A Behavioral Update of Basic Microeconomics

Based on Gabaix (2014).

> Does not depends on the details of the endogenization of attention (i.e. from sparsity or some other procedure).
>
> Effect works for a host of behavioral models, provided they generate some inattention to prices.

8.1 Textbook consumer theory: A behavioral update

8.1.1. Basic consumer theory: Marshallian demand

Consumer's Marshallian demand:

$$c(\boldsymbol{p}, w) := \arg \max_{\boldsymbol{c} \in \mathbb{R}^n} u(\boldsymbol{c}) \text{ subject to } \boldsymbol{p} \cdot \boldsymbol{c} \leq w$$

$\boldsymbol{c}^r(\boldsymbol{p}, w)$ is demand under the traditional rational model

$\boldsymbol{c}^s(\boldsymbol{p}, w)$ is the demand of a behavioral agent

> s stand for: demand given "subjectively perceived prices.

Price of good i: $p_i = p_i^d + x_i$

> p_i^d is the default price (e.g., the average price)
> x_i is an innovation.

Price perceived by a behavioral agent:

$$p_i^s = p_i^d + m_i x_i$$

Or:

$$p_i^s(m) = m_i p_i + (1 - m_i) p_i^d$$

$m_i = 1$: agent fully perceives price p_i

$m_i = 0$: replaces it by the default price.

Proposition 6.1 (Marshallian demand). Given the true price vector $\boldsymbol{p}$ and the perceived price vector $\boldsymbol{p}^s$, the Marshallian demand of a behavioral agent is

$$\boldsymbol{c}^s(\boldsymbol{p}, w) = \boldsymbol{c}^r(\boldsymbol{p}^s, w'),$$

where the as-if budget w' solves $\boldsymbol{p} \cdot \boldsymbol{c}^r(\boldsymbol{p}^s, w') = w$

w' ensures that the budget constraint is hit under the true price

(if there are several such w', take the largest one).

Example 2 (Demand by a behavioral agent with quasi-linear utility)

$u(\boldsymbol{c}) = v(c_1, \ldots, c_{n-1}) + c_n$

v strictly concave

Demand for good $i < n$ is independent of wealth and is: $c_i^s(\boldsymbol{p}) = c_i^r(\boldsymbol{p}^s)$.

Demand of the behavioral agent is the rational demand given the perceived price for all goods but the last one

Residual good n is the "shock absorber" that adjusts to the budget constraint.

In a dynamic context, this good n could be "savings".

Example 3 (Demand proportional to wealth)

When rational demand is proportional to wealth, the demand of a behavioral agent is

$$c_i^s(\boldsymbol{p}, w) = \frac{c_i^r(\boldsymbol{p}^s, w)}{\boldsymbol{p} \cdot \boldsymbol{c}^r(\boldsymbol{p}^s, 1)}$$

Example 4 (Demand by behavioral Cobb-Douglas and CES agents)

$u(\boldsymbol{c}) = \sum_{i=1}^{n} \alpha_i \ln\, c_i,\ \alpha_i \geq 0$

Demand:

$$c_i^s(\boldsymbol{p}, w) = \frac{\alpha_i}{p_i^s} \frac{w}{\sum_j \alpha_j \frac{p_j}{p_j^s}}$$

$u(\boldsymbol{c}) = \sum_{i=1}^{n} c_i^{1-1/\eta}/(1 - 1/\eta)$, with $\eta > 0$

Demand:

$$c_i^s(\boldsymbol{p}, w) = (p_i^s)^{-\eta} \frac{w}{\sum_j p_j \left(p_j^s\right)^{-\eta}}$$

More generally:

> Say that the consumer goes to the supermarket with a budget of $w = \$100$.
>
> Lack of full attention to prices => value of the basket in the cart is actually \$101.
>
> Demand linear in wealth: consumer buys 1% less of all the goods, to hit the budget constraint, and spends exactly \$100.
>
> (this is the adjustment factor $1/\boldsymbol{p} \cdot \boldsymbol{c}^r(\boldsymbol{p}^s, 1) = \frac{100}{101}$)

When demand is not necessarily linear in wealth…

> adjustment is (to the leading order) proportional to the income effect, $\frac{\partial c^r}{\partial w}$,
>
> not to the current basket, $\boldsymbol{c}^r$.

The behavioral agent cuts luxury goods, not necessities.

