

Beyond "Ellison's Matrix": New Directions in Behavioral Industrial Organization*

Kfir Eliaz[†] and Ran Spiegler[‡]

March 19, 2015

Abstract

We discuss new research directions in the field of behavioral industrial organization that we find promising: identification of new unobservables in equilibrium models, strengthening links with abstract choice theory, and integrating behavioral IO models in larger models of the economy.

I. Introduction

When the modern literature on Behavioral Industrial Organization (BIO henceforth) was in its relative infancy, Glenn Ellison wrote a review for the Econometric Society World Congress (Ellison (2006)), in which he said, among other things:

The initial papers in this branch of the behavioral IO literature have tended to focus on how firms will choose prices and product characteristics to exploit behavioral biases and whether competition will eliminate the exploitation. Combining the IO and psychology and economics literatures, however, naturally gives many more than just one paper topic per bias – we can get a whole matrix of paper topics. Think of the set of behavioral biases as the column headings, and put all of the standard models in IO as the row headings: how will a monopolist price, how will a monopolist selling durable goods price; how will a monopolist price discriminate; how will oligopolists selling differentiated goods set prices; how will some action be distorted to deter or accommodate entry, etc. It takes little knowledge or imagination to come up with literally thousands of paper topics...Will this lead to thousands of behavioral IO papers in the next few years? I am pretty sure

*Some of the research this piece is based on was supported by ERC grant no. 230251. We thank two reviewers for useful comments.

[†]Tel Aviv University and University of Michigan. URL: www.kfireliaz.com. E-mail: kfire@post.tau.ac.il.

[‡]Tel Aviv University and University College London. URL: <http://www.tau.ac.il/~rani>. E-mail: rani@post.tau.ac.il.

(and hope) the answer is no. The problem is that most combinations will not produce observations that are sufficiently deep or interesting to warrant their being published.

Indeed, a significant portion of the existing BIO literature can be described as "ticking individual boxes" in "Ellison's matrix". As Ellison himself went on to acknowledge, this practice has generated many valuable insights. Moreover, incorporating decision biases into conventional models of market structure is not a mechanical task - it often generates new modeling challenges and theoretical insights. Nevertheless, like Ellison, we believe that it is important for the BIO literature to evolve beyond the "ticking boxes" mould. Indeed, the recent literature contains works that simply *ask questions of a different kind*: What are the theoretical implications of BIO models for consumer protection (Spiegler (2014a), Grubb (2015))? What determines the incentive to invent new ways to exploit boundedly rational consumers (Heidhues et al. (2012))?

In this modest contribution to this special issue of RIO, we wish to highlight three new theoretical research directions that likewise attempt to expand the methodological scope of BIO: (1) using equilibrium datasets to identify unobservables that characterize non-rational procedures of consumer choice; (2) strengthening the connection between BIO and abstract choice theory; (3) integrating BIO models in larger models of the economy. We wish to draw the reader's attention to these paths because we think they are interesting, and yet few papers (at best) have trodden them. This is emphatically a subjective perspective, and we certainly do not pretend to provide a comprehensive description of current trends in BIO.

II. Eliciting New Unobservables from Equilibrium Datasets

One consequence of recent advances in BIO is the confluence of new types of consumer-choice models and new sources of data on consumer behavior. New models incorporate elements of bounded rationality that require entirely new unobservable primitives that have no counterpart in the standard rational-choice model. At the same time, new datasets provide richer information about the consumer's choice process, and thus potentially enable us to identify the new unobservables. This raises an interesting research direction that lies somewhere between choice theory and "structural" empirical IO: If we assume that a certain (rich) dataset was generated by an *equilibrium market model* with "behavioral" consumers, can we identify the parameters of the unconventional model of consumer choice?

An interesting example of an unconventional unobservable is the subset of feasible alternatives that the consumer *actually considers*, and particularly how this subset responds to firms' marketing strategies. The modern marketplace presents consumers with an overwhelmingly large variety of products. Therefore, consumers often use screening criteria to reduce the number of "relevant" alternatives - thus applying their preferences not to the set of objectively feasible alternatives, but to a potentially smaller set which they construct at

an earlier stage of the decision process. The marketing literature refers to this set as the consumer's "*consideration set*". The basic idea underlying this term is that consumers may be unaware of some of the feasible products, and even when they become aware of a new product, they still need to be *persuaded to consider* it as a potential substitute to their currently consumed product. The "new unobservable" we introduce represents the consumer's susceptibility to attempts to persuade him to consider a new product.

