Title: **Computational Complexity and Finance**

Authors: Peter Bossaerts, Carsten Murawski and Nitin Yadav

(All are from the Brain, Mind and Markets Laboratory, The University of Melbourne)*

Draft: 4 April 2018

Abstract

Central to finance is the problem of determining the values of assets. We refer to this as the *valuation problem*. Traditionally, finance has taken a statistical approach to solving this problem: the value of an asset is computed from past observations of its behaviour, such as prices, using statistical inference. Here, we show that this approach to valuation is (1) often ineffective and (2) sometimes intractable, due to *computational complexity*. This affects both individual investor choices and market behavior, with profound implications for the Efficient Markets Hypothesis: the role of prices is not to provide an accurate value signal, but to impel participants towards better solutions of the valuation problem.

1. Heads or tails?

The canonical approach in finance to resolving ignorance, such as determining the value of an asset, is to associate it with uncertainty. The values of assets are treated as random variables. Thus, the future value of an asset can be estimated based on past observations of its behavior such as prices. Put simply, finance treats valuation of assets like the valuation of a coin flip: in order to the estimate the expected payoff of the flip, we compute the average number of heads of past flips of the coin. The more flips we observe, the better our estimate will be.

In consequence, finance borrows many of its concepts from statistics. For example, the future value of a company is described in terms of a *statistical distribution*; bubbles are price rises that revert with a certain *probability*; and expected returns increase with *beta*, a statistical measure of covariation with the market return (see Fig. 1).

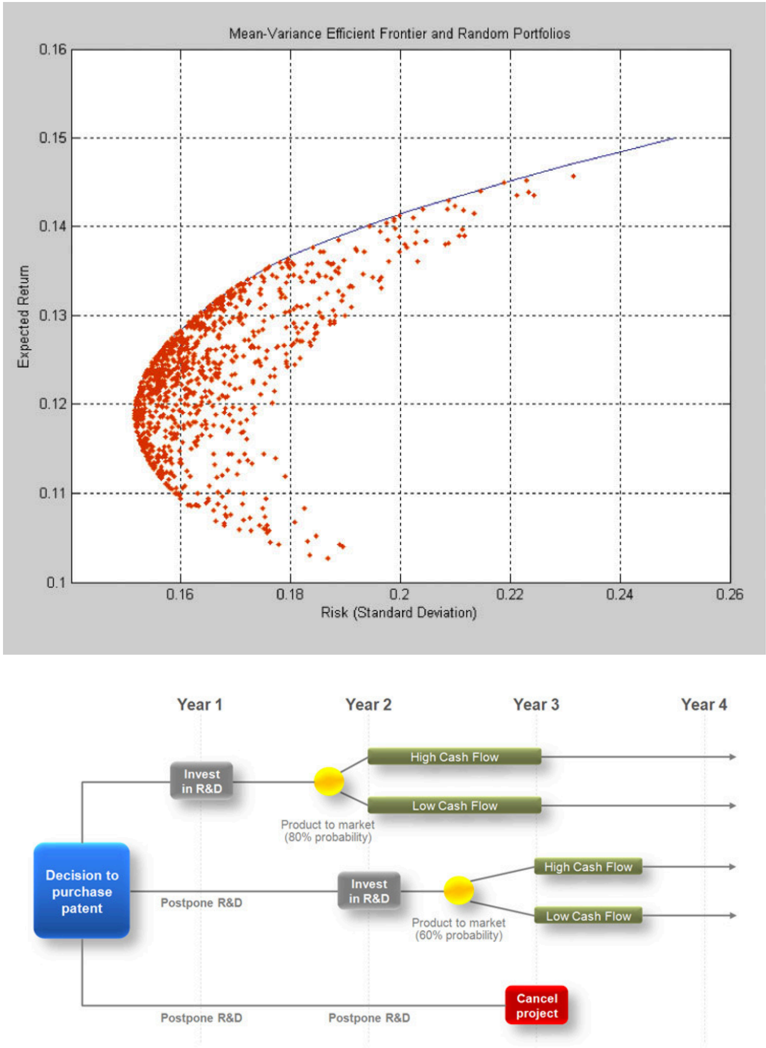

Figure 1. Use of statistics in investments (top) and corporate finance (bottom). Top: Monte-Carlo-based construction of an optimal portfolio, based on maximizing expected return given return variance, and hence assuming that return outcomes can be decribed as proper random variables (source: Mathworks). Bottom: Real options analysis of corporate finance decisions, assuming that probabilities can be associated with possible outcomes (source: Avinash Dixit, Wikimedia Commons)

To an outsider, finance therefore looks like a branch of statistics. The statistical approach has helped tremendously, in providing a disciplined approach to financial analysis and investments, to make sense of historical records of risk and reward, and to identify profit opportunities while controlling for risk. Today, statistics knowledge is a pre-requisite of most standard textbooks of finance [see, e.g., (Ross et al. 2011)].

