

# The Complexity of the Consumer Problem and Mental Accounting\*

Itzhak Gilboa,<sup>†</sup> Andrew Postlewaite,<sup>‡</sup> and David Schmeidler<sup>§</sup>

June 2010

## 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 a human. A simple and intuitive heuristic suggests that the consumer adopt a top-down approach, dividing her budget among main categories, further dividing these amounts to sub-categories and so forth. Such a heuristic may give rise to phenomena of mental accounting.

## 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. Support from Pinhas Sapir Center is gratefully acknowledged, and Postlewaite gratefully acknowledges support from the National Science Foundation.

<sup>†</sup>Tel-Aviv University, and HEC, Paris. [tzachigilboa@gmail.com](mailto:tzachigilboa@gmail.com)

<sup>‡</sup>University of Pennsylvania. [apostlew@econ.sas.upenn.edu](mailto:apostlew@econ.sas.upenn.edu)

<sup>§</sup>Tel-Aviv University and the Ohio State University. [schmeid@tau.ac.il](mailto:schmeid@tau.ac.il)

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). Another implicit assumption that is often dubious is that consumers are aware of all the bundles in their budget set. The following example illustrates.

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. Mary 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 never got to ask himself whether he preferred the coffee or the car. Consequently, it would be inaccurate 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 2 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.

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 idealization 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 used to have with a tight budget constraint. Correspondingly, it is not as obvious that her implicit choices satisfy the behavioral axioms of consumer theory.

The computational difficulties with the neoclassical model suggest that this model does not accurately describe the way consumers make decisions, at least not at the level of specific products. The question then arises, how do consumers make their decisions? Whatever process they use is likely to

lead often to decisions that will appear irrational or anomalous when viewed through the prism of the neoclassical model.

One simple way the consumer may solve her problem is to use a “top-down” approach, going from major categories down to sub-categories, then to sub-sub-categories, and so forth. For example, the consumer may first allocate her income between consumption and savings. Consumption may be split into durable and non-durable goods. Expenses on non-durable goods might be divided into food, transportation, entertainment, etc., and each of these items can be further split. Thus, one may consider a tree whose root represents total income, and every node – an expense on a particular (sub-)category. The consumer can be imagined to make decision regarding expenses in a top-down manner: she begins at the root of the tree and proceeds downwards, where, at every node, the “budget” for the node is the allocation that was decided upon at the node above. The number of sub-nodes relating to each node may be relatively small, and thus, at each step the consumer faces a low-dimension sub-problem, reminiscent of classroom examples in consumer theory. At the end of the process, the consumer arrives at a budget allocation.

An allocation that is arrived at in this fashion may be “locally optimal” in the sense that the consumer will not benefit from any re-allocation of the budget among the sub-nodes of any given node. Yet, such an allocation is not guaranteed to be “globally optimal”, that is, to provide an optimal solution to the original problem, for several reasons. First, the consumer may find it hard to assess the precise utility derived from a given budget allocation at the level of categories, until she knows how she is going to allocate the budget within each category. Even if the consumer does not have any uncertainty about her preferences between any pair of bundles, she cannot imagine which specific bundles will result from different allocations of the budget at the top levels. Second, there are many different ways to allocate products into categories and sub-categories, and different category trees may

result in different budget allocations. Third, certain products may belong to more than one category, and thus the categories may not form a tree in the first place. For example, the consumer may budget some money for clothing and some money for “unexpected opportunities”. A cashmere sweater with a normal price of \$300 may fall in the clothing category, but when it goes on special sale for \$150 it may come under unexpected opportunities. In spite of the difficulties that the consumer may face in allocating her income across categories, the computational complexity discussed above may make this top-down approach an intuitive and useful heuristic.

The top-down heuristic 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. Thus, 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 next section states the complexity result, whose proof is given in an appendix. Section 3 briefly explains how several examples of mental accounting may result from the top-down heuristic to the consumer problem, while Section 4 concludes with a discussion.

## 2 Complexity Results

Many writers have observed that the consumer problem is, intuitively speaking, a complex one. Some (see MacLeod, 1996) 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.<sup>1</sup> 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

---

<sup>1</sup>For 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 has to 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  $I$ .

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. (See further discussion in Appendix B.)

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 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. Aragoes, 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.

## 2.1 Problem 1: Products and characteristics

Consider a consumer who has to choose a bundle composed of  $n$  products. As a leading example, consider electronic products including mobile phones, hand-held computers, laptop computers, etc. The quantity bought of product  $i$  is  $x_i$ . The variable  $x_i$  is naturally a non-negative integer. It may simplify the problem to assume that  $x_i$  is either 0 or 1, but, as we shall see, this simplification will be of little help.