8.1.2. Nominal illusion, asymmetric Slutsky matrix, and inferring attention from choice data

Recall that the consumer "sees" only a part m_j of the price change.

One consequence: nominal illusion.

Proposition 6.2 (Nominal illusion)

Suppose that the agent pays more attention to some goods than others (i.e. the m_i are not all equal). Then, the agent exhibits nominal illusion, i.e. the Marshallian demand $\boldsymbol{c}(\boldsymbol{p}, w)$ is (generically) not homogeneous of degree 0 .

Suppose that the prices and the budget all increase by 10%.

For a rational consumer, nothing changes and he picks the same consumption.

Say a behavioral consumer who pays more attention to good $1 (m_1 > m_2)$.

He perceives that the price of good 1 has increased more than the price of good 2 has

(he perceives that they have respectively increased by $m_1 \cdot 10\%$ vs $m_2 \cdot 10\%$).

So, he perceives that the relative price of good 1 has increased

($\boldsymbol{p}^d$ is kept constant).

Hence, he consumes less of good 1, and more of good 2.

His demand has shifted

In abstract terms:

$\boldsymbol{c}^s(\chi \boldsymbol{p}, \chi w) \neq \boldsymbol{c}^s(\boldsymbol{p}, w)$ for $\chi = 1.1$

i.e. the Marshallian demand is not homogeneous of degree 0

The agent exhibits nominal illusion.

Slutsky matrix

Encodes both elasticities of substitution and welfare losses from distorted prices.

Its element S_{ij} is the (compensated) change in consumption of c_i as price p_j changes:

$$S_{ij}(\boldsymbol{p}, w) := \frac{\partial c_i(\boldsymbol{p}, w)}{\partial p_j} + \frac{\partial c_i(\boldsymbol{p}, w)}{\partial w} c_j(\boldsymbol{p}, w)$$

Rational agent: it is symmetric: $S_{ij}^r = S_{ji}^r$.

Kreps(2012, Chapter 11.6):

"The fact that the partial derivatives are identical and not just similarly signed is quite amazing. Why is it that whenever a $0.01 rise in the price of good i means a fall in (compensated) demand for j of, say, 4.3 units, then a $0.01 rise in the price of good j means a fall in (compensated) demand for i by [...] 4.3 units? [...] I am unable to give a good intuitive explanation."

Varian (1992, p.123):

"This is a rather nonintuitive result."

Mas-Colell, Whinston, and Green (1995, p.70):

"Symmetry is not easy to interpret in plain economic terms. As emphasized by Samuelson (1947), it is a property just beyond what one would derive without the help of mathematics."

If a prediction is non-intuitive to Mas-Colell, Whinston, and Green…

…it might require too much sophistication from the average consumer.

Less rational, psychologically more intuitive prediction:

Proposition 6.3 (Slutsky matrix)

Evaluated at the default price, the Slutsky matrix S^s is, compared to the traditional matrix S^r :

$$S_{ij}^s = S_{ij}^r m_j$$

i.e. the behavioral demand sensitivity to price j is the rational one, times m_j, the salience of price j.

Hence the behavioral Slutsky matrix is not symmetric in general.

Sensitivities corresponding to "non-salient" price changes (low m_j) are dampened.

Instead of looking at the full price change, the consumer just reacts to a fraction m_j of it.

He's typically less responsive than the rational agent.

> Say that $m_i > m_j$, so that the price of i is more salient than price of good j.
>
> Then $|S_{ij}^s|$ is lower than $|S_{ji}^s|$
>
> As good j 's price isn't very salient, quantities don't react much to it.
>
> If $m_j = 0$, the consumer does not react at all to price p_j: substitution effect is zero

Important: asymmetry of the Slutsky matrix indicates that, in general, a behavioral consumer cannot be represented by a rational consumer who simply has different tastes or some adjustment costs!

> Such a consumer would have a symmetric Slutsky matrix.

We may also infer attention from choice data.

Proposition 6.4 (Estimation of limited attention)

Choice data allows one to recover the attention vector m, up to a multiplicative factor $\bar{m}$.

> Suppose that an empirical Slutsky matrix S_{ij}^s is available.

m can be recovered as:

$$m_j = \bar{m} \prod_{i=1}^{n} \left(\frac{S_{ij}^s}{S_{ji}^s} \right)^{\gamma_i}$$

for any $(\gamma_i)_{i=1\ldots n}$ such that $\sum_i \gamma_i = 1$.