We assume that the consumer follows a two-step procedure: given his objectively feasible set, he first constructs a consideration set and then chooses from this set. While a researcher may be able to observe the consumer's objectively feasible set and final choice, he often cannot directly observe the consumer's consideration set and therefore has to infer it from other data. In addition, the procedure's first step makes it harder to elicit the consumer's preferences. The reason is that the consideration set forms a wedge between the consumer's choices and his preferences, in the sense that choosing x over y does not necessarily mean that the consumer prefers x to y ; it may also mean that y was simply not considered.

Understanding *why* a consumer chose x over y has important implications. Imagine that we observe a consumer switching from y to x following a drop in the price of x . If both options were in his consideration set prior to the price change, then we can draw conclusions about the consumer's price elasticity. If, however, only y was originally in the consumer's consideration set, then the move to x may have been the result of a marketing campaign that advertised the price change and hence made x more salient. Thus, identifying the consumer's consideration set from choice data would allow us to draw conclusions regarding his price elasticity as well as the effectiveness of firms' marketing campaigns.

We argue that latent consideration sets may be identified by rich data that includes information beyond the consumer's final choice. For example, recent advances in e-commerce have given rise to datasets that include information about the consumer's search history and his exposure to advertising. Furthermore, if we assume that the data were generated by an *equilibrium* in a suitable market model, then the equilibrium condition imposes restrictions on the data that can be used to make inferences about the formation of the consumer's consideration set.

To illustrate this idea, we present a simple example based on Eliaz and Spiegler (2011). Two firms play a simultaneous-move game with complete information. Each firm chooses an element in $X \times M$, where X is a finite set of products and M is a finite set of marketing messages. An element $m \in M$ may be thought of as advertising image, a slogan, a jingle or a packaging. For simplicity, we allow a marketing message m to be paired with any product x (in reality, a slogan that is appropriate for a product x may be inappropriate for a product y). A strategy (x, m) carries a fixed cost of $c_x + c_m < \frac{1}{2}$, where the costs are measured in terms of market shares.

Firms face a homogenous population of consumers, who are characterized by two primitives. The first primitive is a conventional one, namely a strict preference relation \succ over

X . The second primitive is our "new unobservable": a *consideration function* $f : M \rightarrow 2^X$. Given a strategy profile $((x_i, m_i))_{i=1,2}$, consumers display a clear choice of firm i over firm j if and only if $x_j \in f(m_i)$ and $x_i \succ x_j$. In this case, firm i gets 100% market share and firm j gets zero market share. If consumers do not display a clear choice of one firm over another, each firm i gets 50% market share. Firms' objective is to maximize market share minus fixed costs.

The interpretation of the model is as follows. Each consumer is initially assigned to a random firm and considers buying its product. The set $f(m)$ consists of the products on which marketing message m acts - i.e., the products for which it is effective. When firm i offers the product x_i , and firm j employs the marketing message m_j , then a consumer who is initially assigned to firm i will include firm j 's product in his consideration set if and only if $x_i \in f(m_j)$. The consumer's consideration set is $\{x_i\}$ if $x_i \notin f(m_j)$ and $\{x_i, x_j\}$ if $x_i \in f(m_j)$. In the latter case, the consumer switches away from firm i in favor of firm j if and only if $x_j \succ x_i$. Thus, f captures how marketing messages persuade the consumer to consider a new product. In principle, it enables us to examine how *advertising content* influences consideration sets.

Our model allows marketing to be "targeted", in the sense that the effectiveness of a marketing message is a function of the consumer's initial product. For example, the marketing message may be an advertising banner that highlights a potential need that the consumer's current alternative fails to satisfy. However, if the consumer's current alternative does in fact satisfy this need, the marketing message will fall on deaf ears. Thus, our model incorporates targeted advertising according to the consumer's initial product.

We present the case of a homogenous consumer population because this is a simple case which we have solved. However, it should be clear that for this research agenda to take off, it must allow for consumer heterogeneity in both primitives. In particular, some consumers may be more susceptible to marketing than others, which in our model means that it is easier to attract their attention and trigger their curiosity about new products. Formally, suppose that f and g are two consideration functions, where $f(m) \subset g(m)$ for every m . Then, g captures greater susceptibility to marketing.

We impose the following assumptions on (\succ, f, c) :

No marketing, no consideration: There exists a unique "null" marketing message, $m^0 \in M$, that satisfies $f(m^0) = \emptyset$ and $c_0 = 0$.