There are  $1, \dots, m$  characteristics, where each product has a certain subset of these characteristics. For example, the characteristics may be the ability to (i) place and accept phone calls; (ii) send and receive text messages; (iii) email; (iv) listen to pre-recorded music; (v) surf the internet; (vi) store files and photos; etc. Thus, a simple mobile phone will have characteristics (i) and (ii), but perhaps not (iii)-(vi). An MP3 device will typically have characteristics (iv) and (vi) but not necessarily (i) or (ii), and so forth.

Schematically, the product-characteristic matrix may look as follows:

Products	phone	text	e-mail	music	internet	photos
1	✓	✓				
2				✓		✓
3	✓	✓	✓			
4			✓	✓	✓	✓
5		✓	✓	✓	✓	✓

Product  $i$  has a price  $p_i$ . The consumer's income is  $I > 0$ . The question is, what is the best combination of products that the consumer can afford to buy at the given prices and income. Let us simplify the problem further by restricting attention to a simple class of utility functions: the consumer has a utility of 1 if for each characteristic she has bought at least one product that has this characteristic, and 0 otherwise. We can think of the consumer as insisting on having the ability to communicate by phone, text, e-mail, as

well as to listen to music and store data, etc. According to the matrix in this example, the consumer will be satisfied if she buys products 1 and 4, or, say, 3 and 5, but not if she buys products 2 and 3.

Clearly, one may consider a more general class of utility functions, allowing for more utility levels according to the degree to which the consumer's desires are satisfied. However, even the simple class of functions we consider here suffices for our result.

Let the *Consumer Problem* be: Given natural numbers  $n$  and  $m$ , and a matrix of  $n \times m$  entries  $\delta_{ij} \in \{0, 1\}$ , where  $\delta_{ij}$  denotes whether product  $i$  has characteristic  $j$ , prices  $(p_i)_i$  and income  $I$ , can the consumer obtain a level of utility 1?

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

Notice that with  $n$  products, the consumer might have to consider  $2^n$  different bundles (restricting attention to quantity that is 0 or 1). While the number of bundles is very large even for moderate values of  $n$ , this does not imply that the Consumer Problem is difficult; as we discussed above, there are problems with exponentially many possible alternatives for which there exist efficient algorithms. Claim 1 states more than a mere counting of the possible solutions. It says that, if there were an algorithm that could solve the Consumer Problem efficiently, there would have been such algorithms to each of the thousands of combinatorial problems that are in the class NP, including many well-studied ones. Consequently, it is plausible that actual consumers, whether they enumerate all possible bundles or not, cannot be guaranteed to solve their budget allocation problem optimally.

## 2.2 Problem 2: The classical consumer problem

Given the result above, it should not surprise us that more complicated problems, allowing for a more general class of maximization problems, are also NP-Complete. Yet, it is worth noting that among these more general

problems one can find the neoclassical consumer problem, of maximizing a quasi-concave utility function with a budget constraint. In defining this problem we follow the neoclassical tradition, according to which the various characteristics of the products, as well as the consumer's needs and wants are all encapsulated into the consumer's utility function. Thus, we 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, nondecreasing, and quasi-concave function.

Let the *General 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 General Consumer Problem is NP-Complete.*

### 3 Mental Accounting

In this section we consider several examples of mental accounting that might result from the top-down heuristic for budget allocations.

**Example 1:** (Thaler, 1985): Mr. S admires a \$125 cashmere sweater at the department store. He declines to buy it, feeling that it is too extravagant. Later that month he receives the same sweater from his wife for a birthday present. He is very happy. Mr. and Mrs. S have only joint bank accounts.

As indicated in the introduction, assume that at some node in the budget tree Mr. S divides the category “non-durable consumption” into “repeated expenses” and “special events”. The latter is supposed to be the budget he allocates to birthday and anniversary gifts, Christmas presents and do forth. The sum of \$125 may appear high as compared to the repeated expenses, but may appear reasonable when coming out of the special events budget. Indeed, splitting Mr. S’s annual income, any expense that is categorized as repeated is multiplied by the number of its recurrences per year. As a result, Mr. S has a greater incentive to be frugal when it comes to repeated expenses than when special events are considered. Thus, he may be happy to received the \$125 sweater as a birthday present even if he doesn’t think he can afford such expenses on a regular basis.

**Example 2:** (Thaler, 1985): Mr. and Mrs. L and Mr. and Mrs. H went on a fishing trip in the northwest and caught some salmon. They packed the fish and sent it home on an airplane, but the fish were lost in transit. They received \$300 from the airline. The couples take the money, go out to dinner and spend \$225. They had never spent that much at a restaurant before.