But there are many situations where the statistical approach is ineffective. One such situation is valuation of a unique asset, for example, a start-up that is based on an entirely new technology. In such a situation, there are no historical observations of prices, not even a close analogue, and hence statistical inference is impossible.

Another situation is one in which we have all the necessary information to solve the valuation problem but computing the value is intractable: the computational resources such as time or memory required to solve the valuation problem are beyond those available. The *computational complexity* of the valuation problem is too high. In this situation, statistical inference is possible in principle. Here, we focus on this situation. We will argue that it is encountered much more often than the previous situation (lack of data) and that its consequences are much more severe.

We will illustrate the issues with a canonical optimization problem, the (0-1) *Knapsack Problem* (KP). In this problem, the decision-maker is asked to find, among a set of items with different values and weights, the subset that (i) fits a knapsack with a given weight constraint and (ii) maximizes the total value of the knapsack (see Fig. 2).

The KP permeates finance. It is the problem managers of companies face: which projects, from those available given a budget constraint, should be chosen in order to maximize the value of the company? The traditional portfolio investment problem falls in the same category: which assets to include in a portfolio in order to maximize expected return subject to a risk constraint?

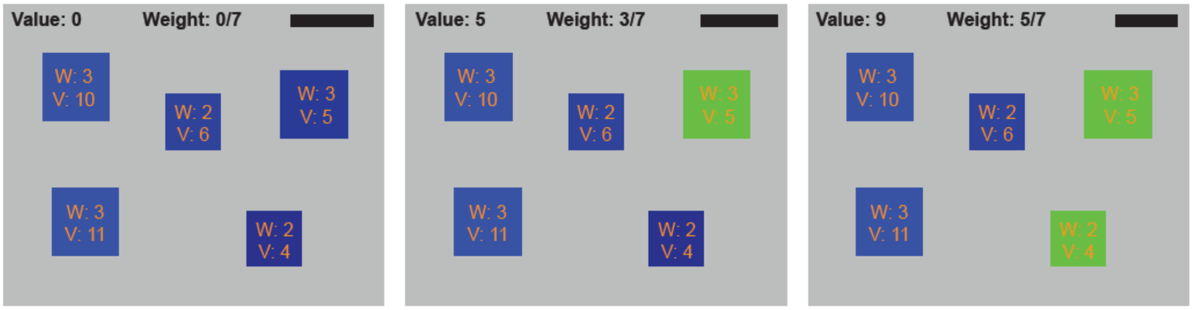

Figure 2. The Knapsack Problem (KP). The player is presented with a number of items, each with a certain value V and weight W. The goal is to find the subset of items with total weight at most equal to the capacity of the knapsack (here: 7) and maximum total value. Participants add one item at a time (Left to Right); the added item is highlighted, capacity use is indicated (on top) and the effect on value (top left) is displayed.

The KP is computationally "hard" (technically, it is part of the complexity class non-deterministic polynomial-time hard, or NP-hard). The number of possible combinations of items increases exponentially in the number of items. Very quickly, finding the optimal solution of an instance requires computational resources that exceed those of even the world's fastest supercomputers – the problem becomes intractable. In practice, electronic computers often resort to approximation algorithms that do not necessarily produce the optimal solution, but at least guarantee a minimum performance (Kellerer et al. 2004).

An important characteristic of computer algorithms for the KP is that they do not sample blindly by randomly trying out different combinations of items. Instead, they provide a disciplined, informed approach, to counter the enormous complexity of the problem. The alternative, using 'trial-and-error' to find the optimal solution, would readily become infeasible. With one thousand items, the number of possible combinations is in the order of 10^{301}, compared to the estimated number of atoms in the Universe, which is about 10^{80}. Importantly, observing a new combination of items does not necessarily provide any information about which items belong to the solution, that is, it does not necessarily 'get one closer' to the solution. One would have to be very lucky indeed to find the optimal solution by chance.

In the following, we compare effectiveness of the statistical approach to valuation, when the valuation problem is similar to estimating the value of a coin flip, against the situation where valuation is closer in nature to the KP. We then discuss empirical evidence that shows how people solve the KP. It turns out that people do far better than random sampling. In fact, they appear to possess the remarkable capacity to recognize whether a given instance of the KP is more difficult, something electronic computers cannot (yet) do.

Subsequently, we look at markets. We discuss how markets value assets when their valuation depends on the solution of a KP instance. Surprisingly, we find that market prices *do* behave "as if" they were the valuations of an individual who solves the instance using trial-and-error.

