The Foerder Institute for Economic Research
at Tel Aviv University

מכון למחקר כלכלי על שם
ד"ר ישעיהו פורדר
על יד אוניברסיטת תל אביב

The Eitan Berglas School of Economics

בית-הספר לכלכלה ע"ש איתן ברגלס

עמותה רשומה

# The complexity of the consumer problem

**Itzhak Gilboa, Andrew Postlewaite, David Schmeidler**

**Working Paper No. 7-2021**

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

מכון פורדר, אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409255 | פקס: 03-6405815
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409255 | Fax 972-3-6405815 | email: foerder@tauex.tau.ac.il

Research in Economics 75 (2021) 96–103

Contents lists available at ScienceDirect

Research in Economics

journal homepage: www.elsevier.com/locate/rie

Research paper

# The complexity of the consumer problem[☆]

Itzhak Gilboa[a,*], Andrew Postlewaite[b], David Schmeidler[c]

[a] Tel-Aviv University, Tel-Aviv, Israel and HEC, Paris, France
[b] University of Pennsylvania, Philadelphia, PA, USA
[c] Tel-Aviv University, Tel-Aviv, Israel

ARTICLE INFO

*Article history:*
Received 21 November 2020
Accepted 2 January 2021
Available online 6 January 2021

*Keywords:*
Consumer theory
Computational complexity
Mental accounting

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. Two implications of this observation are that (i) households may imitate each other's choices; (ii) households may adopt heuristics that give rise to the phenomenon of mental accounting.



## 1. Introduction

Economists seem to be in agreement about two basic facts regarding neoclassical 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 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. And thus, while many writers have commented on the fact that a literal interpretation of the theory does not appear very plausible, economic theory generally adopts the "as if" interpretation of constrained utility maximization, thus rendering the second point largely irrelevant.

However, recent literature in psychology, decision theory, and economics is replete with behavioral counter-examples to the utility maximization paradigm, showing that the theory has many failures even when interpreted as a mere description of behavior. These include direct violations of explicit axioms such as transitivity or Independence of Irrelevant Alternatives, as well as examples that violate implicit assumptions, such as the independence of reference points or the fungibility of money (see Kahneman and Tversky, 1979; Kahneman and Tversky, 1984; Thaler, 1980; Thaler, 1985, and others).

☆ We wish to thank Bill Brainard, Hanming Fang, and Gabi Gayer for discussions and comments. Gilboa gratefully acknowledges ISF Grant 1077/17, the Investissements d'Avenir ANR -11- IDEX-0003 / Labex ECODEC no. ANR - 11-LABX-0047, the AXA Chair for Decision Sciences at HEC and the Foerder Institute at Tel-Aviv University. Postlewaite gratefully acknowledges support from the NSF.

* Corresponding author.
*E-mail addresses:* tzachigilboa@gmail.com (I. Gilboa), apostlew@econ.sas.upenn.edu (A. Postlewaite).

https://doi.org/10.1016/j.rie.2021.01.001
1090-9443/

In this paper we point at a stylized fact that is also incompatible with the classical model of consumer behavior. We provide an example that suggests that consumers are not aware of all the possible bundles they may choose among. We argue that this limited awareness cannot be dismissed as a mere mistake, due to the vast set of possibilities. In order to support this argument we discuss formal complexity results, which suggest that this type of “bounded rationality” is inherent to consumers’ decisions in an affluent society. We then introduce two implications of these results.

### 1.1. A difficulty with optimization

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 caffe 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 caffe 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 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.

### 1.2. 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 is plausible that it would satisfy axioms such as transitivity or the weak axiom of revealed preference. The mathematical model of utility maximization then seems 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 may well have 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 that her choices satisfy the behavioral axioms of consumer theory.

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

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 offer two building blocks that such a theory might employ.

