



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://economics.sas.upenn.edu/pier>

PIER Working Paper 15-013

“Consumer Choice as Constrained Imitation”

by

Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler

<http://ssrn.com/abstract=2575806>

Consumer Choice as Constrained Imitation^{*}

Itzhak Gilboa,[†] Andrew Postlewaite,[‡] and David Schmeidler[§]

February 2015

Abstract

A literal interpretation of neo-classical consumer theory suggests that the consumer solves a very complex problem. In the presence of indivisible goods, the consumer problem is NP-Hard, and it appears unlikely that it can be optimally solved by humans. An alternative approach is suggested, according to which the household chooses how to allocate its budget among product categories without necessarily being compatible with utility maximization. Rather, the household has a set of constraints, and among these it chooses an allocation in a case-based manner, influenced by choices of other, similar households, or of itself in the past. We offer an axiomatization of this model.

1 Introduction

Economists seem to be in agreement about two basic facts regarding neo-classical consumer theory. The first is that the depiction of the consumer as maximizing a utility function given a budget constraint is a very insightful tool. The second is that this model is probably a poor description of the

^{*}We wish to thank Bill Brainard, Hanming Fang, and Gabi Gayer for discussions and comments. We gratefully acknowledge ISF Grant 204/13 (Gilboa and Schmeidler), ERC Grant 269754 (Gilboa), and NSF Grant SES-0961540 (Postlewaite).

[†]HEC, Paris, and Tel-Aviv University. tzachigilboa@gmail.com

[‡]University of Pennsylvania. apostlew@econ.sas.upenn.edu

[§]Tel-Aviv University, the Ohio State University, and the InterDisciplinary Center at Herzliya. davidschmeidler@gmail.com

mental process that consumers go through while making their consumption decisions at the level of specific products.

The first point calls for little elaboration. The neoclassical model of consumer choice is extremely powerful and elegant. It lies at the heart, and is probably the origin of “rational choice theory”, which has been applied to a variety of fields within and beyond economics. Importantly, utility maximization, as a behavioral model, does not assume that a mental process of maximization actually takes place. Behaviorally, utility maximization was shown to be equivalent to highly cogent assumptions regarding consumer choices (see Debreu, 1959).

Yet, many writers have commented on the fact that a literal interpretation of the theory does not appear very plausible. Recent literature in psychology, decision theory, and economics is replete with behavioral counter-examples to the utility maximization paradigm. These include direct violations of explicit axioms such as transitivity, as well as examples that violate implicit assumptions, such as the independence of reference points (see Kahneman and Tversky, 1979, 1984). In this paper we focus on one specific reason for which the neoclassical model does not always appear to be cognitively plausible, namely, computational complexity. One aspect of the latter is emphasized by the example below. It illustrates the implicit and often dubious assumption that consumers are aware of all the bundles in their budget set.

Example 1 Every morning John starts his day in a local coffee place with a *caffè latte grande* and a newspaper. Together, he spends on coffee and newspaper slightly over \$3 a day. He then takes public transportation to get to work. One day Mary joins John for the morning coffee, and he tells her that he dislikes public transportation, but that he can’t afford to buy a car. Mary says that she has just bought a small car, financed at \$99 a month. John sighs and says that he knows that such financing is possible, but that he can’t even afford to spend an extra \$99 a month. Mary replies that if he were to give up on the *caffè latte* and newspaper each morning, he could buy

the car. John decides to buy the car and give up on the morning treat.

What did Mary do to change John’s consumption pattern? She did not provide him with new factual information. John had been aware of the existence of inexpensive financing for small cars before his conversation with Mary. She also did not provide him with new information about the benefits of a car; in fact, it was John who brought up the transportation issue. Rather than telling John of new facts that he had not known before, Mary was pointing out to him certain consumption bundles that were available to him, but that he had failed to consider beforehand. Indeed, the number of possible consumption bundles in John’s budget set is dauntingly large. He cannot possibly be expected to consider each and every one of them. In this case, he failed to ask himself whether he preferred the coffee or the car. Consequently, it would be misleading to depict John as a utility maximizing agent. Such an agent should not change his behavior simply because someone points out to him that a certain bundle is in his budget set.

This example is akin to framing effects (Tversky and Kahneman, 1981) in that it revolves around reorganization of existing knowledge. However, our example differs from common examples of framing effects in one dimension: the ability of the consumer to learn from her mistake and to avoid repeating it. Many framing effects will disappear as soon as the decision problem is stated in a formal model. By contrast, the richness of the budget set poses an inherent difficulty in solving the consumer problem. In our example, John didn’t fail to consider all alternatives due to a suggestive representation of the problem. We argue that he failed to do so due to the inherent complexity of the problem. Specifically, in section ?? we prove that, in the presence of indivisible goods, the consumer problem is NP-Complete. This means that deviations from neoclassical consumer theory cannot be dismissed as “mistakes” that can be avoided should one be careful enough. It is practically impossible to avoid these deviations even if one is equipped with the best software and the fastest computers that are available now or in the foreseeable

future.

1.1 The Affluent Society ¹

There are many problems for which utility maximization can be viewed as a reasonable, if admittedly idealized model of the consumer decisions. Consider, for example, a graduate student in economics, who has to survive on a stipend of \$25,000 a year. This is a rather tight budget constraint. Taking into account minimal expenditure on housing and on food, one finds that very little freedom is left to the student. Given the paucity of the set of feasible bundles, it seems reasonable to suggest that the student considers the possible bundles, compares, for instance, the benefit of another concert versus another pair of jeans, and makes a conscious choice among these bundles. When such a choice among relatively few bundles is consciously made, it stands to reason that it would satisfy axioms such as transitivity or the weak axiom of revealed preference. The mathematical model of utility maximization then appears as a reasonable description of the actual choice process of the student.

Next consider the same student after having obtained a job as an assistant professor. Her tastes have probably changed very little, but her budget is now an order of magnitude larger than it used to be. Housing and food are still important to her, but they are unlikely to constrain her choice in a way that would make her problem computationally easy. In fact, the number of possible bundles she can afford has increased to such an extent that she cannot possibly imagine all alternatives. Should she get box tickets for the opera? Save more money for a Christmas vacation? Buy diamonds? Save for college tuition of her yet-unborn children? For such an individual, it seems that the utility maximization model has lost much of the cognitive appeal it had with a tight budget constraint. Correspondingly, it is far from obvious

¹The title of the subsection is that of the well known book by John Kenneth Galbraith (1958).

that her choices satisfy the behavioral axioms of consumer theory.

Galbraith (1958) suggested that neoclassical consumer theory was developed with poverty in mind, and he pointed out the need to develop alternative theories for affluent societies. Without any pretense to have risen to Galbraith’s challenge, we wish to explore what such a theory might look like.

1.2 Budget Allocation

The computational difficulties with the neoclassical model demonstrate why this model does not accurately describe the way households make decisions, at least not at the level of specific products. The question then arises, how do they make their decisions?

One way to deal with a complicated problem is to decompose it. Looking for an optimal budget allocation among goods, one might use a top-down approach, first dividing the overall budget among a few major categories of goods, then subdividing these amounts among finer categories, and so on. This is a natural heuristic, which may not be guaranteed to produce an optimal solution, but which seems sufficiently reasonable and intuitive as a starting point. In this paper we focus only on the first step of this heuristic, namely, the choice of a budget allocation among a few categories at the top level.