We then point out that this is not the first time experiments generated behavior at the market level that cannot be predicted from that of the individuals who populate the marketplace. Such findings have profound implications for the theory of asset pricing as practiced in academic finance (whether in behavioural or neoclassical finance), where market behavior is *assumed* to reflect the preferences of a "representative" or "marginal" investor (Cochrane 2009).

Since traders use a more effective way to solve the KP, it is therefore not surprising that we find market prices to reveal solutions that are inferior to that of the average trader. Yet we will also report that the very availability of a market improves the solution the average trader obtains.

We conclude by proposing a different role for financial markets than hitherto assumed in (academic) finance. We no longer view markets as providing, through its prices, the correct signals for everyone to become better off (the Efficient Markets Hypothesis, EMH). Instead, we propose that markets merely provide incentives for participants to do better. *Prices reveal little interesting, yet the average participant does better than without markets*.

The remainder of the paper is organized as follows. In the next section, we remind the reader of the essence of statistical inference: sampling. We then illustrate how ineffective standard statistical inference is when computational complexity is high, using the example of KP (Knapsack Problem), while more effective ways to learn exist. In Section 4, we discuss evidence that shows that humans understand this, and do not use standard statistical inference in the KP. In contrast, in Section 5, we show that markets behave “as if” they sample, and hence, prices reflect information that is inferior to that of the average trader. In the concluding Section 6, we argue that this has major implications for the Efficient Markets Hypothesis, and propose that finance de-emphasize the role of prices as signals, while instead looking at them as incentives.

2. Solving the valuation problem using the statistical approach

There are many reasons why we do not know the answer to a question such as ‘How much is company X worth?’. If the question is about the weather, say, in a week’s time, answers can be found using statistics: since weather patterns *repeat themselves* (technically, this is referred to as “ergodicity”), the answer can be determined by drawing inferences based on past observed weather patterns. With a sufficient number of observations, we can quantify the chance that our observation-based prediction is correct. That is, our beliefs can be quantified in the form of a probability distribution: “I believe that it will rain with 60% probability,” meaning that you expect to be correct and it will rain in 60 out of 100 observed cases.

In 1952, Savage showed that this statistical approach is universally “correct” (Savage 1972). That is: even if we are ignorant about something that does not repeat itself, that is, even if the question is unique, like in “What is the chance that the U.K. will finalize Brexit with an accord with the E.U.?”, we are safe to treat the event as one that does repeat itself. This is so because if we formulate beliefs in the form of a probability distribution, as if outcomes were going to repeat themselves, then our choices will be *consistent* (or “rational”). (Whether it is effective is another matter, as we shall see later.)

Choices are called “consistent” if a book-maker cannot bet with the decision-maker (DM) in a way that allows her to make money for sure. Consistency requires, among others, that if the DM is willing to pay 90 cents for you to insure him with one dollar in the case the U.K. will not finalize BREXIT, he ought to bid 10 cents for insurance in case the U.K. will finalize. If the DM bids more on the latter, say 50 cents, then you would be able to accept both insurance contracts, receive 1.40 dollars, with the obligation to only pay 1 dollar.

Finance has been using the term “absence of arbitrage opportunities” (or “no free lunches”) instead of “consistency”. Indeed, if one thinks of the above bids as market prices, it is obvious that it is possible can make money for sure by selling short the insurance contracts, receiving 1.40 dollars, investing these in a risk-free account, and when the insurance payment is due, use the proceeds (1.40 dollars plus interest) to pay the one dollar owed, keeping the remaining 40 cents plus interest. This is a sure profit, so there exists an arbitrage opportunity (or a free lunch).

As time evolves, information may become available that allows the DM to update her beliefs. Here too, Savage requires that updating be done in a consistent manner, using Bayes’ law. This law is a simple rule built around the principle that belief updates require the DM to consider the likelihood of observing the information given all possible outcomes. It does not matter whether this likelihood is true (for example, if history does not repeat itself – we don’t expect the U.K. to exit the E.U. many times, so likelihoods cannot be learned). One merely has to stick to one’s subjective belief, a subjective likelihood, and Bayes’ law.

In our example, if positive news emerges about the U.K. finalizing a Brexit deal, then the DM should decrease her willingness to pay for the insurance in case of failure, to below 90 cents, while increasing the bid for insurance in case U.K. does finalize the deal, to above 10 cents. Bayes’ law ensures that the resulting posterior beliefs continue to satisfy the no-arbitrage condition.

