Behavioral Limits to Complete Markets

Jacopo Magnani Yabin Wang*

September 2025

PRELIMINARY DRAFT, PLEASE DO NOT CITE

Abstract

Standard economic theory predicts that individuals should prefer complete markets to incomplete markets, as the former allow state-contingent claims for every possible outcome. Yet real-world markets remain incomplete, and the demand-side origins of this phenomenon are poorly understood. We develop an experimental framework to examine whether investors may themselves prefer incomplete markets, and highlight two potential mechanisms: preference instability, which exposes agents to greater regret or temptation in complete markets, and complexity costs, which arise because higher dimensionality increases cognitive effort and errors. In our experiment, participants consistently reveal a preference for incomplete markets, contradicting the rational benchmark. Comparing homegrown and induced-preference treatments, we find no evidence that this behavior is driven by preference instability. Instead, utility losses, response times, and subjective ratings indicate that complexity costs drive the preference for incompleteness. Structural estimation confirms that complete markets are several times more complex than incomplete ones, providing a behavioral foundation for market incompleteness.

*Norwegian University of Science and Technology (NTNU), Department of Economics. Email: jacopo.magnani@ntnu.no; yabin.wang@ntnu.no.

1 Introduction

Markets have evolved in many different ways, yet they remain fundamentally distinct from the complete markets envisioned in economic theory, where claims contingent on every possible future state can be traded (Arrow, 1964; Debreu, 1959). Directly assessing market completeness by examining the range of available assets is challenging, but the well-documented evidence of incomplete risk-sharing strongly suggests that real-world markets are indeed incomplete (e.g. Cochrane, 1991). From the perspective of neoclassical theory, this presents a puzzle: if complete markets enable individuals to optimize their choices under uncertainty, why have financial institutions not developed in ways that make them a reality?

Much of the literature on market incompleteness has focused on the transaction costs faced by security issuers or exchanges (Allen and Gale, 1988, 1990, 1991), which limit their ability to introduce new securities and complete the market. Less attention, however, has been given to the investors themselves, specifically, their actual need for and ability to navigate a complete market. In fact, one of the earliest explanations for market incompleteness stems from bounded rationality. For example, in their textbook treatment of incomplete markets, Magill and Quinzii (2002, page 16) make the following argument: "A complete itemization of all possible events at each date, for an extended period into the future would involve far more contingencies than any individual (or computer) could possibly calculate or envision. The individual costs of time and effort involved in the use of contingent contracts become prohibitive." However, these demand-side explanations of market incompleteness have largely remained informal and untested.[1] Do investors actually prefer incomplete markets, and if so, what are the causes of this preference?

We address this question through an experiment aimed at testing demand-side explanations of market incompleteness in a direct way. In the first step of our investigation, we develop a theoretical framework to analyze the demand for complete and incomplete markets. We focus on an individual portfolio choice problem in which each agent faces risk across a finite number S of states and allocates wealth through a set of securities with state-contingent payoffs.[2] In this setting, we consider a complete

[1]This is evident in the fact that models of incomplete markets regularly maintain standard assumptions about agents preferences and cognitive ability. On this point, Leijonhufvud (1993) remarked: "there is by now a sizeable technical literature on missing markets, investigating the conditions under which people can or cannot get around imposed constraints not to transact in parts of the commodity space. What is less clear is why such clever people choose to labor under these constraints."

[2]While the general equilibrium formulation of a complete market involves the possibility to trade

market made of one Arrow-Debreu security for each state, and an incomplete market whose $J < S$ assets induce a coarser partition of the state-space. Critically, the two markets are constructed in such a way that all allocations feasible in the incomplete market are also feasible in the complete market, but not vice versa. Rational agents should therefore always weakly prefer the complete market.

We then explore why agents might instead prefer incomplete markets. We first consider recent models in decision-theory showing that, under some conditions, agents may choose smaller menus over larger ones (Dekel et al., 2001; Gul and Pesendorfer, 2001; Sarver, 2008). This specific ranking of menus obtains either when agents are uncertain about their future preferences and anticipate regret, or when preferences are dynamically inconsistent due to temptation. Thus, we say these explanations are based on preference instability. In our context, preference instability models suggest that complete markets (just like larger menus) may expose agents to higher regret or temptation. Using a tractable example with subjective uncertainty over two possible utility functions (risk-neutral or maximally risk-averse), we demonstrate that the anticipation of regret can overturn the rational ranking, leading to higher welfare in the incomplete market when regret sensitivity is strong. In a second step, we turn to complexity-based explanations. Adapting the model of Gabaix and Graeber (2024), we allow agents to face cognitive costs when processing the parameters of the decision problem, such as prices and state-contingent endowments. Cognitive effort raises decision precision but entails disutility, so agents optimally trade off accuracy against cognitive costs. We show that, under plausible assumptions, complete markets are more complex because they involve a larger number of relevant dimensions, generating higher losses from imprecision and higher effort costs. As a result, when cognitive costs are sufficiently high, agents may prefer incomplete markets despite their lower rational utility frontier. Taken together, preference instability and complexity costs provide distinct theoretical foundations of a preference for incomplete markets.

We then designed an experiment that implements our theoretical portfolio choice problems and allows us to test whether subjects' behavior conforms with the rational model, preference instability models or models of complexity costs. In our experiment, each participant faced a sequence of three portfolio choice tasks, with 16 equiprobable states in every task. In the first two rounds, participants completed one task in a complete market with 16 Arrow–Debreu securities and one task in an incomplete

contingent on states realized at different times, we exclude the time dimension for implementation reasons.

market with only two securities. In the third round, participants chose whether to face a complete or an incomplete market, providing an incentive-compatible measure of revealed preference between the two market types. Our main treatment dimension varied the source of preferences. In the homegrown preferences treatment, payoffs depended on realized states, so that participants' own risk attitudes shaped outcomes. In the induced preferences treatment, payoffs were determined by a deterministic rule that implements a CRRA expected utility function, eliminating scope for preference instability (such as regret or temptation). By comparing market choices across these treatments, we can separate the role of preference instability from complexity-driven mechanisms. In a series of additional treatments, we provided participants with example portfolios (demos) designed to reduce cognitive costs, allowing us to test whether simplifying the task shifts behavior.

We begin by examining revealed preferences over market types and find that less than one third of participants chose the complete market in their third portfolio task, indicating a robust preference for incomplete markets and contradicting the benchmark rational model. Moreover, we find no evidence that instability in risk preferences drives this behavior: choices in the homegrown and induced treatments are statistically indistinguishable. This result point to complexity, rather than preference instability, as the main mechanism behind the preference for incomplete markets. Demos lead to slightly higher rates of complete-market choice, but the effect falls short of statistical significance, suggesting that just providing example portfolios is not enough to reduce cognitive costs.

In the next step of our investigation, we analyze round-level data on portfolio choices and a number of auxiliary variables including response time, cognitive uncertainty and a subjective complexity measure. We find that losses relative to the optimal benchmark are systematically larger in complete markets than in incomplete ones, with differences particularly pronounced in the induced preferences treatments where utility is directly observable. Response times, our proxy for effort, are also substantially higher in complete markets, with participants spending two to three times longer on these tasks. Complementing these measures, subjective ratings confirm that complete markets are perceived as more complex, while cognitive uncertainty is also higher, especially in the induced treatments. These findings jointly suggest that greater complexity generates larger mistakes in spite of higher effort, and reduces net utility.

Finally, we use a structural estimation approach to assess whether the complexity model can quantitatively account for the observed data. We estimate the model using the generalized method of moments, focusing on moments related to effort, relative losses, and market choices in the induced-preferences treatment. While the model somewhat understates the effort differential across markets, it matches the observed patterns of higher losses, higher effort, and lower choice probabilities for the complete market. The resulting parameter estimates are highly significant and imply that the complete market is about five times more complex than the incomplete one, with complexity rising more than proportionally with the number of dimensions to process (which in any market are the prices and state endowments). The model also predicts that subjects allocate attention selectively, focusing on prices but largely ignoring endowments, which helps explain the non-linear scaling of complexity.

Our paper relates to different strands of literature. First, there is a theoretical literature on the welfare properties of complete markets. The main results from this literature are provided in Magill and Quinzii (2002). A seminal contribution is Hart (1975) who showed that while social welfare is maximized by a complete market, introducing a new security without completing the market can lower welfare. Other important works include Malinvaud (1973), who shows that under some conditions a Pareto efficient equilibrium can be reached with fewer securities than those needed for a complete market, and Marin and Rahi (2000), who show that welfare may be higher in an incomplete market under asymmetric information. This literature has largely ignored how behavioral or cognitive factors may interact with market completeness in determining welfare. We take the first step in this direction by using an individual-choice approach rather than a general equilibrium framework.

Our experiment contributes to a growing literature on complexity in economics. The notion of complexity we adopt in this paper is that of task complexity (Oprea, 2024): we focus on descriptive elements of the decision problem, rather analyzing the procedures used by subjects in this task. In particular, our paper relates to other work on task-size complexity. One strand of this literature has looked at the complexity of lotteries, typically measured by the number of states in the lottery (e.g. Puri, 2024). Contrary to this approach, in our experiment the number of states is constant across treatments.

