THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Tractable consumer choice

Citation for published version:

Friedman, D & Sakovics, J 2015, 'Tractable consumer choice', *Theory and Decision*, vol. 79, no. 2, pp. 333-358. <https://doi.org/10.1007/s11238-014-9461-0>

Digital Object Identifier (DOI):

[10.1007/s11238-014-9461-0](https://doi.org/10.1007/s11238-014-9461-0)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Theory and Decision

Publisher Rights Statement:

This is a post-peer-review, pre-copyedit version of an article published in Theory and Decision. The final authenticated version is available online at: <https://link.springer.com/article/10.1007%2Fs11238-014-9461-0>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Tractable Consumer Choice

Daniel Friedman · József Sákovics

Received: date / Accepted: date

Abstract We present a rational model of consumer choice, which can also serve as a behavioral model. The central construct is λ , the marginal utility of money, derived from the consumer's rest-of-life problem. It provides a simple criterion for choosing a consumption bundle in a separable consumption problem. We derive a robust approximation of λ , and show how to incorporate liquidity constraints, indivisibilities, and adaptation to a changing environment. We find connections with numerous historical and recent constructs, both behavioral and neoclassical, and draw contrasts with standard partial equilibrium analysis. The result is a better grounded, more flexible and more intuitive description of consumer choice.

Keywords separable decisions · distributed choice · moneysworth demand · value for money

1 Introduction

John has just realized that it is his wedding anniversary. He enters the only liquor store in town and looks for champagne. There is only one bottle, which costs \$100. How does he decide whether to buy it?

In this paper we argue that a consumer like John should recall his utility value for money, call it λ , and use it to estimate the opportunity cost of spending \$100. He should purchase the champagne if and only if the utility it brings covers that opportunity cost.

D. Friedman
Economics Department, University of California at Santa Cruz

J. Sákovics (corresponding author)
School of Economics, The University of Edinburgh,
30 Buccleuch Place, Edinburgh, EH8 9JT U.K.
Email: jozsef.sakovics@ed.ac.uk,
phone: +44 131 650 3840, fax: + 44 131 650 4514.

Our recommendation might sound like common sense, but it differs sharply from what economics textbooks propose. General equilibrium theory would require John to consider the impact on all possible future consumption plans. As is widely recognized, for actual people, or even *Homo Economicus* facing moderate computation costs or other frictions, this is too complex a problem to solve.

Consequently, micro textbooks (and recent behavioral models of mental accounting) propose an alternative, partial equilibrium solution: John should have a pre-established budget to spend on items including last minute anniversary presents, and should purchase the champagne if and only if that budget exceeds \$100 and, on the margin, the other claims on the budget are less pressing. We shall argue that this alternative solution, although tractable, is twice hobbled. First, it is silent on how to specify the budget set. When and how should John come up with the amount to spend and the range of goods to consider? Second, the budget constraint rules out substitution of purchasing power between present and future, thus preventing John from properly responding to prices that turn out to be higher or lower than expected.

We propose to use λ , the *marginal utility of money*, rather than a budget constraint, to link the present problem to the rest-of-life problem. The numerical value of λ captures the trade-off between current and rest-of-life expenditure, and thus allows the consumer to make optimal saving/borrowing decisions. This approach resonates with the findings of consumer research, and λ itself is a meaningful concept that can be approximated, learned and adjusted in intuitive ways.

Our approach relates to several strands of consumer choice theory, some historical and some more recent. These include *inter alia* Marshall's theory of demand; partial equilibrium analysis; the use of cardinal and quasi-linear utility functions; Frisch demand functions and recursive life-cycle models.

A distinctive aspect of our approach is how it connects present and future choices. Our consumer occasionally adjusts her estimated marginal utility of money looking forward, but when faced with a specific consumption choice she just uses the current value of λ . Existing approaches either ignore the future entirely (partial equilibrium) or else refuse to consider the future as qualitatively distinct from the present (general equilibrium). The simple connection we impose naturally leads to a better understanding of where λ comes from and what makes it useful.

Our free standing theory of consumer choice includes the following innovations:

- Neo-classical foundations for (cardinal) quasi-linear utility and for the use of λ as a rule of thumb.
- Proof of near-optimality of a λ -based response to price surprises.
- Extensions to choice among baskets of indivisible goods, and choice given liquidity constraints.
- An adaptive learning model for λ , complementing the forward looking analysis.

- Examples showing how λ can embody behavioral biases.

Section 2 begins the exposition with a review of the lifetime consumption problem and its textbook solution. We then define separable subproblems, and use the indirect utility function to obtain a quasi-recursive solution to the subproblem and its continuation problem. The first-order condition defines λ as the opportunity cost of subproblem expenditure.

Section 3 analyzes the subproblem solution given an exogenous value of λ . It shows that the consumer moves out along the subproblem’s Income Expansion Path – the locus of points where the ratio of marginal utility to price is the same for each good – until that ratio equals λ . The solution yields the “moneysworth” demand function, whose price effects and λ effects are shown to be quite natural. In particular, demand is always decreasing in own price (hence there are no Giffen goods), while the cross price effects straightforwardly reflect preference-based substitutability or complementarity. Quasilinear preferences emerge as a natural special case.

Section 4 discusses how the consumer can calibrate λ . A Taylor expansion of indirect continuation utility reveals that λ is approximately constant when subproblem expenditures are small relative to lifetime income. An implication is that λ is reusable: the consumer typically need not adjust it between consecutive small purchases. Eventually, however, she does need to update. We posit two complementary methods: one is forward looking and uses relevant news in a life-cycle setting, while the other is experiential and uses previously observed prices to update λ .

Section 5 recalls the textbook partial equilibrium approach, which holds the subproblem budget constant, and yields the Marshallian demand function. Proposition 3 shows that the “true” lifetime optimal demand elasticities are typically well approximated by those of the moneysworth demand function, but only under quite special conditions are they close to the Marshallian elasticities.

Section 6 extends the model to include liquidity constraints, which are shown to correspond to the situation where λ is a step function of available liquidity. Marshallian demand reappears as an extreme case, where $\lambda = 0$ below the budget and $\lambda = \infty$ over the budget. We also show that consumers can handle indivisibilities by comparing λ to an appropriately defined quality-price ratio.

Section 7 reiterates that, compared to textbook approaches, the λ (or moneysworth) approach offers (a) more robust prescriptions for how consumers should react to surprises, (b) a better way to connect partial equilibrium to general equilibrium analysis, and also (c) more plausible descriptions of actual human behavior. That section also notes fruitful avenues for future research.

Appendix A concisely reviews connections to the most relevant strands of historical and recent economic literature, and Appendix B collects mathematical details.

2 Preliminaries

A consumer's lifetime consumption plan $\mathbf{X} = \{(x_1, \dots, x_N) : x_i \geq 0\}$ specifies the quantities of all goods and services consumed at present and at all future dates in all contingencies. The dimension N of the problem is astronomically large, given even moderate numbers of goods, dates and contingencies at each date.¹ We assume that the consumer has cardinal preferences over consumption, represented by the utility function $\tilde{U} : \mathbb{R}_+^N \rightarrow \mathbb{R}$, which also accounts for time preferences and random termination. She takes the (expected) price vector $\mathbf{P} = (p_1, \dots, p_N)$ as given. Initially, we assume that the normalized lifetime purchasing power $L > 0$ is freely transferable across purchases.²

Absent other constraints, the consumer's problem can be written just as in textbooks,

$$\max_{\mathbf{X} \geq 0} \tilde{U}(\mathbf{X}) \text{ s.t. } \mathbf{P} \cdot \mathbf{X} = \sum_{i=1}^N p_i x_i \leq L. \quad (1)$$

The Lagrangian is

$$\max_{(\mathbf{X}, \mu) \geq 0} \left[\tilde{U}(\mathbf{X}) + \mu(L - \mathbf{P} \cdot \mathbf{X}) \right]. \quad (2)$$

Writing $\tilde{U}_i$ for the i^{th} partial derivative (or marginal utility), we have the first-order conditions

$$\begin{aligned} \tilde{U}_i(\mathbf{X}^*) &= \mu p_i, \quad i = 1, \dots, N \\ L &= \mathbf{P} \cdot \mathbf{X}^*. \end{aligned} \quad (3)$$

We have nothing novel to say about multiple, corner or non-differentiable solutions, so to streamline the exposition we henceforth assume that (3) has a unique smooth solution $\mathbf{X}^*(\mathbf{P}, L)$ that solves the original problem (1). Sufficient conditions (e.g., that $\tilde{U}$ is smooth, monotone, Inada and concave) are well known (e.g., Bewley, 2007).

Our main concern is a situation where the consumer is not solving her full lifetime problem, just a small part of it. A *consumer subproblem* is to choose an n -subvector of the lifetime consumption plan $\mathbf{X}$, where $1 \leq n \ll N$. By suitably reindexing, we can write $\mathbf{X} = (\mathbf{x}, \chi)$, where $\mathbf{x} = (x_1, \dots, x_n)$ is the subvector and $\chi = (x_{n+1}, x_{n+2}, \dots, x_N)$ is the rest of life (or continuation) plan. The price subvector of $\mathbf{P}$ for the subproblem is $\mathbf{p} = (p_1, \dots, p_n)$. The idea is to gain tractability by reducing the dimensionality from $N \approx 10^{100}$ in a realistic lifetime problem to something small, perhaps $n = 2$ or 3 .

That reduction is possible if the subproblem is *separable*, that is, if there are subutility functions $u : \mathbb{R}_+^n \rightarrow \mathbb{R}$ and $U : \mathbb{R}_+^{N-n} \rightarrow \mathbb{R}$ such that for every consumption plan $\mathbf{X} = (\mathbf{x}, \chi)$ we can, with negligible error, write

$$\tilde{U}(\mathbf{X}) = u(\mathbf{x}) + U(\chi). \quad (4)$$

¹ Indeed, N could even be infinite and not affect our analysis.