Suppose that the household conceives of a few natural categories of goods such as “housing”, “education”, “transportation”, “food”, etc. (as is customary in empirical work – see, for instance, Deaton and Muellbauer, 1980, Blundell, 1988, and Sabelhaus, 1990.) How does it determine the budget allocation among these categories? Applying the neoclassical consumer optimization model would require that the household have a notion of “utility” derived from vectors of amounts of money allocated to each category. This would entail solving the lower-level optimization problem before one can tackle the top-level one, a process as complex as the initial problem the household started with.

In this paper we offer an alternative model, according to which the top-level budget allocation is not done by maximizing a utility function, whether consciously or not. Rather, it is done in two stages: first, the household has a set of self-imposed constraints, or “rules of thumb”, which simplify the problem by ruling out classes of budget allocations. For example, such rules can be “always save at least 20% of your income” or “do not spend more than 15% of the budget on entertainment”. After considering these rules, which presumably leave a non-empty feasible set of budget allocations, the household may well be left with a collection of feasible, or acceptable allocations. How should it choose among these?

Our model suggests that, at this second stage, the household looks around and observes what other households do. Then it chooses an allocation that is similar to those chosen by similar households. Thus, families would observe what other families, with similar income, size, religion, or background do, and decide to allocate their budgets among categories such as housing, transportation, education and vacations in a similar way to these families. This mode of behavior is compatible both with social learning (see, for instance, Goyal, 2005) and with conformism (as in Bernheim, 1994).

In our model this imitation is done in a “case-based” way, where the household can be thought of as though it chooses the similarity-weighted average of the budget allocations of similar households, provided that these allocations were within the acceptable set. In case an observed allocation is not acceptable to the household, it is replaced by the closest one that is. Thus the household chooses a similarity-weighted average of closest-acceptable allocations. This paper offers an axiomatization of this model.

Our formal model deals with only one step of the budget allocation problem. It is, however, both natural and straightforward to extend the model to multiple levels. In this case one can get a budget tree, where expenditures are allocated to sub-categories, and then to sub-sub-categories, and so forth.

However, when the number of levels grows, one may find that the graph

generated is not a tree: a particular product may be attributed to several possible categories. This may explain phenomena that are referred to as mental accounting (Thaler and Shefrin, 1981; Thaler, 1980, 1985, 2004). To see a simple example, suppose that a sub-category of expenses is split into “standard expenses” and “special events”. In this case, the consumer may decide to buy an item if it is considered a birthday gift, but refrain from buying it if it is not associated with any special event. In other words, the top-down approach implies that the same bundle will be viewed differently depending on the categorization used. Combining this with our basic motivation, we find that computational complexity of the consumer problem may result in mental accounting. Conversely, while mental accounting is certainly a deviation from classical consumer theory, it appears to involve only a very mild form of “bounded rationality”. Treating money as if it came from different accounts is not simply a mistake that can be easily corrected. Rather, it is a by-product of a reasonable heuristic adopted to deal with an otherwise intractable problem.

The rest of this paper is organized as follows. Section 2 states the complexity result. In Section 3 we offer a simple model that captures a household’s budget allocation decision along the lines suggested above. Section 4 provides an axiomatic derivation of the model. Section 5 concludes with a discussion.

2 A Complexity Result

Many writers have observed that the consumer problem is, intuitively speaking, a complex one. Some (see MacLeod, 1996, Arthur, Durlauf, and Lane, 1997) have also made explicit reference to the combinatorial aspects of this problem, and to the fact that, when decisions are discrete, the number of possible bundles grows as an exponential function of the parameters of the

problem.² However, the very fact that there exist exponentially many possible solutions does not mean that a problem is hard. It only means that a brute-force algorithm, enumerating all possible solutions, will be of (worst-case) exponential complexity. But for many combinatorial problems with an exponentially large set of possible solutions there exist efficient algorithms, whose worst-case time complexity is polynomial. Thus, in order to convince ourselves that a problem is inherently difficult, we need to prove more than that the number of possible solutions grows exponentially in the size of the problem.

In this section we show that, when some goods are indivisible, the consumer problem is “hard” in the sense of NP-Completeness. This term is borrowed from the computer science literature, and it refers to a class of combinatorial problems that are deemed to be “hard” in the following sense. For any NP-Complete problem the number of steps in any known algorithm solving the problem grows exponentially in the size of the problem. Consequently, for even moderate size problems, it might take the fastest computers that exist years to solve the problem. Further, if an algorithm were found for which the number of steps in the algorithm was a polynomial in the size of the problem for any NP-Complete problem, the algorithm could be used to construct polynomial algorithms for all NP-Complete problems. Since a variety of these problems have been exhaustively studied for years and no efficient (polynomial) algorithm is known for any of them, proving that a new problem is NP-Complete is taken to imply that it is a hard problem as well.

Thus, the vague intuition that it is hard to maximize a utility function over a large budget set is supported by our complexity result. As rational as consumers can possibly be, it is unlikely that they can solve in their minds

²For example, assume that there are m binary decisions, each regarding the purchase of a product at price p . With income I , the consumer can afford to purchase $\frac{I}{p}$ products. She therefore must consider $\binom{m}{\frac{I}{p}}$ different bundles. If m is relatively large, this expression is of the order of magnitude of $m^{\frac{I}{p}}$, namely, exponential in I .

problems that prove intractable for computer scientists equipped with the latest technology. Correspondingly, it is always possible that a consumer will fail to even consider a bundle that, if pointed out to her, she would consider desirable. It follows that one cannot simply teach consumers to maximize their utility functions. In a sense, this type of violation of utility maximization is more robust than some of the examples of framing effects and related biases. In the example given in the Introduction, John failed to consider a possible bundle that was available to him. After this bundle was pointed out to him by Mary, he could change his behavior and start consuming it. But he had no practical way of considering all consumption bundles, and he could not guarantee himself that in his future consumption decisions he would refrain from making similar omissions.

An NP-Complete problem has the additional feature that, once a solution to it is explicitly proposed, it is easy to verify whether it indeed solves the problem (this is the “NP” part of the definition). Thus, for an NP-Complete problem it is hard to find a solution, but it is easy to verify a solution as legitimate if one is proposed. In this sense, problems that are NP-Complete present examples of “fact-free learning”: asking an individual whether a certain potential solution is indeed a solution may make the individual aware of it, accept it, and change her behavior as a result. Aragones, Gilboa, Postlewaite, and Schmeidler (2005) show that finding a “best” regression model is an NP-Complete problem, and thus that finding regularities in a given database may result in fact-free learning. This section shows that fact-free learning can also occur in the standard consumer problem, arguably the cornerstone of economic theory.

We now turn to show that the neoclassical consumer problem, of maximizing a quasi-concave utility function with a budget constraint is NP-Complete. Consider a problem $P = \langle n, (p_i)_{i \leq n}, I, u \rangle$ whose input is:

- $n \geq 1$ – the number of *products*;
- $p_i \in \mathbb{Z}_+$ is the *price* of product $i \leq n$;

$I \in \mathbb{Z}_+$ is the consumer's *income*; and

$u : \mathbb{Z}_+^n \rightarrow \mathbb{R}$ is the consumer's *utility* function.