Our notion of complexity is closer to other ways of measuring cardinality, like the number of alternatives in a set. This notion has been used, for example, in Ortoleva (2013) and in the large literature on choice overload, partly reviewed in Chernev et al. (2015).[3] While this literature typically focuses on an exogenous and discrete number

[3]Other papers on choice overload include Iyengar and Kamenica (2010); Reutskaja et al. (2011);

of alternatives (such as lotteries), in our experiments participants can construct lotteries over final wealth, subject to feasibility constraints, and we vary the number of available assets. The notion of task-size complexity that seems to best fit our application is provided by the theoretical framework of Gabaix and Graeber (2024), where complexity is related to the number and importance of different problem dimensions (like prices and state-endowments, in our setting).

Finally, our work is related to a number of experimental papers on the effect of complexity and cognitive limitations on portfolio choice, such as Baltussen and Post (2011) and Magnani et al. (2022). Perhaps closest to our work, Carvalho and Silverman (2024) and Halevy and Mayraz (2024) compare behavior across portfolio choice tasks that vary in the number of assets and find the quality of choices (measured by GARP violations) decreases with the number of assets. The main difference from our paper is that all markets in their experiments are complete. Instead, our experiment is designed to understand the demand for incomplete markets.

2 Theoretical framework

2.1 Model setup and rational benchmark

The economy is subject to uncertainty at date 1, represented by a finite number of states of nature $\mathbf{S} = \{1, 2, \ldots, S\}$ with probability measure $Pr(1), Pr(2), \ldots, Pr(S)$. For simplicity we assume states are equiprobable: $Pr(s) = 1/S \forall s$. An agent has preferences over the distribution of final wealth across states, which we denote by: $\mathbf{x} = \{x_1, x_2, \ldots, x_S\}$. The agent has state-dependent endowments $\mathbf{w} = \{w_1, w_2, \ldots, w_S\}$. At date 0 the agent is also endowed with an initial budget w_0, that can be used to purchase assets which pay out at date 1.

The economy contains J securities. The assets prices are exogenous and denoted by: $p_1, p_2, \ldots, p_J$. One unit of asset j yields a payoff in state s denoted by y_s^j and we write the payoff matrix from all J assets as:

$$\mathbf{y}_{S\times J} = \begin{bmatrix} y_1^1 & y_1^2 & y_1^3 & \cdots & y_1^J \\ y_2^1 & y_2^2 & y_2^3 & \cdots & y_2^J \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ y_S^1 & y_S^2 & y_S^3 & \cdots & y_S^J \end{bmatrix}$$

Besedeš et al. (2015); Reutskaja et al. (2018).

The agent makes asset allocation decisions at date 0. We denote the agent's purchases of assets as $q_1, q_2, \ldots, q_J$. Asset purchases can be negative, as long as final wealth is greater than or equal to zero. Final wealth in state s is given by the endowment in that state and the payoffs from the asset holdings: $x_s = w_s + \sum_j y_s^j q_j$. Asset purchases are subject to the budget constraint: $\sum_j p_j q_j \leq w_0$.

The objective of the agent is to maximize his utility from date 1 consumption by choosing his asset holdings at date 0, subject to the constraints:

$$\begin{aligned}
\max_{\{q_j\}_{j=1}^{J}} \quad & U(x_1, ..., x_S) \\
\text{subject to} \quad & x_s = w_s + \sum_j y_s^j q_j \text{ for } s = 1, ..., 16 \\
& \sum_j p_j q_j \leq w_0 \\
& x_s \geq 0 \text{ for } s = 1, ..., 16
\end{aligned}$$

We consider two asset structures: a complete market and an incomplete market. A market is complete if any final wealth profile can be replicated using a portfolio of the J assets (e.g. Ross, 1976). Let $\boldsymbol{q} = (q_1, q_2, ..., q_J)^\top$ be the portfolio quantities and $\mathbf{x} = (x_1, x_2, ..., x_S)^\top$ be an arbitrary final wealth profile. The market is complete if any possible $\mathbf{x}$ can be constructed as a linear combination of the asset payoffs:

$$\mathbf{x} = \mathbf{y}\boldsymbol{q}$$

In other words, the market is complete if and only if the payoff matrix $\mathbf{y}$ has full row rank:

$$\text{rank}(\mathbf{y}) = S$$

A market that is not complete is called an incomplete market.

While there are many asset structures that can be either complete or incomplete, we focus on one simple structure in each class by imposing several assumptions. First, assets have binary payoffs: $y_s^j \in \{0, y\}$, for some $y > 0$. Second, we restrict attention to markets where only one asset can be used to allocate wealth to (or away from) a given state. Third, we require that, for any given state, there is an asset that can be used to allocate wealth to (or away from) that state. The first two assumptions are aimed at simplifying the asset structure, which allows us to focus on the effect of increasing the number of assets, rather than other features such as the correlation of payoffs.[4] By imposing the third assumption we want to rule out cases where there is a state in which all assets become worthless, as this represents an unlikely extreme scenario. To sum up, the asset structure defines a partition of the state space into J subsets, $\{\mathcal{S}_1, ..., \mathcal{S}_J\}$, each representing the set of states where asset j pays out.

Given these assumptions, we have one specific asset structure for the complete market and one for an incomplete market with a given number of assets. First, our complete market has $J = S$ Arrow-Debreu securities, with payoffs:

$$y_s^j = \begin{cases} y, & \text{if } s = j \\ 0, & \text{otherwise} \end{cases}$$

Second, the incomplete market has $J < S$ securities with payoffs:

$$y_s^j = \begin{cases} y, & \text{if } s \in \mathcal{S}_j \\ 0, & \text{if } s \notin \mathcal{S}_j \end{cases}$$

where $\{\mathcal{S}_1, ..., \mathcal{S}_J\}$ is a partition of the state-space, with at least one block containing more than one state. In our experimental design we will require that each block of the partition contains the same number of states.

Finally, we focus on a pair of markets that are linked in the following way. We construct the prices of the assets in the incomplete market by summing up the underlying complete market state-prices:

$$p_j^{incomplete} = \sum_{s \in \mathcal{S}_j} p_s^{complete} \tag{1}$$

The pair of markets we have designed ensures that any final wealth profile that is feasible in the incomplete market can also be achieved in the complete market. However, the converse is not true. For instance, it is easy to see that full insurance is possible in the complete market but not in the incomplete market. Thus, at the optimal portfolio choice, an agent's utility cannot be lower in the complete market than in the incomplete market:

$$U(\mathbf{x}^*_{complete}) \geq U(\mathbf{x}^*_{incomplete})$$

[4] In our markets, assets are perfectly negatively correlated. A more general correlation structure could introduce additional behavioral considerations, since the literature suggests people do not perfectly perceive correlation (e.g. Baltussen and Post, 2011; Ungeheuer and Weber, 2021).

where $\mathbf{x}^*$ denotes the optimal final wealth profile. Rational agents should always (weakly) prefer the complete market over the incomplete market.

What could lead agents to prefer the incomplete market? In the next two sections we consider two potential explanations, based on preferences and complexity respectively.

2.2 Preference-based explanations

Agents may prefer an incomplete market to a complete market because it minimizes regret or temptation. A number of contributions to decision theory have provided models in which agents prefer smaller menus to larger ones (Dekel et al., 2001; Gul and Pesendorfer, 2001; Sarver, 2008). In Gul and Pesendorfer (2001), smaller menus are preferred because they act as a commitment device and reduce self-control costs by eliminating tempting options. In Sarver (2008), smaller menus are preferred because they lower the chance that the agent's choice will be wrong ex post, when the agent is uncertain about his true preferences. In this vein, complete markets offer more flexibility and thus more potential for regret or temptation. To illustrate this idea, we adapt the regret model of Sarver (2008) to our setting.

We assume that when the agent is making his portfolio choice, he faces uncertainty about his risk-preferences. This uncertainty is modeled as a probability measure μ over a set of possible ex-post utility functions $\mathcal{U}$. After making a portfolio choice, and before the state is realized, uncertainty about the risk-preferences of the agent is resolved: ex-post utility is given by $U(\mathbf{x})$, where $\mathbf{x}$ is the profile of final wealth and $U(\cdot) \in \mathcal{U}$. Because the agent has made his portfolio decision before knowing his actual risk-preferences, it is possible that $\mathbf{x}$ is not the optimal choice given U. In this case, the agent experiences ex-post regret, which is is proportional to the difference between the maximum ex post utility attainable in the current market under U and the actual utility attained from the agent's choice. Denoting by M the set of feasible allocations in the current market, regret is:

$$R(M, U, \mathbf{x}) = K \times [U(\mathbf{x}^*) - U(\mathbf{x})] \tag{2}$$

where $\mathbf{x}^* \equiv \max_{\mathbf{x} \in M} U(\mathbf{x})$ and $K \geq 0$ measures the strength of regret.

