant. Because there are almost no people exactly at the kink, the wage distribution that is implied by taking observed income and plugging it into the equation derived above from the first order condition implausibly has a noticeable gap in it with no one (except the people exactly at the kink) at wages between the wage implied by t_a and the wage implied by t_{b}.[35] Our approach is to assume *in our rational model* that taxpayers are uncertain about exactly where the kink is located and so a taxpayer near kink point, *k*, chooses hours of work, *L*, to maximize expected utility where the expectation is over the possible locations of the kink point:

$$(4)\ E[U(L)] = \int_{-\infty}^{wL} p(k)\left[wL - kt_L - t_H(wL - k)\right]dk + \int_{wL}^{\infty} p(k)\left[wL - t_L wL\right]dk - \frac{L^{1+K}}{1+K}.$$

[35] The complication is not simply that this gap is implausible. Because the gap is dependent on where the kinks are, the implied wage distribution will change when the tax schedule changes, a feature that would be inconsistent with the estimation approach we describe below.

In this equation, t_L is the tax rate on the segment of the tax schedule just below the kink point at k, t_H is the tax rate on the segment of the tax schedule that starts at k, and $p(k)$ is the probability distribution function of k. The first term in the equation is after-tax income (consumption) when the kink point is below wL and the second term is after-tax income when the kink point is above wL . If we further assume that the uncertainty is distributed normally with mean at the true kink point and standard deviation $\sigma_{k,}$ the expression can be rewritten as:

$$(5) \qquad E[U(L)] = \Phi(wl)wL(1-t_H) + \int_{-\infty}^{wL} p(k)k(t_H - t_L)dk + (1-\Phi(wl))wL(1-t_L) - \frac{L^{1+K}}{1+K}$$

where Φ is the normal CDF. The first order condition is (surpressing the argument of Φ):

$$(6) \quad 0 = \Phi' w^2 L(1-t_H) + w(1-t_H)\Phi + w(1-t_L)\Phi - \Phi' w^2 L(1-t_L) - w(1-t_L)\Phi - L^K$$

Multiplying through by L and replacing wL with Y (the income observed in the data), we can express L as a function of observed quantities:

$$(7) \qquad L = \left[Y(1-t_H)\Phi + Y(1-t_L)(1-\Phi) + \Phi' Y^2(1-t_H) - \Phi' Y^2(1-t_L)\right]^{\frac{1}{K+1}}.$$

We can then determine w as Y/L. Thus, in our rational model, we have two preference parameters K (the inverse of the labor supply elasticity) and σ_k, *(*the amount of uncertainty around the location of kink points).

It is worth emphasizing that by introducing this uncertainty about the location of the kink point into the rational model we are, in essence, letting the rational taxpayers engage in a bit of schmeduling. In particular, it can be shown that the marginal tax rate perceived by the taxpayer near the kink point between t_a and t_b in this model is a weighted average of t_a and t_b:

(8) $$mtr(Y) = t_a(1 - \Phi(Y)) + t_b\Phi(Y).$$

The implication is that if we find evidence of schmeduling behavior when we test it against this lenient variant of the rational model, it will be even more powerful evidence that schmeduling is occurring.

The Ironing Model

In the ironing model we assume that taxpayers have the same preferences as they do in the rational model (equation 1). However, they respond to average tax rates rather than marginal tax rates. They therefore choose hours, L, to satisfy the following tangency condition

(9) $$L^{*K} = w\left(1 - \frac{T(wL^*)}{wL^*}\right).$$

instead of the conventional one. In this equation, *T(wL*)* is total taxes due at income *wL**. Given K and an observed income level, this equation can be used to find *w*.[36] There is no need to introduce uncertainty about the location of the kinks into the ironing model since the wage distribution implied by ironing does not have a discontinuity at each kink point as it did in the rational model.

[36] In the rational model there is an explicit analytical expression for w. In the ironing model, a simple numerical procedure is needed to solve for w.

Econometric Model

We use the model described above to simultaneously estimate the elasticity, ϵ, and, *s*, the share of taxpayers who are schmedulers. Let S=1 indicate a schmeduler and S=0 indicate a rational taxpayer. Pre-tax income in the two years is given by:

$$
(10)\quad \begin{aligned} m_{97} &= S\mu_S(f_{97},\varepsilon,w) + (1-S)\mu_R(f_{97},\varepsilon,w,\sigma_k) + \nu_{\sigma_O} \\ m_{99} &= S\mu_S(f_{99},\varepsilon,w(1+\gamma)) + (1-S)\mu_R(f_{99},\varepsilon,w(1+\gamma),\sigma_k) + \nu_{\sigma_O} \end{aligned} .
$$

Income, *m*, is determined by the functions μ_S and μ_R which generate an earnings level for a schmeduler and a rational taxpayer respectively as a function of the tax schedule in each year, f_{97} and f_{99}, the elasticity, ϵ, and the taxpayer's wage. In the rational case, there is an extra parameter, σ_k, which reflects the amount of uncertainty about where kinks are located. γ is a parameter that describes the amount of nominal wage growth between the two years. $\nu_{\sigma,}$ is a Gaussian random variable representing optimization error, with standard deviation equal to $\sigma_{O.}$

As we have discussed in the sections on the rational and schmeduling models above, $\mu_S(f,\epsilon,w)$ and $\mu_R(f,\epsilon,w,\sigma_k)$ are invertible in their third arguments, with inverse functions $w_s(f,\epsilon,\mu_S)$ and $w_R(f,\epsilon,\mu_R,\sigma_k)$. Let $w(f,\epsilon,\mu_S,\mu_R,\sigma_k,S)=S\, w_s(f,\epsilon,\mu_S)+(1\text{-}S)w_R(f,\epsilon,\mu_R,\sigma_k)$. Our first key identifying assumption is:

$$
(11)\quad \begin{aligned} & w(f_{97},\varepsilon,m_{S97}-\nu_{\sigma_O 97},m_{R97}-\nu_{\sigma_O 97},\sigma_k,S) = \\ & \frac{1}{1+\gamma} w(f_{99},\varepsilon,m_{S99}-\nu_{\sigma_O 99},m_{R97}-\nu_{\sigma_O 99},\sigma_k,S). \end{aligned}
$$

In other words, we assume that the wage distributions in the two years differ only by the rate of wage growth, γ, and that, apart from wage growth, changes in the income distribution are due to changes in the tax schedule or to random variation in optimization error.[37] Our second key identifying assumption is that each observation in the data set has the same probability, *s*, of being from a schmeduler and a probability of (1-s) of being a non-schmeduler. In other words, the probability of being a schmeduler is independent of income.[38]

We employ a minimum distance method to estimate our parameter values. We take γ, σ_k, and σ_O as known and choose $(\hat{\varepsilon}, \hat{s})$ to minimize:

$$(12) \quad \sup\left|F_{99}(\varepsilon, s) - F_{97}(\varepsilon, s)\right|,$$

the Kolmogorov-Smirnov distance between the wage distributions for the two years. Here $F_{99}(\epsilon,s)$ and $F_{97}(\epsilon,s)$ are the cumulative distribution functions of wages implied by (ϵ,s) for the 1999 and 1997 incomes and tax schedules, respectively. We use a two-dimensional grid search to find the parameter values that minimize (12).[39] For each combination of potential parameter values, we simulate 100 draws from the optimization error distribution for each taxpayer in our data set and include all 100 in estimating the CDF for the wage distribution. This is necessary

---

[37] We thank Alberto Abadie for suggesting this formulation of our first identifying assumption.