Proof: We have $\frac{S_{ij}^s}{S_{ji}^s} = \frac{m_j}{m_i}$, so $\prod_{i=1}^{n} \left(\frac{S_{ij}^s}{S_{ji}^s} \right)^{\gamma_i} = \prod_{i=1}^{n} \left(\frac{m_j}{m_i} \right)^{\gamma_i} = \frac{m_j}{\bar{m}}$, for $\bar{m} := \prod_{i=1}^{n} m_i^{\gamma_i}$.

The underlying "rational" matrix can be recovered as $S_{ij}^r := S_{ij}^s / m_j$, and it should be symmetric, a testable implication. [55] There is a literature estimating Slutsky matrices, which does not yet seem to have explored the role of non-salient prices.

It would be interesting to test Proposition 6.3 directly. The extant evidence is qualitatively encouraging, via the literature on obfuscation and shrouded attributes (Gabaix and Laibson 2006, Ellison and Ellison 2009) and tax salience. [56] Those papers find field evidence that some prices are partially neglected by consumers.

Marginal demand

Proposition 6.5 The Marshallian demand $\boldsymbol{c}^s(\boldsymbol{p}, w)$ has the marginals (evaluated at $\boldsymbol{p} = \boldsymbol{p}^d$):

$$\frac{\partial \boldsymbol{c}^s}{\partial w} = \frac{\partial \boldsymbol{c}^r}{\partial w}$$

and

$$\frac{\partial c_i^s}{\partial p_j} = \frac{\partial c_i^r}{\partial p_j} \times m_j - \frac{\partial c_i^r}{\partial w} c_j^r \times (1 - m_j)$$

Though substitution effects are dampened, income effects $\left(\frac{\partial \boldsymbol{c}}{\partial w} \right)$ are preserved

because w needs to be spent in this one-shot model

9.1. Textbook competitive equilibrium theory: A behavioral update

Notation:

Agent $a \in \{1, \ldots, A\}$ has endowment $\boldsymbol{\omega}^a \in \mathbb{R}^n$ (i.e. he is endowed with ω_i^a units of good i), with $n > 1$.

If the price is $\boldsymbol{p}$, his wealth is $\boldsymbol{p} \cdot \boldsymbol{\omega}^a$, so his demand is $\boldsymbol{D}^a(\boldsymbol{p}) := \boldsymbol{c}^a(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\omega}^a)$.

Economy's excess demand function is $\boldsymbol{Z}(\boldsymbol{p}) := \sum_{a=1}^{A} \boldsymbol{D}^a(\boldsymbol{p}) - \boldsymbol{\omega}^a$.

Set of equilibrium prices is $\mathcal{P}^* := \{\boldsymbol{p} \in \mathbb{R}_{++}^n : \boldsymbol{Z}(\boldsymbol{p}) = 0\}$.

Set of equilibrium allocations for a consumer a is $\mathcal{C}^a := \{\boldsymbol{D}^a(\boldsymbol{p}) : \boldsymbol{p} \in \mathcal{P}^*\}$.

9.1.1. First and second welfare theorems: (In)efficiency of equilibrium

Assume that competitive equilibria are interior, and consumers are locally non-satiated.

Proposition 6.6 (First fundamental theorem of welfare economics revisited: (In)efficiency of competitive equilibrium)

An equilibrium is Pareto efficient if and only if the perception of relative prices is identical across agents.

In that sense, the first welfare theorem generally fails.

Typically the equilibrium is not Pareto efficient when we are not at the default price.

Intuition:

Given two goods i and j, each agent equalizes relative marginal utilities and relative perceived prices:

$$\frac{u^a_{c_i}}{u^a_{c_j}} = \left(\frac{p^s_i}{p^s_j}\right)^a, \quad \frac{u^b_{c_i}}{u^b_{c_j}} = \left(\frac{p^s_i}{p^s_j}\right)^b$$

where $\left(\frac{p^s_i}{p^s_j}\right)^a$ is the relative price perceived by consumer a.

But equilibrium is efficient if and only if the ratio of marginal utilities is equalized across agents, i.e. there are no extra gains from trade, i.e.

$$\frac{u^a_{c_i}}{u^a_{c_j}} = \frac{u^b_{c_i}}{u^b_{c_j}}$$

Hence, the equilibrium if efficient if and only if any consumers a and b have the same perceptions of relative prices $\left(\left(\frac{p^s_i}{p^s_j}\right)^a = \left(\frac{p^s_i}{p^s_j}\right)^b\right)$.