To evaluate a set of feasible allocations M, the agent then computes the expected value of utility minus regret over different realizations of the agent's risk-preferences:[5]

$$\mathcal{W}(M) = \max_{\mathbf{x}\in M} \int_{\mathcal{U}} \left[U(\mathbf{x}) - R(M, U, \mathbf{x})\right] \mu(dU) \tag{3}$$

The final wealth allocation that achieves the maximum in equation (3) is the ex-ante optimal choice.

To show that anticipated regret may cause an agent to prefer an incomplete market to a complete market, we focus on a tractable example. We assume the agent is unsure whether he is risk-neutral or maximally risk-averse, specifically the agent's utility function over final wealth profiles is either $U(\mathbf{x}) = \frac{1}{S}\sum x_s$ or $U(\mathbf{x}) = \min_s x_s$ each with probability 0.5. We then compute the different utility terms using the parameter values of our experiment.

In our parametrization of the complete market, the agent's ex-ante optimal portfolio strategy is to maximize the minimum of the final wealth profile. This policy involves full insurance, that is a constant level of final wealth across states, equal to 266. Thus, this strategy yields an ex-post utility $U = 266$ when the agent is risk-averse and $U = 266$ when he is risk-neutral. The ex-ante optimal strategy is also optimal ex-post if the agent learns that he is in fact risk-averse. If the agent turns out to be risk-neutral, however, the ex-post optimal policy is to sell all securities whose price exceeds the minimum price and use the proceeds (plus the initial budget) to buy the cheapest security(ies). This strategy yields a utility of $U = 484$ to a risk-neutral agent. So, the agent's ex-ante optimal portfolio strategy results in a regret of $K \times (484 - 266) = 218K$ if the agent eventually learns he is risk-neutral. Total utility is then:

$$\mathcal{W}(M_{complete}) = 266 - 0.5 \times 218K = 266 - 109K \tag{4}$$

Similarly, using the numerical values for the parameters of our incomplete market, we find that the agent's ex-ante optimal portfolio strategy is to maximize the minimum final wealth. The agent achieves this by equalizing wealth across states with the lowest endowment in each block of the partition. This strategy yields an ex-post utility $U = 134$ when the agent is risk-averse and $U = 258$ when he is risk-neutral. The difference between the two values is due to the fact that there are several states with final wealth above the minimum. As before, this strategy is optimal also ex-post if the agent turns out to be risk-averse. If instead the agents learns that he is risk-neutral,

[5] The model of Sarver (2008) derives such a representation for an agent's preferences over menus of lotteries constructed from a finite set of prizes, rather than sets of feasible allocations.

the ex-post optimal policy is again to buy only the cheapest security and sell all others. However, sales of any asset j are now bounded by the lowest endowment in the block of states linked to asset j. As a result, the ex-post optimal policy of a risk-neutral agent only yields $U = 288$. Regret in this case is thus $K \times (288 - 258) = 30K$. Total utility is then:

$$\mathcal{W}(M_{incomplete}) = 0.5 \times 134 + 0.5 \times 258 - 0.5 \times 30K = 198 - 15K \tag{5}$$

While the agent prefers the complete market in the absence of regret ($K = 0$), when regret is sufficiently strong the preference is reversed. This is illustrated in Figure 1, which plots $\mathcal{W}(M_{complete})$ and $\mathcal{W}(M_{incomplete})$ over K.

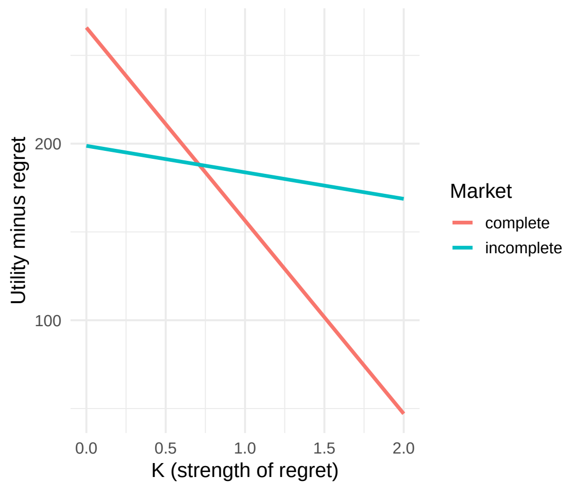

Figure 1: Market completeness and regret

This example illustrates the notion that, when one is uncertain about his own risk-attitude, anticipated regret can generate a preference for incomplete markets. It is also possible to construct a model based on Dekel et al. (2001); Gul and Pesendorfer (2001) that delivers similar results but is interpreted in terms of temptation. For instance in that model an agent who is initially risk-averse may be tempted to choose a risk-neutral portfolio strategy. As a result, the agent may prefer to face an incomplete market where the temptation utility from a risk-neutral portfolio strategy is lower than in a complete market. Overall, both models of anticipated regret and temptation seem to be able to generate a preference for incomplete markets. Both classes of models rely on dynamic changes in the agents' risk-preferences, either due to the resolution of uncertainty about own preferences or to dynamic inconsistency in preferences.

2.3 Complexity-based explanation

An alternative explanation for why agents may prefer the incomplete market is based on complexity. We formalize this notion adapting the model of Gabaix and Graeber (2024). In this model the agent must pay cognitive effort to process relevant parameters of the decision problem. We thus denote by $\mathbf{z}$ the parameters of the portfolio choice problem that require thinking effort: $\mathbf{z}$ includes the asset prices and the state-contingent endowments.[6] Denote by $U(\mathbf{q}, \mathbf{z})$ the expected utility as a function of asset quantities $\mathbf{q}$ and parameters, obtained from the original utility function after substituting the constraints. Denote the rational portfolio choice: $\mathbf{q}^r \equiv \arg\max_{\mathbf{q}} U(\mathbf{q}, \mathbf{z})$.

The problem's parameters (prices and endowments) are assumed to be drawn from a probability distribution: each z_i is is independent and has variance σ_i^2. Moreover, it is assumed that the agent obtains imperfect cognitive cues about each z_i, given by:

$$\tilde{z}_i = m_i z_i + (1 - m_i) z_i^d + \sqrt{m_i(1 - m_i)} K_i \varepsilon_i$$

where $m_i \in [0, 1]$ is the endogenous precision of the cognitive cue, z_i^d is a default, ε_i is a mean-zero noise with variance σ_i^2 and K_i is a scaling constant. While these cognitive cues can be interpreted as a form of imperfect perception, a broader interpretation is that they reflect cognitive uncertainty about how to map parameters into optimal choices, when parameters are different from a familiar default. The default is assumed to be a situation where all assets have the same price and endowments are constant, so that the optimal portfolio is made of equal quantities of each asset $q_j^d = \overline{q} \; \forall j$. While other defaults are possible, this version leads to a particularly easy choice rule. Agents will not simply rely on the default, however, because in general their cognitive cues have some positive precision $m_i > 0$. Using the cognitive cues about the parameters, the agent then chooses a portfolio $\tilde{\mathbf{q}} \equiv \arg\max_{\mathbf{q}} U(\mathbf{q}, \tilde{\mathbf{z}})$

The precision of the cognitive cue m_i is endogenous and is obtained by exerting cognitive effort L_i according to a cognitive production function:

$$m_i = \max\left\{1 - \left(\frac{L_i}{c_i} + 1\right)^{1-\alpha}, 0\right\} \tag{6}$$

[6]We focus on parameters that drive the optimal allocation across assets, rather than the initial budget w_0. As shown by Gabaix (2014), failures in processing the available budget are immaterial when the budget constraint must bind. The budget constraint is binding in our model and it is mechanically enforced in our experiment, as explained later.

where $\alpha > 1$, and $c_i > 0$ is an exogenous measure of the complexity of thinking about dimension i. A higher effort L_i increases precision, while a higher complexity c_i decreases precision.

While effort reduces cognitive uncertainty, it creates disutility. Total utility is assumed to be additively separable in the utility from consuming the future portfolio payoffs and the disutility of cognitive effort, where each unit of effort is assumed to decrease utility by w. As in Gabaix and Graeber (2024), we assume that, when the agent evaluates the effect of cognitive effort on the expected utility, he uses a second-order approximation of $U(\mathbf{q}, \mathbf{z})$ around the default $\mathbf{z}^d$. Then the optimal choice of cognitive effort solves:

$$\max_{L_1,\ldots,L_I} \sum_i V_i m(L_i) - w \sum_i L_i \tag{7}$$

with $V_i = -\frac{1}{2}\mathbf{q}'_{z_i} U_{qq} \mathbf{q}_{z_i} \sigma_i^2$ (all the derivatives are evaluated at the default).