This is an example of unexpected income, which is used with much greater frivolity than expected income.<sup>2</sup> Consider the two couples and imagine that, before going on their trip, they have allocated their budgets along to budget trees according to the top-down heuristic. When their fish are lost, an unexpected \$300 extra income becomes available. Classical consumer theory would suggest that the budget line has been moved “away” from the origin, enlarging the budget set, and calling for a re-calculation of the optimization problem. Indeed, even according to the top-down approach the couples can decide to save some of the \$300 and consume the rest, to spend some of the consumption on durable and some on non-durable goods, etc. However, it is hardly worthwhile to re-calculate all the steps of the allocation problem. In-

---

<sup>2</sup>For a review of the literature on excess sensitivity of consumption see Browning and Lusardi, 1996.

stead, the consumers may decide to add the extra income to one of the leaves of the tree, consuming it in its entirety. Moreover, it makes more sense to add the income to a leaf of the tree where it would be noticeable, and thus adding it to “fancy restaurants” appears to make a bigger difference than adding it to “lifetime savings”. Indeed, additional psychological considerations are needed to explain how the extra money will be spent. Our main point is that the top-down heuristic naturally gives rise to behavior that is typically viewed as mental accounting.

**Example 3:** Revisiting the example from the introduction, John might not have considered the possibility of spending less on coffee and newspaper and more on buying a car because they belonged to different categories, say, “small, daily expenses” vs. “large, one-time expenses”. Once the allocation of the budget between these categories has been decided upon, it is difficult to conceive of trade-off between them further down the budget tree.

## 4 Discussion

### 4.1 Mental accounting

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. Examples 1 and 2 above can be thought of in this way. 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 examples such as 1 and 2 above 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.

## **4.2 Unknown utility**

The complexity result presented in Section 2 should be distinguished from the literature on learning one’s utility function. Indeed, the psychological literature suggests that people do not seem to be particularly 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). In other words, agents may be uncertain about their utility functions, and they may learn them through the experience of consumption. In this sense, a consumer is faced with a familiar trade-off between exploration and exploitation: trying new options in order to gain information, and selecting among known options in an attempt to use this information for maximization of well-being.

By contrast, our formulation of the consumer’s problem ignores this difficulty. We assume that the utility function is given, as an easily applicable formula, and that, given a particular bundle, there is no uncertainty regarding the utility derived from it. In this context, even in the absence of

uncertainty, the consumer’s problem is shown to be dauntingly hard to solve. Behaviorally, the two problems are distinct: a consumer who does not know her utility function needs more factual information to learn it, such as a new experience of consumption. But a consumer who faces a complexity problem may learn from sheer introspection, without any new facts. For example, in the example discussed in the Introduction, John finds new ways to organize his consumption simply because the availability of a bundle is pointed out to him, without learning from new experiences.

Realistic consumer problems are likely to be burdened with both sources of difficulty: first, the utility function may not be known for many bundles that have not been consumed; second, the number of possible bundles of indivisible goods, coupled with complementarity and substitution between them, make the problem hard to solve even under certainty.

## 5 Appendix A: Proofs

### Proof of Claim 1:

We prove the result by reducing the following problem, which is known to be NP-Complete, to the problem CONSUMER:

**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  $\bigcup_{l \leq t} S_{j_l} = S$ ?)

COVER is easily seen to be a special case of our problem. Specifically, given an instance of COVER, we define the characteristics to be the elements of  $S$ . For each subset  $S_j \in \mathfrak{S}$  we define a good  $i$  that has precisely the relevant characteristics. Setting all prices  $p_i$  to 1 and letting the budget be  $I = t$ , the consumer can afford a bundle that obtains utility 1 if and only if  $S$  can be covered by a subset of no more than  $t$  elements of  $\mathfrak{S}$ . Clearly, this reduction is linear in the size of the input.  $\square$

**Proof of Proposition 1:**

The reduction is from COVER again. 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.<sup>3</sup> 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

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

---

<sup>3</sup>One may wish to rule out these instances of COVER as they result in a satiable  $u$ .

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

□

## 6 Appendix B: 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 polynomial 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] Aragonés, 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.
- Browning, M., and A. Lusardi (1996), “Household saving: Micro theories and micro facts,” *Journal of Economic Literature* **34**, 1797 – 1855.
- Debreu, G. (1959). *The Theory of Value*. Cowles Foundation Monographs Series.
- 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.
- 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.
- Kahneman, D., and J. Snell, J. (1990), “Predicting utility”. In R.M. Hogarth (Ed.), *Insights in Decision Making*, University of Chicago Press.
- MacLeod, B. (1996), “Decision, Contract, and Emotion: Some Economics for a Complex and Confusing World”, *The Canadian Journal of Economics*, **29**: 788-810.

- Rubinstein, A. (2000). *Economics and Language*. Cambridge University Press.
- Thaler, R. (1980) “Toward a Positive Theory of Consumer Choice,” *Journal of Economic Behavior and Organization*, **1**: 39-60.
- Thaler, R. (1985) “Mental Accounting and Consumer Choice,” *Marketing Science*, **4**: 199-214.
- Thaler, R. (2004), “Mental Accounting Matters,” in *Advances in Behavioral Economics*, eds. C. Camerer, G. Loewenstein and M. Rabin, Princeton: Princeton University Press.
- Thaler, R.H., Shefrin, H.M., (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.