The function u is assumed to be given by a well-formed arithmetic formula involving the symbols “ x_1 ”, ..., “ x_n ”, “+”, “*”, “−”, “/”, “^”, “(”, “)”, “0”, ..., “9” with the obvious semantics (and where “^” stands for power). As is standard in consumer theory, we assume that this formula, when applied to all of \mathbb{R}_+^n , defines a continuous, non-decreasing, and quasi-concave function.

Let the *Consumer Problem* be: Given a consumer problem $P = \langle n, (p_i)_{i \leq n}, I, u \rangle$ and an integer \bar{u} , can the consumer obtain utility \bar{u} in P ?

(That is, is there a vector $(x_1, \dots, x_n) \in \mathbb{Z}_+^n$ such that $\sum_{i \leq n} p_i x_i \leq I$ and $u(x_1, \dots, x_n) \geq \bar{u}$?)

We can now state:

Proposition 1 *The Consumer Problem is NP-Complete.*

It will be clear from the proof that approximating the solution to the Consumer Problem is also a hard problem; it is not the case that settling for an “almost optimal” choice is less difficult. (The Appendix, where all proofs are gathered, comments on approximations following the proof of the Proposition.)

3 A Model of Case-Based Constrained Imitation

How do households make consumption decisions? We confine our attention to *planned* consumption.³ We imagine the household thinking about its budget and allocating it to product categories. This may be but the first step in an iterative process, where budget is allocated in a top-down way. We seek

³Actual consumption choices might differ from the planned ones as a result of randomness in the process in which consumption opportunities present themselves, as a result of problems of self-control, and so forth.

a model that could be imagined as a description of actual mental decision processes that households go through.

Formally, the household is faced with the problem of dividing income I among the expenditures E_1, \dots, E_n where $i \leq n$ is a category of goods and

$$E_1 + \dots + E_n = I.$$

It will be convenient to think of the budget shares of the categories,

$$z_i = \frac{E_i}{I}$$

so that the vector of budget shares (z_1, \dots, z_n) is a point in the $(n - 1)$ -dimensional simplex.

If $p_i > 0$ is a price index of category i (say, an index of food prices), the quantity of the aggregate good of category i consumed is readily computed from z_i as

$$\frac{z_i I}{p_i}.$$

The context of the problem is a *database of past choices*, made by the household itself and by others, described as

$$D = ((x_{1t}, \dots, x_{mt}), (z_{1t}, \dots, z_{nt}))_{t=1}^T$$

where (x_{1t}, \dots, x_{mt}) describes the characteristics of consumption problem t and (z_{1t}, \dots, z_{nt}) – its solution. We think of (x_{1t}, \dots, x_{mt}) as including (i) income; (ii) demographic variables: the relevant household size, its age distribution, and so forth; (iii) problem variables such as the time of consumption; (iv) the household's identity, which can help us capture personal effects such as habit formation.

It will be convenient to assume that all past problems in a database D were faced given the same vector of price indices (p_1, \dots, p_n) , and that this vector also applies to the current problem. One may use the model more generally, allowing prices to vary across the database and/or between the database and the new problem.

The household in question has a set of characteristics (x_1, \dots, x_m) , among which is its income, I , and has to choose a budget allocation vector $z = (z_1, \dots, z_n)$ in the $(n - 1)$ -dimensional simplex. We assume that it has several *constraints* on the way it splits the budget I among the expenditures E_1, \dots, E_n , and that these are linear constraints in z . That is, we assume that for some set A there exists a collection

$$\{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$$

where f_α is a linear function and $c_\alpha \in \mathbb{R}$.

Typically such constraints could be:

– Minimal quantities, say, the amount spent on food should suffice to cope with hunger:

$$\frac{z_i I}{p_i} \geq \alpha_i$$

or

$$z_i \geq \frac{\alpha_i p_i}{I}.$$

Note that this constraint is linear in z as p_i and I are considered to be fixed for the problem at hand (even if they differ from those appearing in past cases in the database).

– Ratios, say, the amount spent on entertainment cannot exceed 20% of the budget:

$$z_i \leq \alpha_i.$$

Clearly, equality constraints can be described by two opposite inequality constraints

The set of constraints $\{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$ should be expected to define a non-empty feasible set. In some situations, the constraints will uniquely determine the household's choice. For example, consider a household who maximizes a Cobb-Douglas utility function

$$u \left(\left(\frac{z_i I}{p_i} \right)_{i=1}^n \right) = \prod_{i=1}^n \left(\frac{z_i I}{p_i} \right)^{\alpha_i}.$$

As the optimal solution is given by $z_i = \alpha_i$, we can think of the household as picking the unique point z satisfying

$$\begin{aligned} z_i &\leq \alpha_i \\ z_i &\geq \alpha_i \end{aligned}$$

for all i . Indeed, a Cobb-Douglas household can be thought of as if it were operating by “rules of thumb” such as “housing expenditure should always be 30% of the budget”, “entertainment expenditure should always be 15% of the budget” and so forth. Note that in such a case the household’s choice is independent of the database D .

However, more generally one shouldn’t expect rules of thumb to single out a unique point (z_1, \dots, z_n) . For example, it stands to reason that a household would have some general guidelines as “housing expenditure *should not exceed* 40% of the budget” or “savings *should be at least* 25% of the budget” without identifying a unique vector z . Rather, these rules of thumb, or constraints define a *set* of expenditure proportions as *acceptable* to the household.

How does the household select a point in the acceptable set? We propose that it is at this stage that the database of past cases is brought to bear.⁴ Intuitively, one can think of the constraints as the information the household has about its own preferences. It might be aware, say, that spending more than 40% of income on housing will result in undesirable outcomes. Or, it might insist that it save enough. But beyond these constraints the household may well be unsure about its own preferences. Among all the budget allocations that are acceptable, the household does not have a clear ranking or a utility function that can be maximized. In these situations, the household resorts to its available information about the behavior of other households – including its own behavior in the past – and seeks guidance there.

⁴The household may well be affected by a database of past and/or others’ consumption also at the stage of determining the constraints. However, we do not explicitly model this phenomenon here.

Several empirical studies document the causal effect of household consumption decisions on those of other households. (See, for example, Brunnenberg, Dube, and Gentzkow (2012) for packaged goods and Andersen et al. (2013) for automobile brands.) Other studies attempt to go beyond the mere causal effect and show that it is due to social learning, as opposed to conformism. For example, Sorensen (2006) finds social learning in the choice of health plans, and Cai, Chen, and Fang (2009) establish it as a factor in dish selection in restaurants. While there are cases in which the social causal effect is quite clearly a matter of learning, and others in which it is evidently a matter of conformism, often one cannot easily disentangle the effects empirically. Indeed, a person’s own introspection might not be able always to tell them apart.⁵ In this paper we do not attempt to distinguish between the two. We offer a model that we interpret as causal, but that can be read either as social learning or as conformism.

