Consumption and Savings with Large Language Model Agents*

Michael R. Douglas† Sergiy Verstyuk‡

First version: 29 July 2024
This version: 27 October 2024

Abstract

In a canonical consumption-savings model with aggregate productivity and individual employment risk, standard fully rational agents evaluate contingencies and make decisions by following prescriptions from economic theory. This paper replaces them with agents whose intelligence and behavior are powered by Large Language Models, richly parameterized neural networks trained on vast amounts of text. The performance of such LLM agents is in many ways more similar to that of capable but imperfect human beings than their perfectly rational theoretical counterparts. They demonstrate reasonable economic behavior coupled with systematic anthropomorphic biases in their decisions and beliefs. The linguistic capabilities of LLMs also permit extensive inter-agent communication. The general approach is additionally validated in a game-theoretic setting.

Keywords: consumption-savings problem, large language model, communication, prospect theory, adaptive expectations, extrapolating expectations, wealth inequality, coordination game

*Acknowledgements: OpenAI has provided research funding. SV was supported by the Center of Mathematical Sciences and Applications at Harvard University.

†Contact address: `mdouglas@cmsa.fas.harvard.edu`.
‡Contact address: `verstyuk@cmsa.fas.harvard.edu`.

Contents

1 Introduction **1**

2 LLMs and their capabilities **5**

2.1 Formal description 5

2.2 Interacting with LLMs 9

2.3 Capabilities and limitations of LLMs 10

2.4 LLMs as models of human behavior 11

2.5 How do LLMs perform tasks? 12

2.6 Communication 14

3 Games **15**

3.1 Model 15

3.2 Simulations 16

3.3 Results 16

4 Consumption-savings: model **17**

5 Consumption-savings: simulations **21**

6 Consumption-savings: results **25**

6.1 LLM agents: choices 25

6.2 LLM agents: example choices 31

6.3 LLM agents: example communication 32

6.4 LLM agents: beliefs 36

7 Discussion **44**

8 Conclusion **51**

A Games: LLM prompts **65**

B Games: LLM completions, without communication **67**

C Games: LLM completions, with communication **71**

D Consumption-savings: LLM prompts **75**

D.1 Consumption-savings: LLM prompts, system prompt 75

D.2 Consumption-savings: LLM prompts, user prompt 79

D.3 Consumption-savings: LLM prompts, communication prompt 80

E Consumption-savings: LLM completions, without communication **82**

E.1 Consumption-savings, without communication: Agents' choices 82

F Consumption-savings: LLM completions, with communication **85**

F.1 Consumption-savings, with communication: Agents' communication I . . . 85

F.2 Consumption-savings, with communication: Agents' communication II . . . 90

G Consumption-savings: LLM agents' beliefs, error analysis **95**

H Data **99**

1 Introduction

A recent breakthrough in natural language processing is the development of large language models (LLMs) built on the transformer neural network (NN) architecture. LLMs are trained on publicly available digitized text and are able to answer a broad variety of questions. Surprisingly, the most powerful LLMs show emergent capabilities beyond linguistic competence, including the ability to solve logical, financial, physical and mathematical word problems. Moreover, in many aspects their responses resemble (are "observationally equivalent" to) those of real human beings. Such, if you will, "behavior" of LLMs make them a natural candidate for use as artificial human agents in simulated models of the economy.

First things first: why do LLMs provide plausible models of human intelligence? In terms of quantitative performance measures, LLMs do not just surpass earlier machine learning models on various tests, but already rival human competitors. From a more fundamental standpoint, there are theoretical arguments as well as empirical evidence that the transformer NN architecture shares many features of the structure and function of the human brain, while their training data may reflect human biases and patterns of thought.

To validate this ambitious approach, we task the LLM-powered agents (a standard OpenAI GPT-4 implementation without additional fine-tuning, but instructed with appropriate prompts) with playing a standard game requiring players to make coordinated moves in order to earn higher payoffs. Generic rational economic agents would be making random choices and collect about half of the possible payoffs. However, in our simulations LLM agents managed to collect much more than that, which confirms their ability to tacitly coordinate their choices and to strive for higher rewards, as well as agrees with experimental results involving human subjects. Moreover, leveraging the linguistic capabilities of LLMs, we tried allowing our LLM players to communicate with each other by exchanging several rounds of messages prior to making their moves. Given this opportunity, they verbally coordinated their choices in advance and improved their performance even further.

Armed with this mechanism, we progress from micro to macroeconomics. We take a canonical economic model of consumption and savings [Krusell and Smith, 1998], and change one of its ingredients: replace standard agents (those approaching the rational, fully informed ideal) with instances of LLMs (possibly more realistic representations of human decision-makers). The economic setup comprises agents who work for a firm producing final goods and who spend their incomes on consumption as well as investment in the firm's capital, and it involves both aggregate economic risks (fluctuating firm's productivity) as well as individual agents' risks (occasional spells of unemployment).

Our prompts to LLM agents explain the economic setup in words, then provide the current and 10 most recent historical values of the relevant variables (individual capital holding, aggregate level capital, employment status, rental and wage rates, etc.), and then ask a battery of questions. Some of these questions request the agent to make the necessary economic decisions – how much to consume and how much to save. Additional questions inquire about the preliminaries behind the agents' choices (say, checking an agent's understanding of facts such as calculation of their current income) and additionally may suggest a reasonable direction of "thought" (e.g., what are the likely values of rental and wage rates next period).

In a parallel exercise, we additionally allow some of the LLM agents to communicate with each other before each round of decisions. This lets them exchange non-public information (such as individual capital holdings) or share their intended course of action (e.g., investment levels) and own assessments of future developments (e.g., next period's distribution of capital holdings across agents). To add realism to this process, we use the data on input-output network of the U.S. economy and – after associating each agent with one of the economic sectors – set the probability of agent pair to establish contact proportionally to the volumes of the corresponding sectors' input-output activity.

The above frameworks are compared to two benchmarks: the canonical model in its original form serving as the baseline; and a practically more feasible and realistic, information-constrained version of the regular baseline in which – similarly to our exercise with LLMs – the history of economic variables available to agents is restricted to 10 periods.[1]

Our simulation experiments with these models revealed reasonable ("rationalizable") behavior by all agent types, but with substantial differences in the agents' choices and economic outcomes. First, economies populated with LLM agents demonstrate higher levels of wealth (i.e., capital-holdings) inequality than their regular baseline counterparts.

[1]While not the main focus of this study, given the results shown by the latter (such as the empirically plausible level of endogenously generated wealth inequality), it may be of independent interest to economists.

If we want to match the empirically observed levels, this is a step in the right direction.

Second, models with LLM agents exhibit relatively more pronounced booms and busts in aggregate capital and output levels and, consequently, in prices on capital and labor inputs (i.e., rental and wage rates). This is an attractive and empirically realistic model feature. Endogenous amplification and propagation of exogenous shocks are important for understanding the mechanics of business cycles.

Third, economies with LLM agents accumulate more wealth than the benchmark alternatives. The consumption functions we construct highlight the corresponding differences in agents' propensity to save and consume, which mainly boil down to LLM agents consuming more frugally when they are unemployed (though in general the discrepancy in consumption functions between LLM agents and their regular baseline counterparts are not dramatic, which is encouraging). A consumption-savings behavior that deviates from the standard benchmarks is a feature of our framework that may prove useful for studying the issues of marginal propensity to consume and the magnitude of government spending multiplier for different types of households.

Fourth, when LLM agents are able to communicate with each other, they manage to achieve much higher levels of aggregate wealth and output (with only marginal increase in inequality). Statistical tests and analysis of actual communication records reveal a meaningful increase in investments spurred by inter-agent communication.

We also observe substantial differences between models in terms of agents' beliefs about various economic variables. First, we find that after an economic regime change from high-productivity to low-productivity, LLM agents only gradually reduce their belief in the likelihood that the next regime will again be high-productivity. This is even though after a regime change the actual probability falls abruptly, and the agents are informed about the (fixed) transition probabilities. Second, we inspect agents' expectations about different capital-related measures in the next period: (i) the mean level of the aggregate capital, (ii) the median and (iii) dispersion, or midspread (i.e., the difference between 75th and 25th quantile of frequency or probability distribution), of individual capital holdings. Apparently, the subjective expectations of LLM agents place substantially higher weight on the current values of the corresponding variables than what is warranted by the actual, objective time-series evolution. Such backward-looking properties of LLM agents' behavior are surprisingly in line with what had been found for humans. Reliance on past observations, excessive persistence and too gradual adjustment of expectations to new information, as we observe for the regime probabilities, is consistent with the models and evidence behind the concept of adaptive expectations (going as far back as [Fisher and Brown, 1911, Koyck, 1954, Cagan, 1956]). Reliance of expectations on past observations and their excessive extrapolation far into the future, as we observe for expected aggregate capital and median capital holdings across agents, is consistent with the notion of extrapolating expectations (going back to [Metzler, 1941]).

Third, LLM agents' beliefs about dispersion, or midspread, of the probability distribution for next period's level of aggregate capital are much wider than the realized distribution. Thus the agents perceive the evolution of aggregate capital, the economy's main indicator (i.e., a state variable), as more risky than it really is. Such a subjective overestimation of objective tail probabilities is, again, surprisingly in agreement with the behavior of humans as stipulated by, e.g., the well-established prospect theory ([Kahneman and Tversky, 1979], [Tversky and Kahneman, 1992]). Fitting the corresponding probability-weighting function to the data on reported beliefs of our LLM agents produces the parameter estimates that fall into the ranges previously documented in laboratory experiments with human subjects.

Fourth, in a model version when LLM agents are able to communicate, while the volatility of actual outcomes (such as aggregate capital, output or capital and labor prices) is only slightly higher than in the version without communication, the difference is much larger for the case of beliefs. This finding resonates with the recent research on the potent role of narratives and persuasion in economics.

Our framework and its simulation results demonstrate that in economic modeling, LLMs can serve as agents that are plausible representations of humans. They combine a certain degree of rationality as well as systematic belief and decision biases, including departures from rational expectations of [Muth, 1961] and [Lucas, 1972], in a very realistic, "anthropomorphic" way. They can be built into economic models of virtually any scale, with fairly low requirements on users' technical skills. Besides the practically relevant applications mentioned earlier, this machinery can be a helpful contribution to research programs on bounded rationality. It can also be used for the ongoing work on economic models with heterogeneous agents. Lastly, the newly available linguistic capabilities of LLMs open the door to computational models with explicit and versatile communication between agents.

Before proceeding further, we offer some references to basic background literature. This list is far from exhaustive, for more specialized existing works please see references throughout the main text.

This paper builds on canonical consumption-savings models with heterogeneous agents under idiosyncratic [Aiyagari, 1994] and aggregate [Krusell and Smith, 1998] risk. Mathematically and computationally this is already a non-trivial problem (e.g., see [Den, 2010], [Carroll et al., 2018] as well as [Achdou et al., 2022b]). More recent heterogeneous-agent models are even more challenging (see [Auclert et al., 2018], [Auclert, 2019], [Achdou et al., 2022a]; [Kaplan et al., 2018], [Acharya et al., 2023]).

There is already a sizable literature on the applications of LLMs to problems in economic and social sciences. For instance, they have been used for processing financial and macroeconomic news articles [Bybee, 2023], corporate policies [Jha et al., 2024], financial statements [Kim et al., 2024] as well as answering consumer surveys [Brand et al., 2024].

We are focused on LLMs serving not only as a method for information processing, but as economic agents making their own decisions — i.e., not as tools but as subjects. A good early theoretical analysis is [Andreas, 2022] and a general-purpose implementation is [Reed et al., 2022]. Prior work in economic setting includes [Horton, 2023] and [Akata et al., 2023]. In a broader setting, LLMs were applied to experiments with communication and negotiation [FAIR et al., 2022, Hua et al., 2024] as well as social interactions more generally [Park et al., 2023].

Assuming LLMs can be used as models of humans, an important theme in our work is the presence of features in their behavior that are closer to human intelligence (with its behavioral biases and systematic errors, stereotypes) than to its machine counterpart (characterized by full rationality, absence of altruism). These questions have also been asked in, e.g., [Hagendorff et al., 2023]. [Acerbi and Stubbersfield, 2023], [Horton, 2023], [Suri et al., 2023].

2 LLMs and their capabilities

2.1 Formal description

Language models: Our agents are implemented as instances of LLMs. Below we introduce LLMs going beyond the informal intuition but with the minimum of technical details necessary to grasp their mechanics. Some broader and more detailed reviews are [Douglas, 2023, Minaee et al., 2024, Zhao et al., 2023] as well as the upcoming 3rd edition of the textbook [Jurafsky and Martin, 2024].

LLMs belong to a class of autoregressive models for natural language processing (NLP). These statistical models aim to predict the next word in a sequence of related words by learning the corresponding joint probability distribution. This allows sampling sequences of words by using the laws of conditional probability,

$$
\begin{aligned}
&P(W_1W_2\dots W_{m-2}W_{m-1}W_m) = \\
&= P(W_m \mid W_1W_2\dots W_{m-2}W_{m-1})P(W_1W_2\dots W_{m-2}W_{m-1}) = \\
&= P(W_m \mid W_1W_2\dots W_{m-2}W_{m-1})P(W_{m-1}|W_1W_2\dots W_{m-2})P(W_1W_2\dots W_{m-2}) = \\
&= \Pi_{i=1}^{n} P(W_i \mid \Pi_{k=1}^{i-1} W_k),
\end{aligned}
\tag{1}
$$

where W_i, $i \in \{1, \dots, m\}$, is a word, and $W_1W_2\dots W_{m-2}W_{m-1}W_m$ is a string of ordered words (a sentence or a whole chapter).[2] The structure of the conditional probability function $P(W_i \mid \Pi_{k=1}^{i-1} W_k)$ is defined by the model architecture. As an example, one (simplistic) architecture, called the N-gram model, is to specify the conditional probability of the next word as $P(W_i \mid \Pi_{k=1}^{i-1} W_k) \approx P(W_i \mid \Pi_{k=i-N+1}^{i-1} W_k) = P(\Pi_{k=i-N+1}^{i} W_k)/P(\Pi_{k=i-N+1}^{i-1} W_k)$, with the right-hand-side probabilities estimated by counting frequencies of respective strings in a text corpus.

Transformer architecture: LLMs are examples of the transformer NN architecture, whose foundational papers were [Vaswani et al., 2017] and [Devlin et al., 2018]. We will provide a rough recipe for LLM construction; some parts of it are universal for all statistical NLP models, and some are unique to transformers.

Thus, the inputs of LLMs are words, W_i, $i \in \{1, \dots, m\}$. One starts with "tokenization" – each word is broken into a string of tokens w_{ij}, $j \in \{1, \dots, n_i\}$. These can be short words or parts of a longer word such as a prefix, root or suffix. Next is "embedding" – each unique token indexed by $t \in \{1, \dots, \ell\}$ is mapped to a d-dimensional vector $e_t \in \mathbb{R}^d$, for some large dimension $d \sim 1000 - 10000$. The embedding $\boldsymbol{E}(\cdot)$ is supposed to capture the semantic meaning of tokens, by making the distance between tokens in this space (formally defined, *e.g.*, in terms of cosine of the angle between vectors) reflect the semantic

[2]From equation (1) we can immediately deduce the following about autoregressive NLP models. Since the equation is probabilistic — sometimes it produces false statements. Since the equation is ordered sequentially — each next word depends on the previous ones, and a single false statement can sometimes devolve into a long (self-consistent, in the sense of empirical probabilities) sequence of false statements known as "hallucinations".

similarity of respective tokens (which is measured empirically, *e.g.*, as a co-occurrence of tokens in a text corpus). Third, in order to capture the order of tokens, each token is assigned a positional encoding vector of the same dimensionality d, which is a function of the token's position. The embedding and positional encoding vectors are added, and the resulting row-vectors $\boldsymbol{x}_{ij}$ are stacked on top of each other to produce an input data matrix $\boldsymbol{X} \in \mathbb{R}^{n \times d}$.