² Prices that apply to future goods are discounted appropriately. For simplicity, we treat L as a constant, but later note how it can be endogenized.

A sufficient condition for separability is that the cross second partial derivative $\tilde{U}_{ij}(\mathbf{X})$ is zero everywhere for all $i = 1, \dots, n$ and $j = n+1, \dots, N$. For example, even in the same shopping episode the consumer can choose wine and cheese in one subproblem, and choose cereals and milk in a second. Separability seems quite plausible for both subproblems.

Given $\mathbf{P}$, the consumer's lifetime *indirect* utility function is

$$\tilde{V}(L) = \max_{\mathbf{X} \in \mathbb{R}_+^N} \tilde{U}(\mathbf{X}) \text{ s.t. } \mathbf{P} \cdot \mathbf{X} \leq L. \quad (5)$$

By construction, $\tilde{V}(\cdot)$ is homogeneous of degree zero in (L, P) and increasing in L . Given strong classic assumptions of the sort mentioned earlier, we can also assume that it is smooth and concave, i.e., $\tilde{V}'(\cdot) > 0$ and $\tilde{V}''(\cdot) \leq 0$ in the relevant range (c.f. Varian 1992, pp.102ff.).

Let $V(\cdot)$ be the indirect utility function for the continuation plan, obtained by imposing the additional restriction $\mathbf{x} = 0$ in equation (5). Then (4) gives us a helpful expression for indirect utility:

$$\tilde{V}(L) = \max_{\mathbf{x} \in \mathbb{R}_+^n} [u(\mathbf{x}) + V(L - \mathbf{p} \cdot \mathbf{x})], \quad (6)$$

The equation says that if the subproblem is separable, then the only effect that the choice of $\mathbf{x}$ has on rest-of-life utility is pecuniary – the subproblem expenditure $\mathbf{p} \cdot \mathbf{x} = \sum_{i=1}^n p_i x_i$ reduces the consumer's rest-of-life purchasing power.³

3 Moneyworth demand

The first-order conditions for (6), which characterize the principal subvector $\mathbf{x}^*$ of the solution X^* to the lifetime problem, yield $u_i(\mathbf{x}^*) = p_i V'(L - \mathbf{p} \cdot \mathbf{x}^*)$, $i = 1, \dots, n$. Slight reformatting yields

$$\begin{aligned} u_i(\mathbf{x}^*) &= \lambda p_i, \quad i = 1, \dots, n \\ \lambda &= V'(L - \mathbf{p} \cdot \mathbf{x}^*), \end{aligned} \quad (7)$$

where $\lambda = V'(L - \mathbf{p} \cdot \mathbf{x}^*)$ is the consumer's *marginal utility of money*. We see that λ is a sufficient statistic for the continuation problem, and that it tells us how much utility the consumer could gain elsewhere if she cut back subproblem expenditure by a dollar. Otherwise put, λ is the opportunity cost of subproblem expenditure, or the “shadow utility” of purchasing power, or the “conversion rate” between utility and money.

The first line in (7) also has a nice interpretation. It says that the marginal utility vector (or gradient) ∇u is proportional to the price vector $\mathbf{p}$. Varying the

³ In the special case that the same separable subproblem recurs every period, (6) closely resembles the Bellman equation familiar to macroeconomists (where time discounting is built into $V(\cdot)$). Note however, that in general, $\tilde{V}$ differs from V not only in the value of the state variable (remaining wealth) but also in the set of (dated) goods available to purchase.

conversion rate, λ , sweeps out the locus of points satisfying the proportionality condition, a smooth one-dimensional curve often called the Income Expansion Path. The IEP emanates from the origin (where the marginal utilities are maximal) and extends into the interior of the subproblem consumption set $\mathbb{R}_+^n$. The marginal utility of each good decreases smoothly (and each by the same percentage) as we move out along the IEP defined by the price vector $\mathbf{p}$. To find the demand vector satisfying both lines of (7) we simply move out along the IEP until the common proportionality factor falls to the specified value λ .⁴ At that point, (7) holds and the marginal utility of expenditure in the subproblem matches the marginal utility of money λ specified for the rest-of-life.

Compared to equation (3), we see that (7) offers two simplifications: it reduces the apparent dimensionality from $N + 1$ to $n + 1$, and it replaces the budget equation in L by a statement of how λ , the marginal utility of money, is determined.

The complexity of the original problem has not yet disappeared: it has just been concentrated into the endogenous dependence of $\lambda = V'$ on subproblem expenditure $y = \mathbf{p} \cdot \mathbf{x}$. The crucial next step is to show that this dependence is tractable and indeed is often negligible.

3.1 Constant λ

We begin by analyzing demand when λ can be regarded as exogenous. In particular, suppose that the curvature of the indirect utility function is negligible over a range of subproblem expenditure $y = \mathbf{p} \cdot \mathbf{x} \in [0, \bar{y}]$. Then the marginal utility of money is essentially constant at $\hat{\lambda} = V'(L)$, so V is given by its first order Taylor expansion

$$V(L - y) = V(L) - y\hat{\lambda}. \quad (8)$$

The consumer choice problem now becomes quite simple. Substituting (8) into the lifetime optimization problem (6) with $y = \mathbf{p} \cdot \mathbf{x}$ and dropping the

⁴ Of course, expenditure $L - \mathbf{p} \cdot \mathbf{x}$ available for the continuation problem falls as we move out along the IEP so, by concavity, V' increases. Thus a more complete informal description is that the falling proportionality factor meets the rising V' at a unique point $\mathbf{x}^*$ on the IEP; more formally, under present assumptions, the intermediate value theorem and implicit function theorem guarantee a unique smooth interior solution described by (7).

irrelevant constant term $V(L)$ yields the unconstrained optimization problem⁵

$$\max_{\mathbf{x} \geq 0} \left[u(\mathbf{x}) - \hat{\lambda} \mathbf{p} \cdot \mathbf{x} \right]. \quad (9)$$

Streamline notation by dropping the decoration on the now-exogenous parameter $\hat{\lambda}$ and take the first-order conditions for (9) to get

$$u_i(\mathbf{x}^\lambda) = \lambda p_i, \quad i = 1, \dots, n. \quad (10)$$

We will refer to $\mathbf{x}^\lambda(\mathbf{p}) = (x_1^\lambda(\mathbf{p}), \dots, x_n^\lambda(\mathbf{p}))$, the solution of (10), as the *moneysworth* demand function.⁶

3.2 Income and price effects

What are the comparative statics of $\mathbf{x}^\lambda$? A change in subproblem relative prices will shift the IEP, causing both own-price and cross price effects. On the other hand, a move along the original IEP will arise from a change in the subproblem price level, a change in lifetime income L , or anything else that changes λ . In this subsection we analyze both sorts of comparative statics.

Shifts in the IEP are naturally described by second derivatives of the utility function. Let $H = ((u_{ij}))$ denote the $n \times n$ Hessian matrix of second partial derivatives of subproblem utility $u(\mathbf{x})$, and let H_{ij} denote the $(n-1) \times (n-1)$ submatrix with i^{th} row and j^{th} column deleted. Vertical bars, e.g., $|H|$, denote the determinant.

Proposition 1 *Given a separable subproblem of dimension $n \geq 2$ with price vector $\hat{\mathbf{p}} \in \mathbb{R}_{++}^n$ and constant $\lambda > 0$, let $\mathbf{x}^\lambda(\cdot) \gg 0$, be the moneysworth demand function that uniquely solves equation (10). Then its price sensitivities are*

$$\frac{\partial x_j^\lambda}{\partial p_i}(\hat{\mathbf{p}}) = \frac{(-1)^{i+j} \lambda |H_{ij}|}{|H|}, \quad \text{for all } i, j \in \{1, \dots, n\}. \quad (11)$$

A proof appears in Appendix B (as part of the proof of Proposition 3); for related results see also Biswas (1977) and Browning et al. (1985). Note that for the special case of a single good in the subproblem ($n = 1$) we have

$$\frac{\partial x^\lambda}{\partial p}(\hat{p}) = \frac{\lambda}{u'(x^\lambda)}. \quad (12)$$

⁵ The maximand in (9) slightly generalizes quasi-linear utility. Textbook treatments (e.g., Varian 1992, p. 154, 164-7) often assume that $n = 1$ so the variable of interest x is scalar, and that $\hat{\lambda} = 1$. In our notation, the textbook quasi-linear utility function would be written as $u(x) + m$ with budget constraint $m = L - px$. Our approach shows that textbook quasi-linear preferences can be justified for single separable goods, and can be generalized directly for separable bundles of related goods. Given a constant exogenous λ , there is no loss of generality in using a VNM utility function for u (and implicitly, U) that normalizes its value to 1, but our analysis sheds light on the conditions for which the constancy assumption is justified.

⁶ For given constant λ , the Frisch demand function is formally equivalent to $\mathbf{x}^\lambda$, notwithstanding important differences in interpretation (see Section A.2).

Properties of moneysworth demand can be gleaned from the formula (11). The i, j symmetry of the formula (and of the matrix H) imply that the price- i sensitivity of good j is the same as the price- j sensitivity of good i . Since H is negative definite, the determinants $|H_{ii}|$ and $|H|$ have opposite signs, so the formula tells us that $\frac{\partial x_i^\lambda}{\partial p_i} < 0$, i.e., the own price effect is always negative. Moneysworth demand therefore rules out Giffen goods.⁷

The cross-price effect is transparent when there are only two goods in the subproblem, since in that case the formula collapses to $\frac{\partial x_1^\lambda}{\partial p_2} = \frac{\partial x_2^\lambda}{\partial p_1} = \frac{-\lambda u_{21}}{u_{11}u_{22} - u_{12}u_{21}}$. The denominator is positive by concavity, and so the cross price effect simply is a (sign-reversing) rescaling of the cross partial derivative of u at the consumption point. If the goods are substitutes then $u_{21} < 0$ and demand for a good will increase when the other good's price rises, but if they are complements then $u_{21} > 0$ and the same price rise will decrease demand. Thus, relative price effects for moneysworth demand arise naturally from the second partial derivatives of u , and are easy to interpret and explain.