[38] An alternative, possibly more attractive, assumption would be to assume that the schmeduling probability is independent of the wage distribution. We plan on exploring this assumption in a revision.

[39] We intend to estimate the three remaining parameters as well. While doing so is straightforward, we have not yet established that our numerical minimization routine performs adequately in that five parameter case, so we present results only for the simpler case here.

because the optimization error enters nonlinearly in the wage function. We form our 95 percent confidence intervals by bootstrapping. Specifically, we created 500 data sets by randomly sampling from our data with replacement (separately for the two years of the data).

Data and Results

Recent work in public finance emphasizes that taxpayers respond to tax incentives in many ways, not simply by altering their hours of work. For example, they can alter the form of their compensation between cash wages and untaxed benefits, they can adjust their itemized deductions, or they can change their compliance behavior. The literature that has studied these behavioral responses has generally used taxable income as the relevant concept of income. Feldstein (1999) shows that this is the theoretically correct concept for calculating deadweight loss. To meet our purposes, we need a broader concept of income because the taxpayers we study include many with very low taxable incomes whose gross incomes can be large multiples of their taxable incomes. We do not want to simply use adjusted gross income as our broader income concept because that would not allow us to study behavioral responses that occur through itemized deductions. Therefore, we define income as adjusted gross income minus itemized deductions above the standard deduction. In essence, we are treating the sum of the standard deduction and personal and dependent exemptions as a tax bracket with a zero tax rate, rather than subtracting these amounts from adjusted gross income in defining income. Although our model specified in equation (1) and the following equations is specified in terms of a labor supply choice, we view it as a simple structural approximation to methods that apply a constant elasticity to taxable income. Thus, the elasticity estimates we obtain should be seen as comparable to the literature on broader behavioral responses to taxation.

We limit our sample to married couples with at least two children. For these households, the introduction of the child tax credit provided a new tax credit of at least $1000.[40] Our tax model consists of the federal income tax (including the EITC) and the OASDI and HI payroll taxes (modeled as a proportional tax of .0765). We ignore state taxes.

Before turning to results, it is worth asking whether this change in taxes is large enough to allow us to distinguish between the two models. Figures 6A and 6B demonstrate that the two models create noticeably different income distributions. Figure 6A takes the 1997 sample and shows how introducing the child credit in 1997 would have changed the income distribution in that year under the rational model and under the schmeduling model.[41] In particular, it shows the change in income from introducing the child tax credit plotted against 1997 income with an elasticity of 0.40. Under the rational model, taxpayers with incomes between about $10,000 and $25,000 increase their income. But there is no change for taxpayers who are more than a few thousand dollars above the new kink at roughly $25,000 (remember that due to uncertainty about the kink point, some taxpayers above the kink are affected by the change even in the rational model). Under the schmeduling model, taxpayers just above the kink have relatively large responses and the dollar response remains relatively constant for the full range of the income

---

[40] In claiming the child credit, taxpayers could claim only those children age 17 and below. We cannot implement this restriction in our data because we do not observe the age of the children in the 1997 data. Therefore, we simply assume that all taxpayers with dependent children living at home claim the child credit.

[41] For these simulations, the standard deviation of the kink error is set at $3000 and the optimization error is set at $2000.

distribution.[42] Figure 6B shows the CDF of income under the two models. There is a noticeable difference in the schmeduling and rational income distributions between roughly $18,000 and $28,000. In particular, the rational taxpayers who increase their incomes in response to the reduction in marginal tax rates produce a deficit of taxpayers between about $18,000 and $25,000 and the rational CDF is therefore below the ironing CDF over this range. Then the concentration of taxpayers around the kink point in the rational model causes a sudden jump in the CDF around $25,000. Just above this income level is where taxpayers experience the largest reductions in average tax rates. Because of this, the schmeduling CDF falls below the rational CDF as schmedulers who otherwise would have been in this range increase their incomes in response to the reduction in their average tax rates.

As a further check on our model and econometric technique we have simulated data under our model with known values of ϵ and s to see if our estimation technique is capable of recovering the true parameters when the model is correctly specified. We found that roughly 85 percent of the time our technique recovered the true parameters and the remainder it came up with values that were fairly close to the true values.

Table III shows our results from estimating our model on the actual data. We set γ at .05, the average nominal income growth in this income range between the two years studied. We set σ_O=$2000 based on the evidence in Saez (2002). We set σ_k=$4000, which simulations show is the level necessary to produce reasonable looking wage distribution (i.e. with no sharp reduction

[42] In percentage terms, the increase in income falls with income above the kink at the start of the 15 percent bracket. Therefore, if the income growth term, γ, is multiplicative, we can separately identify income growth and schmeduling.

in the density at kink points) for a wide range of elasticities.[43] Given these assumptions, we estimate that the elasticity of income with respect to the after-tax share is 0.22 and that 45 percent of taxpayers are schmedulers. The bootstrapped standard errors suggest a reasonably narrow range around these estimates. Figure 7 uses the 1997 data to predict the 1999 wage distribution using these parameter values under first the schmeduling model (in green) and then the rational model (in blue). It also plots the true distribution of the data (in red). For incomes below $28,000, the predictions of the rational model line up well with the true data. But the sharp jump in the CDF around $28,000 that is predicted by the rational model is completely absent in the true data and at income levels between about $29,000 and $35,000 the schmeduling CDF is closer to the true CDF than the rational CDF is. It is likely that the reason that the share of schmedulers was estimated to be close to 50 percent is that each model fits the data more closely over part of the income range. Moreover, the main prediction of the rational model – the bunching of taxpayers around the new kink point – is not borne out in the data, even though the version of the rational model we estimate allows for uncertainty about the location of the kinks and therefore predicts less bunching than a more extreme version of the rational model would..[44]

B. The Within-month Pattern of Food Consumption by Food Stamp Recipients

Our second empirical test addresses the spotlighting model, looking at food consumption for those on food stamps. Our principal hypothesis is that inframarginal food stamp recipients –

---

[43] The minimum necessary value of σ_k rises with the elasticity. We have also estimated the model under this assumption and we obtained similar results.

[44] The 1979 increase in the level of earnings subject to the OASDI payroll tax provides another natural experiment in which the predictions of the ironing model and the rational model differ. We are working on estimating our econometric model to that policy change as well.

those who spend more than the food stamp amount on food during the month – nonetheless view the cost of spending a dollar of food stamps on food as less than a dollar in terms of lost consumption of other goods. Before they have exhausted their food stamps, they respond, in part, to the local price of spending a dollar of food stamps, which they perceive as far less than one dollar (pure spotlighting would infer a price of zero). That is, during the early period of the month, they fail to realize that the cost of a marginal dollar of food consumption is one dollar. Then, after they have exhausted their food stamps and must spend cash for food, they perceive the true marginal cost of their food. This would imply that actual food consumption would fall through the month. A priori, finding evidence of such behavior might seem unlikely, since food stamp recipients get to play the game many times; every month they get food stamps.