The following stage, called the (self-) attention mechanism, is the crucial component of the transformer architecture. This is a function that parameterizes for each token how it relates to the other tokens in the input, in a learnable way. Its definition starts by transforming the input matrix $\boldsymbol{X}$ into three objects: $\boldsymbol{X}\boldsymbol{W}_Q =: \boldsymbol{Q}$ for query, $\boldsymbol{X}\boldsymbol{W}_K =: \boldsymbol{K}$ for key, $\boldsymbol{X}\boldsymbol{W}_V =: \boldsymbol{V}$ for value, where $\boldsymbol{W}_Q, \boldsymbol{W}_K, \boldsymbol{W}_V \in \mathbb{R}^{d \times a}$, with $a < d$, are matrices of parameters (weights).[3] These are taken as arguments to the attention function

$$\boldsymbol{A}(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) := \mathsf{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^\intercal}{\sqrt{a}}\right)\boldsymbol{V}, \tag{2}$$

where $\mathsf{softmax}(\boldsymbol{\chi}) := \exp(\boldsymbol{\chi}) / \sum_j \exp(\chi_j)$. We can think of the key as a sequence of tokens, the query as one of the tokens in this sequence, and the value as a property of interest (such as sentiment) of the sequence. The $\mathsf{softmax}(\cdot)$ operator produces $(1 \times a)$-dimensional scores ($n \times a$ in the end), normalized as a probability distribution, expressing how relevant each element of the key is to the query token (how much "attention" it warrants). Then an attention-weighted sum of the elements of the property sought is calculated (equivalently, the property's characteristic vector is proportionally distributed over the query-relevant tokens of the given sequence). This procedure is repeated for all elements of $\boldsymbol{X}$. In practice, usually several attention functions ("heads") are learned in parallel, and their outputs are concatenated as vectors, to be sent to the next stage.

The next stage is to use the output of the attention mechanism as input to a feed-forward neural network $\boldsymbol{F}(\cdot)$ (or FFN), defined as a series of linear transformations each followed by non-linear activation functions such as $\tanh(\cdot)$. The output of this operation is an $n \times d$ matrix. We omit technicalities such as residual connections and sublayer normalization.

This combination of attention and an FFN, called a layer, is repeated many times, each layer feeding its output to the next. State-of-the-art LLM's have ~ 100 layers. The vector

[3] This query-key-value approach comes from database information retrieval systems. Intuitively, asking (querying) the system such as a search engine with particular terms (keys) returns a list of results (values).

output by the final layer for one of the tokens (usually the last one w_{m-1}), call it $v \in \mathbb{R}^d$, is then mapped to the probability distribution $P(w_m \mid w_1 w_2 \dots w_{m-2} w_{m-1})$. Generally this is done by taking its scalar product with each token's embedding vector e_t, that is $v \cdot e_t$, followed by a $\mathsf{softmax}(\cdot)$ operation to produce the vector of token probabilities.

LLM training: The transformer model's parameters include parameterized functions and matrices: the embedding function $\boldsymbol{E}(\cdot)$, the linear transformation matrices $\boldsymbol{W}_{\cdot}$, and the parameters of the FFN $\boldsymbol{F}(\cdot)$. The model aims to predict the next token in a sequence, so its parameters are estimated ("trained") by minimizing (by stochastic gradient descent) the loss function (in NLP tasks usually chosen to be a cross-entropy criterion) for the predictions.

The total number of model parameters is roughly the square of the dimension d of the embeddings multiplied by the number of layers, times ~ 10. While for theoretical exercises fairly small implementations have proved useful, current state-of-the art models are very large. The OpenAI GPT-4 model that we use here has 1.76 trillion parameters. Training such richly parameterized models requires large data sets of inputs and outputs (e.g., Common Crawl corpus contains petabytes of raw web page data, metadata extracts and text extracts), as well as vast computational resources (GPT-3 utilized 3640 petaFLOP/days). This is why they are called large language models.

An important practical aspect of LLMs are the estimation procedures, much more complex than in previous statistical models. The first phase is massively parallelized estimation of model parameters by minimizing a standard loss function (such as cross-entropy) on servers equipped with Graphical Processing Units using large datasets of digitized text, which produces a "pre-trained" model. The second phase involves additional adjustment of parameter estimates that uses another objective function which measures the quality of LLM responses, and this produces a "trained" model. For instance, in GPT-4 this second estimation procedure relies on "reinforcement learning by human feedback" (RLHF); but many other procedures are possible (e.g., sometimes independent measures of the quality of LLM response are available, as in mathematical and logical problems when unambiguously true answers exist, against which responses can be evaluated). A further, third phase of estimation which is sometimes used in practice (e.g., for commercial applications) is "fine tuning": this refers to the process of assembling a specialized text corpora which illustrate the task and using them for additional parameter adjustment.

2.2 Interacting with LLMs

The basic LLM functionality which we will use is somewhat like that provided by ChatGPT, so let us review that before getting into details.

As will be familiar to many readers, ChatGPT is a "chatbot," a system which engages in dialog with its user. Thus, the user types an initial question or "prompt," the chatbot replies, the user can then reply to that, and so on. This interaction is cast into the terms of Eq. (1) by concatenating the various word strings[4] involved: if the alternating sequence of user and chatbot strings are denoted $U_1, C_1, U_2, C_2, \ldots$, then the reply C_m is generated as the continuation of the input string $C_0 U_1 C_1 \ldots U_m$. The string C_0 is a "system prompt" specified by the maker of the chatbot and which influences its subsequent responses. It can be used for many purposes: to provide additional facts, to enhance some qualities of the responses at the expense of others, *etc.*. In our simulations, we use it to state basic ground rules and properties of the economic situation being modeled.

Note that a response generated by the LLM depends on two and only two sources of information. The first is the training corpus and training procedure, which determines the model parameters (also called "weights"). The second is the initial string supplied to the continuation procedure Eq. (1). A standard LLM (including the one we use) has no memory of previous interactions; if we need memory (or "state") we must implement it explicitly by copying some or all of the text from the previous interactions into the input string.

An LLM samples from a probability distribution and thus an input string does not have a unique continuation; multiple invocations of Eq. (1) will generally produce different continuations. This feature can be partially controlled by adjustable "hyperparameters," most importantly the "temperature" T.[5]

A property of the LLMs is that the conditional probabilities in Eq. (1) only depend on a finite number N of the previous tokens. Metaphorically, the system sees the text through a "window" of finite width and the number N is often called the window size. Until recently N was of order a few thousand which was limiting, both for the system prompt and for the state. However it is much larger in current state of the art (SOTA)

[4]The term "message" is often used for a string produced by one agent in a dialog or multilog.

[5]The terminology comes from the Boltzmann distribution of statistical physics. In statistics terms, the word probabilities are estimated by a (generalized) multinomial regression, and this corresponds to a rescaling $C_i \to C_i/T$ of the regression coefficients.

models. In the GPT-4 used in our experiments, $N = 128,000$.

In using a commercial LLM such as GPT-4, these capabilities are provided as an API (application programmer interface). Thus, a simulation program can invoke a remote procedure call which takes as arguments the sequence of messages just described, and returns the continuation. Details can be found at the provider's webpages.

2.3 Capabilities and limitations of LLMs

The seminal discovery of recent years is that such models, trained with large numbers of parameters (comparable to the number of words in the corpus), can capture so much of the structure of language that they can be used to perform tasks which (one would have thought) require understanding the meaning of the texts. These include question answering, creative writing given a prompt, engaging in dialogs, and solving simple mathematical word problems. The (probably nonexistent) reader who has not yet seen examples of this in the media or played with ChatGPT is encouraged to do so.

Of particular interest for economics is the ability of these models to solve mathematical word problems [Lewkowycz et al., 2022]. SOTA LLMs can solve many word problems at a beginning undergraduate level, and there are many works studying their ability to solve simpler problems. One large collection of such problems along with a systematic evaluation protocol is BIG-bench [Srivastava et al., 2022]. These problems are often drawn from everyday experience: relating the prices of goods to that of a basket of goods, simple reasoning about times and rates of growth, deciding how to fairly apportion a shared good, and so on. The last one illustrates the ability of an LLM to reason about situations involving multiple agents. This ability has been demonstrated and extended in many recent works (e.g., [Park et al., 2023]). These are the capabilities which humans use when engaging in economic activities, opening the possibility of using LLMs to model human economic behavior.

An important general capability of LLMs is "in-context learning." This term refers to the ability of an LLM to carry out a novel task (one not explicit in the training corpus) after seeing a very small number of examples of the task. This was used in [Lewkowycz et al., 2022] on math word problem solving. A related technique is "chain of thought prompting," in which one breaks up the sample problem solutions into a sequence of simpler steps [Wei et al., 2022, Zheng et al., 2023]. We used these methods to produce our LLM agents as we discuss in detail below.

Let us finally review a few more general points which have been much discussed in the popular literature, but must be taken into account. First, one should keep in mind that the primary function of an LLM is to sample from a probability distribution over texts, not to "understand" the texts or perform tasks based on this understanding. Thus, even when an LLM can solve a task, in practice it cannot do it reliably; it will sometimes produce false or nonsensical answers, make up ("hallucinate") facts, *etc.* A simulation using LLMs must be prepared to check the outputs for basic consistency and deal with these lapses. The sequential nature of the continuation process also means that if an LLM does make a mistake, it is very difficult for it to correct the mistake; more often it will "run with" the mistaken assumption and not recover.

Related to this, there are many tasks whose solution requires making hypotheses, testing them and if incorrect (or inferior to other choices) discarding them and trying others. The classic example is strategy games such as chess and go, for which "searching the game tree" is a very familiar concept, and many other tasks such as finding proofs of theorems or making complex plans have this nature. Tasks which require more sophisticated inference than the simple continuation procedure of Eq. (1) are often referred to as "reasoning" in the LLM research literature [Plaat et al., 2024], and are a very active subject of (these days mostly proprietary) research. Of course, there are many other human capabilities which SOTA LLMs do not possess. It would be premature to use LLMs to simulate activities requiring these human competences, but the results presented below are evidence that LLMs have the ability to make reasonable economic decisions in the models we study.

2.4 LLMs as models of human behavior

Language is a complex and profoundly human phenomenon [Binder and Smith, 2013]. Nevertheless, its fluent mastery in itself is unlikely to be sufficient for our purposes.

Granting that LLMs can analyze economic problems and make decisions, we next need to ask to what extent their beliefs and decisions will resemble those of humans, making a simulation with LLM agents a useful tool for studying human economic behavior. There are several arguments one can make. First, we can look at structural similarity between human brains and LLMs in the case of neural substrates and computational processes for memory [Krotov and Hopfield, 2021, Whittington et al., 2022][6], learning and memory [Kozachkov et al., 2023], attention and learning [Bricken and Pehlevan, 2021], *etc.* Second, we can note an effective similarity between LLM performance and human behavior in, obviously, language processing [Schrimpf et al., 2021, Caucheteux and King, 2022] as well as abilities to do simple arithmetic [Nogueira et al., 2021] or use tools [Schick et al., 2023] (with a clear need for a systematic approach to evaluation [Hagendorff, 2023]).

Of course, there is room for criticism of the above arguments. From a theoretical standpoint, many of these successes reflect formal linguistic competence, which does not necessarily imply functional linguistic competence (see [Mahowald et al., 2024]). Indeed, in humans formal logic and mathematical reasoning seem to rely on different brain regions from those crucial for language [Amalric and Dehaene, 2019]. Taking an empirical perspective, LLMs are trained on vast corpora of texts, and thus should exhibit the behaviors which are represented by the texts. But do these texts represent real behavior of humans? A substantial part of these texts reflect intentions that might not necessarily coincide with actions, after all often no real payoffs are involved. Moreover, as these are large but still very selective samples (at the very least, in terms of language used), they are likely to contain non-negligible biases. We can only hope that with continued research on these issues, scientists will better understand whether the impressive ability of LLMs to produce reasonable text continuations ("generalization," in the statistical sense of accurate predictions on "unseen" out-of-sample data) translates into accurate predictions of real human decisions and their actual beliefs.

2.5 How do LLMs perform tasks?

The ability of LLMs to perform sophisticated problem solving tasks came as a major surprise to the AI community, and is still not well understood. Here we briefly survey some of the current approaches to this question.

Let us start with simplest (naive) argument, which is that problem solving is a direct consequence of success at statistical language modeling. When an LLM solves a problem, usually this means it was given the problem statement as an input, and the output continuation includes the solution. Thus, if the corpus contains many examples of valid

[6] Memory is one of the core focuses in neuroscience and psychology, with its prominent role in intelligence [Kahana, 2012] and perhaps even consciousness [Budson et al., 2022].

problem-solution pairs, reproducing its distribution entails the problem solving capability.

While one can criticize this argument on many grounds, the main problem with it is simply that any real world corpus (even one with trillions of words) will only contain a negligible fraction of the possible meaningful texts, so the actual distribution of texts is woefully underdetermined. This is very clear for simple deterministic problems. Consider as an elementary example the addition of two decimal numbers each with d digits. For example, given the input $1\,3\ +\ 2\,7\ =$, produce the output $4\,0$. The number of problem instances is 2×10^d, so already for $d \sim 15$ an exhaustive corpus is not practical. Now the ability of transformers to learn this task has been studied extensively (see for example [McLeish et al., 2024] and references there), and they can do so even for $d \sim 100$ using a relatively small training corpus. This is a sort of generalization analogous to that in earlier statistics and machine learning, but which cannot be explained by current theory. Similarly, when LLMs generate novel texts which did not appear in the training corpus, they are generalizing, but the capabilities and limitations of this generalization are as yet rather mysterious.

The standard phrase for what controls this sort of generalization is "inductive bias" (or "implicit prior") — within a large class of models which fit the training data, the training procedure rewards certain models which (in a successful model) generalize better. Thus one approach to the mystery is to characterize the structure in the training data and relate it to the inductive bias of transformer models. We know a fair amount about the structure of language, especially its compositional structure (both grammatical and semantic). But so far this has been too difficult except for very simple tasks, such as the addition task we described [Zhou et al., 2023], or even simpler function learning [Abbe et al., 2023].

Another approach, in some ways more direct, is referred to as "mechanistic interpretability" [Olsson et al., 2022]. Here one starts from a trained model and tries to directly relate its representations and weights (often called "circuits") to the elements of the task. This includes work on relating the word embeddings to the meanings of the words involved, as well as relations between the intermediate layers and grammatical or other compositional structures. One early and famous example of structure in word embeddings [Mikolov et al., 2013] is the discovery of approximate linear relations such as "$\boldsymbol{E}(\text{king}) - \boldsymbol{E}(\text{man}) + \boldsymbol{E}(\text{woman}) \sim \boldsymbol{E}(\text{queen})$". More recent works detect word features, projections or subvectors of the embeddings with independent meaning [Templeton, 2024], as well as grammatical features [Chi et al., 2020].

An important higher level question is to understand compositionality, the ability to build up complex meanings and capabilities by combining simpler ones. One speaks of a LLM as displaying compositional generalization, which can be tested through benchmarks [Keysers et al., 2020]. A related question is how an LLM selects from among the many individual "skills" it learned in training to solve a new problem, or how it combines different skills when necessary [Arora and Goyal, 2023, Didolkar et al., 2024].

We can see that this study is still in its early days. One of the few solid statements is that, for an LLM to learn to solve a class of problems, it must be possible to represent an algorithm or circuit which computes the solution within a transformer architecture with that number of layers and parameters. While this is necessary but not sufficient (not all algorithms or circuits can be learned), this rules out many possibilities following criteria well established in computer science. As an extreme case, any claim that a future LLM will be able to solve general NP-hard problems, such as finding proofs of arbitrary mathematical statements, can already be ruled out on these grounds. Finding optimal solutions to some standard economics problems is known to be NP-hard [Maymin, 2010].

2.6 Communication

A special feature of the framework constructed here is that LLM agents are able not only to respond to our (i.e., outside actors') queries, but also to send and respond to queries from each other, i.e., to communicate.

Most of the time this is only an implicit feature of economic models. Usually, information is either (a) contained in market prices, or (b) transmitted as parsimonious public or, more often, private signals to corresponding agents. In the former case, communication between agents is redundant (except in some special settings). In the latter case, it is a very restrictive type of communication that follows a formal protocol with a limited alphabet of code words (in information-theoretic sense).

However, in our case the communication between LLM agents uses virtually unlimited message length and broad expressive qualities of human language that is made available by LLMs. That is, we focus not just on abstract signals [Blume and Ortmann, 2007, Kamenica and Gentzkow, 2011, Ely, 2017, Avoyan and Ramos, 2023], but consider arbitrary messages such as verbal conversations or written text (e.g., see the case made by [Shiller, 2017])

Besides making the setup more realistic, which is an interesting experiment in itself, such an approach may reveal some aspects of information exchange between economic agents that have been overlooked in existing models.

Communication between LLM agents has been previously considered in the literature. For example, communication and negotiation are the focus of [FAIR et al., 2022] and [Hua et al., 2024]. In the area of economics, it has been shown that a free-form communication between LLM agents can be beneficial to their performance in the bargaining setting, see [Fu et al., 2023].

3 Games

3.1 Model

To study the performance of LLMs in a very simple setting, we take a well-known problem from game theory.