It is straightforward to see that a *ceteris paribus* shift in λ can arise from a shift in real lifetime income after subproblem expenditure, or from a change in cardinal preferences, but not from merely rescaling utility nor rescaling (L, P) . The impact of a *ceteris paribus* shift can be expressed in terms of $H^{j(\mathbf{p})}$, the Hessian matrix with its j^{th} column replaced by the price vector $\mathbf{p}$. Differentiating (10) and applying Cramer's Rule we have immediately that

Proposition 2 *The sensitivity of moneysworth demand to the marginal utility of money is*

$$\frac{\partial x_j^\lambda}{\partial \lambda} = \frac{|H^{j(\mathbf{p})}|}{|H|}, \text{ for } j = 1, \dots, n. \quad (13)$$

Biswas (1977) includes the equivalent expression $d\mathbf{x}/d\lambda = H^{-1} \cdot \mathbf{p}$. In the one-good case (13) simplifies to $\frac{\partial x^\lambda}{\partial \lambda} = \frac{p}{u''(x^\lambda)} < 0$, capturing the intuitive notion that demand falls when the (shadow) value of money increases. In the two-good case (13) simplifies to $\frac{\partial x_j^\lambda}{\partial \lambda} = \frac{p_j u_{ii} - p_i u_{ji}}{|H|} = \frac{u_j u_{ii} - u_i u_{ji}}{\lambda |H|}$. Thus an increase in the marginal utility of money can be decomposed into two effects. The first one $\frac{u_j u_{ii}}{\lambda |H|}$ is always negative, since (as noted above) the denominator is positive and the numerator factor u_{ii} is negative by concavity. The second term $-\frac{u_i u_{ji}}{\lambda |H|}$ is positive if (and only if) the goods are substitutes, and it can outweigh the first term. In this last case, the IEP bends backward and a higher λ will lead to a higher consumption of good j .

⁷ This is intuitive, as a constant λ means that there are no income effects. As an empirical matter, it is not clear that Giffen goods exist at all (c.f. Dwyer and Lindsay 1984; Nachbar 1998), with the possible exception of extreme poverty (c.f. Jensen and Miller 2008) when indeed we would not expect consumers to take the future into account and to exhaust whatever purchasing power they have.

4 Calibrating λ

Moneysworth demand is relevant when the consumer can both (a) treat $\lambda = V'(L - y)$ as a constant independent of subproblem expenditure $y = \mathbf{p} \cdot \mathbf{x}$, and (b) obtain a good estimate of that value of λ . Let us deal with each condition in turn.

Constancy is guaranteed when rest-of-life utility is a constant returns to scale CES function, $\tilde{U}(\mathbf{X}) = [\sum_{i=1}^N a_i x_i^r]^{1/r}$ for $0 \neq r < 1$, where $a_i > 0$. In that case, λ is essentially a price index for rest of life consumption and is independent of L and y ; see e.g., Varian (1992, p. 112). The same is true, of course, for Cobb-Douglas utility, which is the $r = 0$ member of the CES family.

More generally, indirect utility is concave, and we need to quantify the rate at which marginal utility declines. The exact Taylor expansion yields

$$V'(L - y) = V'(L) - yV''(L - \alpha y), \text{ for some } \alpha \in [0, 1]. \quad (14)$$

Thus the endogenous marginal utility of money is seen to consist of its zero-expenditure value $\hat{\lambda} = V'(L)$ – a constant exogenous to the subproblem – corrected by a term that is proportional to the unnormalized curvature $V'' = \beta \leq 0$ of the indirect utility function, so the decline in λ is approximately linear in subproblem expenditure $y = \mathbf{p} \cdot \mathbf{x}$.⁸

How good is that linear approximation? The long-time consensus from work on risk aversion (reinforced by recent contributions such as Rabin, 2000) is that the concavity of the indirect utility function should be diminishing: we would expect that the marginal utility for money diminishes more between a wealth of \$100,000 and \$200,000 than between \$1,100,000 and \$1,200,000. For us, then, the worst plausible case is that β is constant. In that case, the exact Taylor expansion $V'(L) = V'(0) + L\beta$ implies $\beta = \frac{V'(L) - V'(0)}{L}$. Hence the error term in (14) is proportional to $\frac{y}{L}$, the size of the expenditure today relative to lifetime income. Even in this worst plausible case, then, the constant λ approximation is good for small subproblems.

A numerical example may help crystalize ideas. Let $V(L) = c \ln L$ as in Bernoulli's classic example, let lifetime income be $L = \$1$ million (e.g., 50k/yr increasing over 20 years at the discount rate), and let subproblem expenditure be $y = \$100$ as for John's champagne. Pick the convenient utility scaling $c = 10^6$ so that John's zero-expenditure marginal utility of money is $\hat{\lambda} = V'(10^6) = c10^{-6} = 1$. Equation (14) tells us that John's marginal utility of money if he spends the \$100 is $V'(10^6 - 100) = \hat{\lambda} - 100\beta = 1 + z$, where the correction $z > 0$ is between $-100\hat{\beta} = 10^2 c [10^6]^{-2} = 10^{2+6-12} = 1.00000 \times 10^{-4}$ (where $\alpha = 0$), and $-100\beta|_{\alpha=1} = 10^2 c [10^6 - 100]^{-2} < 1.00021 \times 10^{-4}$. Notice that John's λ is hardly affected by a purchase in this range – it rises by about one-hundredth of 1% – and that this tiny correction z is well approximated using $\hat{\beta} = V''(L)$.

A general result follows easily from the last Proposition:

⁸ Only approximately linear because β is a second derivative evaluated at a point that can depend on y .

Corollary 1 *The sensitivity of moneysworth demand to lifetime income is given by*

$$\frac{dx_j^\lambda}{dL} = -\widehat{\beta} \frac{|H^{j(\mathbf{p})}|}{|H|}, \quad j = 1, \dots, n, \quad (15)$$

where $\widehat{\beta} = V''(L) \leq 0$.

Proof. Differentiate (14) and evaluate at $y = 0$ to obtain $\frac{d\lambda}{dL} = \widehat{\beta}$. The chain rule tells us that $\frac{dx_j^\lambda}{dL} = \frac{\partial x_j^\lambda}{\partial \lambda} \cdot \frac{d\lambda}{dL}$, so (15) follows from Proposition 2. QED.

Given that $\widehat{\lambda}$ is a good approximation, the remaining question (b) is whether, as a practical matter, it can be estimated reasonably well. The definition of indirect utility is not encouraging, since it refers to the entire lifetime maximization problem. But we only need to estimate the slope of that function, and just for the continuation problem. Of course, at first blush the continuation problem seems almost as complex as the entire problem, but there is a qualitative sense in which it simplifies greatly when we split off the present subproblem. The present is always idiosyncratic, if only because we have observed prices to go on. The future is more nebulous, with enormous amounts of uncertainty, and consumers may prefer to think of prices for broad categories rather than for individual items of future consumption. In other words, while the present subproblem needs to be modeled in fine detail, we can afford to treat the continuation in a more abstract way when estimating λ .⁹ We do not propose a specific model for the continuation, but note that a reasonable forward-looking estimate may be obtained using the stationarity assumptions and aggregation routinely imposed in life-cycle models discussed in Section A.2 below, or using a rule of thumb in the spirit of Love (2013).

It bears emphasizing that, once the consumer has a good estimate of her λ , she can reuse it many times. Unlike a budget constraint, λ changes only slightly from one small subproblem to the next, and doesn't require frequent recalculation. However, eventually λ will require updating, especially as the consumer learns more about herself and her possibly changing circumstances. Some changes in the consumer's view of the future, such as new pension plan, can be dealt with using the Corollary above or minor extensions. Other changes, for example encountering a new product for the first time, concern the current subproblem, but have ramifications for the future – and therefore λ . The consumer should be able to adapt λ as she observes prices over numerous subproblems.

We therefore propose a two step process for updating λ in light of accumulated experience. The first step is to translate a price observation into news about the value of λ , and the second is to determine the magnitude of the update.

⁹ The logic is reminiscent of quasi-hyperbolic discounting: there is a big difference between today and tomorrow, but tomorrow and the day after look similar from the vantage point of today.

Translation is straightforward for indivisible goods. As explained in Section 6.2 below, the quality-price ratio $\frac{u(b^k)}{p \cdot b^k}$ represents a new observation of λ . For divisible goods, that procedure would never get a new observation, since x is *chosen* to satisfy $\frac{u_1(x)}{p_1} = \lambda$. Instead, the consumer evaluates the marginal quality-price ratio at the quantity x^{old} chosen “last time.” That is, the new observation of λ is $\frac{u_1(x^{old})}{p_1^{new}}$.

The second step weights each new observation according to its share in overall consumption, as in simple price indices.¹⁰ Periodically – say, monthly – the consumer collects the new observations, say $\{\lambda_i : i = 1, \dots, m\}$, and computes the share $q_i \in [0, 1]$ of expenditure devoted to good i in the past. The updated value of λ is

$$\lambda' = \left(1 - \sum_{i=1}^m \alpha_i q_i\right) \lambda + \sum_{i=1}^m \alpha_i q_i \lambda_i. \quad (16)$$

Here $\alpha_i \in [0, 1]$ is the parameter measuring how much the consumer weighs new information relative to old. It also captures the consumer’s perception of the permanence of any price changes. Thus a one-off “fire sale” should carry little weight (very low α_i) while a price hike due to a permanent new specific tax levied on a product should have an α_i close to one. Also, we would expect all α_i ’s to be larger for an individual whose marginal utility diminishes more quickly.