Savage’s analysis thus demonstrated that the statistical approach to resolving ignorance is universal. It pays to always think and act like a statistician. If you are asked what the distance is between Paris and London, and you don’t know, treat it as a random variable – even if the distance does not “vibrate,” that is, is not literally random (we are ignoring physical details here). Treat it as a parameter, endow it with a prior, say with mean 250km and standard deviation 50km, fly back and forth a number of times, and each time measure how far you flew. Update your beliefs every time using Bayes’ law. Very quickly, your posterior will be close to the ‘true’ value.

But notice that the goal of Savage’s approach is *not* to make sure that you learn as fast as possible; the goal is to ensure that you make the right choices, and that, if you learn, you keep on making the right choices despite updating your beliefs.

As such, Savage’s approach is not so much about learning. It certainly does not ensure that you learn effectively. Specifically, Savage’s approach is ineffective in the case your ignorance is due to computational complexity, as is the case in the KP. Kolmogorov pointed this out when he advised that probability theory ought not to be used when it is impossible to derive probabilities from random sampling (Kolmogorov 1983).

Kolmogorov came to this conclusion because he was not interested in consistent choices, but in information, that is, in inference. When it is impossible to set up an experiment that could verify one’s beliefs based on random sampling, one should use combinatorics instead, which forms the basis of the theory of computation. The argument is very simple. When random sampling, and the laws of large numbers and central limit theorems associated with them, do not apply, you take away the foundations of statistical inference. Without these foundations, statistical decision theory has little to offer – besides consistent choices.

In conclusion: There are many causes of ignorance, and Savage’s statistical approach is a generic way to address them all. To treat the unknown as a random variable and to learn by sampling guarantees avoidance of inconsistencies in one’s choices. From the point of view of learning, however, sampling may not be effective. There are many situations in which this is the case, among others, when ignorance emerges because of computational complexity. In such situations, another approach is needed.

3. Where The Statistical Approach Is Effective and Where It Is Not

In the following, we briefly consider a situation in which sampling is effective. Imagine there are ten securities that pay a liquidating dividend of 1 or 0 dollars, with equal probabilities. The Decision Maker (DM) collects information to learn these probabilities. Each period ("trial"), our DM receives a signal, which is positive with 70% chance and negative with 30% chance if the final dividend is 1, while positive with 30% chance and negative with 70% chance if the final dividend is 0. Our DM uses Bayes' law to update her beliefs about the final payoff. Assuming a quadratic loss function (as in the CAPM model of finance), her valuation of the securities equals the posterior mean, starting from an unconditional estimate of 0.5.

In this situation, beliefs will quickly converge to their 'true' probabilities (Fig. 3 top). Notice how valuations quickly separate between those that will end up paying 1 and those that expire worthless. The title of the plot lists the final payoffs of the securities, in order, from 1 to 10.