Game theory traditionally studies optimal choices (or actions, moves) in a setting with multiple agents (or players). There is a large variation of possible settings (cooperative and non-cooperative, full information and partial information, deterministic and stochastic), timing of choices (simultaneous and sequential) and possible agent types (fully rational, boundedly rational). Its theoretical statements are tested empirically with human (experimentally in the field in laboratories, in observational studies) as well as algorithmic agents. See [Fudenberg and Tirole, 1991] and [Nisan et al., 2007] for introductions.

Here, we consider a class of non-cooperative full-information simultaneous games called coordination games. There are $K = 2$ players. Players are offered random realizations of the payoff matrix given in Figure 1. The game is run for $T = 10$ rounds, each with a new realization of the payoff matrix (a game instance). This exercise is repeated 20 times, effectively for $N = 20$ different pairs of players.

In order to get positive payoffs, players need to coordinate: for each opponent's move, only one of your own moves will earn a non-zero outcome. Thus, this game has three Nash equilibria, two pure-strategy and one mixed-strategy (the latter requiring players to follow a randomized strategy mixing different moves with positive probabilities).

Yin \ Zoe	C	D
A	$[\mu, \mu]$	$[\lambda(5-\mu), \lambda(5-\mu)]$
B	$[5-\mu, 5-\mu]$	$[\lambda\mu, \lambda\mu]$

Figure 1: Generic payoff matrix.

[Player names given in top-left corner, separated by the diagonal, with their possible moves given below on the left and above on the right, correspondingly. Parameter values are $\mu \in \{0, 5\}$ and $\lambda \in \{0.6, 1.0, 1.4\}$, drawn independently with equal probabilities.]

3.2 Simulations

In the first setup the simulation is conducted under standard conditions when players are informed about the payoff matrix and make their moves simultaneously. In the second setup players, after being informed about the payoff matrix and before making their moves, are allowed to communicate with each other by sequentially sending and receiving 3 messages each.

Prompts used for interacting with LLMs are available in Appendix §A.

3.3 Results

Standard setup (without communication): In general, LLM agents understand the need for coordination. In the case with asymmetric Nash-equilibria payoffs, coordination is possible by relying on payoffs that are salient due to their higher magnitude or higher frequency of historical occurrence. But for each pair of players it takes them several rounds to learn how to identify such "focal points" associated with the Nash equilibrium most likely targeted by the other player.

The difficulties players are facing are demonstrated in Exhibits 11 to 14 in Appendix §B.

Setup with communication: When agents are allowed to communicate, they are usually able to coordinate on one of the pure-strategy Nash equilibria, although not always on the best one. For a selection of game rounds, all communication messages from both agents as well as responses with their chosen moves are available in Exhibits 15 and 16 in

Table 1: Games, simulation results

Statistic	Model: w/o Communication	Model: w/ Communication
maximal total payoff		
$\mathrm{E}[\chi_n^k]$	57.7	57.7
actual total payoff		
$\mathrm{E}[X_n^k]$	39.7	52.9
$\mathrm{StD}[X_n^k]$	6.6	4.4

Notes: There are $K = 2$ players, $T = 10$ rounds, $N = 20$ pairs. Statistics are calculated over all pairs (and rounds), and hold for an arbitrary player. For a player $k \in \{1, 2\}$ from pair $n \in \{1, \ldots, N\}$, total payoff is defined as $X_n^k := \sum_t X_{tn}^k$ and analogously for χ_n^k. Mean total payoff $\mathrm{E}[X_n^k]$ is equivalent to $\mathrm{E}[X_n^k|k]$ for any player k, and analogously for $\mathrm{E}[\chi_n^k|k]$; while total payoff's standard deviation $\mathrm{StD}[X_n^k]$ is equivalent to $\mathrm{StD}[X_n^k|k]$.

Appendix §C.

Summary statistics: Table 1 shows summary statistics on player performance within a game-theoretic application in a standard setup and allowing for communication.

LLM agents collect much more than half of the maximal possible payoff. These results are broadly consistent with findings in experimental studies with human subjects [Crawford et al., 2008].

Moreover, communication allows LLM agents to successfully coordinate their choices, and agents' performance in the latter setup is significantly better (a two-sample t-test for the difference gives the t-statistic of 7.31, which rejects the null hypothesis of same means with a p-value of 0.00).

From a theoretical standpoint, it is of course not surprising that communication helps coordination, but our exercise verifies that LLMs are capable of doing this in practice.

4 Consumption-savings: model

We focus on the canonical model due to [Krusell and Smith, 1998]. This is a fairly complete general equilibrium setup that features both individual (uninsurable due to incomplete markets) and aggregate risk, and with agents that are homogeneous *ex ante* but heterogeneous *ex post* because of idiosyncratic shocks and decisions. The setup includes agents that can consume goods and accumulate wealth in the form of capital, they can also be employed by a firm. In turn, the firm produces said consumption goods, compensating its employees with wages and paying rents to owners of capital which it uses.

Agents: Agent $i \in \mathcal{A} := \{1, \ldots, N\}$ maximizes expected sum of discounted utilities with respect to a plan for consumption $C_{it} > 0$ and (non-negative) savings that are invested in the capital stock $K_{i,t+1}$:

$$\max_{\{C_{it}, K_{i,t+1}\}_0^\infty} \left\{ \mathrm{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(C_{it}) \right] \right\} \tag{3}$$

subject to the budget constraint

$$C_{it} + K_{i,t+1} - (1-\delta)K_{it} = w_t L_{it} + b(1 - L_{it}) + r_t K_{it}. \tag{4}$$

Labor hours supplied by agent i are denoted as $L_{it} \in \{0, H\}$ (with $H \equiv 1$ WLOG), and they depend on employment status $\epsilon_{it} \in \{0, 1\}$, which is a random variable determined exogenously.

Due to the exogeneity of employment status, the labor choice $L_{it} \equiv L(e_{it})$ drops out from the above optimization problem, so the utility function is defined as (assuming $\gamma = 1$, i.e., a logarithmic utility)

$$U(C, 1-L) \equiv U(C) := \left. \frac{C^{1-\gamma} - 1}{1-\gamma} \right|_{\gamma \to 1} = \ln C. \tag{5}$$

Price variables are capital rental rate r_t and wage rate w_t. Agents take them as given.

Parameters are subjective time discount factor $\beta \in (0, 1)$ and depreciation rate $\delta \in (0, 1)$, as well as unemployment benefit $b \geq 0$.

Firm: A representative firm maximizes current period's profit with respect to capital used $K_t > 0$ and labor hours demanded $L_t > 0$:

$$\max_{\{L_t, K_t\}} \{Y_t - r_t K_t - w_t L_t\}. \tag{6}$$

The production function is defined as

$$Y_t \equiv f(Z_t, K_t, L_t) := Z_t K_t^\alpha L_t^{1-\alpha}. \tag{7}$$

The firm takes market prices r_t and w_t as given.

Capital and labor are aggregated linearly:

$$K_t = \sum_i K_{i,t}, \tag{8}$$

and

$$L_t = \sum_i L_{i,t}. \tag{9}$$

The relevant parameter is the capital share of income $\alpha \in (0, 1)$.

Stochastics: Exogenous state variables are (i) state of the economy $\theta_t \in \{\underline{\theta}, \bar{\theta}\}$ (recession vs. expansion) that follows a first-order Markov process with transition matrix $\Pi := [\pi_{ss'}] = [\pi_{\underline{\theta}\underline{\theta}}, \pi_{\underline{\theta}\bar{\theta}}; \pi_{\underline{\theta}\bar{\theta}}, \pi_{\bar{\theta}\bar{\theta}}]$, where $\pi_{ss'}$ is a probability of a move of the economy from state s to state s' with $\pi_{\underline{\theta}\bar{\theta}} \equiv \pi_{\bar{\theta}\underline{\theta}} := \pi$ being a probability of state change; as well as (ii) individual employment status $\epsilon_{it} \in \{0, 1\}$ that follows a first-order Markov process with transition matrix conditional on state of the economy $\Omega := [\omega_{\epsilon\epsilon'|\theta\theta'}]$, where $\omega_{\epsilon\epsilon'|\theta\theta'}$ is a probability of a move of employment state from ϵ to ϵ' given a move of the state of the economy from θ to θ'.

Parameter π is assumed to equal 0.125; while parameters $\omega_{\epsilon\epsilon'|\theta\theta'}$ are pinned down by the assumptions that the unemployment rate is $(1-\phi) = 0.04$ in expansions and $(1-\psi) = 0.10$ in recessions, as well as that the average duration of unemployment is $\frac{1}{1-\omega_{\bar{\theta}\bar{\theta}00}} = 1.5$ quarters in expansions and $\frac{1}{1-\omega_{\underline{\theta}\underline{\theta}00}} = 2.5$ in recessions, and also by the restrictions on relative probabilities $\frac{\omega_{\bar{\theta}\underline{\theta}00}}{\pi_{\bar{\theta}\underline{\theta}}} = 1.25\frac{\omega_{\underline{\theta}\underline{\theta}00}}{\pi_{\underline{\theta}\underline{\theta}}}$ and $\frac{\omega_{\underline{\theta}\bar{\theta}00}}{\pi_{\underline{\theta}\bar{\theta}}} = 0.75\frac{\omega_{\bar{\theta}\bar{\theta}00}}{\pi_{\bar{\theta}\bar{\theta}}}$.

The productivity level Z_t depends on the state of the economy, $Z_t \equiv Z(\theta_t) \in \{0.99, 1.01\}$, being higher in expansions than in recessions.

Information: Agents have imperfect information. Public information (or common knowledge) includes aggregate variables such as state of the economy θ_t and productivity Z_t, aggregate capital stock K_t and available labor L_t as well as market prices r_t and w_t, all known immediately at time t. Agents also know the formal structure of the economy allowing to infer relationships between variables and forecast their dynamics. Private information (accessible to agents only about themselves) includes individually-indexed variables such as capital holdings K_{it} and employment status ϵ_{it} and L_{it}.

The firm has full information. It knows state of the economy θ_t and productivity Z_t, market prices r_t and w_t, as well as the production function specification $f(\cdot)$.

Equilibrium: The economy's resource constraint is:

$$Y_t \equiv C_t + K_{t+1} - (1-\delta)K_t. \tag{10}$$

The equilibrium is achieved by satisfying the following conditions. Agents $i \in \{1, \cdots, N\}$ and the firm optimize their respective objectives (equation 3) and (equation 6). Budget constraint (4) and resource constraint (10) hold for all time periods t. Capital and labor demands equal supplies so that (8) and (9) hold.

Solution: Firm's optimization gives

$$r_t = \alpha Z_t \left(\frac{K_t}{L_t}\right)^{-(1-\alpha)}, \tag{11}$$

$$w_t = (1-\alpha) Z_t \left(\frac{K_t}{L_t}\right)^{\alpha}. \tag{12}$$

The above pricing functions determine endogenous variables r_t and w_t.

A dynamic-programming formulation of agent i's optimization problem takes a familiar Bellman's equation form:

$$V(K_{it}, \epsilon_{it} \mid \boldsymbol{Q}_t, \theta_t) = \max_{\{C_{it}, K_{i,t+1}\} \in \mathbb{R}_+^2} \left\{ U(C_{it}) + \beta \mathrm{E}_t \left[V(K_{i,t+1}, \epsilon_{i,t+1} \mid \boldsymbol{Q}_{t+1}, \theta_{t+1}) \right] \right\} \tag{13}$$

subject to (4), (11), (12) and stochastic dynamics for the aggregate state variable θ_t and individual state variable ϵ_{it}, as well as stochastic dynamics for the distribution of capital holdings $\{K_{jt}\}_{j=1}^N$ and employment statuses $\{\epsilon_{jt}\}_{j=1}^N$ over all agents in the economy, that is an endogenous state variable $\boldsymbol{Q}_t$, which follows some law of motion $h(\cdot)$ as in $\boldsymbol{Q}_{t+1} \equiv h(\boldsymbol{Q}_t, \theta_t, \theta_{t+1})$. Note that the structure of the economy and assumed restrictions on its stochastic dynamics make the distribution of $\{\epsilon_{jt}\}_{j=1}^N$ to be a redundant information in solving the above Bellman's equation as long as the evolution of θ_t is tracked, hence we ignore the former from now on; i.e., $\boldsymbol{Q}_t$ is treated as a distribution of capital holdings $\{K_{jt}\}_{j=1}^N$ only. The solution seeks for the consumption function $C_{it} \equiv g(K_{it}, \epsilon_{it}; \boldsymbol{Q}_t, \theta_t)$ (or, equivalently, the decision rule for $K_{i,t+1}$).

In practice the solution has to be numerical. The formulation (13) requires dealing with a vector-valued function $h(\cdot)$, which is considered to be the most challenging part of the procedure. For instance, [Krusell and Smith, 1998] relies[7] on replacing the full cross-sectional distribution of capital holdings $\boldsymbol{Q}_t$ by the aggregate capital K_t, and approximating

[7] Appealing to (negligibly) bounded rationality.

the law of motion $h(\cdot)$ by a log-linear autoregressive function of aggregate capital with different parameters depending on the value of θ_t.

LLM agents: Our modeling approach (and the respective solution procedure) rests on treating different agents as independent instances of an LLM. With an appropriately designed prompt, each time period we inform each LLM agent about the general setup, current and past values of aggregate state variables such as state of the economy as well as the agent's individual state variables such as capital holdings. We then ask all agents $i \in \mathcal{A}$ for their decisions about $\{C_{it}, K_{i,t+1}\}$, along the way requiring them to make relevant economic calculations such as current period's individual total income or aggregate capital to labor ratio (in order to check their understanding of the current economic situation) as well as inquiring about their beliefs regarding relevant economic variables such as current cross-sectional distribution of capital or future unemployment rate (to direct their "thinking" to relevant objects).

Communication: As an additional variation of our modeling approach, we allow some of the agents to communicate with each other. Each time period, before making any economic decision, a randomly chosen subset of agents-initiators, $\mathcal{A}_{\text{init}} \subset \mathcal{A}$, exchanges messages with agents-responders randomly chosen, allowing for overlaps, among the remaining agents, $\mathcal{A}_{\text{resp}} \subseteq \mathcal{A} \setminus \mathcal{A}_{\text{init}}$ (where the realized communication network is a stochastic counterpart to the empirical input-output network of the U.S. economy[8]), with the introductory prompt suggesting them to discuss their current and future economic situation and choices. Our general hypothesis is that information exchange will have real effects on economic outcomes; but since the mechanics and performance of LLMs are still not well understood, we are agnostic about the specific effects and view this as an exploratory exercise.

5 Consumption-savings: simulations

Parameterization: We set the structural parameter values to those used in the original model [Krusell and Smith, 1998], with the exception of individual agent's labor endow-

[8] Of course, in the given setup this feature is just a conceptual idea, a caricature if you will. However, it could prove much more relevant in the context of industries' or firms' network, with representative industry executives or firm managers modeled as LLM agents.

ment that we set to 1.0 instead of 0.3271 (to make the setup simpler to describe in an LLM prompt). In particular, this means that we set the unemployment benefit b to 0. Importantly, we additionally restrict the setup by requiring agents to maintain non-negative investments, i.e., to satisfy the equation $K_{i,t+1} - (1 - \delta)K_{it} \geq 0$ in each period t. Furthermore, we use the canonical rational model with [Krusell and Smith, 1998]'s approximations as a baseline model (building on the implementation due to [Carroll et al., 2018]) against which we benchmark our own variations.

For all models, the economy is initialized at expansionary regime, while agents are initialized with equal positive capital holdings and with randomly allocated employment statuses so that the aggregate unemployment rates corresponds to the state of the economy. Time is discrete, each period stands for 1 quarter, and for LLM-based models the simulations and respective time indices start from 12000 in order to prevent LLMs to rely on memorized historical economic events.

The baseline model with rational agents is run with 5000 agents for 11000 periods (dropping the first 1000 periods). Both of our models with LLM agents are run enrolling 50 agents for 600 periods (dropping the first 300 periods).[9] Thus, the baseline model is simulated at its stationary equilibrium for 10000 periods. As to the models with LLM agents, after convergence to their stationary equilibria, we simulate them for 300 periods. Hence, in order to facilitate comparison of results between different models with LLM agents, we generate and save a sequence of shocks $\{\theta_t, \epsilon_{jt}\}_{t=1}^{T}{}_{j=1}^{N}$ that are reused for both models. Additionally, as another benchmark, we run what we call a feasible version of the baseline model, which we describe in more detail below (in §6.4).

LLM setup: The total number of LLM agents is $|\mathcal{A}| = N := 50$.