Note that the updating rule also implies that observations of prices of goods that the consumer does not usually purchase (low q_i) have minimal impact on λ . Similarly, if a good gets priced out of a consumer’s reach, she will stop buying it, so q_i will decline and eventually it will have also have minimal impact on λ .

In summary, we think of our consumer as walking around with a constant λ , periodically updating it based on observed prices and occasionally (yearly or when major news arrives) making a forward looking estimate (possibly with the help of a tax accountant).

5 Comparisons

Undergraduate textbooks suggest a rather different tractable subproblem demand function. For an arbitrary budget $B > 0$, write

$$\max_{\mathbf{x} \geq 0} u(\mathbf{x}) \text{ s.t. } \mathbf{p} \cdot \mathbf{x} = B. \quad (17)$$

with first order conditions

$$\begin{aligned} u_i(\mathbf{x}^B) &= \nu p_i, \quad i = 1, \dots, n \\ B &= \mathbf{p} \cdot \mathbf{x}^B. \end{aligned} \quad (18)$$

¹⁰ Here we assume, for simplicity, that the consumer does not try to extrapolate from individual observed prices to changes in the price level. See Deaton (1977) for an exploration of that idea.

Let the solution (assumed unique and smooth) be $\mathbf{x}^B(\mathbf{p}) = (x_1^B(\mathbf{p}), \dots, x_n^B(\mathbf{p}))$. We shall refer to $\mathbf{x}^B$ as the *constant-budget demand function* for the subproblem.¹¹

Observe that (10) differs from (18) only in two respects: it omits the budget constraint, and λ replaces ν . Recall that ν is the marginal consumption utility of a dollar spent over the budget. If B is not chosen optimally (as the optimal subproblem expenditure in the lifetime problem (6)) then ν differs from the true marginal cost of overspending.

Applying a set of fixed budgets to a set of separable subproblems will typically generate a set of different shadow prices ν of expenditure. A standard argument shows that this is inefficient, and the consumer will increase utility by reallocating expenditure from low ν subproblems to those with high ν . In contrast, a consumer who takes the trouble to optimize once can use the same λ over and over without incurring significant efficiency losses.

The case $u_{21} = 0$ of separable goods is particularly instructive. Constant-budget demand asserts that such goods are (gross) substitutes, due to the pecuniary externality incorporated into the income effect. As noted earlier, moneysworth demand has no cross-price effect here, and directly reflects the “want independence” between the goods; the pecuniary externality is fully internalized in λ .

Another difference spotlights unanswered questions about the source of the budget constraint. Relaxing the Inada condition, one still usually gets interior solutions given a constant budget, while moneysworth demand can often be zero. Consider, for example, $u(x) = \ln(x + 1)$, $p = 1$ and $\lambda \equiv 1$. The “bang for buck” that the consumer can get is less than the opportunity cost for any value of x . The textbook consumer would simply spend her budget. For example, if faced with the choice of how much Beluga caviar to buy, an average moneysworth consumer would pass, while a textbook consumer with \$10 in his pocket would buy a gram. The counter argument that the consumer would not consider the \$10 as her budget for the caviar raises the followup question: and how did she arrive at that decision? Every textbook we know is silent here.

We are now ready to ask a crucial question. How good are the two approximations $\mathbf{x}^\lambda(\cdot)$ and $\mathbf{x}^B(\cdot)$ of the exact subproblem solution $\mathbf{x}^*(\hat{\mathbf{P}})$? Recall that the latter expression solves the consumer’s lifetime plan

$$\max_{\mathbf{x} \geq 0} u(\mathbf{x}) + V(L - \hat{\mathbf{p}} \cdot \mathbf{x}) \quad (19)$$

given anticipated prices $\hat{\mathbf{P}}$ (including subvector $\hat{\mathbf{p}}$) and lifetime income L . To make a fair comparison, suppose that B and λ are both chosen optimally, so $B = \hat{\mathbf{p}} \cdot \mathbf{x}^*$ and $\lambda = V'(L - B)$. We’ve set things up so that all three solutions coincide when subproblem prices are exactly as anticipated, but otherwise, of course, the solutions generally diverge. The question then becomes: how do

¹¹ Textbooks often refer to $\mathbf{x}^B$ as the Marshallian demand function, in distinction to the Hicksian demand function which holds constant the utility level rather than B or λ . The literature survey in Section A.1 will note that $\mathbf{x}^B$ actually owes more to Hicks than to Marshall, whose preferred demand function was a special case of $\mathbf{x}^\lambda$.

the subproblem *price sensitivities* of the approximations $\mathbf{x}^\lambda$ and $\mathbf{x}^B$ compare to that of the exact solution $\mathbf{x}^*$?

The next Proposition provides a precise answer, using the Hessian notation introduced earlier.

Proposition 3 *Given a separable subproblem of dimension $n \geq 2$ and a price vector $\hat{\mathbf{p}}$ with subproblem price vector $\hat{\mathbf{p}}$, let the exact and approximate demand functions $\mathbf{x}^*(\cdot)$, $\mathbf{x}^\lambda(\cdot)$ and $\mathbf{x}^B(\cdot)$ be defined as in previous paragraphs. Then there exist a vector $\mathbf{a}(\hat{\mathbf{p}}, H) \in \mathbb{R}^n$ and a constant $b(\hat{\mathbf{p}}, H) \in \mathbb{R}$, such that for all $i, j \in \{1, \dots, n\}$, the price sensitivities are*

$$\frac{\partial x_j^*}{\partial p_i}(\hat{\mathbf{p}}) = \frac{(-1)^{i+j}(-\hat{p}_i)^{n-2}\lambda|H_{ij}| - a_j V''(L-B)}{(-\hat{p}_i)^{n-2}|H| - b V''(L-B)}, \quad (20)$$

$$\frac{\partial x_j^\lambda}{\partial p_i}(\hat{\mathbf{p}}) = \frac{(-1)^{i+j}\lambda|H_{ij}|}{|H|}, \text{ and} \quad (21)$$

$$\frac{\partial x_j^B}{\partial p_i}(\hat{\mathbf{p}}) = \frac{a_j}{b}. \quad (22)$$

Appendix B contains a proof of Proposition 3, including formulae for $\mathbf{a}(\hat{\mathbf{p}}, H)$ and $b(\hat{\mathbf{p}}, H)$.¹² Of course, equation (21) simply recapitulates Proposition 1.

Comparing the first and third equations in Proposition 3, one can see that $\mathbf{x}^B(\cdot)$ is a reliable approximation of the exact demand function $\mathbf{x}^*(\cdot)$ only if $V''(L-B)$ dominates the other factors. That is, the approximation is reliable only if we can accurately determine the appropriate budget $B = \hat{\mathbf{p}} \cdot \mathbf{x}^*$ and we also know that the indirect utility function is tightly curved at just the right point.

By contrast, no fine-tuning is needed for the constant- λ approximation. When $-V''(L-B)$ is small, i.e., the indirect utility function is nearly linear around the optimal expenditure, then the price sensitivity of $\mathbf{x}^\lambda$ is very close to its true lifetime-optimal value given in the first equation. Indeed (see Appendix B for a proof), the quality of this approximation increases monotonically as curvature decreases, and the approximation becomes exact in the “risk-neutral” (locally linear) case:

Corollary 2 $\left| \frac{\partial x_j^*}{\partial p_i}(\hat{\mathbf{p}}) - \frac{\partial x_j^\lambda}{\partial p_i}(\hat{\mathbf{p}}) \right| \downarrow 0$ as $|V''(L-B)| \downarrow 0$, for each $j \in \{1, \dots, n\}$.

¹² For the single good case we have the following simplified expressions

$$\begin{aligned} \frac{\partial x^*}{\partial p}(\hat{p}) &= \frac{\lambda - BV''(L-B)}{u''(x^*) - p^2 V''(L-B)}, \\ \frac{\partial x^\lambda}{\partial p}(\hat{p}) &= \frac{\lambda}{u''(x^*)}, \text{ and} \\ \frac{\partial x^B}{\partial p}(\hat{p}) &= \frac{B}{p^2}. \end{aligned}$$

Note that the sign of the approximation error depends sensitively on the parameters.

Finally, recall from Section 4 that a decent approximation of $\lambda = V'(L - B)$ is readily available for small separable subproblems and, in contrast to the appropriate budget B , it is independent of the anticipated subproblem prices.

6 Extensions

In practice, some goods are indivisible, and most consumers face liquidity constraints. We now show how the analysis extends to cover these important considerations.

6.1 Liquidity constraints

To deal with constraints on transferring purchasing power across time, we need to augment our model with an additional state variable: the freely available liquidity. When the consumer wishes to spend more than that, she will have to bear the cost of borrowing (captured by an interest rate); when she spends less, she benefits from returns on her saving (again captured by an interest rate). To focus sharply on the main point, we assume that interest rates are constant across maturities and quantities borrowed (or saved); extensions to cover term premiums and quantity premiums are conceptually (if not notationally) straightforward.