The second welfare theorem asserts that any desired Pareto efficient allocation $(\boldsymbol{c}^a)_{a=1\ldots A}$ can be reached, after appropriate budget transfers.

It generally fails in this behavioral economy.

Typically, if the first welfare theorem fails, then a fortiori the second welfare theorem fails, as an equilibrium is typically not efficient.

Proposition 6.7 (Second theorem of welfare economics revisited)

The second welfare theorem generically fails, when there are strictly more than two consumers or two goods.

9.1.2. Excess volatility of prices in a behavioral economy

Assume in this section that there is just one representative agent.

Then bounded rationality leads to excess volatility of equilibrium prices.

Suppose that there are two dates, and that there is a supply shock:

Endowment $\boldsymbol{\omega}(t)$ changes between $t = 0$ and $t = 1$.

Let $d\boldsymbol{p} = \boldsymbol{p}(1) - \boldsymbol{p}(0)$ be the price change caused by the supply shock

Consider the case of infinitesimally small changes

Assume that $p_1 = p_1^d$ at $t = 1$).

Assume $m_i > 0$

Proposition 6.8 (Bounded rationality leads to excess volatility of prices).

Let $d\boldsymbol{p}^{[r]}$ and $d\boldsymbol{p}^{[s]}$ be the change in equilibrium price in the rational and behavioral economies, respectively.

Then:

$$dp_i^{[s]} = \frac{dp_i^{[r]}}{m_i}$$

i.e., after a supply shock, the movements of price i in the behavioral economy are like the movements in the rational economy, but amplified by a factor $\frac{1}{m_i} \geq 1$.

Hence, ceteris paribus, the prices of non-salient goods are more volatile.

Denoting by σ_i^k the price volatility in the rational $(k = r)$ or behavioral $(k = s)$ economy, we have:

$$\sigma_i^s = \frac{\sigma_i^r}{m_i}$$

Non-salient prices need to be more volatile to clear the market.

This might explain the high price volatility of many goods, such as commodities.

Consumers are quite price inelastic, because they are inattentive.

In a behavioral world, demand underreacts to shocks…

…but market needs to clear, so prices have to overreact to supply shocks.

9.1.3. Behavioral Edgeworth box: Extra-dimensional offer curve

Take a consumer with endowment $\boldsymbol{\omega} \in \mathbb{R}^n$.

Given a price vector $\boldsymbol{p}$, his wealth is $\boldsymbol{p} \cdot \boldsymbol{\omega}$

His demand is $\boldsymbol{D}(\boldsymbol{p}) := \boldsymbol{c}(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\omega}) \in \mathbb{R}^n$.

The offer curve OC is defined as the set of demands, as prices vary: $OC := \{\boldsymbol{D}(\boldsymbol{p}) : \boldsymbol{p} \in \mathbb{R}^n_{++}\}$.