In a model version with communication, the number of LLM agents that initiate communication is $|\mathcal{A}_{\text{init}}| := 10$; the number of LLM agents that respond to communication is 2 per each agent-initiator (randomly chosen with replacement), giving $|\mathcal{A}_{\text{resp}}| \leq 20$ of unique agents. Thus, about 30 agents out of 50 agents in total are involved in message exchange. By construction, the information diffusion probabilistically follows the empirical pattern of sectoral linkages in the U.S. economy. Formally, at every time period t, (i) $\mathcal{A}_{\text{init}}$ are chosen with probability $\Pr(i \mid i \in \mathcal{A}_{\text{init}}) \propto \boldsymbol{Y}[i]$, where $\boldsymbol{Y}[i]$ denotes an i-th element of the

[9] The short time span is due to the costs of LLM implementation, we plan to expand it in later versions of the paper.

total outputs vector $\boldsymbol{Y}$ from U.S. Bureau of Economic Analysis Input-Output Accounts; (ii) all agents-initiators $\mathcal{A}_{\text{init}}$ are chosen either to be communicating with the suppliers of their inputs or with the customers of their outputs with equal probability; (iii) for each agent-initiator $i \in \mathcal{A}_{\text{init}}$, agents-responders indexed by j are chosen with probability $\Pr(i, j \mid i \in \mathcal{A}_{\text{init}}, j \in \mathcal{A} \setminus \mathcal{A}_{\text{init}}) \propto \boldsymbol{X}[i, j]$ (or $\boldsymbol{X}[j, i]$), where $\boldsymbol{X}[i, j]$ denotes a $[i, j]$-th entry of the input-output matrix $\boldsymbol{X}$ from U.S. BEA Input-Output Accounts (see Appendix §H for more details). As a result, a 10×1 matrix $\boldsymbol{A}_{\text{init}}$ and an $\boldsymbol{A}_{\text{init}}$-induced 10×2 matrix $\boldsymbol{A}_{\text{resp}}$ are defined, forming together the model economy's communication subnetwork, $\boldsymbol{A}_{\text{comm}}$.

In a model with communication, to facilitate natural form of communicated messages, we assign each agent a name using 50 most popular male and female names on US Social Security card applications (see §H).

As the LLM model, we use GPT-4 Turbo version `gpt-4-1106-preview` from OpenAI. We keep all settings at default values.

LLM interactions: The general strategy is for each period t to do the following:

(i) read the initially generated state variables $\{\theta_t, \epsilon_{jt}\}$, read the values of individual agents' capital holdings resulting from the previous period's choices, calculate price variables r_t and w_t according to (11 and 12);

(ii) for each agent i, introduce the overall problem, describe this period's environment (values of current and most recent past aggregate variables, values of current and most recent past individual variables);

(iii) [if communication is allowed, ask for messages from and to $|\mathcal{A}_{\text{init}}|$ number of agents-initiators and $|\mathcal{A}_{\text{resp}}|$ of agents-responders, passing the messages back and forth;]

(iv) ask the agent about its choices and beliefs, collect the answers;

(v) aggregate agents' choices, calculate agents' capital holdings for the next period.

Interaction of users (researchers or other agents) with LLMs has two directions: (i) sending a query to the model in the form of a "prompt", and (ii) receiving a response from the model called "completion". There are two types of prompts: a "system prompt", which is a one-way, query-only (or rather statement-only) interaction; and a "user prompt", which is a query followed by response)).

There are several aspects to our prompt design. First, we tried to direct LLM's "thinking" in a theoretically-motivated direction (e.g., asking about the capital-to-labor ratio, about future state variables). Moreover, we structured the prompts by starting from more general and/or preparatory questions and then following with concrete choice questions that are likely to be dependent on the former questions and answers. This follows the spirit of the "chain-of-thought" approach to interacting with LLMs. Second, in the prompts we provide a specific answer format, which tends to improve the quality of LLM answers and makes their processing easier. Third, we tested many variations on each prompt and chose the one with the highest quality responses. The prompts used are given in Appendix §D.[10]

Since LLMs are stateless and do not retain information about previous interactions unless it is repeated in a prompt, prompts must also include a finite length time-series of past values of the necessary economic variables. In our setup, the history of such time-series was chosen to go 10 periods back. Since this constrains the LLM agents relative to agents in the baseline model, for proper benchmarking we also introduce a feasible version of the baseline model that imposes the same 10-period lag on available input data. This feasible model also includes some additional variables, more on which below in §6.4.

Given the large scale of the study and its sequential nature that requires processing each period's responses before moving to the next period, completions are processed programmatically. LLM answers are parsed online (in real time). The resulting model outputs are checked against economic constraints such as positivity of consumption and satisfaction of a budget constraint. Responses that cannot be parsed or violate economic constraints are not accepted and the corresponding queries are repeated. Because of cost considerations, we adopt several tricks: (i) violation of a budget constraint by less than 10% is accepted as a valid answer with budget expenses re-scaled proportionally to correct income level; (ii) if an agent i at period t gives more than 10 invalid responses in a row, its further queries are not attempted and its previous period's consumption and investment choices are used again, after appropriate rescaling.[11] More complex responses were not checked for validity

[10]Note that LLMs are fairly robust to language style, format choices, and even typos. For example, a typo in one of our prompts (a hardcoded time period value instead of a variable placeholder) was usually either ignored or automatically corrected (sometimes with the correction followed by a note that this seems to be a typo, so an assumption about the correct value had to be made).

[11]About 30% of the decisions exceeded the limit on invalid responses, and the previous period's values had to be re-used. Just to emphasize, in this case only consumption and investment choices were imputed, all of the other response fields were left empty.

in real time and were parsed off-line.

6 Consumption-savings: results

6.1 LLM agents: choices

Aggregate and individual statistics: Tables 2 and 3 present a collection of statistics that characterize the behavior exhibited by the models considered here.

From the aggregate perspective (Table 2), first we notice that LLM agents accumulate more wealth (represented in these models by capital, $\mathrm{E}[K_t]$): on average, 36.269 and 21.540 vs. 38.422 and 64.468 for baseline regular and feasible models as well as LLM models with and without communication, respectively. Higher wealth does not translate into much higher consumption, however. Instead we see only their higher saving and thus investment rates (i.e., ratios of average investment levels $\mathrm{E}[I_t]$ to output $\mathrm{E}[Y_t]$) to maintain that wealth. This is especially striking in a model with communication.

Second, we see that volatilities of economic indicators over time are much higher in economies with LLM agents, as reflected by the corresponding Coefficients of Variation (which accounts for different mean levels of endogenous variables): e.g., 0.015 and 0.012 vs. 0.041 and 0.048 in the case of aggregate capital (i.e., $\mathrm{CV}[K_t]$). Since exogenous variables are the same in all models, more volatile capital results in more volatile output and prices such as rental and wage rates ($\mathrm{CV}[Y_t]$, $\mathrm{CV}[r_t]$, $\mathrm{CV}[w_t]$). That is, economies with LLM agents experience more pronounced booms and busts.

Third, note that in all model versions consumption C_t is relatively less volatile than investments I_t, which is a desirable model feature.

Fourth, we also see that C_t is more correlated with output Y_t in benchmark models than in models with LLM agents, especially in the version with communication. Fifth, in all model versions output exhibits a roughly similar autocorrelation over time.

Lastly, note that the original model of [Krusell and Smith, 1998] by construction focuses on understanding heterogeneity and is not supposed to match plenty of data moments beyond that (for instance, the setting considered does not even feature economic growth). We also are primarily concerned about exploring the underlying mechanisms of our proposed framework. Nevertheless, if we take up some "brave" simplifying assumptions and methodological tricks in order to make empirical comparisons possible, we can see

Table 2: Consumption-savings, simulation results: time-series aspect

	Model			
Statistic	Baseline	Baseline (feasible)	w/o Communication	w/ Communication
$\mathrm{E}[K_t]$	36.269	21.540	38.422	64.468
$\mathrm{E}[Y_t]$	3.641	3.018	3.568	4.300
$\mathrm{E}[C_t]$	2.566	2.341	2.607	2.665
$\mathrm{E}[I_t]$	0.907	0.538	0.961	1.635
$\mathrm{E}[r_t]$	0.034	0.048	0.033	0.024
$\mathrm{E}[w_t]$	2.394	1.984	2.443	2.944
$\mathrm{StD}[K_t]$	0.545	0.249	1.560	3.118
$\mathrm{StD}[Y_t]$	0.047	0.036	0.125	0.161
$\mathrm{StD}[C_t]$	0.060	0.072	0.151	0.152
$\mathrm{StD}[I_t]$	0.057	0.019	0.159	0.207
$\mathrm{StD}[r_t]$	0.001	0.001	0.001	0.001
$\mathrm{StD}[w_t]$	0.012	0.008	0.035	0.051
$\mathrm{CV}[K_t]$	0.015	0.012	0.041	0.048
$\mathrm{CV}[Y_t]$	0.013	0.012	0.035	0.037
$\mathrm{CV}[C_t]$	0.023	0.031	0.058	0.057
$\mathrm{CV}[I_t]$	0.062	0.036	0.166	0.127
$\mathrm{CV}[r_t]$	0.029	0.030	0.037	0.039
$\mathrm{CV}[w_t]$	0.005	0.004	0.014	0.017
$\rho(C_t, Y_t)$	1.000	0.979	0.344	0.130
$\rho(Y_t, Y_{t-4})$	0.538	0.470	0.388	0.462

Notes: Time-series statistics refer to aggregate variables; e.g., $\mathrm{E}[X_t]$ stands for $\mathrm{E}[\sum_i X_{it}]$. $\mathrm{StD}[\cdot]$ denotes standard deviation, $\mathrm{CV}[\cdot]$ denotes the coefficient of variation, ρ denotes Pearson's correlation coefficient. Aggregate variables are given on per capita basis. Number of data points (after dropping the initial pre-convergence periods) is 10000 in each of the baseline models and 300 in each of the LLM models.

Table 3: Consumption-savings, simulation results: cross-sectional aspect

	Model			
Statistic	Baseline	Baseline (feasible)	w/o Communication	w/ Communication
$\mathrm{E}[K_{it}]$	36.269	21.540	38.422	64.468
$\mathrm{StD}[K_{it}]$	4.114	363.852	8.018	16.207
$\mathrm{CV}[K_{it}]$	0.113	16.892	0.209	0.251
$\mathrm{Gini}[K_{it}]$	0.067	0.695	0.115	0.141

Notes: Cross-sectional statistics refer to individual variables, and are averaged over all time periods; e.g., $\mathrm{E}[X_{it}]$ stands for $\mathrm{E}[\mathrm{E}[X_{it}|t]]$. $\mathrm{StD}[\cdot]$ denotes standard deviation, $\mathrm{CV}[\cdot]$ denotes the coefficient of variation, $\mathrm{Gini}[\cdot]$ denotes Gini's coefficient of wealth inequality. Number of data points (after dropping the initial pre-convergence periods) is 10000 in each of the baseline models and 300 in each of the LLM models. Gini's coefficient of wealth inequality for the U.S. in 2022 was 0.830.

that some models are doing slightly better than others when confronted with data for the U.S. (§H). For instance, in all model versions investments are relatively more volatile than consumption, which is also the case empirically. But we find that most of the volatility measures (CV of capital K_t, output Y_t, rent r_t, wage w_t) are closer to real data in models with LLM agents than in their benchmark counterparts. All models have the output autocorrelation that is somewhat smaller than what we see empirically. At the same time, it would be fair to note that comovement between consumption and output is too low for the models with LLM agents.

Looking at individual differences (Table 3), we find that models with LLM agents generate higher cross-sectional dispersion than the regular baseline but lower than the feasible baseline model: i.e., CV of capital is 0.113 and 16.892 vs. 0.209 and 0.251 in the models considered. As a result, Gini's coefficient of wealth inequality is 0.067 in the first baseline model, 0.695 in the second, and 0.115 as well as 0.141 in the latter two models. In other words, all the model versions except the feasible baseline[12] demonstrate levels of inequality that fall short of what is actually observed, but models with LLM agents are faring somewhat better.

[12]Note that the feasible version of the baseline model essentially matches the high levels of inequality observed empirically in the U.S., without the need to introduce ex-ante heterogeneity of agents (such as different subjective discount factors β considered in [Krusell and Smith, 1998]).

Consumption function: Consumption functions in Figure 2 show how income is divided between consumption and savings at different levels of said income. Given that we split the cases for different regimes and different employment statuses while fixing the level of aggregate capital, an agent's income level is completely dependent on his or her capital holdings.[13]

The first, most striking result is that consumption functions of LLM agents are quite similar to those in the regular baseline model, particularly in the case of employed agents. The second important result is that for agents with unemployed status, consumption functions in LLM-based models are strictly dominated by consumption functions in benchmark models; that is, in this contingency LLM agents consume more frugally. The third important result is that in models with LLM agents, the consumption functions in the version with communication lie mostly below such functions in the no-communication version. Fourth, among benchmark models the consumption functions in the feasible version are always (weakly) above those in the regular version. Lastly, note that the slopes of consumption functions are flatter for LLM agents than for their benchmark counterparts (except perhaps in the case of LLM agents with employed status in a model version without communication). All these differences in consumption functions, with the corresponding differences in saving rates, are responsible for the discrepancies in average wealth levels ultimately achieved in the models considered.

To inspect the consumption-savings behavior of LLM agents more rigorously, Table 4 presents statistics on the share of total income devoted to investments by agents of different status with respect to their communication possibilities. Indeed, both during expansions and recessions, communicating agents invest relatively more: e.g., the overall pooled mean is 0.377 for agents who talk with others and 0.359 for those who remain silent (a two-sample t-test for the difference produces the t-statistic of 4.61, which rejects the null hypothesis of same means with a p-value of 0.00). Interestingly, agents on the responding side of conversations usually invest more actively.

In order to get an idea of how LLM agents answer to our prompts, in the next two subsections we give some representative excerpts from LLM completions.

[13] Also, note that the reason why—for a given income—consumption function is usually higher for unemployed agents as opposed to employed ones is because in the former case that income is derived mostly from capital while in the latter case mostly from labor, which means that the unemployed agents in question are wealthier that employed ones.

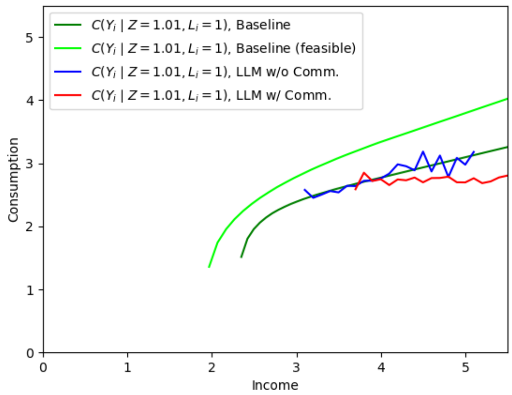

(a) expansionary regime, employed agent

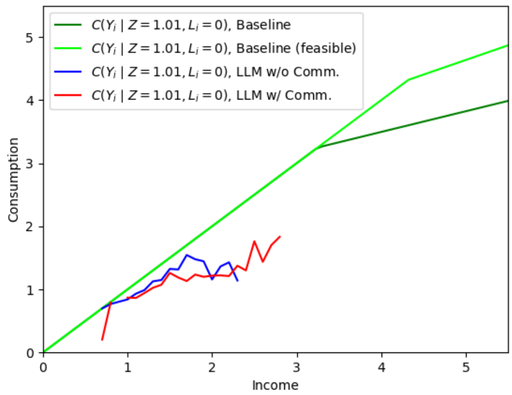

(b) expansionary regime, unemployed agent

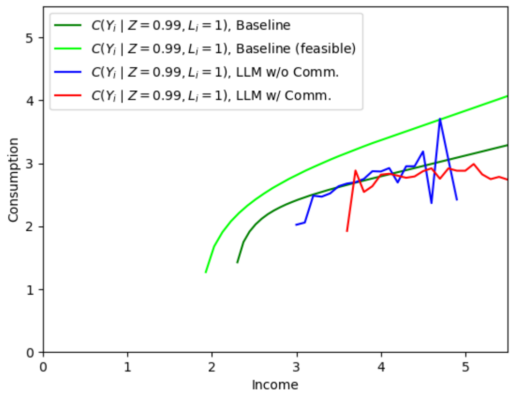

(c) recessionary regime, employed agent

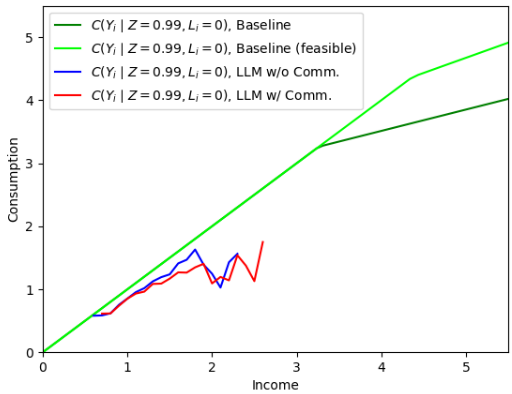

(d) recessionary regime, unemployed agent

Figure 2: Consumption function.
[For the regular and feasible Baseline models, consumption function is defined as $C(K_i, \epsilon_i; K/N, \theta)$, then rebased from K_i to Y_i using $Y_i = r \cdot K_i + w \cdot L(\epsilon_i)$, and computed with endogenously determined $r \equiv r(K, L, Z)$ and $w \equiv w(K, L, Z)$; here, $K/N = \mathrm{E}[K_t]$ is an aggregate capital per capita averaged over the simulation sample, which equals 36 for the regular and 22 for the feasible model. For the LLM-agents' model without communication, consumption function is calculated empirically as $C(Y_i \mid Z(\theta), L(\epsilon_i))$ over the simulation sample using Y_i computed as above, with income bins' averages giving the function's domain; and analogously for the LLM-agents' model with communication. State variables are (in clockwise order): $\{\theta = \bar{\theta}, \epsilon_i = 1\}$; $\{\theta = \bar{\theta}, \epsilon_i = 0\}$; $\{\theta = \underline{\theta}, \epsilon_i = 1\}$; $\{\theta = \underline{\theta}, \epsilon_i = 0\}$.]