Under these assumptions, Gabaix and Graeber (2024) show that it is possible to construct a cognitive production function and a cognitive complexity measure at the level of the whole utility maximization problem. To do this, first define U^r as the expected utility at the rational choice, $U(0)$ as the utility when the agent exerts zero cognitive effort (thus adopting the default choice) and $U(L)$ as utility when the agent exerts total cognitive effort $L = \sum_i L_i$. Then, we can measure the relative utility loss in the portfolio choice problem as $\Lambda(L) \equiv \frac{U^r - U(L)}{U^r - U(0)}$, which is normalized to be between 0 and 1. The main result is the loss function:

$$\Lambda = (\alpha - 1)^{\frac{1-\alpha}{\alpha}} \left(\frac{L}{C} + \Phi\right)^{1-\alpha} \tag{8}$$

with $\Phi \equiv \frac{\sum_i c_i}{C}$, and aggregate complexity C given by:

$$C \equiv \left[\sum_i s_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}}\right]^{-\frac{\alpha}{1-\alpha}} \tag{9}$$

where $s_i = \frac{V_i}{\sum V_j}$ is the relative importance of dimension i.[7] Finally, Gabaix and Graeber (2024) show that under the production function above, higher aggregate complexity

[7] These expressions hold for the case where $L_i > 0 \forall i$. The more general expressions are: $\Lambda = 1 - \sum_{i \in A} s_i + (\alpha - 1)^{\frac{1-\alpha}{\alpha}} \left(\frac{L}{C} + \Phi\right)^{1-\alpha}$, $\Phi \equiv \frac{\sum_{i \in A} c_i}{C}$ and $C \equiv \left[\sum_{i \in A} s_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}}\right]^{-\frac{\alpha}{1-\alpha}}$, where A is the set of dimensions with $L_i > 0$.

C leads to higher losses Λ and, under some conditions, higher cognitive effort L.

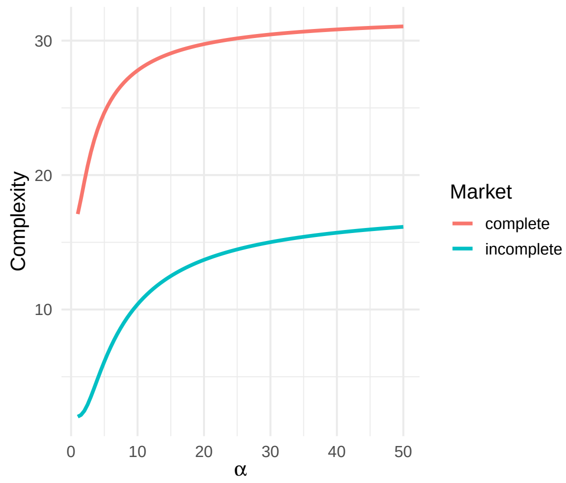

Figure 2: Complexity indices

Does this model predict the complete market is more complex? It is difficult to answer this question in general. If all dimensions of the task have the same importance ($V_i = V \forall i$) and complexity ($c_i = c \forall i$), then complexity C is proportional to the number of dimensions. This suggests that a market with more prices should be more complex. However, different dimensions like prices and endowments are not likely to have the same importance. To obtain a more precise answer, we use the parametrization of our experiment and make an assumption about risk-preferences, namely that they are expected utility preferences with constant relative risk aversion $\gamma = 0.2$ (this is the utility function we induce in one of the treatments of our experiment, as explained below). In this case, under the assumption that $c_i = 1$ for all dimensions and that the agent devotes positive effort in thinking about each dimension, we can compute the C index for the two markets. We plot them in Figure 2 as a function of α. For all α values, the complete market is more complex than the incomplete market. In turn, this implies that losses should be higher in the former.

Moreover, the complexity model suggests this could create a preference for the incomplete market. Denote by $\mathcal{W}$ the expected consumption utility net of cognitive costs at the optimal choice of cognitive effort, obtained from equation (7):

$$\mathcal{W} = U^r - \Lambda[U^r - U(0)] - wL \tag{10}$$

While utility for a rational agent is (weakly) higher under a complete market ($U^r_{complete} \geq$

$U^r_{incomplete}$), losses are larger under a complete market ($\Lambda_{complete} > \Lambda_{incomplete}$). If the losses due to complexity and the additional effort spent in the complete market ($L_{complete} > L_{incomplete}$) are large enough, then this will result in a preference for the incomplete market. To illustrate this point, we analyze a numerical example (based on the actual estimates of the model parameters we obtain from our data, see below). In Figure 3, we plot the utility net of cognitive costs $\mathcal{W}$ for the two markets as a function of the unit cost of effort w. When w is low, the complete market dominates, but for sufficiently high values of w the agent prefers the incomplete market.

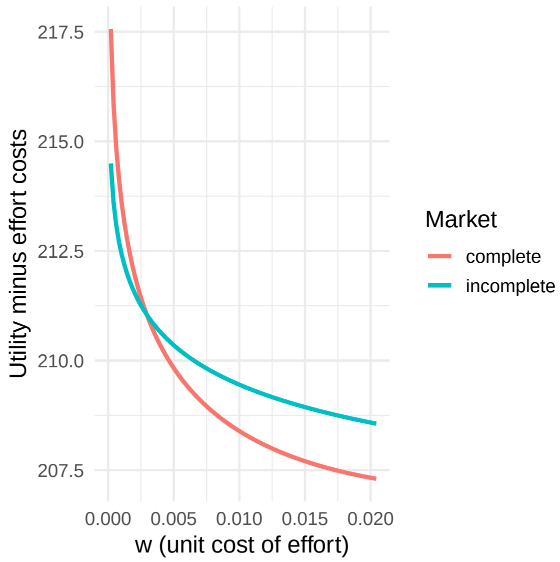

Figure 3: Market completeness and utility under cognitive costs

3 Experimental Design

3.1 Portfolio choice task

Our experiment consists of several variations of a baseline portfolio choice task, both within-subject and between-subject. In each task, the number of states is $S = 16$ and the states are equiprobable. Different portfolio choice tasks vary by the other parameters, like budget, endowments and asset structures.

Each participant makes his portfolio choice using an interface illustrated in Figure 4. The interface lists all the states (called outcomes), state-contingent final wealth levels (called payoffs), the available assets and their prices. The states in which an

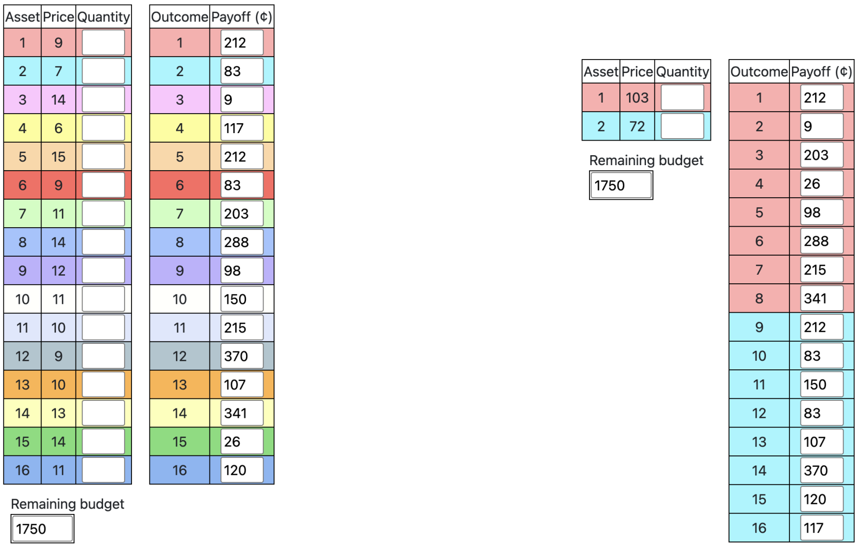

(a) Complete market

(b) Incomplete market

Figure 4: Interface

asset j pays off are displayed in the same color as the asset to help participants easily identify the asset structure. At the start of each task, the payoff column lists the state-endowments, i.e. the state-contingent final wealth participants obtain if they do not engage in any trading. The initial budget is displayed below the asset table.

Participants can choose the quantity of each asset to buy or sell either by typing a number in the corresponding input box or by adjusting the quantity in increments of 0.1 using the up and down arrow keys on their keyboard. As participants make their quantity choices, payoffs and the remaining budget are updated in real time.

Participants can submit the task only if the following conditions are met: 1) all quantities are filled out, 2) no final wealth is below zero, 3) the remaining budget is between 0 and a small value $b > 0$ (specified below). We adopt this modified budget constraint for two reasons. First, although as in the theoretical model, any unspent budget is wasted, it is possible that this kind of mistake could happen in practice. Since we believe this behavior cannot be due to actual cognitive constraints (it can be fixed by simply increasing the purchases of a random asset) and our goal is to focus on mistakes in the portfolio allocation among assets, we require participants to use most of their budget. Second, since asset quantities in the interface can be adjusted only in increments of 0.1, it is not always possible to spend the entire budget and so we allow some slack.

3.2 Rounds and market choice

Each session of our experiment consists of three portfolio tasks, called rounds, as illustrated in Figure 5.