Universal attention grabber: There exists $m^* \in M$ such that $f(m^*) = X$.

Best product is costly: The \succ -maximal product, denoted x^* , is also strictly more costly to produce than any other product, i.e. $c_{x^*} > c_x$ for all $x \prec x^*$. This assumption implies that product differentiation in our model is vertical rather than horizontal. We will use this assumption in Remark 4 below.

Suppose that the economist can observe three things for each data point: the consumer's

initial product, the advertising message to which he was exposed, and his final choice. This is more data than traditionally assumed to be available to the econometrician. However, this is where contemporary e-commerce and (possibly) brick-and-mortar supermarkets offer new opportunities. For instance, an online supermarket records a food brand the shopper regularly buys, a banner ad that appeared on his webpage during a shopping visit, and whether he consequently switched to the advertised new product.

The identification problem arises from the two-stage procedure we assume. Suppose that we observe a consumer who starts with a product x_i , becomes exposed to a marketing action m_j that accompanies a new product x_j , and yet does *not* switch to the new product. This could be either because he simply prefers x_i to x_j , or because the marketing message m_j does not act on x_i - that is, $x_i \notin f(m_j)$. Thus, it may well be that the consumer would have switched from x_i to x_j had he paid serious attention, but he did not because the marketing failed to arouse his interest. Even if we had all three-component data points, we would not be able to identify f and \succ . Indeed, it is easy to show that the most we can identify is the transitive closure of the "beating relation", defined as follows: (x, m) beats (x', m') if $x \succ x'$ and $x' \in f(m)$.

Suppose, however, that we did not observe individual data but *equilibrium* data - namely, all the data points that are consistent with symmetric mixed-strategy Nash equilibrium in the game between the firms. On one hand, since equilibria typically will not have full support, this means that we will have less data than in the ideal case described above. However, assuming Nash equilibrium as an identifying assumption may strengthen our ability to elicit unobservables, as in many examples in empirical IO

If consumers had perfect attention - i.e., if we assumed $f(m) = X$ for all m - then firms would play (x^*, m^0) in Nash equilibrium and earn a payoff of $\frac{1}{2} - c_{x^*}$. The next result describes an implication of symmetric Nash equilibrium under *imperfect* attention (under the three assumptions imposed above).

Proposition 1 *There exists $\bar{c} > 0$ such that if $c_x + c_m < \bar{c}$ for all (x, m) , firms earn $\frac{1}{2} - c_{x^*}$ in any symmetric Nash equilibrium. In particular, (x^*, m^0) is in the support of the mixed equilibrium strategy.*

This result is a consequence of Proposition 2 in Eliaz et al. (2013).¹ If \bar{c} is close to zero, the result is of little interest because the probability with which the rational-benchmark strategy (x^*, m^0) is played in equilibrium is close to one. However, Eliaz and Spiegel (2011) examine consideration functions with an additional "partitional" structure, for which $\bar{c} = \frac{1}{2}$. Thus, a little bit of extra structure on f can deliver a bound \bar{c} which is safely far from zero.

Proposition 1 has an important implication for our purposes, given by a result that Eliaz and Spiegel (2011) dubbed the "**effective marketing property**".

¹This result corrected Proposition 6 in Eliaz and Spiegel (2011).

Corollary 1 (Effective Marketing Property) *Suppose that firms play a Nash equilibrium in which they earn $\frac{1}{2} - c_{x^*}$. Then, for every two realizations $(x, m), (x', m')$ of the equilibrium mixed strategy, the following is true: if $x \neq x^*$ and $x \in f(m')$, then $x' \succ x$.*

Thus, in equilibrium, if a consumer starts with a product which is not the best one, and he is persuaded by marketing to consider the other available product, he will necessarily switch. A priori, this is surprising: off equilibrium, there is no reason to assume that if a consumer is persuaded to consider a product he will end up buying it. In other words, off-equilibrium conversion rates do not have to be 100%. However, in our example, as long as the consumer does not begin with his ideal choice, the conversion rate is 100%! The intuition is that the competitive pressures that bring industry profits to their rational-consumer benchmark also imply that marketing must be highly cost-effective: firms will not waste resources on marketing if it does not fully convert the consumer. Note that in equilibrium, since each firm faces a mixed strategy, it does not know which product the consumer will have as a default, and therefore this cost-effectiveness of marketing is a non-trivial result.

Of course, we do not expect such extreme conversion rates in reality because consumer tastes are heterogeneous. The point of this exercise is to give a simple example in which equilibrium analysis generates a prediction about conversion rates. We will now show how to use this result to elicit unobservables from observed conversion rates, under the equilibrium assumption. Two trivial observations are the following.