The rest of this paper is organized as follows. Section 2 states the complexity results. In Section 3 two implications of these results are discussed. 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.[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. The term is borrowed from the computer science literature, and can be informally described as follows. A problem is considered to be "easy" ("Polynomial") if there exists an algorithm that solves it in a number of computation steps that is bounded by a polynomial of the input size. It is in "NP" if there exists such a "polynomial" algorithm that can *check* whether a suggested solution is indeed correct, though not necessarily to find one on its own. *NP-Complete* problems are a subset of NP, of which the following conditional statement is true: *if* there were a polynomial algorithm for one of them, *then* there would be one for all problems in NP. Because no polynomial algorithm is known for any NP-Complete problem, despite considerable attention that has been devoted to many of them, it is believed that NP-Complete problems do not have polynomial algorithms. 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 et al. (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, in the presence of indivisible goods, fact-free learning can also occur in the standard consumer problem, arguably the cornerstone of economic theory. Indivisibilities introduce a combinatorial aspect to the problem, and with it the possibility of NP-Completeness. Yet, the complexity of the problem depends on the way its input is encoded. For example, if the utility function is given by a matrix, specifying the utility value for each feasible bundle, the optimization problem is trivial (and can be solved in linear time). Thus, the complexity result hinges on a given formulation of the problem, which is to be judged for its psychological plausibility.

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.

### *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 $1, \ldots, m$ 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.

[2] 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 $I$.

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.

We do not prove Claim 1 as it be viewed as a re-interpretation of the problem COVER, which is known to be NP-Complete.[3]

### 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.[4] Thus, we consider a problem $P = \langle n, (p_i)_{i \le n}, I, u \rangle$ whose input is:

$n \ge 1$ – the number of *products*;

$p_i \in \mathbb{Z}_+$ is the *price* of product $i \le n$;

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

$u : \mathbb{Z}_+^n \to \mathbb{R}$ 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 $P = \langle n, (p_i)_{i \le n}, I, u \rangle$ and an integer $\bar{u}$, can the consumer obtain utility $\bar{u}$ in $P$?

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

We can now state:

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

The relevance of this simple proposition depends on the psychological plausibility of the statement of the problem. Is it natural to think of utility as given by an arithmetic formula? Perhaps it would be more natural to think of the function as given by a complete list of all its values, allowing for perfect generality, but resulting in linear complexity? We believe that this is not the case. The fact that the number of possible bundles grows exponentially fast as a function of the budget ($I$)

[3] COVER is defined in the appendix, in the context of the proof of the next result.

[4] See the discussion in Section 4.3 below.

suggests that people do not imagine each and every possible bundle in their feasible set. It seems more plausible to suggest that consumers consider product specifications and use them to "compute" the utility for each bundle. As in the example of electronic gadgets, for each subset of products one can figure out which services (such as phone, texting, e-mail) will be provided. It seems to us that the arithmetic language is relatively intuitive, and, as the proof shows, it can also capture the combinatorial aspects as in Problem 1.

## 3. Implications

### 3.1. Social learning

The notion of NP-Completeness is obviously a theoretical concept. Yet, it has two intuitive features: when a problem is NP-Complete, it is hard to *find* a solution for it, but (being in NP) it is not hard to *verify* that a suggested solution is indeed one. In our context, a household may be at a loss looking for a feasible bundle that obtains a given utility level; but it will find it very easy to ascertain that a possible solution does indeed obtain that level. Put differently, it is hard to answer the question, "Am I doing the best I can?" but easy to answer, "Will I be better off imitating X?"

These two features indicate that we should expect to see social learning in consumption behavior, even in the absence of uncertainty and of peer effects. Observe that, if the consumer problem were computationally easy, then each household could solve its own problem optimally and need not be affected by other households' choices. Conversely, if it were hard to find an optimal bundle, but also hard to check whether a suggested bundle yields a higher utility than the household's current choice, imitation of others could be risky. It is the combination of the two features that makes households particularly prone to imitate others: on the one hand, they cannot imagine all possible bundles; on the other, when they see a consumption pattern they can easily see its merits.