Assume, then, that the consumer has available liquid purchasing power $L_0 \in [0, L]$ and earns interest at rate $q \geq -1$ on unspent liquid balances, but pays interest at rate $r \geq \max\{0, q\}$ on expenditures in excess of L_0 . Using the notation $[y]^+ = \max\{0, y\}$, her continuation purchasing power is

$$L - e(\mathbf{p} \cdot \mathbf{x}; L_0), \text{ where } e(\mathbf{p} \cdot \mathbf{x}; L_0) = \mathbf{p} \cdot \mathbf{x} + r[\mathbf{p} \cdot \mathbf{x} - L_0]^+ - q[L_0 - \mathbf{p} \cdot \mathbf{x}]^+. \quad (23)$$

Rewriting (6) as $\tilde{V}(L, L_0) = \max_{x \in \mathfrak{R}_+^n} [u(x) + V(L - e(\mathbf{p} \cdot \mathbf{x}; L_0))]$, the first-order condition becomes

$$u_i(x) = e_i(\mathbf{p} \cdot \mathbf{x}; L_0) V'(L - e(\mathbf{p} \cdot \mathbf{x}; L_0)), \quad i = 1, \dots, n, \quad (24)$$

where

$$e_i(\mathbf{p} \cdot \mathbf{x}; L_0) = \begin{cases} p_i(1 + q), & \text{if } L_0 > \mathbf{p} \cdot \mathbf{x} \\ p_i(1 + r), & \text{if } L_0 < \mathbf{p} \cdot \mathbf{x}. \end{cases} \quad (25)$$

Approximating $V'(L - e(\mathbf{p} \cdot \mathbf{x}; L_0))$ by $\hat{\lambda} = V'(L)$, as before, we have

$$u_i(x) = p_i \tilde{\lambda}, \quad i = 1, \dots, n, \quad (26)$$

where

$$\tilde{\lambda} = \begin{cases} \hat{\lambda}(1 + q), & \text{if } L_0 > \mathbf{p} \cdot \mathbf{x} \\ \hat{\lambda}(1 + r), & \text{if } L_0 < \mathbf{p} \cdot \mathbf{x}. \end{cases} \quad (27)$$

Figure 1 illustrates the omitted case $L_0 = \mathbf{p} \cdot \mathbf{x}$, where the marginal value of money is not defined in (27), but neither is it necessary for the solution as the optimal expenditure is given by the current liquidity.

The only change relative to the unconstrained case is that λ now is a step function rather than a constant: the consumer continues to move out on her IEP until her price-scaled marginal utility drops to λ , which now is increasing in expenditure (and decreasing in liquidity).

Fig. 1 Optimal consumption with liquidity constraint. Expenditure is along the IEP, where $\frac{u_i}{p_i}$ is the same for all goods $i = 1, \dots, n$. In the case shown, optimal expenditure px^* equals available liquidity L_0 .

Note that this set-up is sufficiently general to capture some important special cases. A “paycheck-to-paycheck” consumer is liquidity constrained in the sense that she is unable to borrow and she cannot earn interest on any saving she might have, so $q = 0$ and $r = \infty$. In Figure 1, the opportunity cost of expenditure simply follows the V' line until L_0 , at which point it goes vertical with no ceiling.¹³

¹³ If such a consumer is faced with several purchasing decisions (subproblems) subject to a unified liquidity constraint, then equation (27) and Figure 1, with $q = 0$ and $r = \infty$, apply only to the final decision. For the prior decisions she should still use the unconstrained λ rule, where the constraint will be built into her continuation (indirect) utility. When the borrowing constraint binds, the consumer’s time horizon effectively shrinks to a single pay period. This would decrease the precision of the estimate of λ , leading to a normative and a positive prediction. First, the final decision should incorporate more goods than usual to decrease the bias in λ . Second, as the estimate is an underestimate – as $V'(L - px) > V'(L)$ – paycheck-to-paycheck consumers using the moneyworth demand should be bad at consumption smoothing and get to the end of the month short of money.

The budget-constrained consumer featured in textbooks is an even more extreme special case. He is supposed to be unable even to store any of the subproblem budget, so $r = \infty$ and $q = -1$. In Figure 1, the bottom step is on the horizontal axis and the other step is in the sky, so the consumer always “chooses” to spend exactly her liquid assets.

It may be worth noting that liquidity constraints do not bring back the textbook income effect. The marginal effect of a change in liquidity continues to be zero, except when subproblem expenditure is close to available liquidity L_0 , and even here the effect is attenuated to the extent that the jump in λ is finite.

A thought experiment offers some quantitative perspective. Assume that at the current price (normalized to 1), our consumer is unconstrained but spends all her liquid wealth on a (composite) good. We now ask the question: by how much does the price have to decrease, so that she is willing to increase her expenditure and borrow? In terms of Figure 1, what proportional price decrease will move u'/p up from the bottom of the step to the top? Inspection of (27) yields the answer of $\frac{1+q}{1+r} \approx q - r$. For a typical consumer the rate difference between borrowing via credit card or depositing in a savings account might be about 2% per month. That is, a mere 2% price decrease would remove her liquidity-constrained inertia. For larger price decreases she would seem identical to an unconstrained consumer with $\lambda = \hat{\lambda}(1 + r)$, and her moneysworth demand would display no income effect.

6.2 Indivisible goods

Suppose that the consumer faces the small separable subproblem of whether or not to buy a single indivisible good (or basket of goods) at price p . Indivisibility is captured in the constraint $x \in \{1, 0\}$, and we normalize $u(0) = 0$. Thus the objective function (9) becomes

$$\max_{x \in \{1, 0\}} [u(x) - \lambda xp] = \max \{0, u(1) - \lambda p\}. \quad (28)$$

Dividing by p , one can say that the consumer calculates the ratio $\frac{u(1)}{p}$ of perceived quality to price and compares it to λ . If the quality-price ratio, interpreted as value for money, exceeds the marginal utility of money, then she will buy, and otherwise not buy.¹⁴

When the consumer has to choose just one of several mutually exclusive varieties or baskets, the quality-price ratio no longer suffices. A very small basket may offer a high value for money, but still provide only a small utility

¹⁴ Hauser and Urban (1986) pose as alternative hypotheses that consumers use “value for money,” u/p , or “net value,” $u - \lambda p$, to prioritize purchases of indivisibles. Our analysis shows that the two rankings are equivalent for yes/no decisions, but we shall now show that “net value” is the appropriate criterion for mutually exclusive alternatives.

gain.¹⁵ Instead, the consumer should rank baskets $\mathbf{b}^k = (x_1^k, \dots, x_n^k)$ of indivisibles (so each $x_i^k = 0$ or 1) at price vector $\mathbf{p}$ according to their net utility gain, $g^k = u(\mathbf{b}^k) - \lambda \mathbf{p} \cdot \mathbf{b}^k$. Then the consumer picks the basket with highest g^k as long as it is positive, and otherwise picks the null basket $\mathbf{b}^0 = (0, \dots, 0)$ at price $\mathbf{p} \cdot \mathbf{b}^0 = 0$ and gain $g^0 = u(0) - \lambda \mathbf{p} \cdot \mathbf{b}^0 = 0$. It follows (after dropping any basket that is dominated by another basket with lower price and higher utility) that basket k will be preferred to basket j if and only if $\frac{u(\mathbf{b}^k) - u(\mathbf{b}^j)}{\mathbf{p} \cdot \mathbf{b}^k - \mathbf{p} \cdot \mathbf{b}^j} \geq \lambda$. The gain in utility by choosing k over j must exceed the shadow utility of the additional expenditure, i.e., the *incremental* quality-price ratio must exceed the marginal utility of money.¹⁶

By contrast, the budget constrained consumer would pick the highest quality item that her budget permits. For example, suppose that the consumer has two baskets available, with $u(\mathbf{b}^1) < u(\mathbf{b}^2)$ and $\mathbf{p} \cdot \mathbf{b}^1 = B < \mathbf{p} \cdot \mathbf{b}^2 = \mathbf{p} \cdot \mathbf{b}^1 + \varepsilon$. If and only if

$$u(\mathbf{b}^2) - u(\mathbf{b}^1) < \lambda \varepsilon \quad (29)$$

we have $0 < g^2 = u(\mathbf{b}^2) - \lambda \mathbf{p} \cdot \mathbf{b}^2 < u(\mathbf{b}^1) - \lambda \mathbf{p} \cdot \mathbf{b}^1 = g^1$, so that both decision rules lead to the purchase of cheaper basket 1. Condition (29) captures how the consumer trades off instantaneous utility gain against the shadow value of money. If the price difference is small enough, she will go for the more expensive basket (unlike the budget constrained consumer). (29) also shows how to respond to the appearance of a new variety, or a change in the valuation of an existing variety. Suppose that the perceived quality difference between the baskets increases sufficiently, keeping the price difference the same. This change would not affect the choice according to the budget rule, but it would again lead to a switch according to the λ rule (and according to lifetime optimization).

The purchase of a single big ticket item, like an automobile, introduces two variations. First, to have a better approximation, the value at which the Taylor expansion is centered should be L minus the representative price of the cars considered. Second, if the price range under consideration is large (e.g., both Fiat and Ferrari are in the relevant choice set) then the marginal utility of money might be higher for more expensive cars. Thus we should set a different marginal utility of money λ_i for each different price(range). The consumer should then choose the variety, q^i , which maximizes her net utility gain $u(q^i) - \lambda_i p_i$.

7 Discussion

The moneysworth approach to consumer choice centers on λ , the marginal utility of money. Defining λ as the opportunity cost of expenditure in the con-

¹⁵ Note that this crucial detail is overlooked in the *melioration* theory of Herrnstein and Prelec (1991), which posits that the option with higher utility per \$ is the one that is chosen (in a distributed choice problem).

¹⁶ Note that for yes/no decisions the quality-price ratio is also incremental but the benchmark is normalized to zero.

tinuation problem that follows a separable subproblem, we obtain the rule that expenditure on a good (or on baskets of goods along the appropriate income expansion path, IEP) should increase until its marginal utility diminishes to λ .