Remark 1 *The best product x^* can be uniquely identified as the one initial product consumers never switch away from.*

Remark 2 *The null marketing message m^0 can be uniquely identified as the one marketing message that accompanies x^* without attracting new clientele.*

We can now apply the effective marketing property to get the following, more interesting identification:

Remark 3 *For every data point in which the initial product is $x_i \neq x^*$, the marketing message to which the consumer was exposed is m_j and the other available product is x_j , if the consumer does not switch from x_i to x_j , then $x_i \notin f(m_j)$.*

Whenever the consumer does not switch away from a non-ideal product x_i , this reveals unambiguously that the marketing message m_j to which the consumer was exposed did not act on x_i . The proof of this result is as follows. By the effective marketing property, for

every data point given by the triple (x_i, x_j, m_j) and drawn from the equilibrium distribution, if $x_i \neq x^*$ and $x_i \in f(m_j)$, then $x_j \succ x_i$. Had m_j acted on x_i , the consumer would have switched to x_j . Therefore, the consumer's failure to switch from x_i to x_j is necessarily due to ineffectiveness of the marketing campaign.

Thus, failure to switch from $x_i \neq x^*$ to x_j when the latter is accompanied by the marketing message m_j reveals that $x_i \notin f(m_j)$. On the other hand, by the very definition of the consumer's choice procedure, if he does switch from x_i to x_j this reveals that $x_i \in f(m_j)$. Therefore, the equilibrium dataset enables a *complete identification of the consideration function, restricted to the set $X^* \times M^*$* , where X^* is the set of products that are offered with positive probability in equilibrium, and M^* is the set of marketing messages that are employed with positive probability in equilibrium.

While the equilibrium assumption enables a relatively powerful identification of the consideration function, it gives us no extra mileage when it comes to identifying preferences. The most we can identify in this regard is the transitive closure of the beating relation. The reason is that we cannot draw any inference about \succ from failure to switch (because as we have seen, it implies that the marketing is ineffective), whereas the inference about \succ from a switch is built into the definition of the consumer's choice procedure, and does not rely on any equilibrium assumption. Thus, we can see that using the equilibrium hypothesis as an identification strategy can be quite useful for one unobservable, and redundant for another.

This simple example suggests a number of interesting extensions. One direction is to allow for consumer heterogeneity, and possibly other types of models of consideration-set formation. Another direction is to introduce prices into the model, such that both fixed costs and mark-ups will be endogenous. Pricing is important not only for the obvious reason that it is a crucial strategic decision in most markets of interests (though not all, e.g. broadcast TV), but also because pricing and marketing decision often interact. For instance, temporary price reductions tend to be advertised, such that consumers' switching behavior indicates not only their price elasticity but also their heightened attention. When we observe switching data, we observe the interplay of two forces: preferences and limited attention, and it is interesting to see if we can use Nash equilibrium in a pricing-cum-marketing game as an identifying assumption for uncovering the role of these two forces, which would potentially enable us to identify the new unobservables.

It is important to point out that the empirical IO literature has recently recognized the importance of consideration sets in interpreting consumer demand, and has produced a few studies that share our interest in using equilibrium models to identify consumers' consideration sets. Sovinsky (2008) estimated a structural market model in which the (independent) probability with which the consumer considers a given product is a function of the intensity of its advertising in various media. A related paper is Clark et al. (2009). Honka (2014) estimates a standard model of rational sequential search, and it shares our idea of using richer data regarding the consumer's search process in order to identify consideration sets.

The research direction we proposed here should also be distinguished from empirical identification of agents' biased prior beliefs (regarding exogenous variables or their own future preferences - see Conlin et al. (2007), Wang (2015) and Grubb and Osborne (2015) for examples), or unconventional utility functions (e.g., (β, δ) intertemporal preferences - see DellaVigna (2009) and Hinnosaar (2014) for examples). In these cases, the unobservable component in question (prior belief, utility function) is entirely conventional, and the novelty of the new behavioral models is that they lift parametric restrictions imposed by standard models.

The novelty of our proposal in relation to all these works is that our model of consumer choice is based on a *non*-rational procedure that involves a genuinely *new unobservable* (the consideration function), and it suggests an *equilibrium* strategy for identifying the new unobservable using suitably *rich* data (conversion rates). Of course, the exercise in this section was a purely theoretical "proof of concept", and considerable effort and ingenuity will be required to make it viable for empirical research with concrete datasets.