Specifically, we assume that the household chooses the similarity-weighted average of the closest points – in the acceptable set – to the consumption proportions in the database. More formally, the household is characterized by

- (i) a similarity function $s : C \rightarrow \mathbb{R}_{++}$
- and
- (ii) a set of constraints

$$F \equiv \{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$$

such that

$$Z \equiv \cap_{\alpha \in A} \{ z \in \Delta(\Omega) \mid f_\alpha(z) \geq c_\alpha \} \neq \emptyset$$

and, for every $T \geq 1$ and every $D = (x_t, z_t)_{t=1}^T$ (where $x_t = (x_{1t}, \dots, x_{mt})$, $z_t = (z_{1t}, \dots, z_{nt})$) the household chooses the expenditure proportion vector z

⁵Indeed, one may argue that the taste for conformity has been evolutionarily selected because it allows for social learning. This would suggest that the distinction between the two is rather tricky. There are situations where one can isolate one of the factors, but there are likely to be many others in which they are intertwined.

that is given by

$$\frac{\sum_{t \leq T} s(x_t) y(z_t)}{\sum_{t \leq T} s(x_t)}$$

where $y(z_t)$ is the closest point to z in Z .

The similarity function s would normally put considerable weight on income I , suggesting that the household will tend to mimic the behavior of households with similar income. Similarly, demographic variables such as the number of people in the household, their ages and genders, are likely to play an important role in similarity judgments. This assumption seems plausible whether we think of the household as imitating others in order to learn what its choices should be based on other households' experience, or whether we tend to view imitation as driven by conformism. We emphasize that our model is compatible with both interpretations, and does not purport to distinguish between them.

This model is axiomatized in Section 4. We assume there that the household's behavior is observable given any database of choices, and is given by a function from databases to points in the simplex of expenditure proportions. We then impose several assumptions on this function, under which it can be represented as above for an appropriately chosen set of constraints F and similarity function s .

3.1 Relation to the Neoclassical Model

Before proceeding to the axiomatic derivation, it is worthwhile to compare our model with the standard one, that is, the neoclassical model of (expected) utility maximization subject to a budget constraint.

The two models differ in some obvious ways: in the neoclassical model the household has a complete ordering over all bundles, so that it is never at a loss when asked which bundle is preferred among several ones. By contrast, in our model the household doesn't have well-defined preferences. It considered some budget allocations acceptable, and the others – as unacceptable.

Thus, one can certainly say that the acceptable allocations are considered to be better than the unacceptable ones. But beyond this the household's constraints are silent on the ranking of allocations. Importantly, the constraints do not provide any guidance for the choice of a bundle among the acceptable ones.

It is here that past cases and the similarity function enter the game in our model. The model thus allows to describe the way that other households' choices and the same household's past choices affect the choice in the new problem. This type of effects cannot be described in the standard neoclassical model.⁶

While the two models seem to be very different, we have noticed that some types of household behavior can be compatible with both, as indicated by the example of Cobb-Douglas preferences, which can be fully captured by constraints in our model. It turns out that this can be generalized:

Proposition 2 *Let there be given a concave utility function u . Then there exists a (typically infinite) set of constraints $F \equiv \{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$ such that, for every similarity function s and every database D every optimal solution to $P(F, s, D)$ defines a maximizer of u (with quantities $z_i I/p_i$).*

However, the construction used in the proof of this proposition is somewhat artificial, and does not correspond to constraints as rules of thumb or otherwise cognitively meaningful concepts. We therefore suggest that the two models are quite different in nature, where ours assumes that preferences are not complete, and that choices are being made partly as a result of the constraints of the specific household, and partly as a result of imitation.

⁶One may enrich the standard model by allowing utility to depend on past choices, as well as on others' choices. But in this case the model will have to be further specified in order to remain meaningful.

4 Axiomatic Foundations

The model suggested in the previous section is motivated mostly by introspection. It attempts to formalize some mental processes that would hopefully sound plausible to the reader, at least in some cases, as does utility maximization in other cases. However, the introspective judgment of plausibility is far from sufficient to consider a model as a basis for consumer theory. One naturally asks, Can the theoretical concepts used in the model be measured by observations? Are there any modes of behavior that are incompatible with the model? What are the testable implications that such a model might have? More precisely, what type of observable data are needed to test the model? What hypotheses could be stated about these data that would be consistent, or inconsistent, with the model? Further, a sympathetic reader might accept the basic tenets of the model, and yet wonder about the particular modeling choices made in it: Why are the constraints linear? Why should we assume the similarity-weighted average, rather than some other formula?

These questions call for setting axiomatic foundations for the model. In doing so, this section will (i) describe the type of in-principle-observable data that are needed to test the model; (ii) characterize which of these data sets are and which aren't compatible with our model; (iii) determine the degree of uniqueness of the theoretical concepts used in the model; (iv) provide some arguments for the modeling choices made, and show that they are not arbitrary.

4.1 General Framework⁷

Let $\Omega = \{1, \dots, n\}$ be a set of expenditure *categories*, $n \geq 3$. Let C be a non-empty set of *cases*. Each case consists of a pair (x, z) where x denotes some economic and demographic variables, and y is a point in the simplex

⁷This sub-section is adapted from Billot, Gilboa, Samet, and Schmeidler (2005). While their model and result deal with probabilities, we re-interpret the vector in the simplex as expenditure proportions.

$\Delta(\Omega)$, denoting the way that the budget was split among the categories in Ω . Thus

$$C = X \times \Delta(\Omega)$$

with X being a subset of \mathbb{R}^k for some $k \geq 0$. A *database* is a sequence of cases, $D \in C^r$ for $r \geq 1$. The set of all databases is denoted $C^* = \cup_{r \geq 1} C^r$. The concatenation of two databases, $D = (c_1, \dots, c_r) \in C^r$ and $E = (c'_1, \dots, c'_t) \in C^t$ is denoted by $D \circ E$ and it is defined by $D \circ E = (c_1, \dots, c_r, c'_1, \dots, c'_t) \in C^{r+t}$.

Observe that the same element of C may appear more than once in a given database. In our model additional observations of the same case would have an impact on the household choices. The reason is that we do not think of the appearance of a case $c = (x, y)$ in the database mostly as a source of information but also as a causal determinant of behavior. If the only pathway in which one household's behavior affects another were awareness – that is, bringing a certain budget allocation to the household's awareness – then repetitions of the same case would have no additional impact beyond that of its first appearance. But if we believe that a given household might imitate others because it takes their choices as implicit evidence that these choices are good ideas, or if the reason for imitation is conformism, then we should expect repeated appearances of the same case to have additional impact.

Household behavior is defined as a choice of a single point in the simplex $\Delta(\Omega)$ for each $D \in C^*$. Thus, we are interested in functions

$$Y : C^* \rightarrow \Delta(\Omega)$$

defined for all non-empty databases. (As will be clarified soon, excluding the empty database simplifies notation.)

For $r \geq 1$, let Π_r be the set of all permutations on $\{1, \dots, r\}$, i.e., all bijections $\pi : \{1, \dots, r\} \rightarrow \{1, \dots, r\}$. For $D \in C^r$ and a permutation $\pi \in \Pi_r$, let πD be the permuted database, that is, $\pi D \in C^r$ is defined by $(\pi D)_i = D_{\pi(i)}$ for $i \leq r$.