The empirical prediction of the spotlighting model for food consumption by food stamp recipients is that food consumption should decline after recipients have exhausted their stamps. No such pattern would occur with rational consumers. If we had data only on food consumption by food stamp recipients, however, it would be hard to distinguish spotlighting from several other hypotheses that could also explain declining food consumption during the month after benefit payment. For example, myopic consumers might also have declining consumption throughout the month, as would consumers who consume a constant minimum level throughout the month but run out of income at the end of the month (and perhaps consume at a higher level at the beginning of the month because they are hungry from running out of income at the end of the previous month).[45]

---

[45]We thank Jesse Shapiro for helping us with these data. Shapiro (2003) provides evidence suggesting that food consumption by food stamp recipients is consistent with hyperbolic discounting.

In order to isolate the pure spotlighting effect, we use data from the 1989-1990 San Diego food stamp cash out experiment. In this experiment, a random sample of the food stamp caseload had their food stamp checks replaced by an equal amount of cash benefits. Mathematica evaluated this experiment for the Department of Agriculture and collected data on food consumption from about 600 randomly selected food stamp recipients and from another 600 who had had their benefits cashed out.[46] Mathematica's evaluation estimated the impact of cash out on average food use at home and concluded that it reduced food use at home by between 5 and 8 percent (Ohls et al 1992) . Whitmore (2002) reevaluated these data, focusing on the difference between marginal and inframarginal consumers as a way to estimate the deadweight loss from paying in food stamps rather than in cash. She found that inframarginal consumers do not alter their food consumption when converted to cash, but that "distorted" food stamp recipients do reduce their consumption, and that they value their stamps at only 80 percent of face value.[47]

These data are useful for testing spotlighting because they allow us to distinguish our theory from other possible explanations for declining food consumption during the month after benefit payment. In particular, for the myopic, hyperbolic discounting, and food insecurity theories, there is no reason why the slope of consumption throughout the month should change if

---

[46] Although we describe the cashed out recipients as receiving "cash," they actually received their payment in a check.

[47] A similar experiment occurred around the same time in Alabama. However, as Whitmore (2002) discusses, the data from this second experiment are less reliable, both because of its limited duration and because caseworkers coached cash recipients not to change their food consumption. In addition, we have been unable to learn the institutional details about AFDC payout dates in Alabama during this period and therefore cannot estimate our model on these data.

stamps were converted to cash. If spotlighting is occurring, in contrast, we would expect to see a greater decline in consumption throughout the month for consumers who are paid in food stamps than for consumers who are paid in cash.

Our base sample consists of all sample members for whom there is complete food consumption data (a total of 541 receiving food stamps and 537 receiving cash). However, we want to restrict the sample to inframarginal consumers – those who would consume more than their food stamps if they received their payment in food stamps.. To identify these consumers we run a probit regression in the food stamp sample with a dependent variable that is one if the household consumed more than their food stamps during the survey month.[48] The independent variables are indicators for household size and a fourth degree polynomial in food stamp benefit level (the food stamp benefit levels in this study are obtained from administrative payment records). We use this estimated equation to predict a probability of being inframarginal for households in both the cash and the food stamp group. Then we limit our sample to households with a predicted probability of 0.95 or above. This results in a sample of 349 in the cash group and 366 in the food stamp group. In this restricted sample, over 98 percent of households in each group spent more on food than they received in food stamps and the percentages are nearly identical in the two groups.

---

[48] Since food data is collected for only one week in this study we scale up the food consumption by # days in month/7.

We use two dependent variables in our regressions testing our hypothesis: total dollars of food consumed during the survey week and total calories of food consumed during the survey week, both measured in natural logarithms.[49] Our OLS regression specification is:

(13) ln(food consumption)=β_0+β_1(paid in cash)+β_2 (days since last AFDC check)+β_3(days since last AFDC check) x (food stamps paid in cash)+β_4(days since last food stamp check) x (paid in food stamps) +γX +ϵ.

AFDC benefits were paid on the first of the month. Cash recipients received their cashed-out food stamp payment as part of the same check. In contrast, food stamp recipients received their food stamp checks at roughly uniform intervals throughout the month. Thus, β_2 should capture the relationship between days since AFDC receipt and food consumption for the entire sample. β_3 should capture any differential relationship between days since AFDC receipt and food consumption for those in the cashed-out group who also received their food stamp payment with their AFDC check. β_4 is the key parameter for testing our hypothesis. In measures the relationship between the number of days since food stamp receipt and food consumption for those receiving food stamps. X is a set of covariates that are assumed not to be affected by the experiment: dummy variables for calendar month and household size and interactions between these variables and treatment group.[50]

---

[49] The Mathematica data set contains an elaborate set of variables measuring the nutritional value of all of the foods consumed by each household.

[50] We have explored the sensitivity of our results to including a more elaborate set of covariates. As would be expected given the independence between covariates and treatment assignment (due to the random assignment), the results are quite insensitive to including additional covariates.

Table 4 shows the results from our regressions. The dependent variable is the log of food consumption, either measured in dollars (column 1) or in calories (column 2). First, consider time since AFDC receipt. The point estimates for days since last AFDC check in the first column represents a 0.27 percent per day reduction in consumption or a total decline of about 8 percent over the month. The coefficients on the interaction between AFDC receipt and payment in cash are positive. However, neither the coefficient on days since AFDC check receipt or on the interaction of days since AFDC receipt with cash payment are statistically different from zero.

Responses to time since food stamp receipt for those paid in food stamps exhibit a quite different pattern. The coefficient on the interaction between "days since food stamp payment" and "payment in stamps rather than checks" is much larger, about 0.8 percent per day, and is statistically significant at the 95 percent level. Thus, individuals receiving payments through food stamps reduce food consumption by about 24 percent over the month. The estimates are similar whether consumption is measured in dollars or in calories. Overall, these results suggest that spotlighting is a quantitatively important aspect of the food consumption behavior of food stamp recipients.

VI. Conclusion

We have argued that schmeduling is likely to be a common form of economic behavior, that it arises in substantively important areas of economic decision making, and that the welfare effects of people responding to schmedules rather than to their true schedules are likely to be large and to have significant policy implications. Moreover, because the conditions that give rise to schmeduling are found in many economic environments, empirical evidence on how people respond to schedules in one environment will help us predict how people will react when faced by schedules in other environments. We have provided two empirical tests of schmeduling. In both cases, the data suggest that a significant amount of schmeduling behavior takes place.

References

Auerbach, Alan J. and Laurence J. Kotlikoff (1985). The Efficiency Gains from Social Security Benefit-Tax Linkage. NBER Working Paper 1645.

Auten, Gerald, and Robert Carroll (1997). The Effects of Income Taxes on Household Behavior. Office of Tax Analysis, U.S. Department of the Treasury.

Bises, Bruno (1990). Income Tax Perception and Labour Supply in a Sample of Industry Workers, *Public Finance*, 1:3-17.

Brannas, Kurt and Niklas Karlsson (1996). Estimating the Perceived Tax Scale within a Labor Supply Model. *Economic Letters*, 52:75-79.

Break, George (1957). "Income Taxes and Incentives to Work: An Empirical Study," *American Economic Review*, 47:529-549.