If all households were identical, social learning could be viewed as akin to computation by parallel processing: each household could be thought of as a processor that is assigned to check out the utility of a given bundle. If all households could observe each other's choices, a single one that happened to find a high-utility bundle could immediately be followed by all the others. Thus, we could imagine a suggested solution being known by all, where each household can improve upon it, to the benefit of all the others. In this sense, society can be viewed as a multi-processor computer (or a non-deterministic Turing machine). Clearly, learning the optimal solution isn't so simple when households differ. Yet, our model assumes that each households knows (with certainty) whether a suggested bundle is better than the one it currently consumes, and thus social learning can occur even without the households being identical. When there are more similar households to a given one, the latter is more likely to find good solutions by social learning.

Importantly, the reason we suggest here for social learning differs from two prominent reasons, namely, uncertainty and conformism. In reality, households obviously face uncertainty about the utility they would derive from given choices, ranging from housing to vacations, clothing to entertainment (see, for instance, Goyal, 2005). Further, conformism is a powerful determinant of consumer choices (as in Bernheim, 1994). Can these factors be told apart?

Our focus is on unawareness: the type of social learning that we discuss is neither along the lines of "if many people do it, it's probably not bad", nor along the lines of "I don't want to be the only one who doesn't have...". Rather, it is closer to "This has never occurred to me". As a result, the kind of social learning we discuss will depend on the number of other households who adopt a given solution discontinuously at zero: if household $i$ does not see any other household $j$ that chooses a given solution, household $i$ may not even think of this bundle. By contrast, when household $i$ observes some other households who consume a given bundle, their exact number does not matter: one is enough to bring the solution into awareness.

Note that the two other reasons for social learning would depend on the number of adopters: if the main reason to imitate others is uncertainty about the utility one derives from a bundle, the more households choose it, the better it is likely to be (though not necessarily in a linear way). Similarly, if peer pressure is the main reason to imitate others, the pressure would increase when there are more of these others who adopt a given solution.

The complexity argument for imitation would therefore be related to discontinuity near zero. In fact, as the example in the Introduction suggests, when unawareness is involved, a household can change its consumption pattern without imitating any other household. It may suffice that the idea is being proposed (as in the case of John and Mary), or that a hypothetical story be told.

### 3.2. Budgets and mental accounting

The consumer may try to simplify her problem by a "divide and conquer" approach, allocating her total budget to major categories, and then going down to sub-categories, 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. (See, for instance, Blundell, 1988; Deaton and Muellbauer, 1980, and Sabelhaus, 1990.) 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 decisions 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.

This process is probably the most common mental process that households who plan their expenditures go through. It is also similar to budgeting in organizations. It begs the question, however, how does one divide the budget at a given level without having decided what to do with it and lower levels? How much money should the household allocate to housing, food, transportation etc.?

Several accounts, which are largely compatible with each other, are possible. First, one may suggest that the household has a general idea of the marginal utility it can derive from each category of products. As often happens in complex systems, it can be easier to estimate aggregates than it is to estimate specific variables. This is certainly true of averages of independent random variables, based on laws of large numbers, but also of phenomena that need not be "truly" random, such as weather forecasts. Thus, a person might have a general idea of "how much better off" she will be if she spends another \$100 on communication as opposed to the same amount on food, without getting into the details of which mobile phone contract she would buy or which restaurant she would go to.

Second, one may use rules of thumb in choosing a budget allocation, which are often not available at the level of specific products. For example, one might recall the rule that "At least 30% of your income should be saved" or "One should spend about 30% of disposal income on housing". By contrast, one can't rely on such general rules when it comes to specific products, of which many may be completely novel. Hence, dividing the budget among categories may be an easier task than among products.

Third, the social learning mentioned in the previous subsection may be applied to categories of products. For example, one may obtain some data on average level of savings, or the average expenses on housing and on education by similar households.