III. Abstraction and Links to Choice Theory

One way for theoretical BIO to transcend "Ellison's matrix" is to formulate more abstract market models that incorporate specific models of consumer behavior as special cases. The model of Eliaz and Spiegler (2011), described in the previous section, is a case in point. The choice procedure does not commit to a particular behavioral phenomenon, and the consideration function is sufficiently general to accommodate a family of consideration-set formation processes. However, as we saw in the previous section, the greater level of abstraction enables results (like the "effective marketing property") that are intelligible and applicable for a large variety of specific models.

Of course, greater generality by itself is not a virtue, and usually leads to weak results. However, the move toward more abstractly formulated BIO models has a number of potential merits. First, it can provide a language for describing specific models in a simple, concise manner, thus facilitating the task of perceiving their differences and commonalities. In addition, sometimes the abstract formulation makes it easier to imagine new specific models that capture concrete behavioral phenomena that were simply not formalized before. Finally, raising the level of abstraction of BIO models facilitates links with choice theory. When a given market model describes consumer choice at the level of abstraction and generality that characterizes the typical choice-theoretic exercise, it becomes easier to examine the implications of general axioms of individual choice on equilibrium outcomes in the market model. In particular, it makes it easier to obey the principle of "revealed preferences" when conducting welfare analysis in the context of the market model.

Let us illustrate the last point using the model of the previous section. Recall the beating relation: (x, m) beats (x', m') if $x' \in f(m)$ and $x \succ x'$. The beating relation is a classical strict revealed preference relation (although it does not coincide with the "true preference relation" \succ that characterizes the consumer). When the consumer is initially assigned to a

firm that plays (x', m') , and the rival firm plays (x, m) , the consumer switches if and only if (x, m) beats (x', m') . This level of abstraction in describing consumer behavior enables us to state general conditions, such as transitivity of the beating relation. This property has a simple implication for the market model of the previous section.

Remark 4 *Suppose that (\succ, f, c) satisfies the properties assumed in Section 2. In addition, assume that the beating relation induced by (\succ, f) is transitive, and that $c_x + c_m < \frac{1}{2}$ for every (x, m) . Then, firms earn the payoff $\frac{1}{2} - c_{x^*}$ in any symmetric Nash equilibrium.*

Proof. Consider a symmetric Nash equilibrium. Let $\sigma \in \Delta(X \times M)$ denote the mixed equilibrium strategy. First, suppose that $\sum_m \sigma(x^*, m) = 0$. Then, a firm can deviate to the pure strategy (x^*, m^*) . By the definitions of x^* and m^* , the firm will get a market share of 1, whereas by the symmetry of the putative equilibrium, the firm's expected market share prior to the deviation is $\frac{1}{2}$. By assumption, $c_x + c_m < \frac{1}{2}$ for every (x, m) . Therefore, the deviation is profitable. Second, suppose that $\sum_m \sigma(x^*, m) = 1$. Then, each firm must play (x^*, m^0) with probability one, because any other strategy (x^*, m) is more costly and yet attains the same market share against σ . However, if both firms play (x^*, m^0) , then by assumption a firm can deviate to (x, m^0) , $x \neq x^*$, and maintain its market share of $\frac{1}{2}$ while lowering its cost. It follows that $\sum_m \sigma(x^*, m) \in (0, 1)$.

Assume that firms' equilibrium payoff is not $\frac{1}{2} - c_{x^*}$. Note that this is the max-min payoff for firms in the game. Therefore, it must be the case that firms earn strictly above $\frac{1}{2} - c_{x^*}$. Therefore (x^*, m^0) is not in the support of σ . Consider an arbitrary strategy (x^*, m) in the support of σ , $m \neq m^0$. Then, (x^*, m) must beat some strategy (x', m') in the support of σ . Note that the strategy (x', m^*) beats *any* strategy (x'', m'') with $x' \succ x''$, by the definition of m^* . By the definition of the beating relation, if (x', m^*) beats (x', m') , then it also beats (x', m^*) . Therefore, by transitivity of the beating relation, (x^*, m) beats (x'', m'') . (Note that for this argument, we do not have to assume that (x', m^*) is in the support of σ .) It follows that there exists a strategy (\tilde{x}, \tilde{m}) in the support of σ which is beaten by all strategies (x^*, m) in the support of σ , and does not beat any strategy in the support of σ . The strategy (\tilde{x}, \tilde{m}) must be a best-reply to σ . However, if a firm switches from this strategy to (x^*, m^*) , it gains $\frac{1}{2}$ in market share (because it avoids being beaten by the opponent whenever it offers x^* and it beats the opponent whenever it offers $x \neq x^*$). Since costs are always strictly below $\frac{1}{2}$, the deviation is profitable, a contradiction. ■