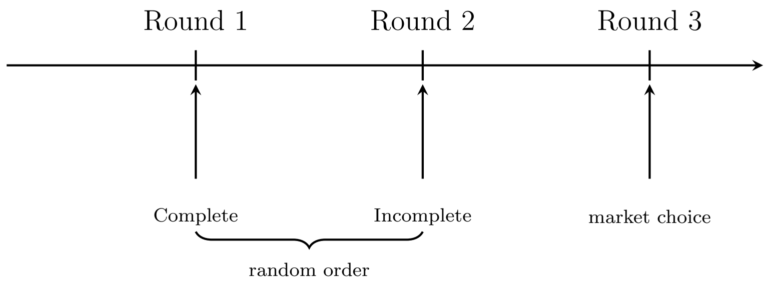

Figure 5: Basic design

The portfolio tasks in the first two rounds have the same budget and state-contingent endowments. The only difference between the first two rounds is that in one the participant faces a complete market and in the other the participant faces an incomplete market, with the order of the two markets randomized at the participant level. The complete market has $J = 16$ Arrow-Debreu securities, while the incomplete market has $J = 2$ securities. Thus the asset structure of the incomplete market induces a partition of the state-space $\{\mathcal{S}_1, \mathcal{S}_2\}$, where each block $\mathcal{S}_j$ contains 8 states. We choose to have two securities in the incomplete market to distinguish it as much as possible from the complete market, while ensuring the participant still faces a meaningful portfolio choice (that would not be the case with $J = 1$). The two markets are linked in the way described in the theory section: the prices of the assets in the incomplete market are obtained by summing up the underlying complete market state-prices within each block of the partition $\{\mathcal{S}_1, \mathcal{S}_2\}$. The assignment of the states to the two blocks was determined randomly (and was the same for all participants).

At the end of each of the first two portfolio choice tasks, we collect the cognitive uncertainty measure of Enke and Graeber (2023) and a subjective measure of complexity (based on the experiment of Gabaix and Graeber, 2024). Specifically, participants were asked to answer the following two questions on a scale from 0 to 100 using a slider: 1) "How certain are you that your chosen budget allocation is the best possible option for you?" and 2) "How complex did you find this task?"

In the third round, participants face another portfolio task. However, they first choose whether it will be a complete market or an incomplete market. This choice is incentive-compatible and thus provides the participant's revealed (weak) preference over market types.[8] At the moment of market choice, the participant are told that the third portfolio choice task will be similar, but not identical, to those encountered before. Specifically, the participant is informed that the number of states and probabilities will remain the same, while prices, budget and initial payoffs will be randomly drawn from the same probability distribution used to generate these parameters in previous rounds (which is further discussed below). This procedure avoids repeating the same portfolio choice problem, which could interfere with the behavioral channels analyzed before, e.g. by reducing uncertainty about one's risk-preferences or reducing the need for cognitive effort. At the same time, this procedure ensures that participants can base their choice on the previous experience with each market type. To facilitate their choice, we also provide participants with a summary of key statistics from their previous two rounds, such as the expected value and range of their wealth profile, their self-reported measures of cognitive uncertainty and complexity and the time spent on each task. No feedback about the realized state in any task is provided before the actual payment, which happens after the end of the experiment.

3.3 Treatments

Our main between-subject treatment design involves varying whether we allow portfolio choices to be driven by the subjects' own risk-preferences or by a preference that we induce in the lab. The goal of the induced-preference treatment is to remove scope for dynamic changes in the agents' risk-preferences that could lead to a preference for incomplete markets, leaving complexity as the residual explanation.

In the homegrown-preferences treatment or *HOM* treatment, the subject's payment in a given task is equal to the final wealth in a randomly selected state (states are equiprobable). In this treatment the participant faces actual risk, as in the original portfolio problem. In this treatment, a preference for the incomplete market may be caused by either dynamic changes in risk-preferences (as in Section 2.2) or complexity (as in Section 2.3).

In the induced-preferences treatment or *IND* treatment, the subject's payment in a given task is given by a deterministic function of the final wealth profile. Specifically,

[8] We chose to not elicit a cardinal measure of the preferences over market types, such as the willingness to pay to switch from an incomplete market to a complete market. This would have involved a multiple-price list and would have made the experiment more complicated and time-consuming for the subjects.

the final payment is equal to:

$$\left[\frac{1}{16}\sum_{s=1}^{16}(x_s)^{0.8}\right]^{1/0.8} \tag{11}$$

where x_s is the final wealth in state s. This payment rule induces CRRA expected utility preferences over profiles of final wealth (with relative risk aversion coefficient $\gamma = 0.2$). Since there is only one way to maximize the deterministic payment in this environment, no changes in the subject's homegrown preferences can lead to regret or temptation. The only remaining cause of a preference for incomplete markets (among those we have identified in the literature) is cognitive costs of complexity.

By comparing the revealed preference for incomplete markets between the *HOM* treatment and the *IND* treatment, we can infer the relative strength of the two explanations put forward in Sections 2.2 and 2.3. If we find that significantly more participants choose the complete market in *IND* compared to *HOM*, this suggests that most of the preference for incomplete markets is driven by anticipated regret or temptation. If instead the rate at which participants choose the complete market is similar across these two treatments, it implies such preference-based explanations are not the main driver, leaving cognitive costs of complexity as a residual mechanism.

Besides our main treatments, we also designed two additional versions of the experiment, called *DEMO* treatments, aimed at lowering complexity costs and providing some robustness for our findings. The only difference between these additional experiments and those described so far is that we let the subjects observe three different example portfolios (demos). One demo is the optimal portfolio for an agent with expected utility preferences and constant relative risk aversion $\gamma = 0.5$ (note that this is different from the utility induced in treatment *IND*). The second demo is the optimal portfolio choice for a maximally risk-averse agent (whose utility is $\min_s x_s$). The third demo is a portfolio randomly selected from the feasible set (the same random portfolio is shown to all subjects in a given task).

Subjects can access the demos through the portfolio choice interface by clicking a button labeled "Click for an example choice." Each click reveals the next example, looping back to the first example after the third. Importantly, participants retain full freedom to adjust asset quantities at any time, regardless of whether or not they have viewed the example choices. We introduce demos in both the homegrown- and induced-preference treatments, leading to two new treatments called *HOM-DEMO* and *IND-DEMO*. To the extent that these demos reduce the cost of cognitive effort, we expect these treatments to decrease effort, losses, complexity, cognitive uncertainty and the proportion of participants choosing the incomplete market.

3.4 Parametrization

In this section, we describe how we set the parameters of the portfolio choice tasks in the first two rounds of the experiment (which constitute the main source of data for our subsequent empirical analysis), in both the *HOM* and *IND* treatment. The parametrization of the third portfolio task is similar and reproduced in the Appendix.

In the *HOM* treatment, we express prices and final wealth in US dollar cents. We set the asset dividend equal to $y = 10$. Asset prices and state endowments were pre-drawn from a random distribution (but kept constant across subjects within the same task). Asset prices in the complete market are integers drawn from a uniform distribution between 6 and 16 (with mean 11 and standard deviation 3.16). State endowments are integers drawn from a uniform distribution between 0 and 400 (with mean 200 and standard deviation 115). The budget was determined after obtaining the prices, to ensure that participants could purchase around 10 units of each asset, so that they could increase their state endowment by a sure payoff of 100 (cents). The budget in each of the first two rounds of the *HOM* treatment was $w_0 = 1750$. The maximum unspent budget allowed was 100 (the opportunity cost of a budget surplus of 100 is equivalent to a sure increase in final wealth of around 5 cents).

Our parametrization has several implications for the distribution of final wealth in the portfolio task (and thus for the subject payments). While we do not know the risk-preferences of the participants, we can compute summary statistics for the distribution of final wealth that results from different portfolio strategies. We collect this information in Table 1. First, consider a risk-neutral subject, who maximizes the expected value of final wealth (strategy MaxEV). This strategy yields an average wealth of 484 cents in the complete market and 288 cents in the incomplete market. Thus, risk-neutral subjects have significant incentives to choose the complete market.[9]

Next, we consider a maximally risk-averse subject who maximizes the minimum wealth (MaxMin strategy). Such a subject would achieve a sure wealth of 266 in

[9] The MaxEV strategy leads also to very high positive skewness in the complete market relatively to the incomplete market, with a maximum wealth of 7750 cents vs. 626 cents. This suggests that the incentives to choose the complete market should be large also for agents with skewness-loving preferences such as prospect theory or rank-dependent utility.

the complete market, and a minimum wealth of only 139 in the incomplete market. Again, utility is significantly higher (almost doubled) in the complete market than in the incomplete market. Finally, we report the outcome of the default portfolio strategy, which we define as purchasing the same quantity of each asset. Although this strategy is not likely to be the optimal strategy for any preference (as it ignores both prices and endowments), it is a heuristic that plays an important role in the complexity model, as discussed above. The default strategy results in a expected wealth of 266 in both markets.