Brown, C.V. (1968). Misconceptions About Income Tax and Incentives. *Scottish Journal of Political Economy*, February: pp1-12.

de Bartolome, Charles A.M. (1995). "Which Tax Rate Do People Use: Average or Marginal?" *Journal of Public Economics*, 56:79-96.

de Villiers, Peter A. and Richard J. Herrnstein. (1976). "Toward a Law of Response Strength," *Psychological Bulletin*, 83:1131-1153. Reprinted in Richard J. Herrnstein, *The Matching Law: Papers in Psychology and Economics* , edited by Howard Rachlin and David I. Laibson, (Cambridge, Harvard University Press), 1997.

Diamond, Peter (1998). "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 88, 83-95.

Dickert, S., S. Hauser, and J.K. Scholz (1994). Taxes and the Poor: A Microsimulation Study of Implict and Explicit Taxes. *National Tax Journal* 47: 76-97.

Eissa, Nada (1995). Taxation and Labor Supply of Married Women: the Tax Reform Act of 1986 as a Natural Experiment. NBER Working Paper No. 5023.

Feldstein, Martin (1995). The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act. *Journal of Political Economy*, 103:551-572.

Feldstein, Martin (1999). Tax Avoidance and the Deadweight Loss of the Income Tax, *Review of Economics and Statistics*, 81: 674-680

Fujii, Edwin T. and Clifford B. Hawley (1988). On the Accuracy of Tax Perceptions, *Review of Economics and Statistics*, 344-347.

Giannarelli, L. and E. Steuerle (1995). It's Not What You Make, It's What You Keep: Tax Rates Faced by AFDC Recipients. Paper presented at the APPAM Annual Research Conference.

Goolsbee, Austin (2000). It's Not About the Money: Why Natural Experiments Don't Work On the Rich. In Slemrod, Joel (editor), *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*. Cambridge University Press, Cambridge.

Grogger, Jeff and Charles Michalopoulous (1999). Welfare Dynamics Under Time Limits. NBER Working Paper 7353.

Gruber, Jon and Emmanuel Saez (2002). "The Elasticity of Taxable Income: Evidence and Implications," *Journal of Public Economics*, 84:1-32.

Hausman, Jerry (1981). "Labor Supply" in Henry Aaron and Joseph Pechman (editors) *How Taxes Affect Economic Activity* (Brookings, Washington, DC).

Heckman, James (1983). "Comment on Stochastic Problems in the Simulation of Labor Supply" in Martin Feldstein (editor) *Behavioral Simulation Methods in Tax Policy Analysis*.

Heim, Bradley T. and Bruce D. Meyer (2003). Work Costs and Nonconvex Preferences in the Estimation of Labor Supply Models. NBER Working Paper 9429.

Barthold, Thomas A., Thomas Koerner, and John F. Navratil (1998). Effective Marginal Tax Rates Under the Federal Individual Income Tax: Death by One Thousand Pin Pricks? *National Tax Journal*, 51: 553-564.

Herrnstein, Richard J. (1961). "Relative and Absolute Strength of Response as a Function of Frequency of Reinforcement." *Journal of the Experimental Analysis of Behavior*, 4: 267-272. Reprinted in Richard J. Herrnstein, *The Matching Law: Papers in Psychology and Economics* , edited by Howard Rachlin and David I. Laibson, (Cambridge, Harvard University Press), 1997.

Herrnstein, Richard J. and Heyman, G. M.(1979). "Is Matching Compatible with Reinforcement Maximization on Concurrent Variable Interval, Variable Ratio?" *Journal of the Experimental Analysis of Behavior*, 31, 209-223.

Herrnstein, Richard J. (1982). "Melioration as Behavioral Dynamism." Reprinted in Richard J. Herrnstein, *The Matching Law: Papers in Psychology and Economics* , edited by Howard Rachlin and David I. Laibson, (Cambridge, Harvard University Press), 1997.

Herrnstein, Richard J. (1990). "Rational Choice Theory: Necessary but Not Sufficient." *American Psychologist*, 45: 356-367. Reprinted in Richard J. Herrnstein, *The Matching*

Law: Papers in Psychology and Economics , edited by Howard Rachlin and David I. Laibson, (Cambridge, Harvard University Press), 1997.

Herrnstein, Richard J. (1991). "Experiments on Stable Suboptimality in Individual Behavor." *American Economic Review* (May), 81: 360-364.

Herrnstein, Richard J. and Drazen Prelec (1991). "Melioration: A Theory of Distributed Choice," *Journal of Economic Perspectives*, 5:137-156.

Herrnstein, Richard J. and Vaughan, w., Jr. (1980). "Melioration and Behavioral Allocation." In J.E.R. Staddon (ed.), *Limits to Action: The Allocation of Individual Behavior* (New York: Academic Press).

Kling, Jeffrey, Jeffrey Liebman, and Larry Katz (2003). Bullets Don't Got No Name: Consequences of Fear in the Ghetto. Forthcoming in Thomas Weisner (editor), *Discovering Successful Pathways*, University of Chicago Press

Kotlikoff, Laurence J. (1995). Privatization of Social Security: How It Works and Why It Matters, *Tax Policy and the Economy*, 10: 1-32.

Liebman, Jeffrey (1996). The Impact of the Earned Income Tax Credit on Labor Supply and Taxpayer Compliance. Harvard University Phd Thesis.

Liebman, Jeffrey (1998). The Impact of the Earned Income Tax Credit on Incentives and the Income Distribution, *Tax Policy and the Economy*, volume 12.

MaCurdy, Thomas, David Green, and Harry Paarsh (1990). "Assessing Empirical Approaches for Analyzing Taxes and Labor Supply," *Journal of Human Resources*, 25: 415-490.

Mazur, J.E. (1981). "Optimization Theory Fails to Predict Performance of Pigeons in a Two-Response Situation." *Science*, 214: 823-825.

Pollack, Harold and Richard Zeckhauser (1996). Budgets as Dynamic Gatekeepers, *Management Science*, 42: 642-658.

Rachlin, Howard and David I. Laibson (1997). "Introduction," in Richard J. Herrnstein, *The Matching Law: Papers in Psychology and Economics* , edited by Howard Rachlin and David I. Laibson, (Cambridge, Harvard University Press).

Romich, Jennifer and Thomas Weisner (2002). How Families View and Use the Earned Income Tax Credit, in Bruce Meyer and Douglas Holtz-Eakin editors, *Making Work Pay*. Russell Sage Foundation, New York.

Rosen, Harvey S. (1976). Taxes in a Labor Supply Model with Joint Wage-Hours Determination, *Econometrica*, 44:485-507.

Saez, Emmanuel (2002). Do Taxpayers Bunch at Kink Points. Mimeo UC Berkeley.

Saez, Emmanuel (1999) The Effect of Marginal Tax Rates on Income: A Panel Study of Bracket Creep, *Journal of Public Economics*.

Slemrod, Joel (1990). *Do Taxes Matter?*. MIT Press, Cambridge, MA.

Vaughn, W., Jr. (1981). "Melioration, matching, and maximization." *Journal of the Experimental Analysis of Behavior*, 36: 141-149.