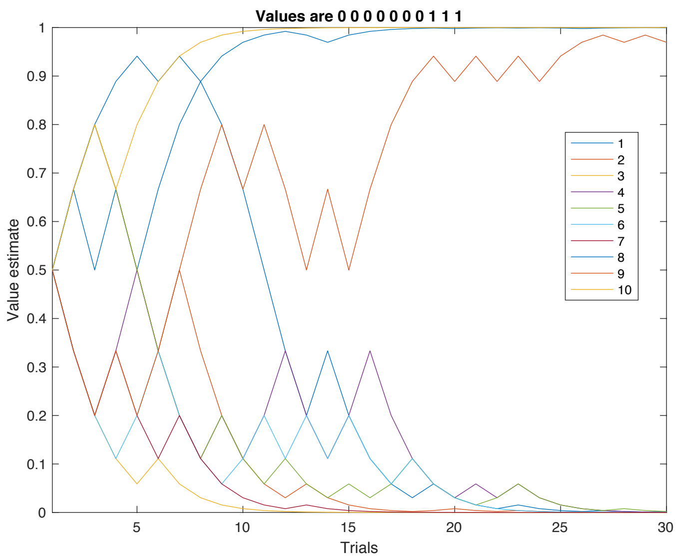

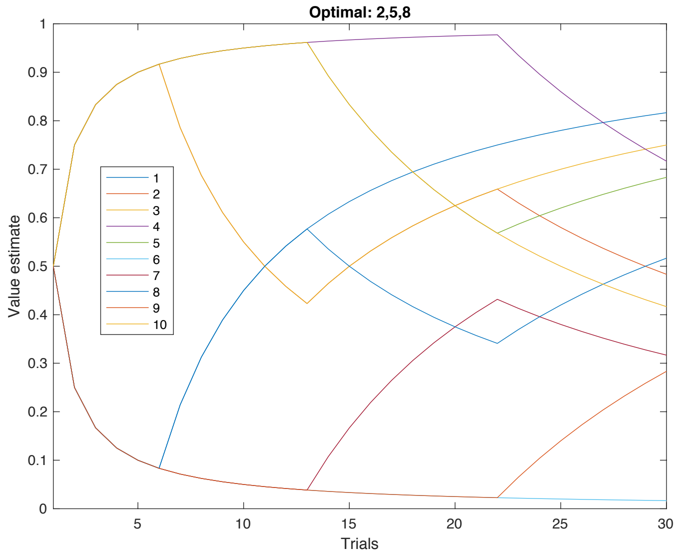

Figure 3. Belief Evolution Through Random Sampling when Determining the Composition of an Urn (Top) against when Determining Which Samples Matter (Bottom). Top: Ten belief (value) estimates; true value is either 0 or 1 (listed on top of plot, in order). Bottom: Ten belief (value) estimates; true value is either 1 (for assets in optimal solution, namely, assets 2, 5 and 8) or 0 (all other assets).

As such, repeated sampling (quickly) gets our DM closer to the truth over time. Specifically, the chance that she deviates too far from knowing the truth decreases with sample size. As time progresses, the chance of one valuation to veer off (say, from close to 1 down to 0) is drastically reduced. Even with a finite (and small number of) samples (signals), the valuation estimate is "probably approximately correct" (PAC) (Valiant 1984).

We now consider a different problem. Consider a set of securities whose terminal values depend on the solution of the following instance of the 0-1 KP. There are 10 items, with the following values (first number) and weights (second number): (31, 21), (141, 97), (46, 32), (30, 21), (74, 52), (105, 75), (119,86), (160, 116), (59, 43), (71, 54). The total weight limit is 265, which means that it can contain items only if their combined weight is at most 265.[†] The problem is to find the subset of items that maximizes the total value of items without exceeding the weight constraint.

There is a security for each available item, that is, there are 10 securities in total. If the item is in the optimal solution, then the corresponding security's terminal value equals 1; otherwise it expires worthless.

The problem our DM faces is how to value the securities and how to update her valuations as she samples potential solutions of the KP.

In accordance with Savage's framework, the DM could assign a prior value to each security, which can be completely arbitrary. She might assign the value 0.5 to each one, effectively revealing that she thinks each item has 50% chance of being part of the solution. She could then gradually adjust her beliefs by sampling, as follows. In each trial, she randomly tries a subset of the items that fills the KS to capacity. She then computes the value of this knapsack and compares it to the maximum value obtained in previous trials. She then constructs "signals" of the likelihood that an item is in the optimal solution, as follows. If the new value of the knapsack is less than the previous maximal value, then the signal for an item equals 1 if it was in the earlier (better) solution, and 0 otherwise. If the new knapsack value is higher than the trailing maximum, then the signal equals 1 if the item is in the current solution; otherwise the signal is 0. The DM then updates security valuations in trial t (t = 1, …, T) by weighing the valuation in the previous trial (t-1) by (t-1)/t and the signal in the present trial by 1/t.

This may not sound like Bayesian because it does not use the true likelihoods (of observing a particular knapsack value given the optimal solution). However, this is without consequence, because Savage only requires the DM to use some prior, not the true distribution. That is, our DM is allowed to start with the wrong beliefs.

The bottom panel of Fig. 3 illustrates how our DM's securities valuation would change over time. In each trial, she considered a randomly chosen subset of the items that filled the knapsack to capacity, and she updated her beliefs as explained in the previous paragraph. The title of the plot indicates which items are in the optimal solution (items 3, 5 and 8), and hence, which securities end up paying one dollar; all others expire worthless.

Even after 30 trials, it is still unclear which items are in the optimal solution, and hence, which securities will pay one dollar. To make matters worse, beliefs, and hence, valuations evolved arbitrarily, that is, they moved from high to low and *v.v.* even in later trials. This behavior of beliefs is in stark contrast with the previous case, where uncertainty was reduced by random sampling (top panel in Fig. 3).

Since there are only a finite number of possible capacity-filled knapsacks, our DM should eventually come across the optimal solution. This will happen when, just by chance, the best knapsack is drawn. But this takes time. In the present situation, there are (only!) 82 possible capacity-filled knapsacks, so the chance of drawing the best knapsack in any trial equals 1/82; this implies, among others, that the chance that the algorithm finds the optimal knapsack within 30 trials is about 5%.

To demonstrate how ineffective sampling is, notice that it takes approximately four hundred trials to increase to 95% the chance of coming across the optimal knapsack. But there are only 82 capacity-filled knapsacks, so in this case it would have been much better to list all 82 possibilities and pick the optimum. This simple, deterministic procedure, finds the optimum for sure in 82 steps, while the sampling approach may not find it in 5% of the cases even after 5 times as many steps!