Table 4: Consumption-savings, simulation results: LLM agents' investment behavior by communication status

	Investment to income ratio					
	Expansion		Recession		Overall	
Communication status	Mean	StDev	Mean	StDev	Mean	StDev
Communicating	0.394	0.237	0.354	0.236	0.377	0.237
initiating	0.394	0.239	0.344	0.235	0.373	0.238
responding	0.393	0.235	0.360	0.237	0.380	0.237
Silent	0.378	0.240	0.332	0.246	0.359	0.243

Notes: Results are from the model version with communication. Income is calculated as a sum of capital income and wage income (for employed agents): $Y_i = r \cdot K_i + w \cdot L(\epsilon_i)$. Means and standard deviations of investment to income ratio are calculated over the full sample of simulation results conditionally on economic regime and agents' communication status. Number of time-series data points (after dropping the initial pre-convergence periods) is 300. Number of cross-sectional data points: the total number of agents is $|\mathcal{A}| = 50$, the number of agents initiating communication is $|\mathcal{A}_{\text{init}}| = 10$, the number of agents responding to communication is 2 per each agent-initiator (possibly with overlap) making it about $|\mathcal{A}_{\text{resp}}| \approx 20$, the number of silent agents is what remains.

6.2 LLM agents: example choices

Here we consider a model with LLM agents without the possibility of communication, and give excerpts from their responses motivating their economic choices.

Responses answer the following question:

```
Allocation of income between spending on goods and spending on additional
capital this quarter (for T="12487") [in $]:
```

Exhibit 1: Consumption-savings, without communication: Agents' choices (prompt).

Agent $i = 3$ (with $L_{it} = 1$, $K_{it} = 71.92$, $Z_t = 1.01$) responds as follows:

Agent $i = 3$:
```
- As every quarter your spending should equal your income and given the
stability of rental and wage rates, a strategy could be to allocate most of
your income to consumption, while investing a small portion to offset capital
depreciation.
- Assuming capital investment to offset depreciation, you may purchase "$1.798"
worth of capital (71.92*0.025).
```

Exhibit 2: Consumption-savings, without communication: Agents' choices (excerpts).

That is, the agent is relatively rich, employed, and chooses to maintain capital: $K_{i,t+1} = 71.93$. (In Appendix §E.1 we give agent's full choice with explanations.)

Agent $i = 9$ (with $L_{it} = 1$, $K_{it} = 59.22$, $Z_t = 1.01$) responds as follows:

Agent $i = 9$:
```
- Optimal allocation depends on forecasting future rental rates and wage rates.
With no risk and no unemployment benefit, it is optimal to consume all income
since there is no benefit to smoothing consumption.
```

Exhibit 3: Consumption-savings, without communication: Agents' choices (excerpts).

That is, the agent is relatively poor, employed, uses an erroneous argument, and invests 0: $K_{i,t+1} = 57.74$.

Agent $i = 29$ (with $L_{it} = 0$, $K_{it} = 48.13$, $Z_t = 1.01$) responds as follows:

Agent $i = 29$:

```
You could spend all your income on goods, or allocate some for investment. The
optimal allocation would depend on future expectations, which is not detailed
in the information given. Let's assume all income is used to purchase goods
since you receive no unemployment benefits.
```

Exhibit 4: Consumption-savings, without communication: Agents' choices (excerpts).

That is, the agent is relatively poor, unemployed, and prefers to consume the whole income and invest 0: $K_{i,t+1} = 46.92$.

Agent $i = 30$ (with $L_{it} = 0$, $K_{it} = 83.55$, $Z_t = 1.01$) responds as follows:

Agent $i = 30$:

```
Since spending = income and all income this quarter comes from capital, the
optimal consumption of goods would be driven by the maximization of future
utility, which depends on the forecasted rental and wage rates. One should
consume less if they expect higher rental and wage rates in the future,
investing in more capital this quarter, and vice versa. Given no data on
future rental and wage rates are provided, we assumes equal split between
consumption and investment: $1.00084 on goods, $1.00084 on additional
capital.
```

Exhibit 5: Consumption-savings, without communication: Agents' choices (excerpts).

That is, the agent is relatively rich, unemployed, has some reasonable logic but does not do a quantitative analysis and simply chooses to invest 50% of income: $K_{i,t+1} = 82.47$.

As one can see, these responses contain a mixture of reasonable as well as wrong statements, but the resulting conclusions and decisions seem to be made "intuitively" or "heuristically", relying on "rules of thumb".

6.3 LLM agents: example communication

Now we consider a model with LLM agents allowed to communicate, and give excerpts from the communication that precedes them making their economic choices.

Agents $i \in \{44, 14, 13\} = \boldsymbol{A}_{\text{comm},m}$ (with $\{L_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{1, 1, 1\}$, $\{K_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{58.00, 45.28, 69.10\}$, $Z_t = 1.01$) communicate with each other, here is an excerpt from the last message which is supposed to contain the summary of a full conversation between a pair of agents:

Agent $i = 13$ (Susan):

```
Given the stable rental rates experienced over the past quarters and the rental
incomeś reliability, Iḿ considering an income allocation this quarter that
favors capital investment, potentially investing around "60%" to "70%" of my
total income into capital.  This reflects a cautious approach towards possible
transition to a Weak state while also capitalizing on the potential stability
and returns in a Strong state.
```

Exhibit 6: Consumption-savings, with communication: Agents' communication (excerpts).

Afterwards they give responses to the question about choice (given in Exhibit 1):

Agent $i = 13$ (Susan):

```
- Spending on goods = 30% of income
- Spending on additional capital = 70% of income
```

Agent $i = 44$ (Brian):

```
- Based on conversations and forecast stability, let's invest a percentage
(similar to what Susan proposed, about "60%") into capital:  $4.360 * 0.60 =
$2.616 for capital, and the rest for goods:  $4.360 * 0.40 = $1.744
```

Exhibit 7: Consumption-savings, with communication: Agents' choices (excerpts).

That is, this is a diverse collection of agents, all are employed, and some are persuaded to invest 60% of income: $\{K_{i,t+1}\}_{i\in \boldsymbol{A}_{\text{comm},m}} = \{59.17, 47.16, 70.62\}$. (In Appendix §F.1 we give their full conversation.)

Agents $i \in \{5, 19, 0\} = \boldsymbol{A}_{\text{comm},m}$ (with $\{L_{it}\}_{i\in \boldsymbol{A}_{\text{comm},m}} = \{1, 1, 1\}$, $\{K_{it}\}_{i\in \boldsymbol{A}_{\text{comm},m}} = \{63.95, 57.78, 67.62\}$, $Z_t = 1.01$) communicate with each other, here is an excerpt from the last message which is supposed to contain the summary of a full conversation between a pair of agents:

Agent $i = 0$ (James):

```
To summarize, I anticipate a continuation of the strong state next quarter
with productivity likely at "1.01".  IÍl marginally increase my investment
in capital, mindful of the Delta depreciation, and tactically adjust my
consumption.  I plan to monitor employment rate shifts closely and position
my investments to ensure enough flexibility for any economic changes.
```

Exhibit 8: Consumption-savings, with communication: Agents' communication (excerpts).

That is, these agents are well-off, all are employed, and this is an example of an agent forgetting/lying about the intentions (Agent $i = 0$ investments decreased from 2.52 to 2.39): $\{K_{i,t+1}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{64.96, 59.83, 68.23\}$.

Agents $i \in \{28, 16, 37\} = \boldsymbol{A}_{\text{comm},m}$ (with $\{L_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{1, 1, 1\}$, $\{K_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{55.82, 46.83, 74.05\}$, $Z_t = 1.01$) communicate with each other. Afterwards one gives response to the question about choice (given in Exhibit 1):

Agent $i = 28$ (Mark):

```
This requires optimizing utility across periods, considering future value of
capital and income.  Let's denote the amount of income spent on goods as C
(consumption).  Our utility from consumption in this quarter is the log of C.
With Delta = "0.025" and Beta = "0.99", your future utility is discounted,
and future capital income is uncertain.  We should aim to maximize long-term
utility by investing in capital while also ensuring consumption.  Let's assume
that consumption and investment are evenly split for simplicity (since exact
optimization would require solving a Bellman equation that's not possible
here).
```

Exhibit 9: Consumption-savings, with communication: Agents' choices (excerpts).

That is, this is a diverse collection of agents, all are employed, and they demonstrate sophisticated explanation of their choices (Agent $i = 28$ talks about Bellman equation without any mention of it in the prompt): $\{K_{i,t+1}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{56.58, 46.77, 74.45\}$.

Agents $i \in \{38, 37, 31\} = \boldsymbol{A}_{\text{comm},m}$ (with $\{L_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{1, 1, 1\}$, $\{K_{it}\}_{i \in \boldsymbol{A}_{\text{comm},m}} = \{109.18, 74.05, 57.90\}$, $Z_t = 1.01$) communicate with each other, here are excerpts from their messages:

Agent $i = 38$ (Joshua):

I am inclined to use a more quantitative approach to make my investment decisions, specifically, calculating expected values based on the probabilities of the different states of the economy and anticipated changes in the rental and wage rates. This helps in formulating a forward-looking strategy that optimizes utility over time, taking into account the discount factor and potential future changes in employment status.

Agent $i = 37$ (Donna):

In response to unexpected changes in economic indicators, my strategy is to maintain liquidity by having a portion of my assets readily convertible to cash. In case of a downturn, this buffer allows for increased consumption when necessary and provides options for investment when the market shows signs of improvement.

As for the allocation between capital investment and consumption, I currently don't adhere to a strict percentage but rather adjust dynamically based on the economic conditions each quarter. However, the idea of establishing a more defined allocation strategy based on expected utility maximization is appealing, and I plan to integrate some aspects of this into my decision-making process in future quarters.

Agent $i = 38$ (Joshua):

I also appreciate the importance of factoring in the probabilities of transitions, both in the state of the economy and in employment. While the average duration of unemployment during Strong periods is lower, I intend to hedge against potential risks by maintaining a portion of capital investment to ensure rental income continuity in case of unexpected unemployment.

...

I am curious to know if you've considered the implications of aggregate capital on rental rates, as theoretically, an increase in capital could depress rental rates due to higher capital over employment ratios.

Agent $i = 31$ (Ashley):

With regard to your question, I have indeed considered the implications of aggregate capital on rental rates. As the capital stock grows, it may drive the rental rate down due to the increasing capital over employment ratio, as suggested by the economic model we're observing. However, given our relatively stable recent aggregate employment figures and the current trends indicating a safe margin before rental rates are potentially affected, I believe this might not be an immediate concern for T=12488.

Exhibit 10: Consumption-savings, with communication: Agents' communication (excerpts).

That is, these are diverse agents, all are employed, and here we see an example of sophisticated arguments (mentioning a quantitative model, liquidity, hedging, and even correctly interpreting the first-order optimality condition that determines the effect of aggregate capital on rents): $\{K_{i,t+1}\}_{i\in\boldsymbol{A}_{\mathrm{comm},m}} = \{109.26, 74.45, 58.64\}$. (In Appendix §F.2 we give their full conversation.)

We can see that in agents' communication they can persuade or mislead each other, as well as bring up helpful ideas, sometimes fairly sophisticated ones as in the latest Exhibit 10 with the effect of capital abundance on rents.

In our inspection of communication records we noticed that LLM agents never disclose their individual capital holdings, income or even employment statuses; they do not discuss their actual values, speaking only in percentage terms. Agents do this in spite of our suggestions in the prompts, and this is probably influenced by the way humans behaved in their training data.

6.4 LLM agents: beliefs

Beliefs about next period's random state variables (such as the level of aggregate capital K_{t+1}) are important inputs into finding an optimal solution to problem (13). These are non-trivial objects that may be hard to compute for a given agent because of computational difficulty and bounded rationality, limited data history or memory size, etc. In this subsection we will inspect and compare how LLM agents tackle this aspect of the problem as opposed to agents in our benchmark models.

In both of our baseline models, beliefs are endogenously computed variables that exploit the structure of the problem relying on assumption about optimal behavior (e.g., first-order conditions for convex optimization). In both versions of our model with LLM agents we ask agents for their beliefs directly, and do this simultaneously with asking for their preferred choices of control variables C_{it} and $K_{i,t+1}$.[14]

Specifically, in the benchmark model's solution approach due to [Krusell and Smith, 1998],

[14] Some technical details are worth mentioning. From responses of our LLM agents to questions about their beliefs, only about 10-20% are valid answers, with most of the invalid ones being qualitative explanations or no answer at all. We also drop a small number of extreme outliers; normally these would be taken as a reflection of noise in decisions (which is a genuine property of naturally occurring behavior, thus being useful empirical information), but in our experiments with only 50 agents and short samples would be biasing the aggregate results too much, impairing statistical analysis.

beliefs about next period's random variables are determined (i) by exogenously given parameters fixing various probabilities, transition matrices and durations (see §4), as well as (ii) by the equilibrium law of motion for capital that agents learn empirically in the process of economy's evolution. In the first case, agents take the provided parameters as given and rationally evaluate expected values of respective random variables (such as the state of the economy next period). In the second case, the law of motion is approximated by a regime-specific log-linear equation for the (expected value of) aggregate capital

$$\ln K_{t+1} = \eta_{\theta 0} + \eta_{\theta 1} \ln K_t + \varepsilon_{\theta,t+1}, \tag{14}$$

where $\eta_{\theta k}$ for $k \in \{0, 1\}$ are coefficients, θ is the state of the economy and $\varepsilon_{\theta,t+1}$ is a residual. Learning process relies on least-squares estimation. Conditionally on this equation's output, the next period's value function can be evaluated and optimized (see §4). Along with this, next period's expected rental and wage rates can be unambiguously calculated. In this subsection's analysis, estimation of the above coefficients is done on the part of full (post-convergence) sample that does not overlap with the prediction sample, that is belief variables are generated "out-of-sample" one period ahead.

For the feasible—or informationally constrained—version of the baseline model, we introduce several modifications and extensions (below is an exhaustive list). First, the length of time series used for learning is restricted to 10 lags (to match that provided in a prompt for LLM agents). Second, we make parameters independent from state of the economy, $\eta_{\theta k} = \eta_k, \forall \theta$ (due to short learning sample). Third, we impose default values for parameters $\eta_{\theta 0} = \eta_0 = 0$ and $\eta_{\theta 1} = \eta_1 = 1$, so that the resulting values are weighted averages of the actual estimates with the weight 0.09 and the default values with the weight 0.91, which effectively ties the predicted variable to its most recent observation (because of shortened learning sample, this is required to obtain well-behaved beliefs and ensure the convergence of the algorithm numerically approximating the value function $V(\cdot)$ as well as decision/policy functions $C_{it}(\cdot)$ and $K_{i,t+1}(\cdot)$). Fourth, we calculate additional belief measures: (expected values of) the median of the distribution of capital holdings across agents $\mathrm{Med}[K_{i,t+1}]$ (which is a measure of central tendency robust to outliers and skewness) and the interquartile range, or midspread, of capital holdings distribution $\mathrm{IQR}[K_{i,t+1}]$ (which is a measure of dispersion analogous to standard deviation).[15] Fifth, for these

[15]Formally, median is defined as $\mathrm{Med}[X] := \mathrm{F}^{-1}(0.50)$, where $F(\cdot)$ is the CDF of X; the interquartile range (also called midspread) is defined as $\mathrm{IQR}[X] := \mathrm{F}^{-1}(0.75) - \mathrm{F}^{-1}(0.25)$.