Table 1: Final wealth in the *HOM* treatment portfolio tasks

	Complete market	Incomplete market
MaxEV max wealth	7750	626
MaxEV expected wealth	484	288
MaxEV min wealth	0	0
MaxMin min wealth	266	139
MaxMin expected wealth	266	258
Default max wealth	470	470
Default min wealth	109	109
Default expected wealth	265	265

Note. All values are rounded to the nearest integer. Default is purchasing the same amount of each asset (around 10 units of each asset).

Our parametrization of the first two rounds in the *IND* treatment aims at achieving average subject payments similar to those in the *HOM* treatment, so as to keep incentives comparable across the two treatments. This is achieved partly by choosing an appropriate induced utility function, jointly with the other parameters. In the *IND* treatment, prices, payoffs and induced utility are measured in dollar cents. The asset payoff is equal to $y = 10$. Asset prices in the complete market are integers drawn from a uniform distribution between 1 and 10 (with mean 5.5 and standard deviation 2.87). State endowments are integers drawn from a uniform distribution between 0 and 200 (with mean 100 and standard deviation 58). Again, the budget was determined after obtaining the prices, to ensure that participants could purchase around 10 units of each asset. The budget in each of the first two rounds of the *IND* treatment was $w_0 = 870$. The maximum unspent budget allowed was 50 (with a similar opportunity cost as in the *HOM* treatment).

In the *IND* treatment, there is an objective optimal portfolio strategy. As summarized in Table 2, the optimal portfolio strategy leads to a utility of 412 cents in the complete market, and 288 cents in the incomplete market. These seem significant incentives to choose the complete market, similar in magnitude to those present in the *HOM* treatment. The default portfolio strategy, consisting in buying the same quantity of each asset, leads instead to a utility of 207 cents.

Table 2: Induced utility in the *IND* treatment portfolio tasks

	Complete market	Incomplete market
Optimal	412	288
Default	207	207

Note. These are payoffs computed using the induced utility function. All values are rounded to the nearest integer. Default is purchasing the same amount of each asset (around 10 units of each asset).

3.5 Other details

Our experiment was conducted on Prolific in June 2025. 120 subjects participated in each treatment, for a total of 480 subjects. The experiment lasted 20 minutes on average. After reading the instructions, participants completed a quiz. All participants received a base payment of $4. In addition, a bonus payment was implemented: for one quarter of the participants, a randomly selected task was chosen for payment. The final wealth or utility realized in that task was converted directly in dollar cents and added to the base payment. The experiment was preregistered at AsPredicted, with id 232198.

4 Results

4.1 Reduced form results

Our empirical analysis focuses on the behavior of participants who answered at least three out of four comprehension questions correctly, as pre-registered. This leaves us with around 80% of the original participants: 106 subjects in the *HOM* treatment, 86 subjects in the *IND* treatment, 96 subjects in the *HOM-DEMO* treatment and 96 subjects in the *IND-DEMO* treatment. Moreover, we restrict attention to the first two rounds of the experiment, with the exception of data on market choices, which are made at the beginning of the third round.

We begin by examining the subjects' revealed preference for complete markets. Table 3 reports the proportion of participants who choose the complete market for their third portfolio task. This proportion ranges from 23% to 32%. Thus, a majority of participants prefer the incomplete market, contrary to the predictions of the rational model put forward in Section 2.1. For each treatment, we test whether the proportion is statistically different from 50% and we reject the null hypothesis in each case.

Table 3: Proportion choosing complete market across treatments

Treatment	**% Choice = Complete**	**p-value**
Homegrown	0.2830	0.0000
Induced	0.2326	0.0000
Homegrown_Demo	0.3125	0.0003
Induced_Demo	0.3229	0.0007

Note: P-values are from two-sided binomial tests assessing whether the proportion choosing market 16 is significantly different from 50%. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Next, we examine the differences between market choices across treatments through a number of two-sample proportion tests, reported in Table 4. First, we look at the difference between the *HOM* treatment and *IND* treatment. The difference is not statistically significant. This suggests that the dynamic changes in risk-preferences underlying preference-based explanations have a limited role. In fact, participants seem to be more likely to choose the complete market in *IND* treatment (28%) than in the *HOM* treatment (23%). Similarly, the proportion of subjects choosing the complete market is essentially indistinguishable between treatments *IND-DEMO* and *HOM-DEMO*.

Second, we look at the effect of the demo treatments. Overall the demo treatments seem to increase the preference for the complete market as predicted, but the change is not significant. The change is close to the 10% significance level when comparing *IND* and *IND-DEMO*, or when comparing the pooled baseline data with the data pooled from the two demo treatments.

So far, the evidence suggests that: 1) there is a strong preference for incomplete markets, 2) dynamic changes in preferences like those assumed in models of regret and temptation are not the main drivers of this preference. Although the demo treatments seem to have failed to significantly change market choices, this could be due to the ineffectiveness of demos in lowering cognitive costs. Thus, complexity remains a plausible mechanism for the observed preference for incomplete markets. Next, we examine whether data from the portfolio choice tasks is consistent with this explanation.

The main prediction of the complexity model (under the assumption that the

Table 4: Comparison of proportion choosing complete market across treatments

Treatment 1	Treatment 2	% Complete 1	% Complete 2	p-value
Homegrown	Induced	0.2830	0.2326	0.7346
Homegrown	Homegrown_Demo	0.2830	0.3125	0.3807
Induced	Induced_Demo	0.2326	0.3229	0.1171
Homegrown_Demo	Induced_Demo	0.3125	0.3229	0.5000
Baseline	Demo	0.2604	0.3177	0.1301

Note: P-values are from one-sided proportional tests assessing whether the proportion choosing market 16 is significantly greater in Treatment 2 than in Treatment 1. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

complete market is more complex) is that mistakes will be larger in the complete market than in the incomplete market. We thus need a measure of mistakes. In the *IND* treatment we can easily quantify the deviation from optimality in the utility space as:

$$loss_{IND} = U(\mathbf{x}^{*}) - U(\tilde{\mathbf{x}}) \tag{12}$$

where $\tilde{\mathbf{x}}$ is the final wealth profile resulting from the actual portfolio choice of the subject and $\mathbf{x}^{*}$ is the optimal wealth profile for the induced utility function $U(\cdot)$. In the *HOM* treatment, however, we do not know the utility of the subjects nor the optimal choice. We thus resort to a measure of first-order stochastic dominance efficiency loss proposed by Kuosmanen (2004) for generic portfolio problems (see also Kopa and Post (2009)). This measure is given by:

$$loss_{HOM} = E(\hat{\mathbf{x}}) - E(\tilde{\mathbf{x}}) \tag{13}$$

where $\hat{\mathbf{x}}$ is the wealth profile that maximizes the expected value of final wealth among the feasible profiles that first order stochastic dominate $\tilde{\mathbf{x}}$. This measure can be computed solving a mixed integer linear programming problem reproduced in the Appendix. $loss_{HOM}$ represents the maximum increase in mean wealth over $\tilde{\mathbf{x}}$ obtainable without aggravating the risk exposure of the portfolio, where risk is measured in terms of first-order stochastic dominance. This is a lower bound on the utility loss from mistakes (expressed in monetary terms).[10] Clearly, our loss measures for the *HOM* and *IND* treatments are not comparable. However this is not a problem because we

[10]The actual utility loss may be larger that this in some cases. For example, in the complete market, full insurance is not first order stochastically dominated by any other feasible allocation, and so it results in $loss_{HOM} = 0$, but it leads to a utility loss of 218 if the subject is risk-neutral, see Table 1.

are interested in comparing losses between markets within each treatment.

Together with our loss measure, we analyze how market completeness affects effort, cognitive uncertainty and subjective complexity. These variables are all motivated by our theoretical framework. The complexity model of section 2.3 predicts higher complexity will lead to higher effort under some parametric restrictions. Moreover, in that model, higher effort contributes to lower net utility and can thus explain a preference for the less complex environment. We measure effort as the time spent on a portfolio task before it is submitted, and we refer to it as response time for short.

Cognitive uncertainty measures the participants' subjective uncertainty over their ex ante utility-maximizing decision. Both preference-based models and the complexity-model suggest cognitive uncertainty should be higher in the complete market than in the incomplete market. In the preference-based models of section 2.2, agents are uncertain about their risk-preferences and when they face a larger set of options, the ex-ante optimal choice is less likely to be optimal ex-post. However, according to this model, there should be no residual uncertainty about the optimal choice in the *IND* treatment. On the other hand, in the complexity model of section 2.3 cognitive uncertainty arises from imprecision in finding an optimal choice, rather than true uncertainty about preferences. This model predicts higher complexity will lead to a lower precision of cognitive cues and thus higher cognitive uncertainty, independently of whether preferences are known or not. We measure cognitive uncertainty as 100 minus the value chosen by participants to indicate how certain they were about the optimality of their portfolio choice.

Finally, the subjective measure of complexity we collect in the experiment can provide an observable counterpart to the objective complexity index defined in the model of section 2.3. A similar subjective measure of complexity has been shown to correlate with model-implied complexity by Gabaix and Graeber (2024), although in a different experimental setting.