Thus, when the conventional strict revealed preference relation satisfies the basic choice-theoretic property of transitivity, this implies that firms earn $\frac{1}{2} - c_{x^*}$ in any symmetric Nash equilibrium. And as we saw in the previous section, a corollary of this result is the effective marketing property, which has implications for observed conversion rates (and therefore for identification of the consumer's consideration function). The lesson is that once the model is described at a level of abstraction that is suitable for choice-theoretic characterizations, we

can use general axioms of the choice model (stated in terms of a revealed-preference relation) to obtain relevant characterizations of equilibrium in the market model.

Piccione and Spiegler (2012) and Spiegler (2014b) are further steps along this direction. The latter paper brings the theoretical literature on market competition with imperfectly rational consumers into contact with the choice-theoretic literature on "choices with frames" (Salant and Rubinstein (2008), Bernheim and Rangel (2009)). By "frames" we mean different ways of describing or presenting an option, which is payoff irrelevant but nevertheless may affect choice. For example, a particular product may come in many different types of packages. Similarly, the interest on a loan may be expressed in nominal terms or in real terms.

The framework is simple. A market consists of n firms and one consumer. Each firm chooses a pair (x, m) as in the model of the previous section. The profile of marketing messages induces a probability distribution over "frames". The consumer chooses a firm according to a probabilistic choice function, which is a function of the profile of products and the frame. Spiegler (2014b) shows that the framework incorporates a number of prominent models from the BIO literature. In addition, the framework facilitates the construction of new market models that capture phenomena such as bracketing of financial risk, or manipulating of similarity judgments. Finally, the abstraction enables the formulation of a general property of the choice model - specifically, the mapping from profiles of marketing messages to frames - that is satisfied in many specific cases, and has useful implications for the analysis of Nash equilibrium in the game played among the firms. It remains to be seen whether similarly general and abstract properties of the consumer's frame-dependent choice function are both interesting at the level of individual choice and fruitful for the equilibrium analysis in the market model.²

IV. Integrating BIO Models in Larger Models of the Economy

IO models are typically "partial equilibrium" exercises, focusing on a single market. However, sometimes economists are interested in the connection between the market behavior in the single market and phenomena that are relevant for behavior in a collection of markets, or for the economy as a whole. In this section we describe two promising research directions that fit this description.

Macroeconomic applications

Price rigidity is an example of a friction that has long captured the interest of macroeconomists. In many macro models, price rigidity is modeled mechanically, without any theoretical account of its origin. At the same time, empirical studies have documented rich patterns of price rigidity in a variety of datasets (Klenow and Malin (2010)). BIO models can give rise to similarly rich patterns of price rigidity from very simple settings. For example, Koszegi and Heidhues (2014) study a model of monopoly pricing when consumer preferences display loss

²For another attempt to marry choice-theoretic analysis of consumer behavior with an IO application, see Papi (2014).

aversion with expectational reference points. They show that the optimal pricing strategy is random, with an atom on a high "regular price" and a continuous density on an interval of lower "sales prices" that are bounded away from the regular price. Spiegler (2012) obtains a very similar pattern from a multi-firm market model, in which firms compete in price distributions, and consumers (who also have a sufficiently attractive outside option) choose according to a naive sampling procedure. Given the utter simplicity of the settings that generate this rich pattern, and the relative simplicity of the equilibrium characterizations, we believe it is worthwhile to try embedding the mechanisms that generate it in larger macroeconomic models.

Eliasz and Spiegler (2013) is an example of such an attempt, in the different context of labor markets. The basic ingredient of the model is a repeated, finite-horizon principal-agent interaction between an employer and a worker. It is built around the idea that the labor contract is highly incomplete, and therefore worker effort relies on his "intrinsic motivation", rather than on high-powered material incentives. At every period, the employer offers a flat "spot" wage contract. Conditional on accepting the offer, the worker shirks whenever the wage falls below the worker's "reference wage", which is the lagged-expected wage for that period (calculated according to the information available to the parties at the end of the previous period). The worker's outside option fluctuates according to some stochastic process. This model generates patterns of wage rigidity (in the sense of being unresponsive to the current fluctuations) and endogenous job destruction. Eliasz and Spiegler (2013) incorporate this account of the labor relation into a search-matching model of the labor market with a business cycle, and show that the unique (subgame-perfect) equilibrium of our model exhibits the following properties: existing workers experience downward wage rigidity, as well as destruction of output following negative shocks due to layoffs or loss of morale; newly hired workers earn relatively flexible wages, but not as much as in the benchmark without reference dependence; market tightness is more volatile than under this benchmark. Thus, this paper sheds some light on a puzzle that has attracted much attention in the macroeconomic literature ("the Shimer puzzle"): how can we explain large fluctuations in unemployment alongside downward rigidity in wages? In a similar vein, we believe that BIO models of price rigidity may be fruitfully incorporated in larger macroeconomic models.