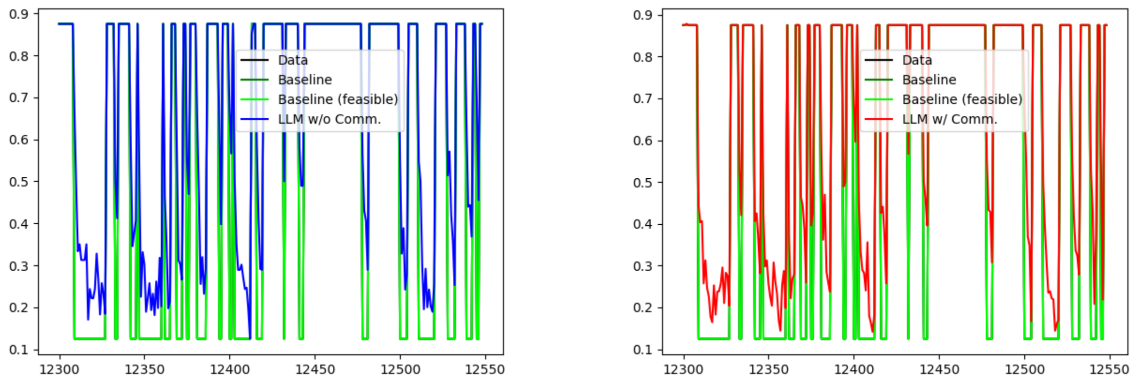

(a) Without Communication

(b) With Communication

Figure 3: Beliefs vs. true probability of expansionary regime next period.
[True probability of state change is $\pi = 0.125$. Therefore, true probability of expansionary regime next period is 0.875 when the current regime is expansionary, and 0.125 when the current regime is recessionary. Data are true probabilities. Beliefs about the probability of expansionary regime next period for baseline and feasible baseline models are true probabilities. Beliefs about the probability of expansionary regime next period for LLM agents are collected in each period, with reported values averaged across all agents with valid answers.]

additional measures we consider a similar law of motion as (14), except that the logarithmic transformation for midspread is dropped. Estimation is done on the sample going 10 periods back, and belief variables are generated one period ahead.

Levels and dynamics: The most transparent example is beliefs about fixed parameters. For instance, true probability of a change in state of the economy, or regime, is $\pi = 0.125$ and the probability of staying in the same state is $(1 - \pi) = 0.875$. Agents in both versions of the baseline model are able to use this information correctly and recognize the non-stochastic nature of this parameter as long as the current state is known, see Figure 3.

LLM agents are also fully informed about the true probabilities as well as about the current state of the economy. Nevertheless, as can be seen from Figure 3, after a change from an expansionary to recessionary regime, when the probability of the next period's regime being an expansionary one abruptly falls from 0.875 to 0.125, it takes LLM agents many periods to gradually adjust their assessment of that probability from 0.875 to 0.125 (both in a regular version and in a version with communication).

A puzzling accompanying result, however, is that the reverse adjustment happens without any delay: after a move from recessionary to expansionary regime, LLM agents immediately recognize that the probability of staying in expansionary regime next period jumps from 0.125 to 0.875. It is unclear why LLMs exhibit such a behavior; our best guess is that it is related to the fact that when our prompts to LLMs mention different regimes, these are always ordered as expansionary regime first and recessionary second[16], which may inadvertently suggest the presence of some asymmetry between them.

Let us now turn to beliefs about random variables. Figures 4 and 5 present the time series of beliefs implied by the baseline model and its feasible alternative, as well as by LLM models in regular formulation and with communication.

Visually, LLM agents' beliefs match next periods' actual realizations of the corresponding variables not much worse than their counterparts in the feasible baseline model. Note that even the magnitude of noise in generated series are usually comparable. One mild exception is expected mean of the aggregate capital, in which case expectations of LLM agents consistently fall short of the actual realizations (but the average magnitude is only about 1% of the target). Formally, the goodness of fit measured in terms of R^2 for most cases is lower for LLMs than for the feasible baseline model by several percentage points, although sometimes the gap is larger: see Tables 9, 10, 11, 12 in Appendix §G.

Conducting a statistical analysis of expectation errors for the variables covered by Figures 3, 4 and 5 suggests that LLM agents' beliefs use available information fairly efficiently. As can be seen from Tables 9, 10, 11 and 12 in Appendix §G, both in the regular version and in the version with communication regressing errors on available to LLM agents variables (or their transformations) are rarely statistically significant under conventional levels of confidence, which means that information contained in them has already been taken into account and can not improve predictions further. But there are some exceptions: e.g., a higher moving average $\mathrm{MA}[\mathrm{Med}[K_{it}]]$ is associated with lower errors $w_{t+1} - \mathrm{E}_t[w_{t+1}]$ when these errors are negative making them even worse, which means that agents' errors could be reduced in absolute magnitude by decreasing the expectation $\mathrm{E}_t[w_{t+1}]$ when $\mathrm{MA}[\mathrm{Med}[K_{it}]]$ is high in the setup without communication.[17]

[16]Strictly speaking, our prompts use the terms "Strong" and "Weak" state of the economy.

[17]In Tables 9, 10, 11, 12, $\mathrm{MA}[X]$ denotes an unweighted moving (rolling) average of X over 10 periods.

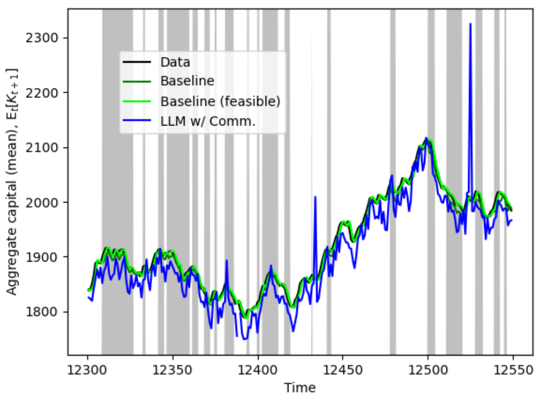

(a) Aggregate capital, mean

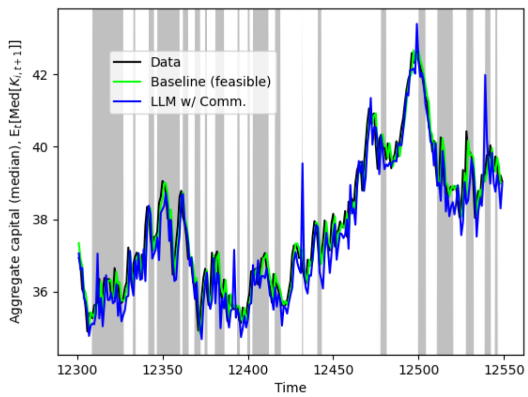

(b) Aggregate capital, median

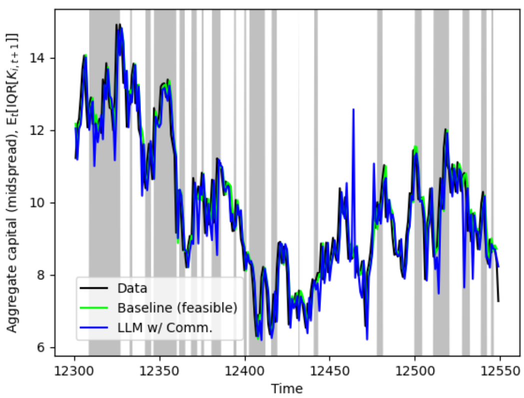

(c) Aggregate capital, midspread

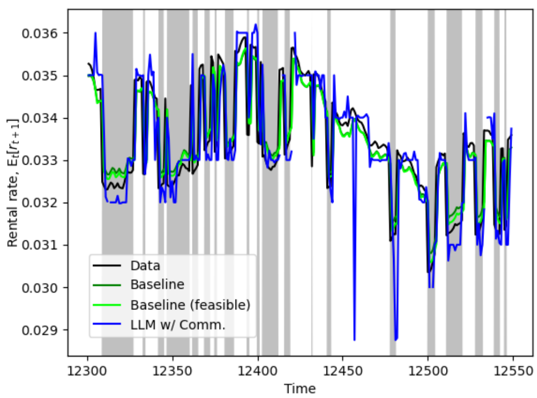

(d) Rental rate

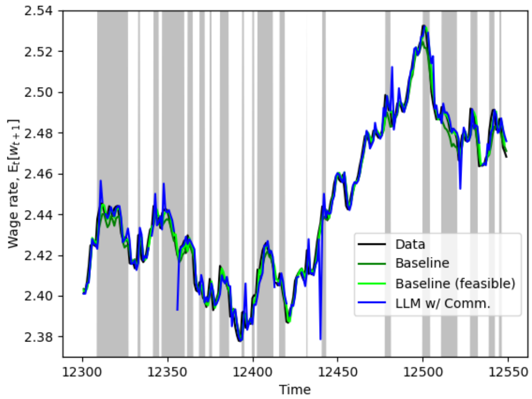

(e) Wage rate

Figure 4: Beliefs vs. actual realizations of economic variables, without Communication.
[Data are actually realized values of the corresponding variables. Beliefs implied by baseline, feasible baseline and LLM models are one-period-ahead predictions of the corresponding variables conditionally on available state variables following the procedure described in the text. Shaded areas indicate recessions.]

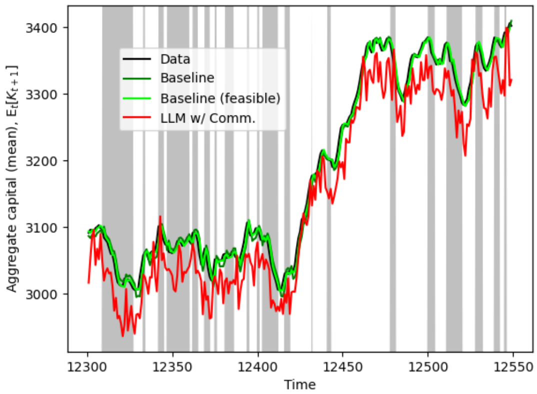

(a) Aggregate capital, mean

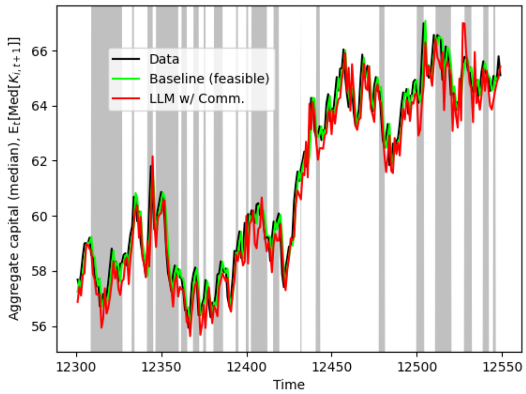

(b) Aggregate capital, median

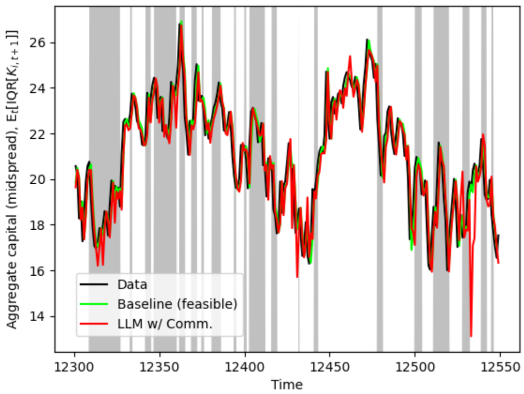

(c) Aggregate capital, midspread

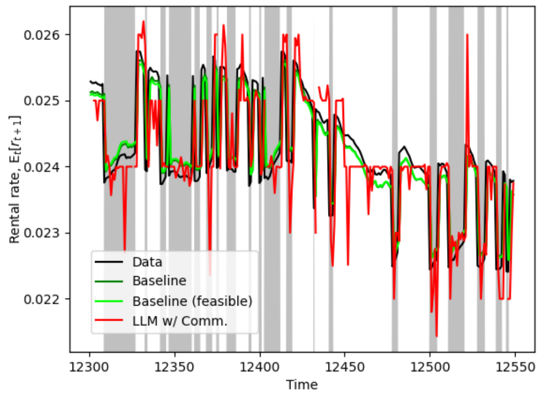

(d) Rental rate

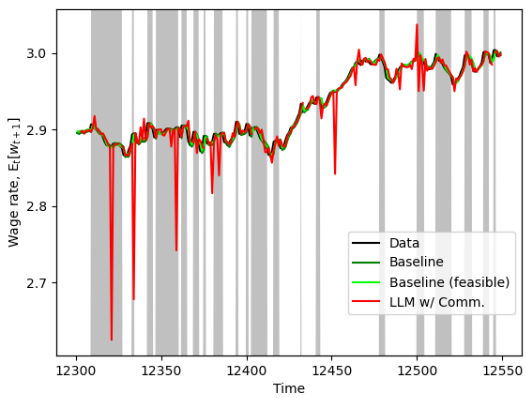

(e) Wage rate

Figure 5: Beliefs vs. actual realizations of economic variables, with Communication.
[Data are actually realized values of the corresponding variables. Beliefs implied by baseline, feasible baseline and LLM models are one-period-ahead predictions of the corresponding variables conditionally on available state variables following the procedure described in the text. Shaded areas indicate recessions.]

Consistency: Next, we study the properties of beliefs and investigate their consistency with the actual variables they refer to. We start with temporal aspects. Table 5 estimates the dependence of next period's expected level of aggregate capital as well as expected median and midspread of individual capital holdings on their currently observed values, and compares the estimates to empirical laws of motion for the actual realizations of these variables.

We see that LLM agents universally put higher weights on the most recent values of the variables they form beliefs about (although because of short samples our estimates are not always sharp enough to formally reject the equality of the corresponding coefficients at conventional confidence levels). The difference is particularly striking for the midspread, in which case LLM agents place an order of magnitude larger weight than what is optimal. As a result, LLM agents' beliefs are less accurate than their counterparts in both versions of the baseline models, as can be seen from goodness-of-fit measures in the top rows of Tables 9, 10, 11, 12 in Appendix §G (of course, as shown in Table 5, R^2 for the fitted equations of expectation determinants themselves may still be high).

Another consistency check focuses on the scale of variables agents deal with: specifically, on cross-sectional heterogeneity and measures of risk. Table 6 presents the calculations on dispersions, or midspreads, of the frequency and probability distributions of capital next period.

On aggregate, LLM agents' beliefs about next period's expected dispersion of capital holdings across agents are well-calibrated relatively to actual realizations both in the regular setup, with mean dispersions of 9.881 in the data vs. 9.785 in agents' beliefs, and in a setup with communication, with 21.055 vs. 20.908 (two-sample t-tests for the difference give the t-statistic of 0.56 with a p-value of 0.58, and the t-statistic of 0.68 with a p-value of 0.50, respectively). This does not seem to be the case for LLM agents' beliefs about the probability distribution for next period's aggregate capital level realizations and specifically its dispersion in either setup, with mean dispersions being 11.216 vs. 16.576 in a version without communication, and 13.735 vs. 36.430 in a version with communication (although the structure of our simulations does not allow us to conduct a formal test in this case, the differences are clearly substantial).