Be that as it may, households may find it easier to divide their budgets among a few major categories of expenses than to make specific decisions about each product. And, given such a categories-budget, it is natural to proceed and divide the budget in each category among sub-categories and so forth.

The top-down heuristic may explain phenomena that are referred to as mental accounting (Thaler, 1980; 1985; 2004; Thaler and Shefrin, 1981). 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. This would not be the case if the budget allocation graph generated a tree, where the root represents total income, and the leaves – specific products. But if the graph is not a tree, so that a given product can be related to the root by multiple paths, mental accounting can occur.

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.

## 4. Discussion

### 4.1. 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 et al., 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 simple 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 utilities 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.

### 4.2. Mental accounting and bounded rationality

Ever since mental accounting was introduced (Thaler, 1980; 1985; Thaler and Shefrin, 1981), it has had several intuitive explanations, of which complexity is but one. People may divide their income into sub-budgets to cope with self-control problems, namely, with the risk of spending too much on some goods resulting in not having enough for others. (See a recent contribution by Galperti, 2019.) Relatedly, a multiple-selves approach, where different agents are responsible for different decisions, can also result in mental accounting (see "narrow thinking" in Lian (2018)). Another explanation has to do with over-generalization, where the consumer suspects that the decision she would make at a given point may be repeated by her in the future, not being able to recall past consumption decisions. (See Gilboa and Gilboa-Schechtman, 2003, who relate mental accounting to the Absent-Minded Driver Paradox of Piccione and Rubinstein, 1997.)

We argue that the computational complexity account of mental accounting provided here is "more rational" than the others. Suppose that a consumer could relegate her future consumption decisions to a computer program. Relegating decisions to a computer program allows her to deal with various self-control problems that may arise in the future. Further, she can ensure that the program keeps track of the large picture, and eliminates memory failures. Yet, the consumer cannot assume that the computer program will solve NP-Hard problems optimally. Hence, even if the consumer chooses to relegate her decisions to a machine, she may need to program various heuristics (such as sequential budgeting) and mental accounting may be one of the heuristics.

### 4.3. Predetermined utility

Our formulation of the consumer problem takes the utility function as part of the input. One may wonder whether this is a plausible model of the problem consumers face. After all, each consumer has a specific utility function, and our focus is on deterministic preferences. Why would a consumer have to think about other utility functions? The question is particularly relevant as specific functions need not be hard to maximize. For example, if $u(x) = x_1$ it is trivial to find an optimal bundle. Further, evolution, which may have shaped consumer preferences to begin with, might have also equipped consumers with the tools to find such an optimal bundle.

However, as mentioned above, in the neoclassical model the utility function encodes all the information about the products characteristics, the way and degree they satisfy needs etc. Considering Problem 1 again, the consumer considers subsets of five products, each of which is only denoted by an index. The fact that, for instance, Product 1 can satisfy the needs "phone" and "text", but not the other needs, is given only in the utility function. When new products are introduced into the market, the number of variables grows and the utility function has to be re-defined to describe the consumer's preferences given the new available bundles. It thus seems to us reasonable to assume that a single consumer faces an optimization problem where the utility function is part of the input.

## Declaration of Competing Interest

Authors declare that they have no conflict of interest.

## Appendix A. Proof

**Proof of Proposition 1:**

We prove the result by reducing the following problem, which is known to be NP-Complete (Gary, Johnson, 1979) to the General Consumer 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 $t \le q$, are there $t$ subsets in $\mathfrak{S}$ whose union contains $S$?

(That is, are there indices $1 \le j_1 \le \dots \le j_t \le q$ such that $\bigcup_{l \le t} S_{j_l} = S$ ?)