We formulate the following axioms.

Invariance: For every $r \geq 1$, every $D \in C^r$, and every permutation $\pi \in \Pi_r$, $y(D) = y(\pi D)$.

Concatenation: For every $D, E \in C^*$, $y(D \circ E) = \lambda y(D) + (1 - \lambda)y(E)$ for some $\lambda \in (0, 1)$.

The Invariance axiom might appear rather restrictive, as it does not allow cases that appear later in D to have a greater impact on behavior than do cases that appear earlier. But this does not mean that cases that are chronologically more recent cannot have a greater weight than less recent ones: should one include time as one of the variables in X , all permutations of a sequence of cases would contain the same information. In general, cases that are not judged to be exchangeable differ in values of some variables. Once these variables are brought forth, the Invariance axiom seems quite plausible.

The Concatenation axiom states that the expenditure behavior induced by the concatenation of two databases cannot lie outside the interval connecting the expenditure behavior induced by each database separately.

The following result appeared in Billot, Gilboa, Samet, and Schmeidler (2005):

Theorem 3 *Let there be given a function $Y : C^* \rightarrow \Delta(\Omega)$. The following are equivalent:*

- (i) *Y satisfies the Invariance axiom, the Concatenation axiom, and not all $\{Y(D)\}_{D \in C^*}$ are collinear;*
- (ii) *There exists a function $y : C \rightarrow \Delta(\Omega)$, where not all $\{y(c)\}_{c \in C}$ are collinear, and a function $s : C \rightarrow \mathbb{R}_{++}$ such that, for every $r \geq 1$ and every $D = (c_1, \dots, c_r) \in C^r$,*

$$Y(D) = \frac{\sum_{j \leq r} s(c_j) y(c_j)}{\sum_{j \leq r} s(c_j)}. \quad (*)$$

Moreover, in this case the function y is unique, and the function s is unique up to multiplication by a positive number.

4.2 Constrained Case-Based Behavior

Theorem 3 deals with the way that accumulation of cases results in behavior. However, it remains silent on the behavior that corresponds to any single case. It focuses on the function $y : C \rightarrow \Delta(\Omega)$ which is basically the function Y restricted to databases of length 1, but it doesn't restrict this function in any way. In particular, it is possible that for two cases, $c = (x, z)$ and $c' = (x, z')$ (with $z \neq z'$) we will have $y(c) = z'$ and $y(c') = z$. This type of behavior might make sense if, for instance, the household in question is non-conformist, and wishes to behave differently from those it observes around it. Yet, this is not the type of behavior we seek to model. We devote this sub-section to additional assumptions, which will restrict the behavior of the function y .

Recall that, for any case $c = (x, z)$, $y(c) = y((x, z))$ is the expenditure proportion chosen by the household if it only observed a single case $c = (x, z)$. In light of Concatenation, this would also be the expenditure chosen by the household if it observed a very large database consisting of cases of type $c = (x, z)$ alone. It will often be more productive to think of large databases. In this context, a “single case c ” should be thought of as a large database consisting only of repeated appearances of c .

We will impose the following assumptions on the function y :

A3 Independence: For all $x, x' \in X$, and all $z \in \Delta(\Omega)$, $y((x, z)) = y((x', z))$.

Independence says the following: imagine that the household has observed only cases of type (x, z) . Thus, asking itself “what do others do?” it gets an unequivocal answer: households choose the budget allocation z . This might suggest that the household would choose z , too. Notice that the households observed had characteristics x , which might well differ from the household in

question. Yet, in the absence of any other households, there is no meaningful way in which more similar households can be more relevant.

If, having observed only z , the household also chooses z , Independence would hold: in this case we would have, for all $x, x' \in X$, and all $z \in \Delta(\Omega)$, $y((x, z)) = z = y((x', z))$. However, this might be too strong: the household in question might have some constraints that render z unacceptable. In this case, even if all the others have chosen z , our household might opt for another allocation. The Independence axiom only requires that this other allocation will not depend on the (only) x observed in the database. That is, it retains the requirement that $y((x, z)) = y((x', z))$ but allows it to differ from z .

Under Independence we can simplify notation and write $y(z)$ for $y((x, z))$.

Our main assumption is the following:

A4 Distance: For all $z \in \Delta(\Omega)$ and all $z', z'' \in \text{Im}(y)$, if $z' = y(z)$ and $z'' \neq z'$ then $\|z' - z\| < \|z'' - z\|$.

(Where Im denotes image of a function.)

The intuition behind this assumption is the following: when the household observes that everyone is selecting the allocation z , it tends to do the same. However, it is possible that the point z is not feasible for the household. For example, the household may have a constraint on the minimal amount of food it requires, and even if everyone in its sample consumes less, it would find it unacceptable to do the same. However, if two points, z', z'' are acceptable to the household – as evidenced by the fact that they are in the image of y , that is, that they can be chosen for some databases – then the choice among them is based on their distance from the point z . That is, given that everyone else chooses z , the household chooses the closest acceptable point to z .

We can now state the following:

Theorem 4 *The following are equivalent:*

(i) *The function y satisfies A3 and A4;*

(ii) *There exists a set of constraints*

$$\{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$$

where f_α are linear functions and $c_\alpha \in \mathbb{R}$ such that

$$Z \equiv \cap_{\alpha \in A} \{ z \in \Delta(\Omega) \mid f_\alpha(z) \geq c_\alpha \} \neq \emptyset$$

and, for all $x \in X$ and all $z \in \Delta(\Omega)$, $y(z) = y((x, z))$ is the closest point to z in Z .

Further, in this case the set Z is unique and it is the image of y .

Putting two theorems together yields the desired result. It is an immediate corollary of these theorems, and we dub it “theorem” only to indicate its conceptual import:

Theorem 5 *The following are equivalent:*

(i) *Y is not collinear, it satisfies A1, A2, and the resulting y satisfied A3 and A4;*

(ii) *There exists a $s : C \rightarrow \mathbb{R}_{++}$ and a set of constraints*

$$\{ f_\alpha(z) \geq c_\alpha \mid \alpha \in A \}$$

such that

$$Z \equiv \cap_{\alpha \in A} \{ z \in \Delta(\Omega) \mid f_\alpha(z) \geq c_\alpha \}$$

is a non-empty set which is not contained in an interval, and, for every $r \geq 1$ and every $D = (c_1, \dots, c_r) \in C^r$, $Y(D)$ is

$$\frac{\sum_{j \leq r} s(c_j) y(c_j)}{\sum_{j \leq r} s(c_j)}$$

where, for each $c = (x, z)$, $y(c)$ is the closest point to z in Z .

Furthermore, in this case the set Z is unique and the similarity function is unique up to multiplication by a positive constant.

Observe that the formula derived in the Theorem 5 differs from ?? in two ways: first, in the Theorem it is derived only for the case in which Y isn't collinear, whereas in the description of the model it was assumed to hold more generally. Second, the Theorem yields a similarity function $s(c_j)$ that depends on the entire case $c_j = (x_j, z_j)$ whereas in the previous section we assumed that this function is based solely on the similarity between the vector of characteristics in the past case, x_j , and the corresponding vector in the present one (whose notation is suppressed). One may formally add another axiom that would render $s((x_j, z_j))$ independent of z_j .