Multi-tasking

BIO models examine the limitations on the rationality of consumer behavior in the context of a single industry. However, many of these limitations arise because the consumer is simultaneously involved in many market and non-market interactions. As consumers we must juggle our consumption decisions together with work-, family- and health-related problems. Since there is only a limited amount of time and attention resources we can devote to all these problems, we are constantly faced with the daunting task of prioritizing and scheduling our decision problems. In other words, consumers are "reluctant multi-taskers".

This observation has often been used by critics of behavioral economics to claim that

consumers' biases and errors in a given choice situation can be "rationalized away" as part of an optimal second-best solution to the multi-tasking problem. There seems to be a schism between "behavioral" theorists, who take consumers' procedures in a single market as given, and "rational-choice" theorists who claim that these choice procedures are admissible only in-so-far as they can be justified as second-best optima in a larger, multi-market model. At the current stage of the BIO literature, we see no reason for this dichotomy. We believe it is important to examine consumers' "multi-tasking" problems, using either "behavioral" or "rational choice" modeling approaches; interest in the multi-tasking problem should not undermine the "partial equilibrium" modeling practice of taking the consumer's choice procedure as given in a single-industry model.

Nevertheless, the papers we describe below do adopt rational-choice modeling strategies to analyze the consumer's multi-tasking problem. de Clippel, Eliaz and Rozen (2014) consider a consumer who would like to buy one unit from each of K different markets (say, one gym membership, one internet provider, one cell phone provider, etc.). The consumer starts with a default provider in each of the K markets such that the consumer knows how much he currently pays in each market. In order to switch to a different provider in some market, the consumer needs to search for an alternative in that market and to learn its price. However, the consumer has only a limited amount of time that he can devote for searching an alternative. Specifically, de Clippel et al. assume that a consumer of type k can inspect up to k markets, where $k = 0, \dots, K$ follows some distribution. When $k = 0$ the consumer is fully inattentive and never switches. When $k = K$ the consumer is fully attentive and can inspect all markets.

de Clippel, Eliaz and Rozen (2014) consider a simple case in which every market has only two firms, the default firm and a challenger, that compete in prices by simultaneously choosing a distribution over prices. All consumers face the same default firm in each market and they only observe the price realization of the default firms. Given the realized prices of the default firms, each consumer decides how best to allocate his budget of k units of attention across the markets. Upon inspecting a market, the consumer optimally chooses which firm to buy from. All consumers have unit demands and are willing to pay at most 1 for each unit. de Clippel et al. (2014) establish the existence of a unique partially symmetric equilibrium (all default firms use the same strategy and all challenger firms use the same strategy). They show that this equilibrium has the property that an increase in partial attention (in the sense that the distribution of partially attentive consumer undergoes a first-order stochastic shift) decreases consumer welfare.

Coviello et al. (2014) introduce the dynamic aspect of the consumer's multi-tasking problem. Consumption problems arrive randomly over time. When the consumer encounters a new problem, he first needs to decide how to schedule it in the queue of problems. Therefore, the consumer's dynamic problem involves sequencing the consumption problems subject to his time and attention constraints, and solving each problem at a time. Coviello et al. (2014) consider a single decision maker (DM) who faces a stream of decision problems that arrive

over time at some constant rate α . Each decision problem requires the same amount of time T to complete. The DM cannot solve all problems simultaneously, and each problem requires some time to solve. The DM therefore creates a queue of problems and draws problems to solve from this queue at some constant rate β . The DM uniformly allocates his attention over all the problems drawn from the queue.

Coviello et al. (2014) characterize the DM's "production function", which links the parameters α and β to the rate with which problems are solved. Coviello et al. (2014) then embed this production function in a model where decision problems are generated by strategic agents who can, by incurring a cost, manipulate the arrival rate of their problems. It is shown that there exists a unique equilibrium in which the arrival rate of decision problems is constant. This model has potentially interesting implications for market behavior: firms bombard consumers with consumption problems, and the consumer's delay in reaching decisions can manifest itself in "default bias".