Figure 1
Pigeons Spotlighting

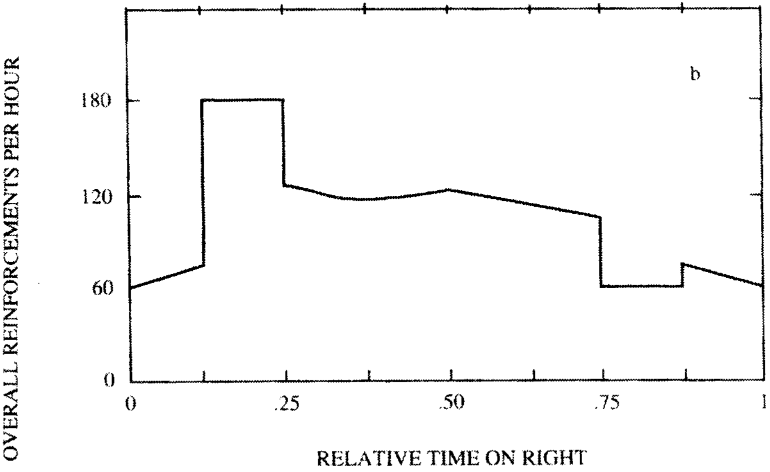

Source: Vaughn (1981) as adapted by Herrnstein (1982).