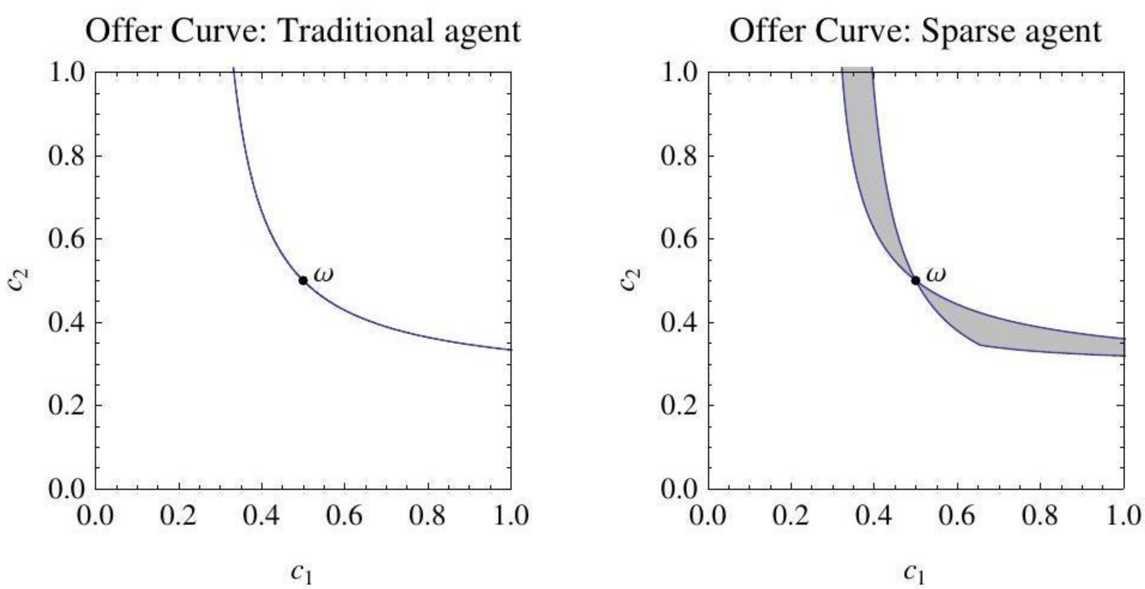

Figure 3: This Figure shows the agent's offer curve: the set of demanded consumptions $\boldsymbol{c}(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\omega})$, as the price vector $\boldsymbol{p}$ varies. The left panel is the traditional (rational) agent's offer curve. The right panel is the behavioral agent's offer curve (in gray): it is a 2-dimensional surface.

Start with two goods ($n = 2$).

Left panel of Figure 3 is the offer curve of the rational consumer:

it has the traditional shape.

The right panel plots the offer curve of a behavioral consumer with the same basic preferences:

the offer curve is the gray area.

Offer curve has acquired an extra dimension, compared to the one-dimensional curve of the rational consumer.

The OC is a now two-dimensional "ribbon", with a pinch at the endowment; if mistakes are unbounded, the OC is the union of quadrants north-west or south-east of $\boldsymbol{\omega}$

In the traditional model, the offer curve is one-dimensional:

demand $\boldsymbol{D}(\boldsymbol{p}) = \boldsymbol{c}(\boldsymbol{p}, \boldsymbol{p} \cdot \boldsymbol{\omega})$ is homogeneous of degree 0 in $\boldsymbol{p} = (p_1, p_2)$...

...then only relative price p_1/p_2 matters.

Behavioral model:

demand $\boldsymbol{D}(\boldsymbol{p})$ is not homogeneous of degree 0 in $\boldsymbol{p}$ any more: nominal illusion

Hence, offer curve is effectively described by two parameters (p_1, p_2)

(rather than just their ratio),

so it is 2-dimensional

This holds even though the Marshallian demand is a single-valued function.

In the traditional model, equilibria are the intersection of offer curves.

However, this is typically not the case here, as we shall now see.

9.1.4. A Phillips curve in the Edgeworth box

Traditional model with one equilibrium allocation:

set of equilibrium prices $\mathcal{P}^*$ is one-dimensional ($\mathcal{P}^* = \{\chi\overline{\boldsymbol{p}}: \chi \in \mathbb{R}_{++}\}$)

and $\mathcal{C}^a$ is just a point, $\boldsymbol{D}^a(\overline{\boldsymbol{p}})$

Behavioral setup, $\mathcal{P}^*$ is still one-dimensional.

However, to each equilibrium price level corresponds a different real equilibrium.

This is analogous to a "Phillips curve": $\mathcal{C}^a$ has dimension 1.

Consider the case of one rational consumer and one behavioral consumer:

Proposition 6.9

Suppose agent a is rational, and the other agent is behavioral with $m_1 = 1, m_2 = 0$

There are two goods.

The set $\mathcal{C}^a$ of a's equilibrium allocations is one-dimensional: it is equal to a's offer curve.

Suppose we start at a middle point of the curve in Figure 4, right panel.

Suppose that consumer b is a worker, good 2 is food, and good 1 is "leisure,"

When he consumes less of good 1 , he works more.

Let us say that $m_1 > m_2$:

he pays keen attention to his nominal wage, p_1, and less to the price of food, p_2.

Suppose now that the central bank raises the price level.

Then, consumer b sees that his nominal wage has increased, and sees less clearly the increase in the price of good 2.

So he perceives that his real wage $\left(\frac{p_1}{p_2}\right)$ has increased.

Hence (under weak assumptions) he supplies more labor:

he consumes less of good 1 (leisure) and more of good 2.

Hence, the central bank, by raising the price level, has shifted the equilibrium to a different point.