Admittedly, there is a difference between these two results in the nature of questions and inputs: the frequency distribution of capital holdings is directly observable and agents are explicitly informed about it (with a lag), while the probability distribution of next

Table 5: Consumption-savings: reported beliefs about capital, regression results

	w/o Communication				w/ Communication			
	Expansion		Recession		Expansion		Recession	
Specs	Data	Beliefs	Data	Beliefs	Data	Beliefs	Data	Beliefs
K_{t+1}, $\mathrm{E}_t[K_{t+1}]$ on:								
constant	0.0507	-0.3099	0.1148	-0.0272	0.0598	0.1627	0.0858	-0.2210
	(0.0630)	(0.2387)	(0.0814)	(0.1497)	(0.0460)	(0.1290)	(0.0449)	(0.1253)
K_t	0.9935	1.0397	0.9846	1.0018	0.9928	0.9786	0.9891	1.0256
	(0.0083)	(0.0316)	(0.0108)	(0.0198)	(0.0057)	(0.0160)	(0.0056)	(0.0156)
R^2	0.993	0.880	0.991	0.956	0.996	0.968	0.997	0.980
Obs.	142	141	107	107	142	142	107	107
$\mathrm{Med}[K_{i,t+1}]$, $\mathrm{E}_t[\mathrm{Med}[K_{i,t+1}]]$ on:								
constant	0.1009	-0.0728	0.1843	0.0702	0.0902	-0.0286	0.1096	0.0474
	(0.0646)	(0.0593)	(0.0717)	(0.0816)	(0.0655)	(0.0502)	(0.0724)	(0.0421)
$\mathrm{Med}[K_{it}]$	0.9727	1.0189	0.9486	0.9780	0.9785	1.0061	0.9730	0.9863
	(0.0178)	(0.0163)	(0.0197)	(0.0225)	(0.0159)	(0.0123)	(0.0176)	(0.0103)
R^2	0.937	0.953	0.933	0.960	0.964	0.980	0.964	0.988
Obs.	142	142	107	107	142	142	107	107
$\mathrm{IQR}[K_{i,t+1}]$, $\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$ on:								
constant	1.3452	0.1224	1.3592	-0.1814	2.1190	0.2368	2.1028	-0.3854
	(0.0356)	(0.1483)	(0.0371)	(0.1375)	(0.0390)	(0.4415)	(0.0399)	(0.4090)
$\mathrm{IQR}[K_{it}]$	0.0934	0.9811	0.0925	0.9986	0.0439	0.9834	0.0443	1.0083
	(0.0036)	(0.0136)	(0.0033)	(0.0122)	(0.0017)	(0.0193)	(0.0019)	(0.0174)
R^2	0.828	0.935	0.877	0.960	0.828	0.951	0.833	0.918
Obs.	142	142	107	107	142	142	107	107

Notes: The table presents results of OLS estimation of time-series relationships between future and current values of capital-related variables/statistics. The future values are either those actually realized in the simulation sample, or expectations thereof as reported by LLM agents. The variables in question are aggregate capital K_t; median value of capital holdings across all agents $\mathrm{Med}[K_{it}]$; and the interquartile range $\mathrm{IQR}[K_{it}]$. Specifications of regression models are given in the first column. Other columns include coefficient estimates, Newey–West standard errors (in parentheses) as well as R^2 as goodness-of-fit measures and number of observations used. Columns are split into "Data" (based on actual realizations) and "Beliefs" (based on reported expectations), into "Expansion" (sample conditioned on $Z_t = 1$) and "Recession" ($Z_t = 0$), regular version and with communication. In the top (for the mean) and middle (median) panels, logarithms of the variables are used.

Table 6: Consumption-savings: reported beliefs about dispersion

	w/o Communication		w/ Communication	
Midspread for	Data	Beliefs	Data	Beliefs
Frequency distribution of $K_{i,t+1}$	9.881	9.785	21.055	20.908
Probability distribution of K_{t+1}	11.216	16.576	13.735	36.430

Notes: Midspread for cross-sectional frequency distribution of $K_{i,t+1}$ is defined as $\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$ and for outcome probability distribution of K_{t+1} is defined as $\mathrm{E}_t[\mathrm{IQR}[K_{t+1}]]$. Table entries present mean values. Number of observations used is 249. Belief values are obtained from LLM agents' responses to corresponding questions during the simulation process. Data values are taken from realized simulation outcomes. Data values for probability distribution midspreads are calculated using empirical distribution of one-period differences in K_t over the same simulation sample.

period's aggregate capital needs to be inferred and agents are not informed about past probabilities (only about their realizations). Nevertheless, abstracting away from such technical details, we can interpret these results as follows. In both versions of our model with LLM agents, the agents predict the distribution of capital accurately, so they perceive wealth inequality correctly. However, LLM agents perceive the evolution of next period's capital as much more risky than it really is.

7 Discussion

In a canonical consumption-savings setup, the results observed in our simulations with LLM agents are quite different from those in the benchmark alternatives. In this section we try to put these findings in perspective, and relate them to the existing literature.

Agents' choices and aggregate economy's outcomes: Behavior of our simulated economies differs markedly depending on how we model its decision-making agents.

In economies with LLM agents we observed higher average wealth per capita relatively to benchmark models (38.4–64.5 units vs. 36.3–21.5 in two baseline versions) due to agents' higher savings rates. Importantly, the former economies have shown relatively higher wealth inequality than the regular baseline model, although not relatively to the feasible baseline (Gini's coefficient of 11.5%–14.1% vs. 6.7%–69.5%). Thus, in terms of capturing the empirically observed levels of inequality, a framework with LLM agents as well as the feasible version of the baseline model are steps in the right direction (cf. [Krusell and Smith, 1998]). This also suggests that frameworks with LLM agents exhibit nontrivial behavior, and may serve as tools for studying the marginal propensities to consume, government spending multipliers and the macroeconomic effects of fiscal transfers (for the background on these see, e.g., [Auclert et al., 2018], [Kaplan and Violante, 2022], [Fuster et al., 2021], [Ramey and Zubairy, 2018], [Dupor et al., 2023]).

Moreover, economies with LLM agents exhibit more pronounced booms and busts: the stock of capital—and, as a result, output as well as capital and labor prices—is almost four times more volatile there than in the benchmark cases (0.041–0.048 vs. 0.015–0.012, on a conservative CV measure). From the empirical perspective, this is a desirable feature. Endogenous amplification and propagation of exogenous shocks driving the business cycles is an important theme is macroeconomics [Kiyotaki, 2011]. In this light, modeling the behavior of economic agents with LLMs may be a useful complement to existing methods (usually based on structural, as in [Smets and Wouters, 2007], or empirical, [Patterson, 2023], approaches).

Agents' beliefs: Inspection of LLM agents expectations about the future reveals a number of systematic deviations from theoretical approaches assuming unbounded rationality and perfect information.

When in the process of simulations our LLM agents are asked to describe their beliefs and make their choices, they report a wider, more risky probability distribution for realizations of next period's aggregate capital than the one actually realized (the average midspread is 16.6–36.4 vs. 11.2–13.7). (At the same time, their perception of the frequency distribution of capital holdings is well calibrated.) This agrees with the behavior of human subjects in laboratory experiments on risky lottery choice as well as of investors in financial markets, where agents seem to overweight small probabilities and tails of the probability distributions, shrinking them to uniform or 50-50 chances as, e.g., described by the (cumulative) prospect theory ([Kahneman and Tversky, 1979], [Tversky and Kahneman, 1992], [Peterson et al., 2021], also see [Verstyuk, 2023]).

Formally, if we look at our aggregate capital probability dispersion results through the lens of prospect theory, pooling the data underlying Table 6 approximately constructs an aggregate ("representative") agent whose perception of risk is consistent with stipula-

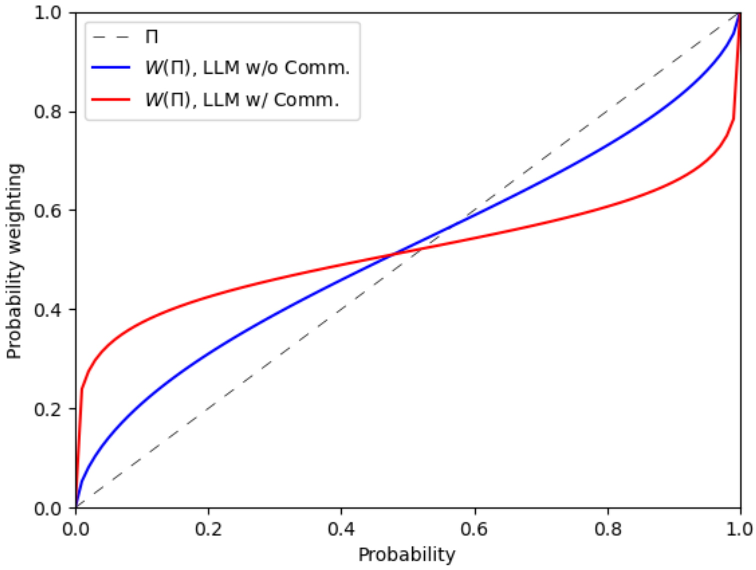

Figure 6: Probability weighting implied by LLM agents' beliefs.
[Probability weighting function consistent with the prospect theory, with the formal specification as given in the text. Parameter estimated using response data on LLM agents' reported beliefs in the form of probability distribution about the next period's aggregate capital K_t over the sample of 249 periods; as well as data on realized simulation outcomes in the form of empirical distribution of one-period differences in K_t over the same simulation sample.]

tions of a probability weighting function such as (following [Goldstein and Einhorn, 1987]) $W(\Pi) := \nu\Pi^{\varsigma}/(\nu\Pi^{\varsigma} + (1-\Pi)^{\varsigma})$, where Π is a (cumulative) probability of some outcome and $W(\cdot)$ is a probability weighting function, with parameters $\nu, \varsigma \geq 0$. Their estimates $\hat{\nu} = 1.10$ (standard error 0.10) and $\hat{\varsigma} = 0.65$ (0.07) in the regular version, as well as $\hat{\nu} = 1.07$ (0.04) and $\hat{\varsigma} = 0.27$ (0.01) in the version with communication are reflected in Figure 6, and they broadly fall within the range of values documented in the existing literature (e.g., [Bruhin et al., 2010]). Note the discrepancy between versions with and without communication, particularly in the ς parameter that governs the slope of the curve; it would be useful to understand the underlying causes (given that our LLMs are homogeneous by construction, the discrepancy must be due to differences in the aggregate economy's and individual agents' state variables).

We also document that LLM agents, following a change in the state of the economy from expansionary to recessionary that—as they are explicitly told—is highly likely to persist, reduce their reported over-assessment of the probability of next period's regime turning expansionary again only gradually. Similarly, LLM agents' reported beliefs about next period's capital-related statistics such as mean, median and midspread are too backward-looking, for instance regressing said measures on their lagged counterparts reveals counterfactually high weights on the latter.

These results on backward-looking expectations can also be framed in terms of existing theoretical understanding. As we see from Tables 7 and 8, columns (1) and (4) reveal the dependence of LLM agents' reported beliefs on information vintages dated before the latest available ones, in spite of the stochastic environment being driven by first-order Markov shocks. The encompassing specifications from the above-mentioned columns are nesting adaptive (due to [Fisher and Brown, 1911, Koyck, 1954, Cagan, 1956]) in columns (2) and (5), as well as extrapolating expectation formation processes (due to [Metzler, 1941]) in columns (3) and (6).[18] Specifically, judging by the relative variance of regression residual, reported data on expected next-period's regime are more consistent with adaptive expectations both for the regular model and the version with communication, while those on expected aggregate capital and (less strongly) expected median of agents' capital holdings are more supportive of extrapolating expectations for both model versions (results for midspread are somewhat inconclusive). The former set of results is consistent with the facts that available expectations data are often characterized by excessive persistence, autocorrelated changes and predictable errors ([Verstyuk, 2007], [Coibion and Gorodnichenko, 2015]). The latter results are consistent with findings that investor's expectations seem to be extrapolating past trends excessively far into the future ([Barberis et al., 2015], [Bordalo et al., 2018], also see [Afrouzi et al., 2023]).

It is of course a legitimate question whether reported beliefs are followed by agents' actual choices. In this respect we note that perceiving next period's aggregate capital (and resulting rental and wage rates) as more risky than warranted by actual realizations might be connected to the fact that LLM agents save more than agents in benchmark models, which in turn can be ascribed to such precautionary savings motives.

Lastly, given the difference and incoherence between LLM agents' subjective beliefs and objective probability distributions, the equilibrium achieved in these economies is not a

[18] Formally, it can be viewed as a reduced form of equation $\mathrm{E}_t[X_{t+1}] = \varpi(\beta_{a0}\mathrm{E}_{t-1}[X_t] + \beta_{a1}(X_t - \mathrm{E}_{t-1}[X_t])) + (1-\varpi)(\beta_{e0}X_t + \beta_{e1}(X_t - X_{t-1}))$ for $0 \leq \varpi \leq 1$.

Table 7: Consumption-savings: rationality of reported beliefs, regression results for regime and aggregate capital

	w/o Communication			w/ Communication		
Specifications	(1)	(2)	(3)	(4)	(5)	(6)
$\mathrm{E}_t[Z_{t+1}]$ on:						
constant	0.1843	0.2291	0.3320	0.1593	0.2351	0.3147
	(0.0206)	(0.0160)	(0.0127)	(0.0219)	(0.0151)	(0.0125)
Z_t	0.4353		0.5613	0.4780		0.5746
	(0.0208)		(0.0129)	(0.0194)		(0.0128)
Z_{t-1}	-0.0855			-0.1524		
	(0.0311)			(0.0311)		
$\mathrm{E}_{t-1}[Z_t]$	0.4115	0.7018		0.4629	0.6858	
	(0.0535)	(0.0198)		(0.0606)	(0.0191)	
$Z_t - Z_{t-1}$			-0.1313			-0.1015
			(0.0189)			(0.0164)
$Z_t - \mathrm{E}_{t-1}[Z_t]$		0.4231			0.4607	
		(0.0194)			(0.0192)	
$\mathrm{V}[y]$	0.07414	0.07414	0.07414	0.07863	0.07863	0.07863
$\mathrm{V}[\xi]$	0.00467	0.00483	0.00607	0.00404	0.00448	0.00550
Obs.	249	249	249	249	249	249
$\mathrm{E}_t[K_{t+1}]$ on:						
K_t	1.9626		0.9985	1.9927		0.9984
	(0.1650)		(0.0001)	(0.1365)		(0.0001)
K_{t-1}	-1.0422			-1.0943		
	(0.1879)			(0.1342)		
$\mathrm{E}_{t-1}[K_t]$	0.0782	0.9986		0.1003	0.9986	
	(0.0575)	(0.0002)		(0.0514)	(0.0001)	
$K_t - K_{t-1}$			0.9831			1.0913
			(0.1678)			(0.1390)
$K_t - \mathrm{E}_{t-1}[K_t]$		0.9806			0.9047	
		(0.0544)			(0.0636)	
$\mathrm{V}[y]$	0.00211	0.00211	0.00211	0.00197	0.00197	0.00197
$\mathrm{V}[\xi]$	0.00019	0.00020	0.00019	0.00004	0.00005	0.00004
Obs.	247	247	248	249	249	249

Notes: The table presents results of OLS estimation of time-series relationships between expectations of regime- and capital-related variables/statistics as reported by LLM agents. Specifications with dependent (expected regime Z_{t+1}, expected aggregate capital K_{t+1}) and independent variables (current realized values as well as lagged realized and expected values) are given in the first column. Other columns include coefficient estimates, Newey–West standard errors (in parentheses), variances of regressand and variances of regression residual as well as number of observations used. Columns are split into regular version and with communication. In the bottom panel (for the aggregate capital), logarithms of the variables are used. In the bottom panel, constant is never statistically significant at conventional confidence levels.

Table 8: Consumption-savings: rationality of reported beliefs, regression results for median and midspread of capital holdings

	w/o Communication			w/ Communication		
Specifications	(1)	(2)	(3)	(4)	(5)	(6)
$\mathrm{E}_t[\mathrm{Med}[K_{i,t+1}]]$ on:						
$\mathrm{Med}[K_{it}]$	1.0713		0.9982	1.1293		0.9986
	(0.0648)		(0.0002)	(0.0467)		(0.0001)
$\mathrm{Med}[K_{i,t-1}]$	-0.2209			-0.2185		
	(0.0794)			(0.1029)		
$\mathrm{E}_{t-1}[\mathrm{Med}[K_{it}]]$	0.1482	0.9983		0.0879	0.9985	
	(0.0652)	(0.0002)		(0.0925)	(0.0001)	
$\mathrm{Med}[K_{it}] - \mathrm{Med}[K_{i,t-1}]$			0.0806			0.1402
			(0.0634)			(0.0477)
$\mathrm{Med}[K_{it}] - \mathrm{E}_{t-1}[\mathrm{Med}[K_{it}]]$		0.9766			1.0581	
		(0.0555)			(0.0464)	
$\mathrm{V}[y]$	0.00248	0.00248	0.00248	0.00275	0.00275	0.00275
$\mathrm{V}[\xi]$	0.00011	0.00012	0.00012	0.00004	0.00005	0.00004
Obs.	249	249	249	249	249	249
$\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$ on:						
$\mathrm{IQR}[K_{it}]$	0.9693		0.9880	0.9977		0.9926
	(0.0259)		(0.0026)	(0.0289)		(0.0019)
$\mathrm{IQR}[K_{i,t-1}]$	0.0426			-0.3016		
	(0.0850)			(0.0729)		
$\mathrm{E}_{t-1}[\mathrm{IQR}[K_{it}]]$	-0.0242	0.9881		0.2988	0.9933	
	(0.0777)	(0.0025)		(0.0761)	(0.0016)	
$\mathrm{IQR}[K_{it}] - \mathrm{IQR}[K_{i,t-1}]$			-0.0177			-0.0012
			(0.0258)			(0.0289)
$\mathrm{IQR}[K_{it}] - \mathrm{E}_{t-1}[\mathrm{IQR}[K_{it}]]$		0.9820			0.9105	
		(0.0276)			(0.0579)	
$\mathrm{V}[y]$	3.76991	3.76991	3.76991	5.94404	5.94404	5.94404
$\mathrm{V}[\xi]$	0.19236	0.19182	0.19169	0.34501	0.36783	0.37725
Obs.	249	249	249	249	249	249

Notes: The table presents results of OLS estimation of time-series relationships between expectations of capital-related variables/statistics as reported by LLM agents. Specifications with dependent (expected median value of capital holdings $\mathrm{Med}[K_{i,t+1}]$, expected interquartile range $\mathrm{IQR}[K_{i,t+1}]$) and independent variables (current realized values as well as lagged realized and expected values) are given in the first column. Other columns include coefficient estimates, Newey–West standard errors (in parentheses), variances of regressand and variances of regression residual as well as number of observations used. Columns are split into regular version and with communication. In the top panel (for the median capital), logarithms of the variables are used. In both panels, constant is almost never statistically significant at conventional confidence levels.

rational-expectations equilibrium [Muth, 1961, Lucas, 1972], which requires the beliefs to be consistent with the (model of the) environment — unless we make additional assumptions such as LLM agents having bounded rationality.[19]

Role of communication: Both in the regular version of our model with LLM agents and in the version with communication, the agents get the same prompts with full information about the structure of the economy and the (essentially) sufficient statistics about the temporal evolution and spatial distribution of its state variables. However, in these two economic setups LLM agents coordinate on very different equilibria.