Figure 2a

Consumers' Budget Constraint in Monopolist Case

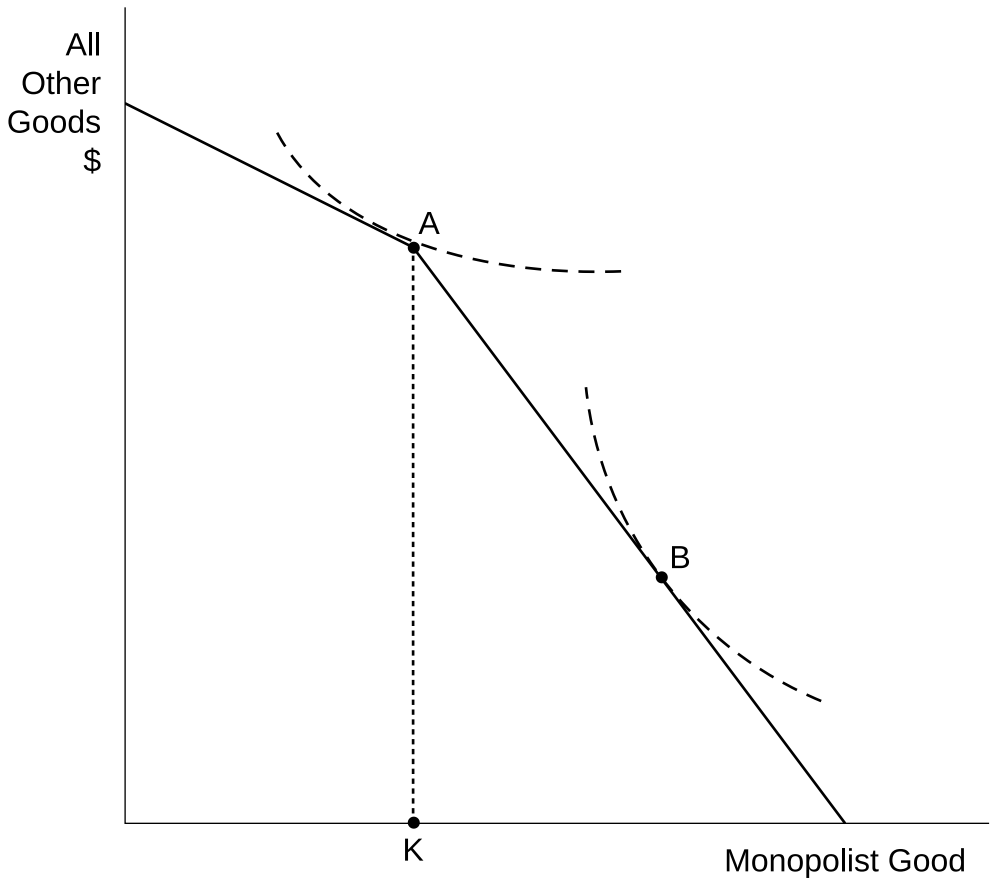

Figure 2b

Monopolist Case

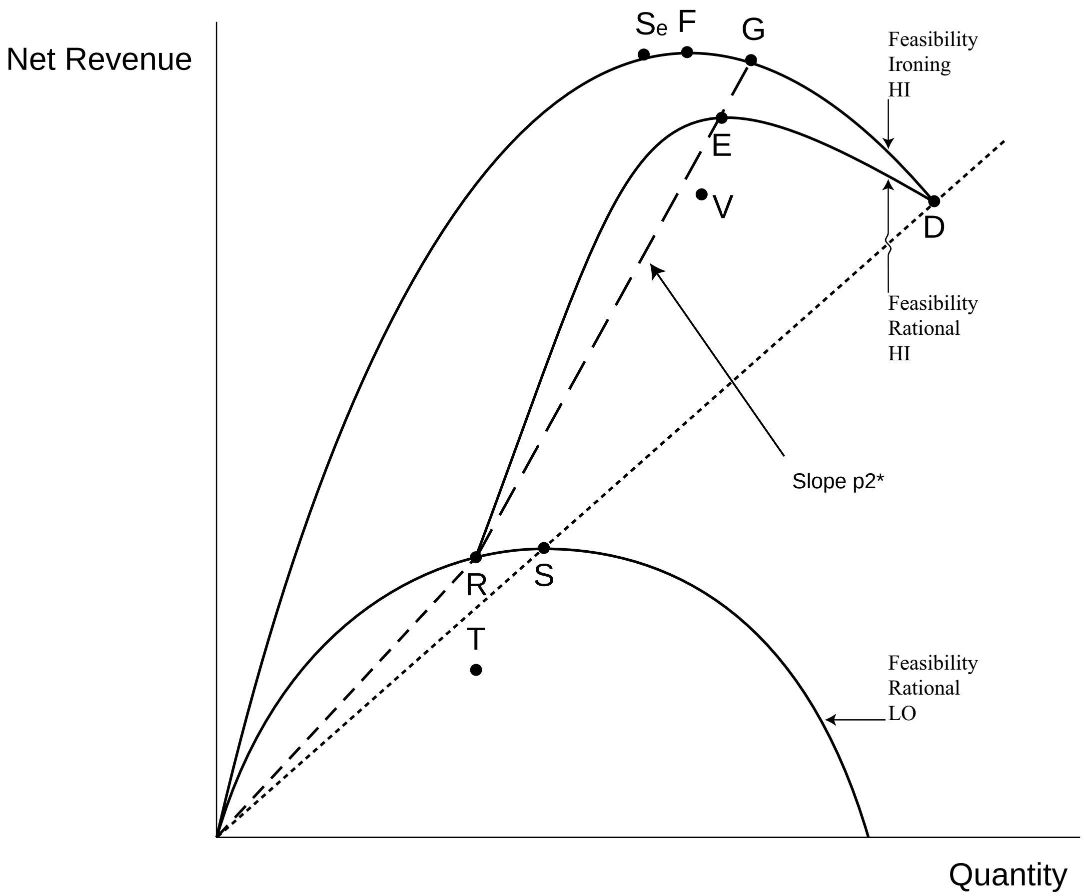

Figure 3
Ramsey Pricing

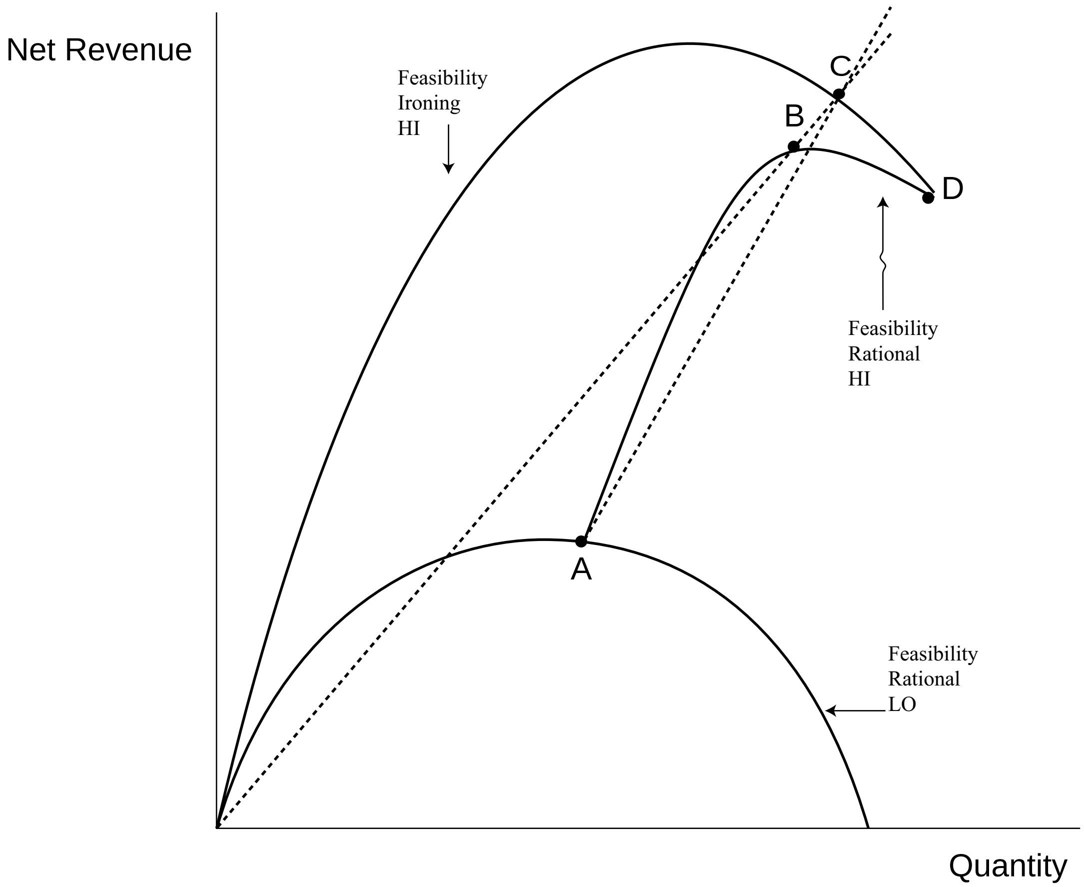

Figure 4

Schmeduling in the Optimal Income Tax

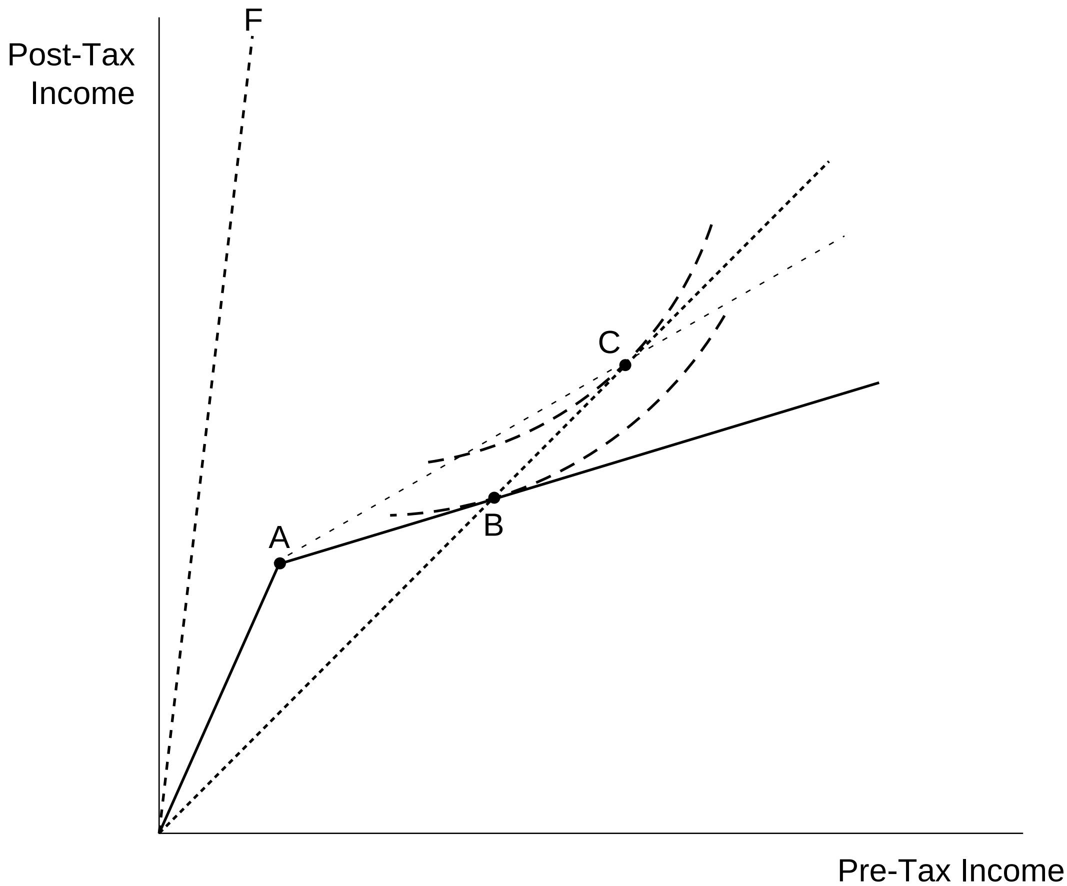

Figure 5
Change in Average and Marginal Tax Rates from Introduction of Child Credit