Table 5 reports the medians of our round-level variables by market type and treatment. First, we find positive losses on average in almost all treatments and markets, with the exception of the *HOM*-treatment incomplete market. More importantly, we find that losses are systematically larger in the complete market than in the incomplete market. In the *IND* and *IND-DEMO* treatments, the complete market leads to losses more than twice as large as in the incomplete market. Even in the *HOM* and *HOM-DEMO* treatments, where our loss measure provides only a lower bound, the difference is significant: while the median portfolio seems close to first-order stochastic

Table 5: Median values of loss, response time, complexity and uncertainty.

Variable	**Incomplete**	**Complete**	**p-value**
	HOM		
loss	0	52	0.0000***
response time	42	159	0.0000***
complexity	53	66	0.0000***
uncertainty	28	28	0.0524*
	IND		
loss	81	205	0.0000***
response time	59	155	0.0000***
complexity	55	66	0.0000***
uncertainty	21	29	0.0003***
	HOM-DEMO		
loss	2	43	0.0000***
response time	34	97	0.0000***
complexity	40	59	0.0000***
uncertainty	23	28	0.0007***
	IND-DEMO		
loss	81	202	0.0000***
response time	50	101	0.0000***
complexity	50	57	0.0206**
uncertainty	19	18	0.2520

Note: Median of each variable is reported (rounded to the nearest integer). P-values from Wilcoxon signed-rank tests comparing num_asset = 2 and num_asset = 16 within participants. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

dominance in the incomplete markets (loss close to zero), participants incur a loss of around 50 in the complete markets (meaning they could switch to a portfolio that guarantees them 50 more cents in expectation without worsening the payment distribution). These patterns support the predictions of the cognitive costs model when the complete market is more complex than the incomplete market (see section 2.3).[11]

Second, we find that response time is also systematically higher in the complete markets than in the incomplete markets. Without demos, subjects spend around 155-159 second in the complete market task, but only 42-59 seconds in the incomplete market task. Interestingly, the response time is slightly lower in our *DEMO* treatments. However, even when subjects can utilize the demos, they spend significantly more time in the complete market task than in the incomplete market task. Interpreting response

[11] While these are absolute losses, the model's predictions are formulated in terms of relative losses in section 2.3. Our results are robust to measuring losses in relative terms, as shown below in section 4.2.

time as a proxy for effort, these patterns are consistent with the cognitive costs model, which under some parametric restrictions predicts effort increases with complexity .

Third, subjective complexity is significantly higher in the complete market than the incomplete market across treatments. This validates our standing assumption that the complete market is more complex. The cognitive costs model can also generate a higher complexity index for the complete market in some parameterizations.

Finally, we find that cognitive uncertainty is significantly higher in the complete market for treatments *IND* and *HOM-DEMO*. We also find that cognitive uncertainty is roughly similar across our homegrown-preferences and induced-preferences treatments, suggesting that instability in risk-preferences is not the main cause.

On average, participants put more effort in the complete market task than the incomplete market task, but end up making worse decisions and with stronger feelings of complexity. These facts and the overwhelming preference for incomplete markets we find in our data are jointly consistent with the cognitive costs model.

4.2 Structural estimation

The previous section has documented a number of facts that are qualitatively consistent with the cognitive costs model of section 2.3. To provide a better test of whether such model can quantitatively explain the data, in this section we structurally estimate the model. Structural estimation is also useful for deriving implications from the current experiment for other settings: once deep parameters (such as micro-complexities and cognitive costs) are estimated, they can potentially be used to calibrate models of financial markets and general equilibrium.

Our structural approach focuses on estimating the parameters that are directly related to the production of cognitive effort. However, solving the model requires knowledge of a number of auxiliary parameters including the agent's risk-preferences. To facilitate the estimation, we exploit the fact that we know the utility function in the *IND* treatment, and thus we restrict attention to this subset of our data. Another set of auxiliary parameters are the budget, asset prices and endowments. We calibrate these values to the actual parameters used in the experiment. Similarly, the variance of prices and endowments is set equal to the value we used to generate them in the lab (see our earlier section on implementation). Finally the model also requires specifying a default. We assume that at the default all securities have the same price and all states have the same endowment. In this setting, given the CRRA utility we induce in the lab, the default action is to purchase the same quantity of each asset.

The model described in section 2.3 provides predictions about effort and losses that we can try to match to our data. The model also allows computation of utility net of the costs of cognitive effort. While we do not observe utility directly, we observe revealed preferences over complete and incomplete markets. To use revealed preference data, we augment the model with a logit decision rule that maps net utility into stochastic market choices. We assume that the probability with which an individual chooses the complete market is:

$$\frac{1}{1+e^{\lambda\times(\mathcal{W}_{incomplete}-\mathcal{W}_{complete})}} \tag{14}$$

where the $\mathcal{W}$ terms are net utilities given by equation (10) and λ measures the responsiveness of choice to utility. Further we assume that all dimensions of the problem that require attention, namely prices and endowments, have the same micro-complexity $c_i = c\forall i$ (though we allow them to have different importance weights). Thus, the parameters of the model are $\Theta = \{\alpha, c, \lambda, w\}$, where α is the curvature of the cognitive production function, c is micro-complexity, λ is the sensitivity of market choice to net utility and w is the unit cost of effort.

We estimate the parameters Θ by the generalized method of moments. This technique selects parameter values that minimize the distance between the moments produced by the model ($\Psi^M(\Theta)$) and their corresponding empirical values (Ψ^E):

$$\hat{\Theta} = \arg\min_{\Theta} \left[\Psi^E - \Psi^M(\Theta)\right]' \hat{W} \left[\Psi^E - \Psi^M(\Theta)\right] \tag{15}$$

in which $\hat{W}$ is a positive definite matrix that converges in probability to a deterministic positive definite matrix W. We set the weighting matrix $\hat{W}$ equal to the inverse of the diagonal elements of the variance-covariance matrix of the data moments.

To estimate the parameters of the model, we use five moments. The first two moments are the average effort in each of the two markets, $L_{complete}, L_{incomplete}$. The empirical counterparts are the average response times in those two tasks (measured in seconds). The second two moments are the expected value of the relative losses $\Lambda_{complete}, \Lambda_{incomplete}$. The empirical counterparts are average of relative losses, computed using the formula $\Lambda(L) = \frac{U^r-U}{U^r-U(0)}$, where U is the actual utility realized in the task, U^r is the maximum utility attainable in the task and $U(0)$ is the utility at the default. The last moment is the expected value of the market choice, measured empirically as the proportion of subjects who choose the complete market.

Table 6: GMM estimates

	α	c	λ	w
Estimate	1.0216	4.9970	1.0055	0.0102
Std. Error	0.0015	0.3252	0.0854	0.0006

Table 7: Moments and moment conditions.

	L_{complete}	$L_{\text{incomplete}}$	$\Lambda_{\text{complete}}$	$\Lambda_{\text{incomplete}}$	$Pr(\text{choose complete})$
Empirical	230.02	90.72	0.96	0.92	0.23
Theoretical	181.48	125.22	0.98	0.95	0.25
p-value	0.03	1.00	0.96	0.85	0.68

Tables 6 and 7 show the estimation results. All parameters are statistically significant. The model fits the data fairly well, as shown in Table 7. With the exception of effort in the complete market, we cannot reject that the predicted moments are statistically indistinguishable from the empirical moments. Although the model predicts a smaller difference in effort levels across the two markets than in reality, it correctly predicts the complete market will lead to a higher effort level and higher losses than the incomplete market. It also matches the high loss levels and the preference for the incomplete market observed in the experiment.

Using the estimated model, we can make a number of interesting observations. First, the model provides us with a numerical measure of complexity: $C_{\text{complete}} = 45, C_{\text{incomplete}} = 9$. Thus, the complete market is around five times more complex than the incomplete market. The increase in complexity is larger than the increase in the number of prices and endowments, which raises from 18 in the incomplete market to 32 in the complete market. Thus, our example cautions against proxying complexity by the number of dimensions that the agent must process.

Second, the model also allows us to estimate the effort paid to process different dimensions, $\{L_i\}_{i=1}^{N}$. We find that the agent pays positive effort to thinking about prices in both markets. The precision m_i with which each single price is processed, i.e. the weight placed on the actual price vs. the default price, is around 5% in the incomplete market and falls to 2.5% in the complete market. More importantly, in both markets the model predicts the agent devotes no effort to thinking about endowments, treating the portfolio problem as one where endowments are constant. This feature helps explaining the non-linearity of complexity in the total number of dimensions: when all the agent focuses on are prices, moving from 2 prices to 16 prices represents a dramatic increase in complexity.

Finally, since induced utility is measured in dollar cents in our experiment, our estimate of w yields a monetary cost per unit of time spent in thinking about prices and endowments: we find that one hour of effort costs around 36 cents.