For small separable subproblems, we obtain moneysworth demand functions $\mathbf{x}^\lambda$ that share some features with their standard “Marshallian” counterparts $\mathbf{x}^B$ – notably, both lie on the same subproblem IEP – but $\mathbf{x}^\lambda$ has several distinct advantages.

- It is very simple and specific – the single number λ is a sufficient statistic for the hugely complex rest-of-life problem. By contrast, each subproblem requires its own budget B in the standard approach, and short of re-solving the lifetime problem, it is unclear how B might be determined.
- It is robust to changes in subproblem prices; no change in λ is required when p changes. By contrast, the appropriate B depends sensitively on p .
- Its elasticities are a first-order approximation of the true (lifetime-optimal) elasticities. Standard Marshallian elasticities are close to true elasticities only in the special case that they coincide with $\mathbf{x}^\lambda$ ’s.
- It is quite intuitive – the consumer tries directly to get her money’s worth.¹⁷

Our moneysworth approach offers a fresh perspective on the traditional distinction between non-pecuniary consumption externalities and pecuniary externalities. The former recognizes that how much the consumer values a certain quantity of a good may depend on what else is in her consumption basket. Often this kind of externality extends only to a small set of goods, e.g., a few complements and close substitutes, and so can be internalized in a low-dimensional separable subproblem of the lifetime problem. On the other hand, pecuniary externalities are pervasive because expenditures are mutually exclusive: money spent on one good is not available to purchase *any* other good. But λ precisely captures this kind of externality.

By contrast, it is hard to specify the appropriate subproblem to which to apply a budget B in the standard approach. The larger the set of goods, the better internalized is the pecuniary externality, and the closer one gets to the true lifetime problem. But at the same time, one loses the tractability of low-dimensional partial equilibrium analysis.

Standard partial equilibrium analysis routinely conflates the available liquidity L with the expenditure B targeted at a subproblem. Otherwise put, the natural subproblem boundaries need not coincide with the boundaries of binding liquidity constraints. Given access to perfectly liquid financial markets, it is erroneous to specify B other than the lifetime budget constraint. When there are additional liquidity constraints, they are best dealt with as in section 6.1, where the textbook analysis is shown to be an extreme and unrealistic special case.

A key insight from the moneysworth approach is that budgets should not be applied piecemeal to subproblems, but instead all subproblems should be

¹⁷ We suspect that it is also quite descriptive of actual consumer behavior, but we are not aware of any empirical studies so far that compare the predictive power of $\mathbf{x}^\lambda$ and $\mathbf{x}^B$.

solved consistently using a single sufficient statistic, the marginal utility of money. The point is that funds are fungible, as has long been recognized in other branches of economics.¹⁸

The moneysworth approach is intrinsically cardinal, and thus goes against textbook orthodoxy favoring ordinal preferences. But that orthodoxy seems increasingly anomalous. After modernization by Von Neumann and Morgenstern (1944) and especially following Friedman and Savage (1948), cardinal preferences have been routinely assumed in risky choice theory. They are ubiquitous in applied work (e.g., in industrial organization) and in game theory, and of course production functions remain cardinal. Why should consumer choice in the absence of risk be the only topic left that insists on using ordinal preferences?¹⁹

To streamline our presentation, we focused on consumer choice, but similar reasoning applies to work decisions. For example, a piece rate worker (or consultant) optimally works until the marginal disutility equals the wage times λ . Likewise the “big ticket” version of the λ rule tells a worker whether to accept a salaried job, or one with fixed hours. We also streamlined the presentation by neglecting corner solutions, multiple optima and non-differentiable solutions. We leave all these generalizations for future research.

One final thought. For several generations, economics instructors have tortured undergraduates with various decompositions of income and substitution effects, with picky distinctions between gross and net substitutes and complements, and with ordinal versus cardinal utility. An important benefit of the moneysworth approach is that it sweeps away such dross, and brings consumer choice theory closer to common sense.

A Appendix: Connections

Our approach to consumer choice connects numerous topics that economists previously have considered in isolation. Here are some that have caught our attention.

A.1 Historical perspectives

The idea of a cardinal utility function defined over purchasing power goes back at least to Bernoulli (1738) and Cramer (1728), but Bentham (1802) was apparently the first to

¹⁸ In public finance, for example, it is well known that a uniform income tax is more efficient for raising a given amount of revenue than a collection of specific taxes on individual items.

¹⁹ Arguably, revealed preference theory is the real contribution of ordinality. A companion paper (Sákovics 2013) shows that revealed preference theory has a cardinal counterpart, with the Generalized Axiom of Revealed Preference replaced by the stronger Axiom of Revealed Valuation (ARV). Formally, assume that we have T observations of a consumer's chosen bundles, $x^t \in \mathbb{R}_+^n$, at price vectors, $p^t \in \mathbb{R}_{++}^n$. Let $a_{tj} = p^t \cdot (x^j - x^t)$ denote the pecuniary advantage of the chosen bundle x^t relative to an arbitrary bundle x^j . Then ARV states: *For every ordered subset $\{i, j, k, \dots, s\} \subseteq \{1, 2, \dots, T\}$, $a_{ij} + a_{jk} + \dots + a_{si} \geq 0$. The main theorem states: If and only if the observed choices satisfy ARV, they are rationalizable by a cardinal utility function of the form $u(x) - \lambda p \cdot x$.*

develop the idea to explain choice in the absence of risk. These economists, and their successors, presumed that the utility function was concave, if only because the more attractive consumption opportunities would be selected first.

Marshall's (1890, 1920) theory of consumer demand centered on λ . In the special case that utility is additively separable in each good, he obtained the crucial first order condition that marginal utility for each good equals its price times the marginal utility of money. Given the market price p_i for some good i (tea was Marshall's favorite example), the consumer increases or decreases his consumption x_i until the marginal utility it brings is equal to the market price scaled by λ , the marginal utility of money, assumed exogenous and constant.

"Edgeworth destroyed this pleasant simplicity and specificity when he wrote the total utility function as $f(x_1, x_2, x_3, \dots)$," says Stigler (1950, p. 322). It fell to Hicks and Allen (1934) to show how to impose a budget set to derive demand functions and cross-price elasticities when goods might have complements and substitutes. Their analysis, developed further by Paul Samuelson and a host of other economists, ultimately became textbook orthodoxy.

Biswas (1977) eventually showed that Marshall's approach could be extended to interdependent goods without the imposition of a budget constraint. He noted that the marginal utility of money is constant only in the short run and thus its adjustment is of relevance. Biswas' purpose was to explain Marshall's approach, so he did not develop a full-fledged alternative theory of the consumer; e.g., as far as we know he did not show the relation to the IEP or explore the basis of the constant λ assumption. He was content to establish the existence of a downward-sloping demand function satisfying Marshall's first-order conditions, and labored to salvage Giffen goods by writing out an adjustment process for λ based on the – what seems to us peculiar – assumption that the consumer observes the *ex post* optimal expenditure (E^*) while she adjust her λ only partially.

A.2 Macro connections

Frisch (1932, 1959) was among the first economists to pose a representative agent's intertemporal consumption problem as maximizing the discounted sum of a fixed concave felicity (i.e., subutility) function v for composite consumption c each period t , subject only to a lifetime budget constraint. Of course, relative to the general equilibrium problem posed in Section 2, Frisch's problem is a special case in which the continuation is an evenly spaced sequence of subproblems identical to the current subproblem, and separation is built in. Assuming that the agent's personal discount rate is equal to the interest rate in a frictionless financial market, the first order (i.e., Euler) condition for this optimization problem states that marginal felicity $v'(c_t)$ is constant across periods and equal to the Lagrange multiplier on the budget constraint (e.g., Deaton 1992, Chapter 1, equations 5-9).²⁰ Here λ is not a simple approximation, but rather is the exact GE opportunity cost of expenditure, albeit in a highly stylized model.

Macroeconomists exploring consumption smoothing and the permanent income hypothesis adapted and extended Frisch's approach.²¹ The papers in this tradition that come closest to our own work were written in the 1980s by Deaton, Browning, and various coauthors. Browning (1985, 2005) re-expresses the Euler equation in terms of the marginal utility of expenditure, and obtains what we would call a moneyworth demand function. Referring to this function as Frisch demand, he argues that it is more relevant to fiscal policy analysis than the usual constant-budget ("Marshallian") or constant-utility ("Hicksian") demand

²⁰ Allen (1933) interpreted Frisch's constant λ as an unjustified imposition of additional restrictions on the utility function. Eventually, Brown and Calsamiglia (2003) noted that Frisch demand functions are better understood as related to Marshall's theory of consumer demand.

²¹ Heckman (1974) was apparently the first to use λ -constant comparative statics in such models. Bewley (1977) proposed a consumer theory for the life-time problem using the constancy of marginal utility of money. See Bewley (2007, Chapter 8.3) for a recent synthesis.

functions, and shows how it can be estimated from aggregate data despite the unobservability of λ . Browning et al. (1985) consider the somewhat more abstract problem where consumption is additively separable in all goods (or at least “block additive,” i.e., that it consists of a finite number of separable subproblems) and is subject to a single overall budget constraint. They construct the consumer “profit function,” a renormalized version of our objective function (9), and use variants on textbook duality identities to obtain properties of the corresponding Frisch demand function. The main properties, some of which help to identify econometric specifications, are (a) degree 0 homogeneity in p and $r = 1/\lambda$ (which also is transparent from our equation (10) or even from (9)); (b) symmetry of cross partials (as noted after our Prop 1); (c) substitution matrix is proportional to H^{-1} (as can also be seen from Prop 1); and (d) downward sloping demand (again noted after Prop 1). The article also gives formulas for moving among Marshallian, Hicksian and Frisch demand.