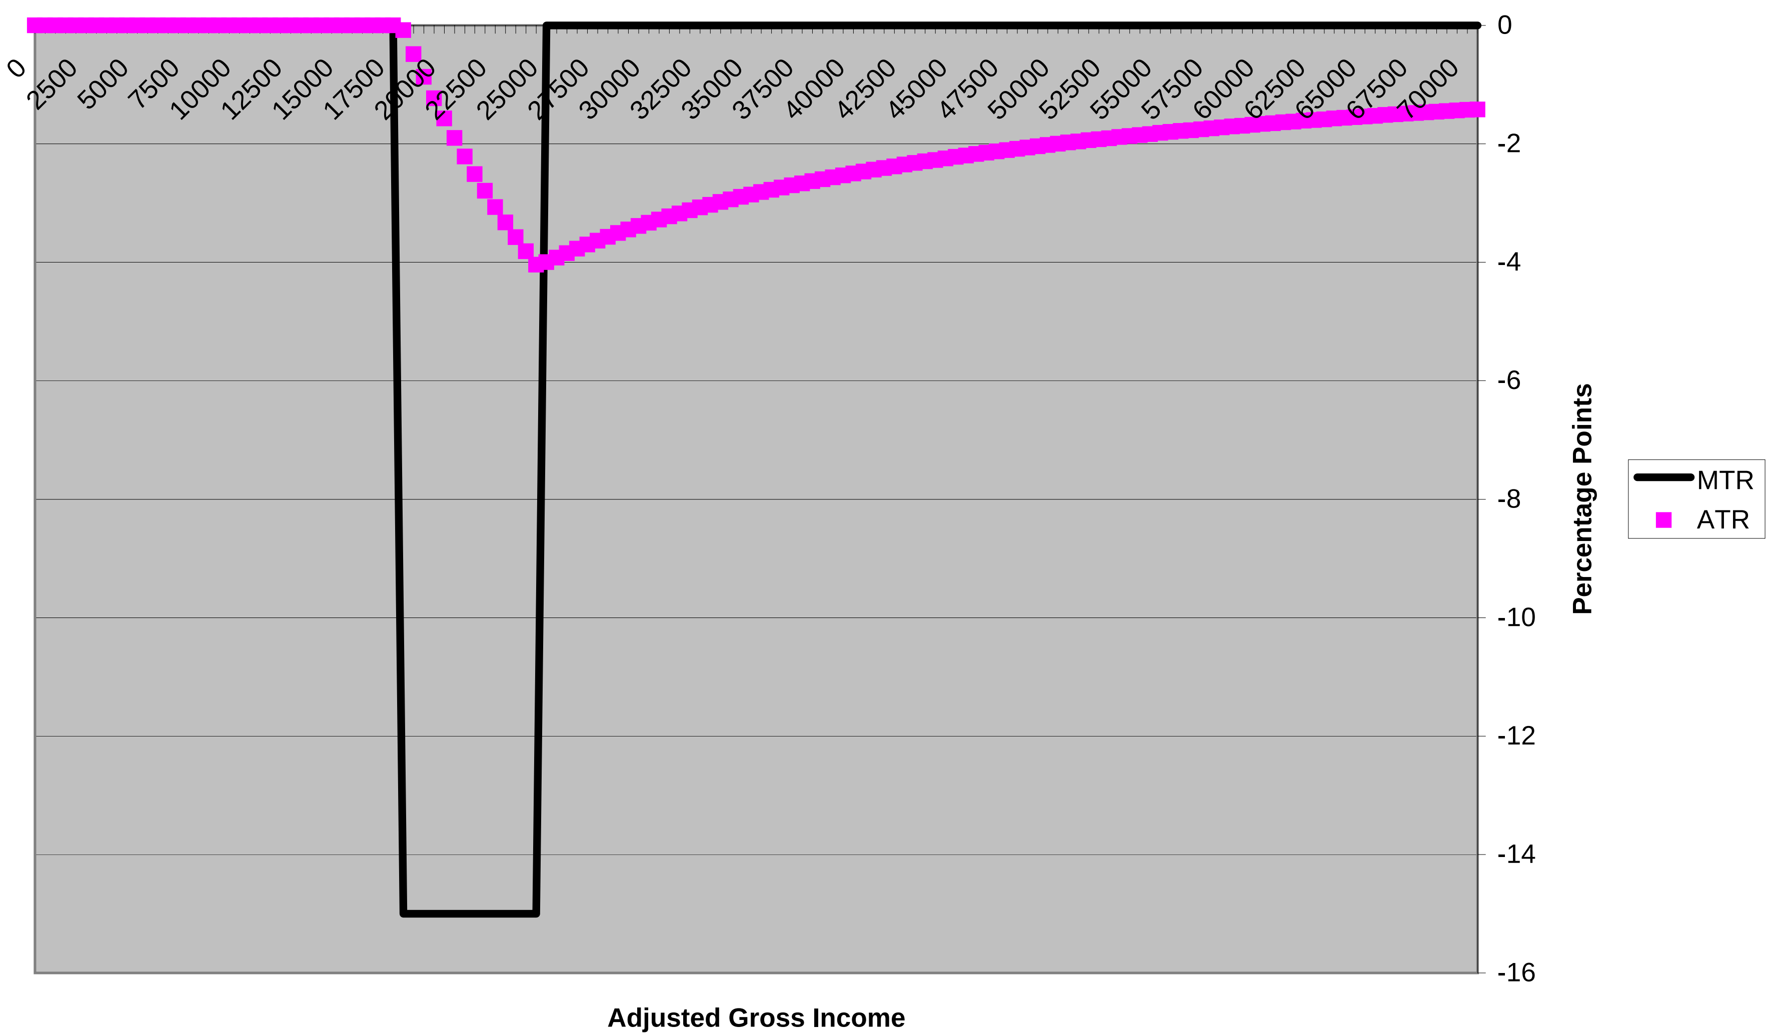

Figure 6A: The Different Predictions of the Two Models

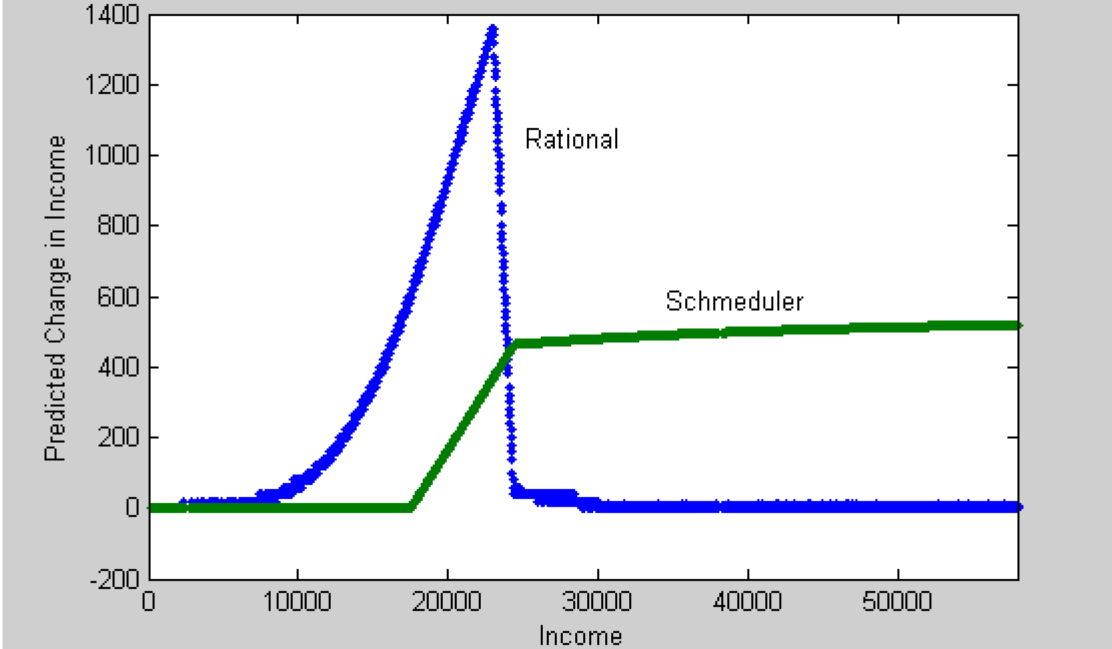

Figure 6B: CDFs of the Predicted Income Distribution Under Two Models

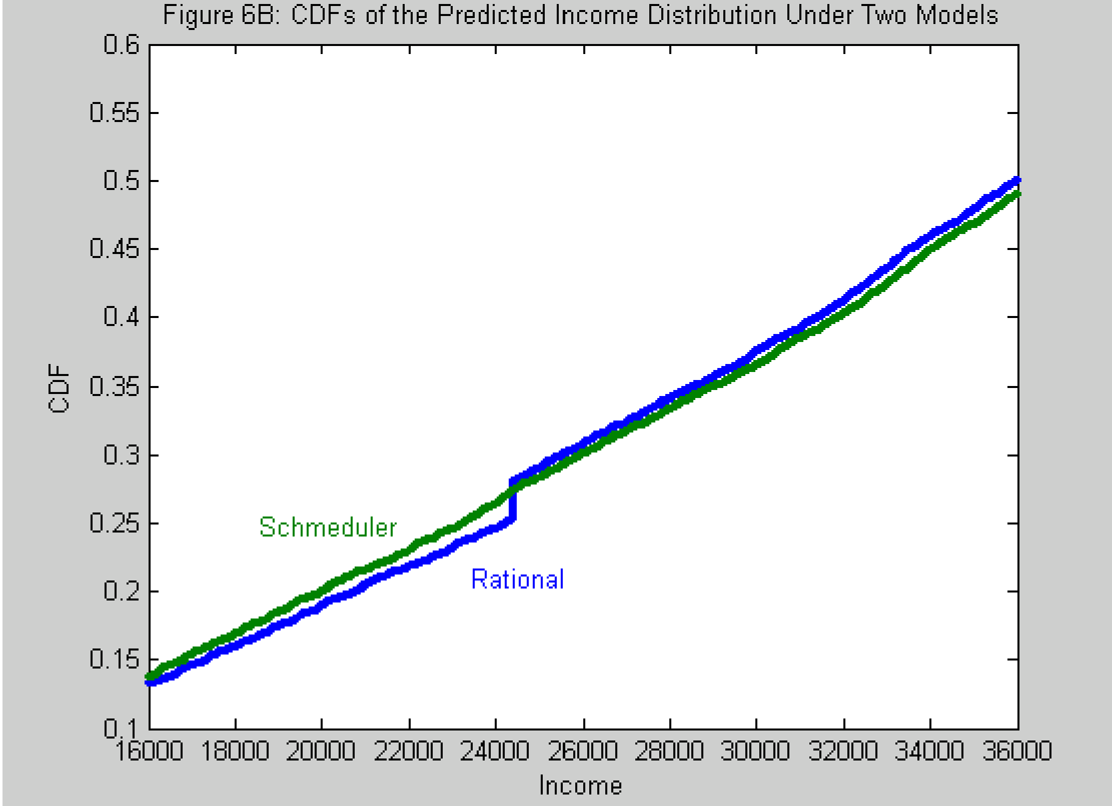

Figure 7: Predicted and Actual Distributions of Income

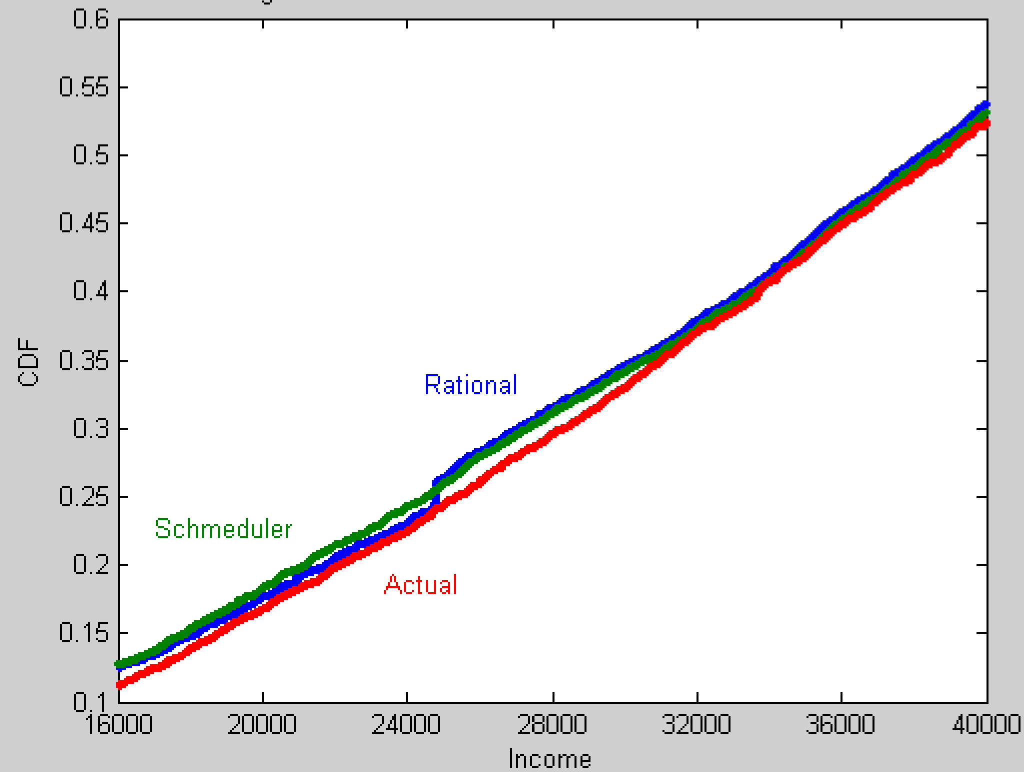

Table I
Five Examples of Conditions that Give Rise to Schmeduling

	Non-linear pricing	Complexity	Frequent revisions	Delayed Payoffs	Bundled Consumption	Nonstationary environment	Schedule heterogeneity	Obscure units	False signals
Tax schedules	X	X	X	X	?	X	X		X
Public assistance benefit formulas	X	X	X	?	?	X	?		
Utility pricing	X	X	X	X	X	X	X	X	?
Richard's parking tickets	X			X	X				X
Non-linear pricing of consumer goods	X				X	X	?	?	X

Note: X means that the condition is usually present for that example. ? means that this condition is sometimes present and sometimes not in that example.

Table II
Deadweight Loss in the Two Models with Elasticity of 0.4
(billions of dollars)

	Taxable Income	Revenue	Deadweight Loss
Schmeduling	4233.3	974.7	56.7
Rational	4019.9	913.4	109.0

Table III
Parameter Estimates from the Econometric Model of the Introduction of the Child Credit

	Point Estimate	Bootstrapped 95 percent confidence interval
Elasticity	0.22	[0.18, 0.26]
Share of sample that are schmedulers	0.45	[0.39, 0.51]

Notes. Standard deviation of optimization error is set to $2000. Standard deviation of uncertainty around the kink is set at $4000. Nominal wage growth is set at 5 percent between 1997 and 1999.

Table IV: Regression Results for the Within-Month Pattern of Food Consumption by Food Stamp Recipients

	Dependent Variable	
	Log of food consumption in dollars (1)	Log of food consumption in calories (2)
Days since last AFDC check	-.00266 (.00363)	-.00458 (.00360)
Days since last AFDC check interacted with Food Stamps paid in cash	.000676 (.005208)	.00531 (.00502)
Days since last Food Stamp check (for those paid in stamps)	-.00784 (.00295)	-.00755 (.00326)
R^2	0.20	0.27
Sample size	715	715

Notes: Regressions also include indicator variables for calendar month, household size, and interactions between these variables and experimental group. Robust standard errors in parentheses.