4. Do Individuals Always Follow the Statistical Approach?

Savage has shown that the statistical approach provides a generic, but possibly ineffective, way to resolve uncertainty. As early as the 1960s, Ellsberg showed that in at least one circumstance, humans do not follow Savage's prescriptions. Indeed, when people have difficulty determining whether one prior is better than another, that is, in situations of *ambiguity*, humans stop being statisticians, to the point that their choices become inconsistent (Ellsberg 1961). It has been shown conclusively in experiments that this phenomenon also affects financial markets (P. Bossaerts et al. 2010; Asparouhova 2015).

When faced with combinatorial problems such as the KP, people likewise do not follow Savage's approach. In particular, they do not use random sampling to solve problems. Instead, they follow strategies that are best described in terms of the deterministic algorithms that computer scientists use to get to a solution as fast as possible without compromising too much accuracy. The latter means that these algorithms are fast (faster than random sampling), and guarantee a minimum performance (Kellerer et al. 2004).

In an experiment where people were paid to solve eight instances of KP, individuals performed far better than if they had used random sampling of potential solutions, that their initial choices reflect a well-known algorithm, known as the “greedy algorithm” (whereby one fills the knapsack in the order of the ratio of value over weight of the items), but that they subsequently deviated from it, similar to branch-and-bound algorithms (Murawski & P. Bossaerts 2016a).

In essence, humans solve the KP, not like statisticians, but like computer scientists. Critically, participants spent more time on instances that take computer algorithms longer to solve (Yadav et al. 2018), and their performance deteriorated rapidly as computational resource requirements increased – just as computer algorithms (Murawski & P. Bossaerts 2016a). Human and electronic computers appear to take similar approaches, and hence, computers are truly “artificially intelligent” (if one calls humans intelligent). While electronic computers are a lot faster than humans in the KP, they cannot, however, identify when a given instance will require more time to solve, while humans appear to be able to tell (Murawski & P. Bossaerts 2016a).

5. Behavior at the Market Level

We pointed out earlier that individuals do not appear to follow the statistical approach in case there is ambiguity about which prior to apply, and that this has major implications for market prices.

We now ask whether the same is true when valuation requires market participants to solve a combinatorial problem, namely, the KP. Do markets behave in ways that are fundamentally different from those predicted in asset pricing models, which assume that, at least implicitly, that all investors behave as statisticians?

Surprisingly, the answer is NO.

To address this question, we conducted an experiment where terminal values of securities depended on the solution of instances of the KP, as in the example in Section 3 (P. Bossaerts et al. 2018). Each item in the instance mapped into a security that participants could trade in an online marketplace. Exchange was organized using a continuous double-sided open book system, similar to the mechanisms used by electronic stock markets globally. Trading was done on the online experimental markets platform Flex-E-Markets.[‡] Participants traded for 15 minutes, or in later rounds, less. At the end of each round, securities were liquidated, and those that corresponded to items in the optimal solution paid one dollar per unit held, while the others expired worthless. Each round corresponded to a different instance of the KP.

Importantly, participants could earn money (beyond a small fixed sign-up reward and the terminal values of the securities they happened to be endowed with) <u>only</u> by trading in the marketplace. Even those who could quickly figure out the optimal solution to the KP instance at hand could not monetize their knowledge without trading. Participants had to take positions in the items they believed to be in the optimal knapsack by buying shares of the corresponding security. They could also sell shares corresponding to items they believed not to be in the solution, though short-sales were not permitted.

We discovered that prices behaved very much like the valuations of an individual who solves the KP instance using sampling. If one compares Figure 3, Bottom (valuations of an individual attempting to solve the KP using sampling) with Figure 4 (market prices), it is apparent that in both cases beliefs/prices continue to change abruptly, even after a lengthy adjustment period. At the end, neither the individual (Figure 3, Bottom) nor the market (Figure 4) can be confident that they are close to the solution.