5 Discussion

5.1 Mental accounting

As briefly mentioned in the introduction, our model can be used to explain some mental accounting phenomena: given the consumer's (planned) budget allocation, and given a new consumption opportunity, the consumer might find that she cannot afford it if it belongs to one budget, but she can if the good is differently categorized. For example, a consumer who admires a cashmere sweater (Thaler, 1985) might think that she has spent all the money allocated to the clothes category, but still has funds in the birthdays category. Likewise, in the example we started out with, John might not have money in the category of "large, one-time expenses" in order to buy a car, but he may have money in the category of "small, daily expenses" to finance a car, at least if is willing to give up on *caffe latte* from that category.

It is sometimes suggested that mental accounting is a tool a consumer subject to self-control problems might use to control spending. An individual might use a mental accounting system to "... keep spending under control" (Thaler, 2004). Roughly the idea is that an individual can be thought of as consisting of multiple selves, with the current self setting out rules and budgets to discipline future selves and to limit their deviations from the plans

that are optimal from the current self’s point of view. In the present paper consumers might employ mental accounting, but for fundamentally different reasons. In our framework, consumers are as coherent as in the neoclassical model, but face complexity constraints in making decisions. In particular, they can be made better off if someone were to point out alternatives that they hadn’t considered, as in the example in the introduction, but there is no difficulty in determining whether changing their consumption choices is beneficial, as there is in multiple selves models. In this sense our “explanation” of mental accounting is conceptually a smaller deviation from the neoclassical model. At the same time, while multiple selves model are usually analyzed by standard game theoretic techniques, it is less obvious how consumer theory should be expanded to deal with the complexity challenges we discuss here.

We should make clear that we are not arguing that the account above for why a consumer might employ mental accounting is the only, or even the best, foundation for doing so. The suggestion is only that the complexity of the consumer’s problem can lead to mental accounting by consumers with completely standard neoclassical preferences.

5.2 Causal Accounts

Our model suggests that a household’s budget allocation is causally affected by budget allocations of other, similar households, as well as the same household’s past choices. Such a causal relationship may come in different flavors: is it driven by the fact that a household implicitly believes that others around it have learnt something that would be relevant to itself, or it is a result of conformism? That is, are we discussing a type of social learning, where others’ experiences are merely a source of information in a highly uncertain environment,⁸ or is it the case that these experiences are direct determinants

⁸Note that even in deterministic set-ups there exists uncertainty about one’s own preferences. Indeed, the psychological literature suggests that people do not seem to be par-

of one's preferences?

As stated above, our model does not formally distinguish between these explanations. Admittedly, the additive nature of the concatenation axiom may be more easily reconciled with conformism than with learning: if learning were the key factor to imitate others, then the accumulation of data should exhibit a decreasing marginal value of information. Yet, the model can be interpreted in both ways.

5.3 Normative Questions

One of the obvious differences between the Neoclassical model and ours is that the former features a notion of a utility function, which allows one to address normative questions, whereas the latter does not, at least not explicitly. If consumers behave according to the model proposed here, how can we tell if they are better off as a result of trade? How can we judge allocations for something along the lines of Pareto optimality?

These questions are of paramount importance. However, that fact that welfare economics exercises might be conceptually more challenging in the present model than in the classical one is not a reason to reject the former in favor of the latter. We hold that we should not choose a model only because it facilitates welfare economics exercises; rather, we should pose the normative questions within a model that is cognitively plausible.

Should one accept the constrained imitation model, there are several issues to be addressed regarding its interpretation. Consider first the set of constraints. If these are thought of as basic needs, such as minimal amounts spent on food and shelter, then the “size” of the feasible set may be taken as a measure of welfare, where a household who has more leeway beyond these constraints may be considered to be better off. But, as we saw in the Cobb-Douglas example, the constraints may also represent peculiarities of

ticularly successful in predicting their own well-being as a result of future consumption. Consumers do not excel in “affective forecasting” (see Kahneman and Snell, 1990, and Gilbert, Pinel, Wilson, Blumberg, and Wheatley, 1998).

tastes rather than basic needs. A household who has a very clear idea of its “utility” function may have constraints that leave no room for imitation, without necessarily being any worse off than without more lax constraints.

Similarly, the mechanism of imitation can be differently interpreted. If the main reason for imitation is social learning, then imitation is, overall, a desirable phenomenon, whereby households use each other’s experience to solve an otherwise unwieldy problem. If, however, imitation is largely due to conformism and social pressure, household who imitate others may not necessarily be thereby maximizing their independently-defined well-being.

The descriptive model proposed in this paper is insufficient to resolve these interpretational questions. It should probably be augmented by additional data that would help us determine whether constraints are an expression of needs or of tastes, and whether imitation is driven by learning or by conformism, before we could address the major normative questions.

6 Appendix: Proofs

Proof of Proposition 1:

It is straightforward that the Consumer Problem is in NP. Indeed, verifying that a proposed solution obtains the utility level \bar{u} is done in a number of algebraic operations that is linear in the length of the formula describing u .

To see that the problem is NP-Complete, we reduce to it the classical minimal set cover problem:

Problem COVER: Given a natural number r , a set of q subsets of $S \equiv \{1, \dots, r\}$, $\mathfrak{S} = \{S_1, \dots, S_q\}$, and a natural number $t \leq q$, are there t subsets in \mathfrak{S} whose union contains S ?

(That is, are there indices $1 \leq j_1 \leq \dots \leq j_t \leq q$ such that $\bigcup_{i=1}^t S_{j_i} = S$?)

Let there be given an instance of COVER: a natural number r , a set of subsets of $S \equiv \{1, \dots, r\}$, $\mathfrak{S} = \{S_1, \dots, S_q\}$, and a natural number t . Let

$(y_{ij})_{i \leq q, j \leq r}$ be the incidence matrix, namely $y_{ij} = 1$ if $j \in S_i$ and $y_{ij} = 0$ if $j \notin S_i$.

We now define the associated consumer problem. Let $n = q$. For $i \leq n$, let $p_i = 1$, and define $I = t$. Next, define u by

$$u(x_1, \dots, x_n) = \prod_{j \leq r} \sum_{i \leq n} y_{ij} x_i.$$

Finally, set $\bar{u} = 1$.

A bundle $(x_1, \dots, x_n) \in \mathbb{Z}_+^n$ satisfies $\sum_{i \leq n} p_i x_i \leq I$ and $u(x_1, \dots, x_n) \geq \bar{u}$ iff $\sum_{i \leq n} x_i \leq t$ and $\sum_{i \leq n} y_{ij} x_i \geq 1$ for every $j \leq r$. In other words, the consumer has a feasible bundle $x \equiv (x_1, \dots, x_n)$ obtaining the utility of 1 iff (i) no more than t products of $\{1, \dots, n\}$ are purchased at a positive quantity at x , and (ii) the subsets S_i corresponding to the positive x_i form a cover of $S = \{1, \dots, r\}$. Observe that the construction above can be performed in linear time.