We should point out that the model of Coviello et al. (2014) does not have an explicit IO context. However, we chose to describe it here because it has obvious implications for IO settings. An important question it raises is how the biases and errors consumers commit in individual markets might be affected by aspects of the stochastic process that generates the multiple tasks, and how this process is affected by the firms' own activities.

References

- [1] Bernheim, D., & Rangel, D. (2009). Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. *Quarterly Journal of Economics*, 124, 51-104.
- [2] Conlin, M., T. O'Donoghue, T. & Vogelsang, T. (2007). Projection Bias in Catalog Orders. *American Economic Review*, 97, 1217-1249.
- [3] Coviello, D., A. Ichino, & Persico, N. (2014). Time Allocation and Task Juggling. *American Economic Review*, 104, 609-623.
- [4] Clark, R., U. Doraszelski, & Draganska, M. (2009). The Effect of Advertising on Brand Awareness and Perceived Quality: An Empirical Investigation using Panel Data. *Quantitative Marketing and Economics*, 7, 207-236.
- [5] de Clippel, G., K. Eliaz, & Rozen, K. (2014). Competing for Consumer Inattention. *Journal of Political Economy*, 122, 1203-1234.
- [6] DellaVigna, S. (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature*, 47, 315-372.

- [7] Eliaz, K., T. Gamp, & Spiegler, R. (2013). Consideration Sets and Competitive Marketing: Corrigendum. Working Paper. Tel-Aviv University.
- [8] Eliaz, K., & Spiegler, R. (2011). Consideration Sets and Competitive Marketing. *Review of Economic Studies*, 78, 235-262.
- [9] Eliaz, K., & Spiegler, R. (2013). Reference-Dependent Preferences and Labor Market Fluctuations. *NBER Macroeconomic Annual*, 28, 159-200.
- [10] Ellison, G. (2006). Bounded Rationality in Industrial Organization. In Blundell, R., Newey, W. & Persson, T. (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Ninth World Congress (pp. 142-174). Cambridge: Cambridge University Press.
- [11] Grubb, M. (2015). Consumer Inattention and Bill-Shock Regulation. *Review of Economic Studies*, 82, 219-257.
- [12] Grubb, M. & Osborne, M. (2015). Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock. *American Economic Review*, 105, 234-271.
- [13] Hinnoosaar, M. (2014). Time Inconsistency and Alcohol Sales Restrictions. Working Paper. Collegio Carlo Alberto.
- [14] Klenow, P., & Malin, B. (2010). Microeconomic Evidence on Price-Setting. In Friedman, B.M. & Woodford, M. (Eds.), *Handbook of Monetary Economics Vol. 3* (pp. 231-284). Oxford: Elsevier Science.
- [15] Köszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. *Quarterly Journal of Economics*, 121, 1133-1166.
- [16] Heidhues, P., & Köszegi, B. (2014). Regular Prices and Sales. *Theoretical Economics*, 9, 217-251.
- [17] Heidhues, P., B. Köszegi & Murooka, T. (2012). Exploitative Innovation. *American Economic Journal: Microeconomics*. (in press).
- [18] Papi, M. (2014). Noncompensatory Consideration and Compensatory Choice: an Application to Stackelberg Competition. *Economic Theory Bulletin*, 2, 53-63.
- [19] Piccione, M., & Spiegler, R. (2012). Price Competition Under Limited Comparability. *Quarterly Journal of Economics*, 127, 97-135.
- [20] Salant, Y., & Rubinstein, A. (2008). (A,f): Choice with Frames. *Review of Economic Studies*, 75, 1287-1296.
- [21] Sovinsky, M. (2008). Limited Information and Advertising in the US Personal Computer Industry. *Econometrica*, 76, 1017-1074.

- [22] Spiegel, R. (2012). Consumer Bounded Rationality and Rigidity/Flexibility Retail Price Patterns. *Economics Letters*, 116, 335-338.
- [23] Spiegel, R. (2014a). On the Equilibrium Effects of Nudging. *Journal of Legal Studies*. (in press).
- [24] Spiegel, R. (2014b). Competitive Framing. *American Economic Journal: Microeconomics*, 6, 35-58.
- [25] Wang, J. (2015). Why Hire Loan Officers? Examining Delegated Expertise. Working Paper. University of Michigan.