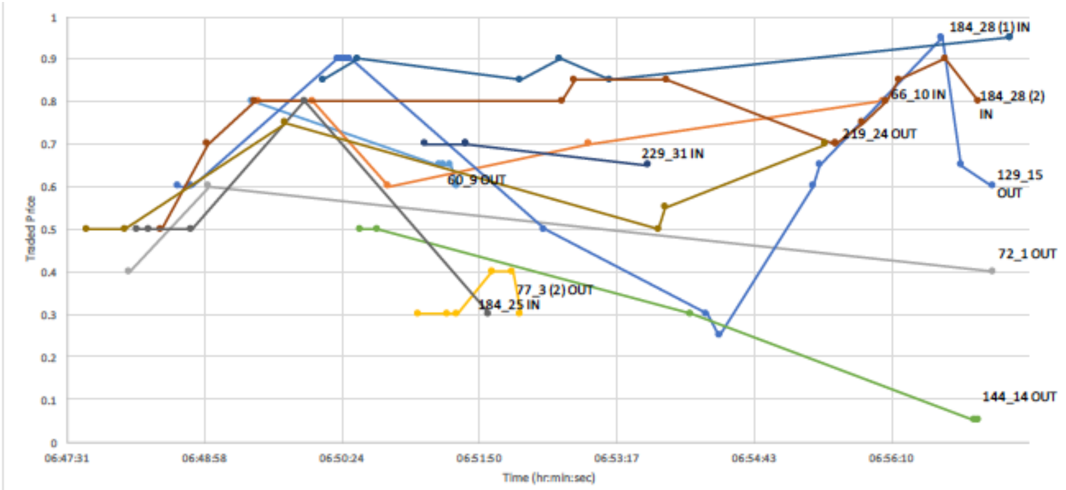

Figure 4.

Evolution of transaction prices in a market where values depended on the solution to an instance of the KP. Final payoffs equalled 1 if the solution was in the optimal knapsack, and zero otherwise. Securities are named by weight and value (weight_value) of the corresponding item, and whether the item was in (IN) or out (OUT) of the optimal knapsack.

Yet, market participants did not behave individually as if they were randomly sampling solutions. They searched systematically for the solution, as document in previous studies (Murawski & P. Bossaerts 2016b) & (P. Bossaerts et al. 2018).

In our experiments, we have repeatedly found markets to exhibit features that deviate markedly from those of the individual participants. In experiments with multiple single-period securities, individual participants care little about diversification, yet tight CAPM pricing obtains (P. Bossaerts & Plott 2004; P. Bossaerts et al. 2007; Asparouhova et al. 2003). In experiments with multi-period securities, individual participants seem to care little about hedging changes in the price of insurance, yet prices fluctuate in a pro-cyclical way, as in the core model of macro-finance, the “Lucas” model (often referred to as “DSGE,” Dynamic Stochastic General Equilibrium) (Asparouhova et al. 2016).

The message to take home from these markets experiments is that *one cannot predict market outcomes from individual behavior*. Markets have their own “laws,” and these laws are the results of interactions between many individuals all of whom behave in idiosyncratic ways. Unlike what is assumed in traditional financial modelling (Cochrane 2009), the behavior of the market cannot simply be predicted from that of a “representative” agent whose behaviour is like that of an individual market participant.

6. Implications for the Efficient Markets Hypothesis

When correct valuation depends on the solution of a problem of high computational complexity such as KP, markets prices evolve erratically (Fig. 4). As such, they hardly provide any information about the correct solution. This may be a good thing: “noise” keeps prices from fully revealing the attempts of individual participants, and thus provide incentives for participants to spend more effort. Indeed, noise is generally considered to be important for the good functioning of markets (Black 1986). Without noise, everyone can see what others know, and nobody has an incentive to spend any effort to seek a better solution. With noise, in contrast, everyone may eventually become better.

This prediction is also borne out in our experiments. More participants find the solution to a KP instance than in the absence of markets (Meloso et al. 2009). This is perhaps not surprising in view of the evidence depicted in Figure 4: if market prices behave “as if” the marginal investor uses an inferior trial-and-error approach to solving a KP instance, while individuals use the more effective, deterministic algorithms from computer science, individuals should do better.

To directly compare the two, we can construct candidate knapsack solutions by solely using transaction prices, and compare those against the solutions that traders individually suggested to us after market close. This direct comparison shows that the market does worse than the average trader (P. Bossaerts et al. 2018).

Closer inspection of how traders change their suggested solutions during trading reveals that they adjust their candidate solutions in response to a combination of prices and volume: if the market assigns little value to an item in an individual’s solution, and this valuation is accompanied by large trading volume in the corresponding security, then individuals tends to remove this item from their candidate solutions. The tactic generally improves the solution.

Altogether, a new picture of the role of a financial market emerges: the market provides incentives for everyone to do better. Prices reveal little interesting information, yet the little information they do reveal allows participants to end up knowing better than in a situation without a market.