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 $\left(y_{ij}\right)_{i \le q, j \le r}$ 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 $i \le n$, let $p_i = 1$, and define $I = t$. Next, define $u$ by

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

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

A bundle $(x_1, \dots, x_n) \in \mathbb{Z}_+^n$ satisfies $\sum_{i \le n} p_i x_i \le I$ and $u(x_1, \dots, x_n) \ge \bar{u}$ iff $\sum_{i \le n} x_i \le t$ and $\sum_{i \le n} y_{ij} x_i \ge 1$ for every $j \le r$. 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 $y_{ij} \ge 0$, $u$ is non-decreasing in the $x_i$'s. We turn to prove that it is quasi-concave.

If there exists $j \le r$ such that $y_{ij} = 0$ for all $i \le n$, $u(x_1, \ldots, x_n) = 0$, and $u$ is quasi-concave.[5] 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 $\{x_1, \ldots, x_n\}$. On the domain $\{x \mid u(x) > 0\}$, define $v = \log(u)$. It is obviously sufficient to show that

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

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

## References


Aragones, E., Gilboa, I., Postlewaite, A., Schmeidler, D., 2005. Fact-free learning. Am. Econ. Rev. 95, 1355–1368.
Bernheim, B.D., 1994. A theory of conformity. J. Polit. Econ. 102, 841–877.
Blundell, R., 1988. Consumer behaviour: theory and empirical evidence—A survey. Econ. J. 98, 16–65.
Deaton, A., Muellbauer, J., 1980. An almost ideal demand system. Am. Econ. Rev. 70, 312–326.
Galbraith, J.K., 1958. The Affluent Society. Penguin Group USA.
Galperti, S., 2019. A theory of personal budgeting. Theor. Econ. 14, 173–210.
Gary, M., Johnson, D.S., 1979. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. Freeman and Co., San-Francisco, CA.
Gilbert, D.T., Pinel, E.C., Wilson, T.D., Blumberg, S.J., Wheatley, T., 1998. Immune neglect: a source of durability bias in affective forecasting. J. Personal. Soc. Psychol. 75, 617–638.
Gilboa, I., Gilboa-Schechtman, E., 2003. Mental accounting and the absentminded driver. In: Brocas, I., Carrillo, J.D. (Eds.), The Psychology of Economic Decisions. In: Rationality and Well-being, Vol. 1. Oxford University, Oxford, pp. 127–136.
Goyal, S., 2005. Learning in networks: a survey. In: Demange, G., Wooders, M. (Eds.), Group Formation in Economics: Networks, Clubs and Coalitions. Cambridge University Press, Cambridge.
Kahneman, D., Snell, J., 1990. Predicting utility. In: Hogarth, R.M. (Ed.), Insights in Decision Making. University of Chicago Press.
Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk,. Econometrica 47, 263–291.
Kahneman, D., Tversky, A., 1984. Choices, values and frames. Am. Psychol. 39, 341–350.
Lian, C., 2018. A theory of narrow thinking. Working paper.
MacLeod, B., 1996. Decision, contract, and emotion: Some economics for a complex and confusing world. Can. J. Econ. 29, 788–810.
Piccione, M., Rubinstein, A., 1997. On the interpretation of decision problems with imperfect recall,. Games Econ. Behav. 20, 3–24.
Sabelhaus, J., 1990. Testing neoclassical consumer theory with aggregate and household data. Appl. Econ. 22, 1471–1478.
Thaler, R., 1980. Toward a positive theory of consumer choice,. J. Econ. Behav. Organ. 1, 39–60.
Thaler, R., 1985. Mental accounting and consumer choice,. Mark. Sci. 4, 199–214.
Thaler, R., 2004. Mental accounting matters. In: Camerer, C., Loewenstein, G., Rabin, M. (Eds.), Advances in Behavioral Economics. Princeton University Press, Princeton.
Thaler, R.H., Shefrin, H.M., 1981. An economic theory of self-control,. J. Polit. Econ. 89, 392–406.
Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. Science 211, 453–458.


[5] One may wish to rule out these instances of COVER as they result in a satiable $u$.