5 Conclusion

Our study set out to investigate why real-world markets remain incomplete despite the theoretical advantages of complete markets. We developed a portfolio choice experiment in which participants faced both complete and incomplete markets and were asked to reveal their preferences between them. Only a minority of participants preferred the complete market, contradicting the benchmark rational model. Moreover, by comparing treatments with homegrown and induced preferences, we found no evidence that regret or temptation drive this behavior. Instead, our results point to cognitive costs as the primary mechanism: participants made larger mistakes, spent more time, and reported higher perceived complexity in complete markets. Structural estimation of a complexity model corroborates these findings, indicating that the complete market is several times more complex than the incomplete one. Together, these results provide direct experimental evidence that cognitive limitations can generate a robust demand for market incompleteness. In the remaining of this section we outline several implications of this finding that can be explored in future research.

First, while we provide only indirect evidence that individual cognitive limitations have constrained the evolution of actual markets, it seems possible to develop theoretical arguments for this mechanism. In a model where each security exchange chooses how many assets to offer for trading, competition among exchanges can internalize individual preference for incomplete markets, similarly to how firm competition internalizes the preference of complexity-averse consumers for sticky prices in Gabaix (2025).

Second, even if a complete set of securities is offered by the market at any given time, e.g. consisting of bonds, stocks and options, our results suggest that many investors will refrain from using the entire set of securities to finely allocate wealth across states. Instead, many individual investors are likely to focus only on a set of securities that generate a coarse partition of the state-space (e.g. only bonds or stocks). This implies that markets may be segmented along cognitive costs of different investors groups.

Third, complexity costs may also explain other features of real-world financial markets. Most of real-world securities are not binary state-contingent claims traded in a static market. Although from a theoretical point of view a set of state-contingent securities can be replicated by a different asset structure, involving for example sequential markets for risky assets, people may find one of these structures more complex than the other. Magnani et al. (2022) show that subjects incur greater losses in a portfolio task with Arrow-Debreu securities than in a formally equivalent dynamic portfolio task with a risky asset and a safe asset. While Magnani et al. (2022) did not elicit revealed preferences for different formats or measure complexity as in this paper, the results suggest dynamic trading of stocks and bonds is less complex than trading an equivalent set of state-contingent securities. In turn, this could explain why there are few state-contingent securities in real-world markets.

Finally, our paper has implications not only for the evolution of existing markets but also for the design of new markets, for example the design of prediction markets. There has been a recent interest in designing combinatorial prediction markets that allow participants to bet on a rich set of outcomes, which can be thought of as a joint probability distribution over many related states (e.g. Ledyard et al., 2009). One issue in designing such markets is how to partition the state space into assets. While complete markets are theoretically better at revealing information (as supported by recent experiments like Corgnet et al. (2023), though see Bossaerts et al. (2024) for a counterexample), our experiment suggests that there may be a trade-off between participation (and thus market thickness) and the number of assets. Just like simplicity is important in designing auctions and matching markets (Li, 2024), our paper suggests complexity costs should be taken into account when designing prediction markets.

References

Allen, Franklin and Douglas Gale, "Optimal security design," *The Review of Financial Studies*, 1988, *1* (3), 229–263.

_ **and** _ , "Incomplete markets and incentives to set up an options exchange," *The geneva papers on risk and insurance theory*, 1990, *15*, 17–46.

_ **and** _ , "Arbitrage, short sales, and financial innovation," *Econometrica: Journal of the Econometric Society*, 1991, pp. 1041–1068.

Arrow, Kenneth J, "The role of securities in the optimal allocation of risk-bearing," *The review of economic studies*, 1964, *31* (2), 91–96.

Baltussen, Guido and Gerrit T Post, "Irrational diversification: An examination of individual portfolio choice," *Journal of Financial and Quantitative Analysis*, 2011, *46* (5), 1463–1491.

Besedeš, Tibor, Cary Deck, Sudipta Sarangi, and Mikhael Shor, "Reducing choice overload without reducing choices," *Review of Economics and Statistics*, 2015, *97* (4), 793–802.

Bossaerts, Peter, Elizabeth Bowman, Felix Fattinger, Harvey Huang, Michelle Lee, Carsten Murawski, Anirudh Suthakar, Shireen Tang, and Nitin Yadav, "Resource allocation, computational complexity, and market design," *Journal of Behavioral and Experimental Finance*, 2024, *42*, 100906.

Carvalho, Leandro and Dan Silverman, "Complexity and sophistication," *Journal of political economy microeconomics*, 2024, *2* (1), 43–76.

Chernev, Alexander, Ulf Böckenholt, and Joseph Goodman, "Choice overload: A conceptual review and meta-analysis," *Journal of Consumer Psychology*, 2015, *25* (2), 333–358.

Cochrane, John H, "A simple test of consumption insurance," *Journal of political economy*, 1991, *99* (5), 957–976.

Corgnet, Brice, Cary Deck, Mark DeSantis, Kyle Hampton, and Erik O Kimbrough, "When do security markets aggregate dispersed information?," *Management Science*, 2023, *69* (6), 3697–3729.

Debreu, Gerard, *Theory of value: An axiomatic analysis of economic equilibrium*, Vol. 17, Yale University Press, 1959.

Dekel, Eddie, Barton L Lipman, and Aldo Rustichini, "Representing preferences with a unique subjective state space," *Econometrica*, 2001, *69* (4), 891–934.

Enke, Benjamin and Thomas Graeber, "Cognitive uncertainty," *The Quarterly Journal of Economics*, 2023, *138* (4), 2021–2067.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

_ , "A Theory of Complexity Aversion," Technical Report, Harvard University, Department of Economics; National Bureau of Economic Research (NBER); Centre for Economic Policy Research (CEPR); European Corporate Governance Institute (ECGI) March 2025. Available at SSRN: https://ssrn.com/abstract=5185671 or http://dx.doi.org/10.2139/ssrn.5185671.

_ **and Thomas Graeber**, "The complexity of economic decisions," Technical Report, National Bureau of Economic Research 2024.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and self-control," *Econometrica*, 2001, *69* (6), 1403–1435.

Halevy, Yoram and Guy Mayraz, "Identifying rule-based rationality," *Review of Economics and Statistics*, 2024, *106* (5), 1369–1380.

Hart, Oliver D, "On the optimality of equilibrium when the market structure is incomplete," *Journal of economic theory*, 1975, *11* (3), 418–443.

Iyengar, Sheena S and Emir Kamenica, "Choice proliferation, simplicity seeking, and asset allocation," *Journal of Public Economics*, 2010, *94* (7-8), 530–539.

Kopa, Miloš and Thierry Post, "A portfolio optimality test based on the first-order stochastic dominance criterion," *Journal of Financial and Quantitative Analysis*, 2009, *44* (5), 1103–1124.

Kuosmanen, Timo, "Efficient diversification according to stochastic dominance criteria," *Management Science*, 2004, *50* (10), 1390–1406.

Ledyard, John, Robin Hanson, and Takashi Ishikida, "An experimental test of combinatorial information markets," *Journal of Economic Behavior & Organization*, 2009, *69* (2), 182–189.

Leijonhufvud, Axel, "Towards a not-too-rational macroeconomics," *Southern Economic Journal*, 1993, pp. 1–13.

Li, Shengwu, "Designing simple mechanisms," *Journal of Economic Perspectives*, 2024, *38* (4), 175–192.

Magill, Michael and Martine Quinzii, *Theory of incomplete markets*, Vol. 1, Mit press, 2002.

Magnani, Jacopo, Jean Paul Rabanal, Olga A Rud, and Yabin Wang, "Efficiency of dynamic portfolio choices: An experiment," *The Review of Financial Studies*, 2022, *35* (3), 1279–1309.

Malinvaud, Edmond, "Markets for an exchange economy with individual risks," *Econometrica: Journal of the Econometric Society*, 1973, pp. 383–410.

Marin, Jose M and Rohit Rahi, "Information revelation and market incompleteness," *The Review of Economic Studies*, 2000, *67* (3), 563–579.

Oprea, Ryan, "Complexity and Its Measurement," 2024. Working paper.

Ortoleva, Pietro, "The price of flexibility: Towards a theory of thinking aversion," *Journal of Economic Theory*, 2013, *148* (3), 903–934.

Puri, Indira, "Simplicity and Risk," *Journal of Finance, Forthcoming*, 2024.

Reutskaja, Elena, Axel Lindner, Rosemarie Nagel, Richard A Andersen, and Colin F Camerer, "Choice overload reduces neural signatures of choice set value in dorsal striatum and anterior cingulate cortex," *Nature Human Behaviour*, 2018, *2* (12), 925–935.

_ , Rosemarie Nagel, Colin F Camerer, and Antonio Rangel, "Search dynamics in consumer choice under time pressure: An eye-tracking study," *American Economic Review*, 2011, *101* (2), 900–926.

Ross, Stephen A, "Options and efficiency," *The Quarterly Journal of Economics*, 1976, *90* (1), 75–89.

Sarver, Todd, "Anticipating regret: Why fewer options may be better," *Econometrica*, 2008, *76* (2), 263–305.

Ungeheuer, Michael and Martin Weber, "The perception of dependence, investment decisions, and stock prices," *The Journal of Finance*, 2021, *76* (2), 797–844.