This is in sharp contrast with the role of a financial market in the Efficient Markets Hypothesis (Fama 1991). There, the market is always superior, in that prices always reveal all the available information, so that individuals can blindly rely on those prices to do better. Not only is there plenty of evidence that this view is generally wrong, both from the field (D. Hirshleifer 2001) and from experiments (such as the one discussed before). A strict reading of the hypothesis leads to multiple internal inconsistencies, such as lack of incentives for individuals to pay effort to seek out information in the first place (Grossman & Stiglitz 1980), or, even if the information is free, to ignore it since the price is superior anyway (Hellwig 1980; Diamond & Verrecchia 1981). In addition, prices may reveal too much information, leading to inefficient allocation of risks (J. Hirshleifer 1971).

A field study on prediction markets confirms this novel view on the role of financial markets as merely providing incentives for individuals to do better: in a volatile market, more better-informed individuals are active, and changes in their positions provide more accurate forecasts of events than market prices do (F. Bossaerts et al. 2018)

Thus, prices ought to no longer be considered accurate signals. Instead, prices should be viewed as the means with which participants are incentivized to know better, and through these incentives, make everyone better off.

The price is dead. Long live the price.

References

Asparouhova, E., 2015. Asset pricing and asymmetric reasoning. *Journal of Political Economy*, 123(1), pp.66–122.

Asparouhova, E. et al., 2016. `Lucas' in the Laboratory. *Journal of Finance*, 71, pp.2727–2780.

Asparouhova, E., Bossaerts, P. & Plott, C., 2003. Excess Demand and Equilibration in Multi-Security Financial Markets: The Empirical Evidence. *Journal of Financial Markets*, 6, pp.1–21.

Black, F., 1986. Noise. *The Journal of Finance*, 41(3), pp.528–543.

Bossaerts, F. et al., 2018. When Prediction Markets Work: The Wisdom in the Crowd.

Bossaerts, P. & Plott, C., 2004. Basic Principles of Asset Pricing Theory: Evidence from Large-scale Experimental Financial Markets. *Review of Finance*, 8(2), pp.135–169.

Bossaerts, P. et al., 2010. Ambiguity in asset markets: Theory and experiment. *The Review of Financial Studies*, 23(4), pp.1325–1359.

Bossaerts, P. et al., 2018. Information aggregation under varying levels of computational complexity: Experiments.

Bossaerts, P., Plott, C. & Zame, W., 2007. Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiment. *Econometrica*, 75, pp.993–1038.

Cochrane, J.H., 2009. *Asset Pricing*, Princeton University Press.

Diamond, D.W. & Verrecchia, R.E., 1981. Information aggregation in a noisy rational expectations economy. *Journal of Financial Economics*, 9(3), pp.221–235.

Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. *The quarterly journal of economics*, pp.643–669.

Fama, E.F., 1991. Efficient Capital Markets: II. *The Journal of Finance*, 46(5), pp.pp. 1575–1617.

Grossman, S. & Stiglitz, J., 1980. On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3), pp.393–408.

Hellwig, M.F., 1980. On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3), pp.477–498.

Hirshleifer, D., 2001. Investor Psychology and Asset Pricing. *The Journal of Finance*, 56(4), pp.1533–1597.

Hirshleifer, J., 1971. The private and social value of information and the reward to inventive activity. *American Economic Review*, 61(4), pp.561–574.

Kellerer, H., Pferschy, U. & Pisinger, D.D., 2004. *Knapsack problems*, Berlin ; New York : Springer.

Kolmogorov, A.N., 1983. Combinatorial foundations of information theory and the calculus of probabilities. *Russian Mathematical Surveys*, 38(4), pp.29–40.

Meloso, D., Copic, J. & Bossaerts, P., 2009. Promoting intellectual discovery: patents versus markets. *Science*, 323(5919), pp.1335–1339.

Murawski, C. & Bossaerts, P., 2016a. How Humans Solve Complex Problems: The Case of the Knapsack Problem. *Scientific Reports*, 6, p.34851.

Murawski, C. & Bossaerts, P., 2016b. How Humans Solve Complex Problems: The Case of the Knapsack Problem. *Scientific Reports*, 6.

Ross, S.A., Westerfield, R. & Jordan, B.D., 2011. *Essentials of corporate finance*, New York : McGraw-Hill Irwin.

Savage, L.J., 1972. *The foundations of statistics*, New York, NY: Dover.

Valiant, L.G., 1984. A theory of the learnable. *Communications of the ACM*, 27(11), pp.1134–1142.

Yadav, N. et al., 2018. Phase transition in the knapsack problem: Computational and human perspectives.

---

[*] Address: 198 Berkeley Street, Carlton VIC 3010 (Australia). Email: peter.bossaerts@unimelb.edu.au. We thank participants at the 7th Behavioural Finance and Capital Markets Conference [Melbourne (Australia) 2017] for comments.

[†] This instance is number 8 in (Murawski & P. Bossaerts 2016b).

[‡] See http://www.flexemarkets.com.