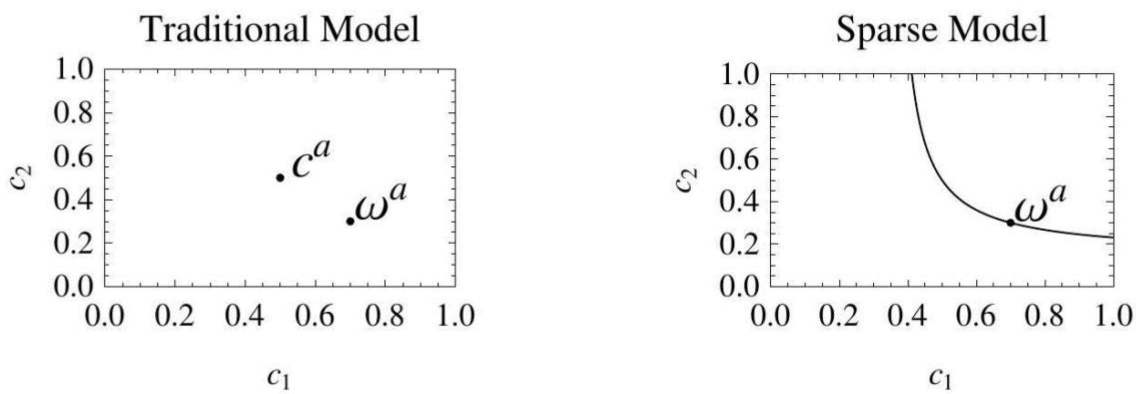

Figure 4: These Edgeworth boxes show competitive equilibria when both agents have CobbDouglas preferences. The left panel illustrates the traditional model with rational agents: there is just one equilibrium, $\boldsymbol{c}^a$. The right panel illustrates the situation when type a is rational, and type b is boundedly rational: there is a one-dimensional continuum of competitive equilibria (one for each price level) - a "Phillips curve." Agent a 's share of the total endowment (ω^a) is the same in both cases.

Is this Phillips curve something real and important?

This question is debated in macroeconomics, with an affirmative answer from New Keynesian analyses (Gali 2011).

Standard macro deals with one equilibrium, conditioning on the price level (and its expectations).

To some extent, this is what we have here.

Given a price level, there is (locally) only one equilibrium (as in Debreu 1970), but changes in the price level change the equilibrium (when there are some frictions in the perception or posting of prices).

This is akin to a (temporary) Phillips curve:

when the price level goes up, the perceived wage goes up, and people supply more labor.

Hence, we observe here the price-level dependent equilibria long theorized in macro, but in the pristine and general universe of basic microeconomics.

One criticism of the influential Lucas (1972) view is that inflation numbers are in practice very easy to obtain, contrary to Lucas' postulate.

This criticism does not apply here: behavioral agents actively neglect inflation numbers, which means the Phillips curve effect is valid even when information is readily obtainable.

9.2. What is robust in basic microeconomics?

Use the sparsity benchmark not as "the truth," of course, but as a plausible extension of the traditional model, when agents are less than fully rational.

Propositions that are not robust

Tradition: There is no money illusion.

Behavioral model: There is money illusion.

When the budget and prices are increased by 5%, the agent consumes less of goods with a salient price (which he perceives to be relatively more expensive);

Marshallian demand $\boldsymbol{c}(\boldsymbol{p}, w)$ is not homogeneous of degree 0.

Tradition: The Slutsky matrix is symmetric.

Behavioral model: It is asymmetric.

Elasticities to non-salient prices are attenuated by inattention.

Tradition: The offer curve is one-dimensional in the Edgeworth box.

Behavioral model: It is typically a two-dimensional pinched ribbon.

Tradition: The competitive equilibrium allocation is independent of the price level.

Behavioral model: Different aggregate price levels lead to materially different equilibrium allocations.

It's like in a Phillips curve.

Tradition: The Slutsky matrix is the second derivative of the expenditure function.

Behavioral model: They are linked in a different way.

Tradition: The Slutsky matrix is negative semi-definite. The weak axiom of revealed preference holds.

Behavioral model: These properties generally fail in a psychologically interpretable way.