It is left to show that we have obtained a legitimate utility function u . Continuity holds because this is a well-defined function that is described by an algebraic formula. Since $y_{ij} \geq 0$, u is non-decreasing in the x_i 's. We turn to prove that it is quasi-concave.

If there exists $j \leq r$ such that $y_{ij} = 0$ for all $i \leq n$, $u(x_1, \dots, x_n) = 0$, and u is quasi-concave.⁹ Let us therefore assume that this is not the case. Hence u is the product of r expressions, each of which is a simple summation of a non-empty subset of $\{x_1, \dots, x_n\}$. On the domain $\{x \mid u(x) > 0\}$, define $v = \log(u)$. It is obviously sufficient to show that

$$v(x_1, \dots, x_n) = \sum_{j \leq r} \log \left(\sum_{i \leq n} y_{ij} x_i \right)$$

is quasi-concave. But it is not hard to see that v is concave, hence quasi-concave: for every $j \leq r$, $\log \left(\sum_{i \leq n} y_{ij} x_i \right)$ is a concave function, and the sum

⁹One may wish to rule out these instances of COVER as they result in a satiable u .

of concave functions is concave. This completes the proof of the proposition.

□

A comment on approximations is in order. When we think of the consumer problem as maximization of the function u , approximation is naturally defined relative to this function. However, in the construction above the question is whether u can obtain the value 1, and no value can approximate it from below. (To be more precise, one may add max and min operations to the language, and define u above as $\min \sum_{i \leq n} y_{ij} x_i$, so that the approximation will be limited to $\{0, 1\}$.)

One can also look at the dual problem, and ask what is the minimal cost needed to obtain the level of utility \bar{u} . In our construction this would be equivalent to the minimal size of the set cover, for which approximations are unimpressive. (The problem is approximable only up to a factor $1 + \log(r)$.)

Proof of Proposition 2:

Let there be given a concave u . A maximizer of u is a vector $w = (w_1, \dots, w_n)$ (of quantities of the aggregate goods in each category) such that, for every (other) w' ,

$$u(w) \geq u(w').$$

Noticing that $w_i = z_i I / p_i$, we observe that $z_i = p_i w_i / I$ defines an optimal solution (w) if and only if, for every w' ,

$$u(z_1 I / p_1, \dots, z_n I / p_n) \geq u(w').$$

As prices p_i and income are fixed, u is a concave function of (z_1, \dots, z_n) . Hence it is the minimum of a collection of linear functions. That is, there exists a set

$$\{ f_\alpha(z) \mid f_\alpha \text{ is linear, } \alpha \in A \}$$

such that

$$u(z_1 I / p_1, \dots, z_n I / p_n) = \min_{\alpha \in A} f_\alpha(z).$$

It follows that z defines a maximizer of u if and only if, for every $w' \in \Delta(\Omega)$, and every $\alpha \in A$,

$$f_{\alpha}(z) \geq u(w').$$

It remains to define the set F as the set of all such linear constraints, where its index would be the cross product of A and $\Delta(\Omega)$. \square

Proof of Theorem 4: (i) implies (ii): Consider $\text{Im}(y)$ which is a non-empty subset of $\Delta(\Omega)$. Clearly, A4 implies that for every $z \in \text{Im}(y)$ is a fixed point of y , that is, $f(z) = z$. Conversely, every fixed point of y is in $\text{Im}(y)$. Let us denote the image of y , which is also the set of its fixed points, by Z .

Since A4 says that the point $y(z)$ is the closest one to z in Z .

Since the function y , selecting the closest point in Z , is well-defined, by Motzkin's Theorem (Motzkin, 1935, see also Phelps, 1957) it follows that Z is closed and convex. Hence Z can be written as the intersection of weak linear inequalities, proving the representation in (ii).

That (ii) implies (i) is immediate, as is the uniqueness of Z . \square

7 Appendix (Not Intended for Publication): Computational Complexity

A **problem** can be thought of as a set of legitimate inputs, and a correspondence from it into a set of legitimate outputs. For instance, consider the problem “Given a graph, and two nodes in it, s and t , find a minimal path from s to t ”. An input would be a graph and two nodes in it. These are assumed to be appropriately encoded into finite strings over a given alphabet. The corresponding encoding of a shortest path between the two nodes would be an appropriate output.

An **algorithm** is a method of solution that specifies what the solver should do at each stage. **Church’s thesis** maintains that algorithms are those methods of solution that can be implemented by **Turing machines**. It is neither a theorem nor a conjecture, because the term “algorithm” has no formal definition. In fact, Church’s thesis may be viewed as defining an “algorithm” to be a Turing machine. It has been proved that Turing machines are equivalent, in terms of the algorithms they can implement, to various other computational models. In particular, a PASCAL program run on a modern computer with an infinite memory is also equivalent to a Turing machine and can therefore be viewed as a definition of an “algorithm”.

It is convenient to restrict attention to **YES/NO problems**. Such problems are formally defined as subsets of the legitimate inputs, interpreted as the inputs for which the answer is YES. Many problems naturally define corresponding YES/NO problems. For instance, the previous problem may be represented as “Given a graph, two nodes in it s and t , and a number k , is there a path of length k between s and t in the graph?” It is usually the case that if one can solve all such YES/NO problems, one can solve the corresponding optimization problem. For example, an algorithm that can solve the YES/NO problem above for any given k can find the minimal k for which the answer is YES (it can also do so efficiently). Moreover, such an

algorithm will typically also find a path that is no longer than the specified k .

Much of the literature on computational complexity focuses on **time complexity**: how many operations will an algorithm need to perform in order to obtain the solution and halt. It is customary to count input/output operations, as well as logical and algebraic operations as taking a single unit of time each. Taking into account the amount of time these operations actually take (for instance, the number of actual operations needed to add two numbers of, say, 10 digits) typically yields qualitatively similar results.

The literature focuses on **asymptotic** analysis: how does the number of operations grow with the size of the input. It is customary to conduct **worst-case** analyses, though attention is also given to average-case performance. Obviously, the latter requires some assumptions on the distribution of inputs, whereas worst-case analysis is free from distributional assumptions. Hence the complexity of an algorithm is generally defined as the order of magnitude of the number of operations it needs to perform, in the worst case, to obtain a solution, as a function of the input size. The complexity of a problem is the minimal complexity of an algorithm that solves it. Thus, a problem is **polynomial** if there exists an algorithm that always solves it correctly within a number of operations that is bounded by a polynomial of the input size. A problem is **exponential** if all the algorithms that solve it may require a number of operations that is exponential in the size of the input, and so forth.

Polynomial problems are generally considered relatively “easy”, even though they may still be hard to solve in practice, especially if the degree of the polynomial is high. By contrast, exponential problems become intractable already for inputs of moderate sizes. To prove that a problem is polynomial, one typically points to a polynomial algorithm that solves it. Proving that a YES/NO problem is exponential, however, is a very hard task, because it is generally hard to show that there does *not* exist an algorithm that solves

the problem in a number of steps that is, say, $O(n^{17})$ or even $O(2^{\sqrt{n}})$.