The concerns of this strand of literature – consumption smoothing and estimating elasticities from aggregate national consumption data – are quite different from ours, and consequently so are the contributions. This literature does not consider the quality of approximations to GE demand as in our Proposition 3 (they do not need to approximate in their stylized GE model), nor investigate how λ might be determined and adjusted at the disaggregated individual level, nor discuss separability and the microfoundations of individual consumer behavior.

A.3 Behavioral connections

Miscalibration. Some consumers may have a persistently biased view of their true marginal utility for money. A *miser* is someone who (perhaps because of an impoverished childhood and stalled adjustment) maintains an unreasonably high λ , relative to typical preferences and to his actual lifetime opportunities. Likewise, a *spendthrift* maintains an unreasonably low λ . Thus, if λ is badly calibrated, the moneysworth approach will describe some prominent forms of suboptimal behavior.

Money illusion. Lifetime utility $U(X)$ depends only on actual consumption, and not directly on nominal quantities such as L and P ; recall that the indirect utility function is homogeneous of degree 0 in (L, P) . Hence a proportional change in all prices (and income) will have no effect on the consumption plans or utility of *Homo Economicus*. The same is true in the textbook treatment of a subproblem with the budget B adjusted in the same proportion as the price vector.

However, inflation in the real world is much messier than a simultaneous proportional price change. An important empirical regularity is that relative prices become more volatile as the measured rate of inflation rises (see, for example, Heymann and Leijonhufvud 1995). Actual people updating λ , therefore, are unlikely to immediately adjust to a change in the price level. According to (16) they will lag in reacting to observed nominal increases in prices of purchased goods (and, by extension, overreact to nominal increases in wages; see, for example, Genesove and Mayer 2001). That is, in the parlance of macroeconomic theory, they will suffer from *money illusion*.

If prices settle down, an adaptive moneysworth consumer will, after some lag, find a new λ appropriate to the new price level. The illusion eventually fades, as it should.

System 1 vs 2. Our consumer has two operational phases: normally she just goes around with her λ and makes quasi-automatic decisions, while occasionally she updates her λ , using a more thoughtful procedure. This dual process fits well with the two-system approach originated by Stanovich and West (2000), surveyed by Evans and Frankish (2009), and popularized by Kahneman (2011). System 1 works like a reflex and has no significant cost to operate (a bit like the body’s vegetative system), while System 2 ponders decisions but requires time and attention, which are in limited supply. There are certain cues in the environment which give the control over to System 2, otherwise the default decision maker is System 1. Our model goes beyond this switching scheme by introducing an element of communication between the systems: λ .

By tracking neural responses to the observation of the price of an item, Knutson et al. (2007) show that the price of an item generates a significant cue about whether the item

will be purchased. This implies that the decision to buy is mostly based on the price (and the utility value) of the good, fitting well with moneysworth demand, and its System 1 interpretation.

Mental accounting. Read et al. (1999) and Thaler (1999) launched a sizeable empirical literature called narrow bracketing or mental accounting. In our terminology, the idea is that consumers sometimes treat non-separable subproblems as if they were separable, resulting in inefficient choices. For example, a consumer may regard the health risk as negligible when considering smoking a single cigarette. The addictive properties of nicotine warrant a broader definition of the subproblem, and the consumer might come to a different conclusion when considering smoking a pack a day for a year, or for a decade.

We see the moneysworth approach as sharpening the empirical questions to be investigated. First, when are humans more likely to treat non-separable consumption problems as separable? Second, when faced with a subproblem, separable or otherwise, when do humans use a fixed budget to make their choice, and when do they use the marginal utility of money?

The empirical literature has made progress on the first question, especially in highlighting self-control issues. For example, a liquidity-constrained gambler may bring exactly \$2500 to Las Vegas in an attempt to ensure that he can pay the rent next month. (The ubiquitous presence of cash machines in casinos suggests that this self-control device is less than 100% successful!) We believe that the question can be asked more broadly; perhaps there are cognitive issues as well as self-control issues in defining separable subproblems. The answer to this empirical question is at least as important to our moneysworth approach as it is to the standard approach.

In our reading of the mental accounting literature, the second question has seldom been asked, even though to us it seems pivotal. The question seems amenable to laboratory investigation, especially since payments received in the lab are normally quite separate from the rest of the subject's life. One could set up various moderately complicated lab tasks, and in one treatment frame the tasks in terms of budgets and in another treatment frame them in terms of the marginal utility of money. Then one could check which treatment encouraged more efficient choices. A more direct test is simply to provide access to budgeting tools as well as λ -oriented tools, and to see which the subjects prefer to use.

Other biases. Many sorts of behavioral biases can be captured by assigning plausible but non-optimal values to λ . Non-standard time preferences illustrate the procedure. An agent wishing to pay debt early and delay payment for work done as in Prelec and Loewenstein (1998) can be accommodated by positing a higher λ for consumption and a lower one for work income.

Considering our model as the rational benchmark, it can “suffer” from yet other behavioral biases. For example, as in Gennaioli and Shleifer (2010)²², the consumer may engage in “local thinking” by not following through with a full-blown updating of λ , rather recalling the small subset of situations that first come to her mind.

Appendix B: Mathematical Details

Proof of Proposition 3. Note that all three setups share the $n - 1$ equations defining the IEP:

$$p_i u_k = p_k u_i \quad \text{for } k \neq i. \quad (30)$$

The n^{th} equations are²³

$$u_i(\mathbf{x}) = p_i V'(L - \mathbf{p} \cdot \mathbf{x}), \quad u_i(\mathbf{x}) = p_i \lambda, \quad \text{and } \mathbf{p} \cdot \mathbf{x} = B, \quad (31)$$

²² See also Bordalo et al. 2013, for a related model based on the salience of past observations.

²³ We have chosen the i^{th} coordinates as we will be checking sensitivity to the i^{th} price.

for the global, the λ and the budget solutions, respectively. Differentiation of the common equations with respect to p_i yields

$$u_k + p_i \sum_m u_{km} \frac{\partial x_m}{\partial p_i} = p_k \sum_m u_{im} \frac{\partial x_m}{\partial p_i}, \quad (32)$$

which simplifies to

$$\sum_m (p_k u_{im} - p_i u_{km}) \frac{\partial x_m}{\partial p_i} = u_k. \quad (33)$$

Differentiating the other three equations, we have

$$\sum_m u_{im} \frac{\partial x_m}{\partial p_i} = V'(L - \mathbf{p} \cdot \mathbf{x}) - p_i \left(x_i + \sum_m p_m \frac{\partial x_m}{\partial p_i} \right) V''(L - \mathbf{p} \cdot \mathbf{x}), \quad (34)$$

$$\sum_m u_{im} \frac{\partial x_m}{\partial p_i} = \lambda, \quad (35)$$

and

$$x_i + \sum_m p_m \frac{\partial x_m}{\partial p_i} = 0. \quad (36)$$

If $V''(L - B) = 0$, then it is immediate that the price sensitivities of the lambda rule will coincide with the optimal ones at $\mathbf{p} = \hat{\mathbf{p}}$, while the budget rule would only be a good approximation by coincidence. Otherwise, denoting the i^{th} row of the Hessian of $u(\mathbf{x}^*(\hat{\mathbf{p}}))$ by H_i , we obtain the following three equation systems:

$$\begin{pmatrix} \hat{p}_1 H_i - \hat{p}_i H_1 \\ \hat{p}_2 H_i - \hat{p}_i H_2 \\ \vdots \\ H_i + \hat{\mathbf{p}} \hat{p}_i V'' \\ \vdots \\ \hat{p}_n H_i - \hat{p}_i H_n \end{pmatrix} \begin{pmatrix} \frac{\partial x_1^*}{\partial p_i} \\ \frac{\partial x_2^*}{\partial p_i} \\ \vdots \\ \frac{\partial x_i^*}{\partial p_i} \\ \vdots \\ \frac{\partial x_n^*}{\partial p_i} \end{pmatrix} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ V' - x_i^* \hat{p}_i V'' \\ \vdots \\ u_n \end{pmatrix} \quad (37)$$

$$\begin{pmatrix} \hat{p}_1 H_i - \hat{p}_i H_1 \\ \hat{p}_2 H_i - \hat{p}_i H_2 \\ \vdots \\ H_i \\ \vdots \\ \hat{p}_n H_i - \hat{p}_i H_n \end{pmatrix} \begin{pmatrix} \frac{\partial x_1^\lambda}{\partial p_i} \\ \frac{\partial x_2^\lambda}{\partial p_i} \\ \vdots \\ \frac{\partial x_i^\lambda}{\partial p_i} \\ \vdots \\ \frac{\partial x_n^\lambda}{\partial p_i} \end{pmatrix} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ \lambda \\ \vdots \\ u_n \end{pmatrix} \quad (38)$$

$$\begin{pmatrix} \hat{p}_1 H_i - \hat{p}_i H_1 \\ \hat{p}_2 H_i - \hat{p}_i H_2 \\ \vdots \\ \hat{\mathbf{p}} \\ \vdots \\ \hat{p}_n H_i - \hat{p}_i H_n \end{pmatrix} \begin{pmatrix} \frac{\partial x_1^B}{\partial p_i} \\ \frac{\partial x_2^B}{\partial p_i} \\ \vdots \\ \frac{\partial x_i^B}{\partial p_i} \\ \vdots \\ \frac{\partial x_n^B}{\partial p_i} \end{pmatrix} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ -x_i^B \\ \vdots \\ u_n \end{pmatrix} \quad (39)$$

Denoting the matrices by G, L and D , respectively, we can apply Cramer's rule to solve the systems of equations:

$$\frac{\partial x_j^*}{\partial p_i} = \frac{|G^j|}{|G|}, \frac{\partial x_j^\lambda}{\partial p_i} = \frac{|L^j|}{|L|}, \text{ and } \frac{\partial x_j^B}{\partial p_i} = \frac{|D^j|}{|D|}, \quad (40)$$

where a superindexed matrix means that its j^{th} column is replaced by the vector on the right-hand side of its corresponding system. Since when the realized prices equal the expected ones $V'(L - \mathbf{p} \cdot \mathbf{x}) = \lambda$ and $x_i^* = x_i^B$ it is immediate that

$$|G| = |L| + |D|\widehat{p}_i V'' \text{ and } |G^j| = |L^j| + |D^j|\widehat{p}_i V''. \quad (41)$$

By subtracting multiples of row i in L it is easy to see that $|L| = (-\widehat{p}_i)^{n-1} |H|$. By the same method – using the fact that $u_k = p_k \lambda$ for all k – we can ensure that the j^{th} column of $|L^j|$ has zeros everywhere except in the i^{th} place where it has λ . Writing out the Laplace expansion it is immediate that $|L^j| = (-1)^{i+j} \lambda (-\widehat{p}_i)^{n-1} |H_{ij}|$. Finally, let $a_j = |D^j|$ and $b = |D|$. Q.E.D.