Small robustness: Propositions that hold at the default price, but not away from it, to the first order

Marshallian and Hicksian demands, Shephard's lemma and Roy's identity:

values of the underlying objects are the same in the traditional and behavioral model at the default price

but differ (to the first order in $\boldsymbol{p} - \boldsymbol{p}^d$) away from the default price.

> This leads to a U-shape of errors in welfare assessment (in an analysis that does not take into account bounded rationality) as a function of consumer sophistication…
>
> the econometrician would mistake a low elasticity due to inattention for a fundamentally low elasticity.

Greater robustness: Objects are very close around the default price, up to second order terms

Tradition: People maximize their "objective" welfare.

Behavioral model: people maximize in default situations, but there are losses away from it.

Tradition: Competitive equilibrium is efficient, and the two Arrow-Debreu welfare theorems hold.

Behavioral model: Competitive equilibrium is efficient if it happens at the default price.

> Away from the default price, competitive equilibrium has inefficiencies, unless all agents have the same misperceptions.
>
> Welfare theorems do not hold in general.
>
> Values of the expenditure function $e(\boldsymbol{p}, u)$ and indirect utility function $v(\boldsymbol{p}, w)$ are the same, under the traditional and behavioral models, up to second order terms in the price deviation from the default $(\boldsymbol{p} - \boldsymbol{p}^d)$. [66]

Traditional economics gets the signs right

> More prudently put, the signs predicted by the rational model (e.g. Becker-style price theory) are robust under a sparsity variant.

Those predictions are of the type "if the price of good 1 does down, demand for it goes up", or more generally "if there's a good incentive to do X, people will indeed tend to do X," Those sign predictions make intuitive sense, and, not coincidentally, they hold in the behavioral model:

> those sign predictions (unlike quantitative predictions) remained unchanged even when the agent has a limited, qualitative understanding of his situation.

Indeed, when economists think about the world, or in much applied microeconomic work, it is often the sign predictions that are used and trusted, rather than the detailed quantitative predictions.

Open Questions and Conclusion

We need more measures of inattention

> This survey showed a number of measure of attention
>
> Currently, to produce one good measure of attention m, we need a full paper.
>
> It would be nice to scale up production - in particular, to always attempt to provide a quantitative measure of attention, rather than a demonstration that it is not full.
>
> In particular, can we relate the "physical measures of inputs to attention" (e.g. eye-tracking) to attention itself?

Investigating Varian in the lab

In physics textbooks, assertions and results (e.g. force = mass times acceleration) have been verified exquisitely in the lab.

Not so in economics.

You open, say, Varian (1992) or Mas-Colell, Whinston, and Green (1995), and see many assertions and predictions, with very few experimental counterparts

And indeed, one suspects that the assertions will actually be wrong if they are to be tested.

It would be great to make economics more like physics.

To do so, it seems important to experimentally investigate basic and behavioral microeconomics in the lab.

Challenges:

(i) to implement a notion of "clearly perceived" and "more opaque" prices
(ii) measure attention m
(iii) implement in a roughly naturalistic way the basic problem.

When you look at Mas-Colell, Whinston, and Green (1995)…

a few chapters have been extensively investigated: expected utility, with prospect theory as a benchmark, or basic game theory.

Other chapters, such as basic microeconomics of the consumer-theory / ArrowDebreu style, have been investigated very little

Such a study would be drier than, say, work on discrimination or fairness, but useful for economics.

Hopefully that imbalance will be corrected.

We need more experimental evidence on the determinants of attention

There are now several theories of attention, but measurement is somewhat lagging in refinement.

What's the cost of inattention?

Could we get some sense of the shape of the cost, and of the attention function?

More structural estimation

The early papers found evidence for imperfect attention, with large economic effects.

A more recent wave of papers has estimated inattention - its mean, variance, and how it varies with income, education and the like.

A third generation of papers might estimate more structurally models of inattention, to see if the predictions do fit, and perhaps suggest newer models.

Using this to do better policy: enerating attention

All this work may lead to progress in how to generate attention, e.g. for policy.

Making consumers more rational is difficult even when the right incentives are in place - for example, consumers overwhelmingly fail to minimize fees in allocating their portfolios (Choi, Laibson, and Madrian 2009).

The work on nudges (Thaler and Sunstein 2008) is based on psychological intuition rather than quantified principles.

Also, knowing better "best practices" for disclosure would be helpful.

Firms are good at screening for consumer biases (Ru and Schoar 2016), but public institutions less so, and debiasing is quite hard.