A **non-deterministic Turing machine** is a Turing machine that allows multiple transitions at each stage of the computation. It can be thought of as a parallel processing modern computer with an unbounded number of processors. It is assumed that these processors can work simultaneously, and, should one of them find a solution, the machine halts. Consider, for instance, the Hamiltonian path problem: given a graph, is there a path that visits each node precisely once? A straightforward algorithm for this problem would be exponential: given n nodes, one needs to check all the $n!$ permutations to see if any of them defines a path in the graph. A non-deterministic Turing machine can solve this problem in linear time. Roughly, one can imagine that $n!$ processors work on this problem in parallel, each checking a different permutation. Each processor will therefore need no more than $O(n)$ operations. In a sense, the difficulty of the Hamiltonian path problem arises from the multitude of possible solutions, and not from the inherent complexity of each of them.

The class **NP** is the class of all YES/NO problems that can be solved in **P**olynomial time by a **N**on-deterministic Turing machine. Equivalently, it can be defined as the class of YES/NO problems for which the validity of a suggested solution can be verified in polynomial time (by a regular, deterministic algorithm). The class of problems that can be solved in polynomial time (by a deterministic Turing machine) is denoted **P** and it is obviously a subset of NP. Whether $P=NP$ is considered to be the most important open problem in computer science. While the common belief is that the answer is negative, there is no proof of this fact.

A problem A is **NP-Hard** if the following statement is true (“the conditional solution property”): if there were a polynomial algorithm for A , there would be a polynomial algorithm for any problem B in NP. There may be many ways in which such a conditional statement can be proved. For instance, one may show that using the polynomial algorithm for A a poly-

nomial number of times would result in an algorithm for B . Alternatively, one may show a polynomial algorithm that translates an input for B to an input for A , in such a way that the B -answer on its input is YES iff so is the A -answer of its own input. In this case we say that B is **reduced** to A .

A problem is **NP-Complete** if it is in NP, and any other problem in NP can be reduced to it. It was shown that the **SATISFIABILITY** problem (whether a Boolean expression is not identically zero) is such a problem by a direct construction. That is, there exists an algorithm that accepts as input an NP problem B and input for that problem, z , and generates in polynomial time a Boolean expression that can be satisfied iff the B -answer on z is YES. With the help of one problem that is known to be NP-Complete (**NPC**), one may show that other problems, to which the NPC problem can be reduced, are also NPC. Indeed, it has been shown that many combinatorial problems are NPC.

NPC problem are NP-Hard, but the converse is false. First, NP-Hard problems need not be in NP themselves, and they may not be YES/NO problems. Second, NPC problems are also defined by a particular way in which the conditional solution property is proved, namely, by reduction.

There are by now hundreds of problems that are known to be NPC. Had we known one polynomial algorithm for one of them, we would have a polynomial algorithm for each problem in NP. As mentioned above, it is believed that no such algorithm exists.

References

- [1] Anderson, S. T., R. Kellogg, A. Langer, and J. M. Saltee (2013), “The Intergenerational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy,” NBER Working Paper Series 19535.
- Aragones, E., I. Gilboa, A. Postlewaite, and D. Schmeidler (2005), “Fact-Free Learning”, *American Economic Review*, **95**: 1355-1368.
- Arthur, B., S. Durlauf, and D. Lane, Eds. (1997). *The Economy as an Evolving Complex System II*. Addison-Wesley, Reading, Mass., Series in the Sciences of Complexity.
- Billot, A., I. Gilboa, D. Samet, and D. Schmeidler (2005), “Probabilities as Similarity-Weighted Frequencies”, *Econometrica*, **73**: 1125-1136.
- Bernheim, B. D. (1994), “A Theory of Conformity”, *Journal of Political Economy*, **102**: 841-877
- Blundell, R. (1988), “Consumer Behaviour: Theory and Empirical Evidence – A Survey”, *The Economic Journal*, **98**: 16-65.
- Bronnenberg, B. J., J.-P. H. Dube, and M. Gentzkow (2012), “The Evolution of Brand Preferences: Evidence from Consumer Migration”, *American Economic Review*, **102**: 2472–2508.
- Cai, H., Ch. Yuyu, and H. Fang (2009), “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, **99**: 864-882.
- Debreu, G. (1959). *The Theory of Value*. Cowles Foundation Monographs Series.
- Deaton, A. and J. Muellbauer (1980), “An Almost Ideal Demand System”, *AER*, **70**: 312-326.
- Galbraith, J. K. (1958). *The Affluent Society*. Penguin Group USA.

- Gary, M. and D. S. Johnson (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. San-Francisco, CA: W. Freeman and Co.
- Gilbert, D. T., E. C. Pinel, T. D. Wilson, S. J. Blumberg, and T. Wheatley (1998), “Immune Neglect: A Source of Durability Bias in Affective Forecasting”, *Journal of Personality and Social Psychology*, **75**: 617-638.
- Goyal, S. (2005), “Learning in Networks: A Survey”. In G. Demange and M. Wooders (Eds.), *Group Formation in Economics: Networks, Clubs and Coalitions*, Cambridge: Cambridge University Press.
- Kahneman, D., and J. Snell (1990), “Predicting utility”. In R. M. Hogarth (Ed.), *Insights in Decision Making*, University of Chicago Press.
- Kahneman, D. and A. Tversky (1979), “Prospect Theory: An Analysis of Decision Under Risk”, *Econometrica*, **47**: 263-291.
- Kahneman, D. and A. Tversky (1984), “Choices, Values and Frames”, *American Psychologist*, **39**: 341-350.
- MacLeod, B. (1996), “Decision, Contract, and Emotion: Some Economics for a Complex and Confusing World”, *The Canadian Journal of Economics*, **29**: 788-810.
- Motzkin, Th. (1935), “Sur quelques proprietes caracteristiques des ensembles convexes”, *Atti Acad. Naz. Lincei Rend* 6, **21**: 562-567.
- Phelps, R. R. (1957), “Convex Sets and Nearest Points”, *Proceedings of the American Mathematical Society*, **8**: 790-797.
- Sabelhaus, J. (1990), “Testing Neoclassical Consumer Theory with Aggregate and Household Data”, *Applied Economics*, **22**: 1471-1478.
- Sorensen, A. T. (2006), “Social Learning and Health Plan Choice”, *RAND Journal of Economics*, **37**: 929-945.
- Thaler, R. H. (1980) “Toward a Positive Theory of Consumer Choice”, *Journal of Economic Behavior and Organization*, **1**: 39-60.

- Thaler, R. H. (1985) “Mental Accounting and Consumer Choice”, *Marketing Science*, **4**: 199-214.
- Thaler, R. H. (2004), “Mental Accounting Matters,” in *Advances in Behavioral Economics*, eds. C. Camerer, G. Loewenstein and M. Rabin, Princeton: Princeton University Press.
- Thaler, R. H. and H. M. Shefrin (1981), “An Economic Theory of Self-Control”, *Journal of Political Economy*, **89**: 392–406.
- Tversky, A. and D. Kahneman (1974), “Judgment under Uncertainty: Heuristics and Biases”, *Science*, **185**: 1124-1131.
- Tversky, A. and D. Kahneman (1981), “The Framing of Decisions and the Psychology of Choice”, *Science*, **211**: 453-458.