Proof of Corollary 2: $\left| \frac{|L^j| + a_j \widehat{p}_i V''}{(-\widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V''} - \frac{|L^j|}{(-\widehat{p}_i)^{n-1} |H|} \right| = \left| \frac{-(a_j (-\widehat{p}_i)^n |H| + |L^j| b \widehat{p}_i) V''}{((-\widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} \right| =$
 $\frac{-|a_j (-\widehat{p}_i)^n |H| + |L^j| b \widehat{p}_i V''}{|((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|}|$, that is clearly equal to zero when $V'' = 0$. To see that the convergence is monotone, note that a_j and b are independent of V'' :

$$\frac{d}{dV''} \left| \frac{V''}{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} \right| = \begin{cases} \frac{\left| \frac{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|}{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} - V'' \right| b \widehat{p}_i (-\widehat{p}_i)^{n-1} |H|}{\left| \frac{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|}{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} \right|^2} \\ \frac{\left| \frac{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|}{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} + V'' \right| b \widehat{p}_i (-\widehat{p}_i)^{n-1} |H|}{\left| \frac{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|}{((- \widehat{p}_i)^{n-1} |H| + b \widehat{p}_i V'') (-\widehat{p}_i)^{n-1} |H|} \right|^2} \end{cases}, \quad (42)$$

where we obtain the upper line if the denominator on the left-hand side is increasing in V'' and the bottom line otherwise. As $V'' < 0$, the upper line is clearly positive. The bottom line as well, as the denominator is decreasing in V'' implying that $b \widehat{p}_i V'' (-\widehat{p}_i)^{n-1} |H| > 0$, while $((-\widehat{p}_i)^{n-1} |H|)^2 > 0$ as well. Q.E.D.

Acknowledgements This paper supersedes the previously circulated “The marginal utility of money: a new Marshallian approach to consumer choice”. We thank Olga Rud Rabanal for research assistance. For helpful suggestions we thank two anonymous referees of this Journal, Luciano Andreozzi, Buz Brock, Roberto Burguet, Andrew Clausen, Joan-Maria Esteban, Steffen Huck, Axel Leijonhufvud, Youn Kim, Michael Mandler, Albert Marcet, Carmen Matutes, David de Meza, Ryan Oprea, Martin Shubik, Peter Simmons, Nirvikar Singh, Hal Varian, Donald Wittman, and seminar audiences at Barcelona GSE, Cambridge, CEU (Budapest), Edinburgh, Glasgow, Roy Seminar (Paris), UBC-UHK Summer Micro Seminars (Hong Kong), UCIII (Madrid), UC at Santa Cruz, and York (England).

References

1. Allen, R.G.D. (1933) “On the Marginal Utility of Money and Its Application” *Economica*, No. 40, 186-209.
2. Bentham, J. (1802) *Theory of legislation*. Edited by E. Dumont, translated from the original French by C.M. Atkinson in 1914. Oxford University Press.
3. Bernoulli, D. (1738) “Specimen theoriae novae de mensura sortis.” English translation, “Exposition of a New Theory on the Measurement of Risk,” *Econometrica*, Jan. 1954, 22, 23-26.
4. Bewley, T. (1977) “The Permanent Income Hypothesis: A Theoretical Formulation,” *Journal of Economic Theory*, 16, 252-292.
5. Bewley, T. (2007) *General Equilibrium, Overlapping Generations Models, and Optimal Growth Theory*, Harvard University Press.
6. Biswas, T. (1977) “The Marshallian Consumer,” *Economica*, 44(173), 47-56.
7. Bordalo, P., Gennaioli, N. and Shleifer, A. (2013) “Salience and Consumer Choice,” *Journal of Political Economy*, 121(5), 803-843.
8. Brown, D.J. and Calsamiglia, C. (2003) “The Strong Law of Demand”, Cowles Foundation Discussion Paper No. 1399.
9. Browning, M. (2005) “A working paper from April 1985: Which demand elasticities do we know and which do we need to know for policy analysis?” *Research in Economics*, 59, 293-320.
10. Browning, M., Deaton, A. and Irish, M. (1985) “A profitable approach to labor supply and commodity demands over the life-cycle,” *Econometrica*, 53(3), 503-544.
11. Cramer, G. (1728) *Letter to Nicolas Bernoulli*. Reprinted in Bernoulli (1738).
12. Deaton, A. (1977) “Involuntary Saving Through Unanticipated Inflation,” *American Economic Review*, 67(5), 899-910.
13. Deaton, A. (1992) *Understanding Consumption*, Oxford: Clarendon Press.
14. Dwyer, G.P. Jr. and Lindsay, C.M. (1984) “Robert Giffen and the Irish Potato,” *American Economic Review*, 44(1), 188-192.
15. Evans, J.S.B.T. and Frankish, K. eds. (2009) *In Two Minds: Dual Processes and Beyond*, OUP, New York.
16. Friedman, M., and Savage, L. J. (1948). “The utility analysis of choices involving risk,” *The Journal of Political Economy* 56(4), 279-304.
17. Frisch, R. (1932) *New Methods of Measuring Marginal Utility*. J.C.B. Mohr (Paul Siebeck). Tubingen.
18. Frisch, R. (1959) “A Complete Scheme for Computing All Direct and Cross Demand Elasticities in a Model with Many Sectors,” *Econometrica* 27, 177-196.
19. Genesove, D. and Mayer, C. (2001) “Loss Aversion and Seller Behavior: Evidence from the Housing Market,” *Quarterly Journal of Economics*, 116(4), 1233-1260.
20. Gennaioli, N. and Shleifer, A. (2010) “What Comes to Mind,” *Quarterly Journal of Economics*, CXXV(4), 1399-1433.
21. Hauser, J.R. and Urban, G.L. (1986) “The Value Priority Hypotheses for Consumer Budget Plans” *Journal of Consumer Research*, 12(4), 446-462.
22. Heckman, J. (1974) “Life Cycle Consumption and Labor Supply: An Explanation of the Relationship between Income and Consumption Over the Life Cycle,” *American Economic Review*, 64(1), 88-94.

23. Herrnstein, R.J. and Prelec, D. (1991) "Melioration: A Theory of Distributed Choice," *Journal of Economic Perspectives*, 5(3), 137-156.
24. Heymann, D. and Leijonhufvud, A. (1995) *High Inflation*, The Arne Ryde Memorial Lectures, Oxford University Press.
25. Hicks, J.R. and Allen, R.G.D. (1934) "A Reconsideration of the Theory of Value," *Economica*, Part I: 1(1), 52-76; Part II: 1(2), 196-219.
26. Jensen, R.T. and N.H. Miller. (2008) "Giffen Behavior and Subsistence Consumption," *American Economic Review*, 98(4): 1553-77.
27. Kahneman, D. (2011) *Thinking, Fast and Slow*, Farrar, Straus and Giroux.
28. Knutson, B., Rick, S., Wimmer, G.E., Prelec, D. and Loewenstein, G. (2007) "Neural Predictors of Purchases," *Neuron*, 53, 147-156.
29. Love, D.A. (2013) "Optimal Rules of Thumb for Consumption and Portfolio Choice," *The Economic Journal*, 123 (September), 932-961.
30. Marshall, A. (1890/1920) *Principles of economics*. London: Macmillan.
31. Nachbar, J.H. (1998) "The last word on Giffen goods?" *Economic Theory*, 11(2), 403-412.
32. von Neumann, J. and Morgenstern, O. (1944) *Theory of Games and Economic Behavior*. Princeton University Press, New Jersey.
33. Prelec, D. and Loewenstein, G. (1998) "The red and the black: mental accounting of savings and debt" *Marketing Science*, 17(1), 4-28.
34. Rabin, M. (2000) "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68(5), 1281-1292.
35. Read, D., Loewenstein, G. and Rabin, M. (1999) "Choice bracketing," *Journal of Risk and Uncertainty*, 19(1-3), 171-197.
36. Sákovic, J. (2013) "Revealed Cardinal Preference," *Economic Theory Bulletin*, 1(1), 39-45.
37. Stanovich, K.E. and West, R.F. (2000) "Individual Differences in Reasoning: Implications for the Rationality Debate," *Behavioral and Brain Sciences*, 23, 645-665.
38. Stigler, G. J. (1950) "The Development of Utility Theory. I," *Journal of Political Economy*, 58, 307-27.
39. Thaler, R.H. (1999) "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 12(3), 183-206.
40. Varian, H.R. (1992) *Microeconomic Analysis*, Third Edition, W.W. Norton and Co., New York.