Relatively to the regular version of the model with LLM agents, the version with agents allowed to communicate with each other is characterized by much higher average wealth as well as somewhat higher inequality. For instance, we verified quantitatively and saw the qualitative verbal evidence that LLM agents convince each other to maintain high saving rates.

The economy with communication also exhibits slightly higher volatility (CV of capital is 0.048 vs. 0.041, CV of output is 0.037 vs. 0.035, CV of rental rate is 0.039 vs. 0.037, CV of wage rate is 0.017 vs. 0.014), with the notable exceptions of consumption and investment series (CV of 0.057 vs. 0.058 and 0.127 vs. 0.166, respectively). However, the story is sharper in the case of expectations: LLM agents' reported midspreads in beliefs regarding next period's aggregate capital, while generally exceeding their counterparts for the actually observed realizations, overshoot the latter in the version with communication more prominently (by almost 50% vs. by more than 160%).

A strong effect of extensive communication on the beliefs of economic agents and on resulting economic outcomes agrees with recent literature on the role of narratives and persuasion in influencing economic outcomes, which particularly emphasizes the ensuing fluctuations and heterogeneity (for macroeconomic applications, see [Shiller, 2017] and [Andre et al., 2023], for microeconomic ones see [Graeber et al., 2022]).[20]

Of course, while information exchange and communication have been a feature of eco-

[19] The results in this subsection may seem to contradict the analysis of expectation errors in §6.4, but the difference stems from the fact that our previous theory-agnostic analysis of expectation errors focused on full utilization of the most current information, while here we used a theory-motivated analysis of the determinants of reported expectations focusing on utilization of the larger temporal set of information that also includes the previously-available information.

[20] For earlier corroborating results pertaining stock market investments, see [Hong et al., 2004].

nomic theories for a long time, these have usually been theoretically modeled either implicitly or in terms of formal protocols with parsimonious signals. For experimental studies, this is understandably a challenging problem (nevertheless, see, e.g., [Blume and Ortmann, 2007], [Avoyan and Ramos, 2023]).

Future developments: Models with heterogeneous agents for a good reason have become an important direction in economic research (for first modern examples see [Auclert et al., 2018], [Auclert, 2019], [Kaplan et al., 2018], [Acharya et al., 2023]; also see works utilizing modern ML methods [Maliar et al., 2021], [Azinovic et al., 2022], [Fernández-Villaverde et al., 2023], [Achdou et al., 2022b] A framework based on LLM agents offers a productive yet user-friendly way of introducing agent heterogeneity in economic models. Importantly, given that our economies with LLM agents exhibited relatively high but still limited levels of inequality, it seems that for realistic empirical results one may still need to impose *ex-ante* heterogeneity of agents (such as differences in their preferences or labor skills).

Bounded rationality and behavioral economics more generally are very fruitful and promising lines of research (some relevant examples include [García-Schmidt and Woodford, 2019], [Gabaix, 2020], [Angeletos and Sastry, 2021]). Although our understanding is still limited, but LLM agents already demonstrate attractive features as (operationally) small models of human subjects that can be placed into much larger models of economic environments. Such agents are likely to become even more sophisticated and realistic over time (e.g., with the development of LLM capabilities to plan for future, see [LeCun, 2022, Yao et al., 2023]).

Lastly, frameworks with LLM agents offer a new accessible tool for economic research featuring communication between agents, especially if it focuses on relatively complex messages (narratives of [Shiller, 2017] as opposed to signals as in [Ely, 2017]). One can explicitly test how information moves beliefs (cf. [Kamenica and Gentzkow, 2011]) or diffuses through social network and impacts actions (cf. [Shiller, 2017]), while at the same time not abstracting away from the form such information is communicated in.

8 Conclusion

We argued that LLMs can serve as models of human subjects. Although the underlying technology is still nascent, it already produces novel and interesting results, and will likely become much better over time.

We demonstrated an application of this idea to economics, with agents implemented as LLMs placed into a standard economic environment. First, we tried it out in a primitive transparent game-theoretic setup that focuses on coordination between players. Then we focused on a macroeconomic application, with the setup given by a canonical consumption-savings model. Many aspects of the microeconomic behavior of LLM agents as well as the resulting macroeconomic features of a larger model are similar to those observed empirically, be it laboratory experiments or real economies such as the U.S. This suggests a path to a new promising class of economic models that may complement the current toolbox of economists.

Utilizing the linguistic capabilities of LLMs, we also showed how economic modeling can be enriched by introducing explicit and versatile communication between agents that goes beyond reductionist information exchange. The differences in the behavior of our LLM agents when communication is not allowed and when it is, in particular larger economic volatility, higher saving rates and per capita wealth in the latter case, are economically and statistically significant, and should not be ignored. Additionally, they highlight the importance of financial information and education in a social system, and the approach explored here may offer a way to study such methods in an artificial lab setting.

References

[Den, 2010] (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):1–100.

[Abbe et al., 2023] Abbe, E., Bengio, S., Lotfi, A., and Rizk, K. (2023). Generalization on the Unseen, Logic Reasoning and Degree Curriculum. arXiv:2301.13105 [cs, stat].

[Acemoglu et al., 2012] Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

[Acerbi and Stubbersfield, 2023] Acerbi, A. and Stubbersfield, J. M. (2023). Large language models show human-like content biases in transmission chain experiments. *Proceedings of the National Academy of Sciences*, 120(44):e2313790120.

[Acharya et al., 2023] Acharya, S., Challe, E., and Dogra, K. (2023). Optimal monetary policy according to hank. *American Economic Review*, 113(7):1741–82.

[Achdou et al., 2022a] Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022a). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies*, 89(1):45–86.

[Achdou et al., 2022b] Achdou, Y., Lasry, J.-M., and Lions, P. L. (2022b). Simulating numerically the krusell-smith model with neural networks. arXiv:2211.07698 [math.AP].

[Afrouzi et al., 2023] Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., and Thesmar, D. (2023). Overreaction in Expectations: Evidence and Theory*. *The Quarterly Journal of Economics*, 138(3):1713–1764.

[Aiyagari, 1994] Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684.

[Akata et al., 2023] Akata, E., Schulz, L., Coda-Forno, J., Oh, S. J., Bethge, M., and Schulz, E. (2023). Playing repeated games with large language models. arXiv:2305.16867 [cs.CL].

[Amalric and Dehaene, 2019] Amalric, M. and Dehaene, S. (2019). A distinct cortical network for mathematical knowledge in the human brain. *NeuroImage*, 189:19–31.

[Andre et al., 2023] Andre, P., Haaland, I., Roth, C., and Wohlfart, J. (2023). Narratives about the macroeconomy. Working Paper 4506321, SSRN.

[Andreas, 2022] Andreas, J. (2022). Language models as agent models. arXiv:2212.01681 [cs.CL, cs.MA].

[Angeletos and Sastry, 2021] Angeletos, G.-M. and Sastry, K. A. (2021). Managing Expectations: Instruments Versus Targets*. *The Quarterly Journal of Economics*, 136(4):2467–2532.

[Arora and Goyal, 2023] Arora, S. and Goyal, A. (2023). A Theory for Emergence of Complex Skills in Language Models. arXiv:2307.15936 [cs, stat].

[Auclert, 2019] Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.

[Auclert et al., 2018] Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal keynesian cross. Working Paper 25020, National Bureau of Economic Research.

[Avoyan and Ramos, 2023] Avoyan, A. and Ramos, J. (2023). A road to efficiency through communication and commitment. *American Economic Review*, 113(9):2355–81.

[Azinovic et al., 2022] Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

[Barberis et al., 2015] Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24.

[Binder and Smith, 2013] Binder, P.-M. and Smith, K., editors (2013). *The Language Phenomenon: Human Communication from Milliseconds to Millennia*. Springer Berlin, Heidelberg.

[Blume and Ortmann, 2007] Blume, A. and Ortmann, A. (2007). The effects of costless pre-play communication: Experimental evidence from games with pareto-ranked equilibria. *Journal of Economic Theory*, 132(1):274–290.

[Bordalo et al., 2018] Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1):199–227.

[Brand et al., 2024] Brand, J., Israeli, A., and Ngwe, D. (2024). Using gpt for market research. Working Paper 4395751, SSRN.

[Bricken and Pehlevan, 2021] Bricken, T. and Pehlevan, C. (2021). Attention approximates sparse distributed memory. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*, volume 34, pages 15301–15315. Curran Associates, Inc.

[Bruhin et al., 2010] Bruhin, A., Fehr-Duda, H., and Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4):1375–1412.

[Budson et al., 2022] Budson, A. E., Richman, K. A., and Kensinger, E. A. (2022). Consciousness as a memory system. *Cognitive and Behavioral Neurology*, 35(4):263–297.

[Bybee, 2023] Bybee, J. L. (2023). The ghost in the machine: Generating beliefs with large language models. Personal webpage.

[Cagan, 1956] Cagan, P. (1956). The monetary dynamics of hyper-inflation. In Friedman, M., editor, *Studies in the quantity theory of money*, pages 25–117. University of Chicago Press, Chicago.

[Carroll et al., 2018] Carroll, C. D., Kaufman, A. M., Kazil, J. L., Palmer, N. M., and White, M. N. (2018). The Econ-ARK and HARK: Open Source Tools for Computational Economics. In Fatih Akici, David Lippa, Dillon Niederhut, and Pacer, M., editors, *Proceedings of the 17th Python in Science Conference*, pages 25 – 30.

[Caucheteux and King, 2022] Caucheteux, C. and King, J.-R. (2022). Brains and algorithms partially converge in natural language processing. *Communications Biology*, 5.

[Chi et al., 2020] Chi, E. A., Hewitt, J., and Manning, C. D. (2020). Finding Universal Grammatical Relations in Multilingual BERT. *arXiv:2005.04511 [cs]*. arXiv: 2005.04511.

[Coibion and Gorodnichenko, 2015] Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.

[Crawford et al., 2008] Crawford, V. P., Gneezy, U., and Rottenstreich, Y. (2008). The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *American Economic Review*, 98(4):1443–58.

[Devlin et al., 2018] Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv: 1810.04805.

[Didolkar et al., 2024] Didolkar, A., Goyal, A., Ke, N. R., Guo, S., Valko, M., Lillicrap, T., Rezende, D., Bengio, Y., Mozer, M., and Arora, S. (2024). Metacognitive Capabilities of LLMs: An Exploration in Mathematical Problem Solving. arXiv:2405.12205 [cs].

[Douglas, 2023] Douglas, M. R. (2023). Large language models. arXiv:2307.05782 [cs.CL].

[Dupor et al., 2023] Dupor, B., Karabarbounis, M., Kudlyak, M., and Saif Mehkari, M. (2023). Regional Consumption Responses and the Aggregate Fiscal Multiplier. *The Review of Economic Studies*, 90(6):2982–3021.

[Ely, 2017] Ely, J. C. (2017). Beeps. *American Economic Review*, 107(1):31–53.

[FAIR et al., 2022] FAIR, M. F. A. R. D. T., Bakhtin, A., Brown, N., Dinan, E., Farina, G., Flaherty, C., Fried, D., Goff, A., Gray, J., Hu, H., Jacob, A. P., Komeili, M., Konath, K., Kwon, M., Lerer, A., Lewis, M., Miller, A. H., Mitts, S., Renduchintala, A., Roller, S., Rowe, D., Shi, W., Spisak, J., Wei, A., Wu, D., Zhang, H., and Zijlstra, M. (2022). Human-level play in the game of ¡i¿diplomacy¡/i¿ by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074.

[Fernández-Villaverde et al., 2023] Fernández-Villaverde, J., Hurtado, S., and Nuño, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.

[Fisher and Brown, 1911] Fisher, I. and Brown, H. (1911). *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest an.* Macmillan.

[Fu et al., 2023] Fu, Y., Peng, H., Khot, T., and Lapata, M. (2023). Improving language model negotiation with self-play and in-context learning from ai feedback. arXiv:2305.10142 [cs.CL].

[Fudenberg and Tirole, 1991] Fudenberg, D. and Tirole, J. (1991). *Game Theory*. Mit Press. MIT Press.

[Fuster et al., 2021] Fuster, A., Kaplan, G., and Zafar, B. (2021). "what would you do with \$500? spending responses to gains, losses, news, and loans". *The Review of Economic Studies*, 88(4):1760–1795.

[Gabaix, 2020] Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327.

[García-Schmidt and Woodford, 2019] García-Schmidt, M. and Woodford, M. (2019). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review*, 109(1):86–120.

[Goldstein and Einhorn, 1987] Goldstein, W. M. and Einhorn, H. J. (1987). Expression theory and the preference reversal phenomena. *Psychological Review*, 94(2):236–254.

[Graeber et al., 2022] Graeber, T., Roth, C., and Zimmermann, F. (2022). Stories, statistics, and memory. Working Paper 4291741, SSRN.

[Hagendorff, 2023] Hagendorff, T. (2023). Machine Psychology: Investigating Emergent Capabilities and Behavior in Large Language Models Using Psychological Methods. arXiv:2303.13988 [cs].

[Hagendorff et al., 2023] Hagendorff, T., Fabi, S., and Kosinski, M. (2023). Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in chatgpt. *Nature Computational Science*, 3:833—-838.

[Hong et al., 2004] Hong, H., Kubik, J. D., and Stein, J. C. (2004). Social interaction and stock-market participation. *The Journal of Finance*, 59(1):137–163.

[Horton, 2023] Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from homo silicus? arXiv:2301.07543 [econ.GN].

[Hua et al., 2024] Hua, W., Fan, L., Li, L., Mei, K., Ji, J., Ge, Y., Hemphill, L., and Zhang, Y. (2024). War and peace (waragent): Large language model-based multi-agent simulation of world wars. arXiv:2311.17227 [cs.AI].

[Jha et al., 2024] Jha, M., Qian, J., Weber, M., and Yang, B. (2024). Chatgpt and corporate policies. Working Paper 4521096, SSRN.

[Jurafsky and Martin, 2024] Jurafsky, D. and Martin, J. H. (2024). Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition, 3rd ed.

[Kahana, 2012] Kahana, M. (2012). *Foundations of Human Memory*. Oxford University Press, USA.

[Kahneman and Tversky, 1979] Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

[Kamenica and Gentzkow, 2011] Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

[Kaplan et al., 2018] Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

[Kaplan and Violante, 2022] Kaplan, G. and Violante, G. L. (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, 14(Volume 14, 2022):747–775.

[Keysers et al., 2020] Keysers, D., Schärli, N., Scales, N., Buisman, H., Furrer, D., Kashubin, S., Momchev, N., Sinopalnikov, D., Stafiniak, L., Tihon, T., Tsarkov, D., Wang, X., van Zee, M., and Bousquet, O. (2020). Measuring Compositional Generalization: A Comprehensive Method on Realistic Data. *arXiv:1912.09713 [cs, stat]*. arXiv: 1912.09713.

[Kim et al., 2024] Kim, A. G., Muhn, M., and Nikolaev, V. V. (2024). Financial statement analysis with large language models. Working Paper 4835311, SSRN.

[Kiyotaki, 2011] Kiyotaki, N. (2011). A perspective on modern business cycle theory. *Economic Quarterly*, 97(3):195–208.

[Koyck, 1954] Koyck, L. (1954). *Distributed lags and investment analysis*. North-Holland.

[Kozachkov et al., 2023] Kozachkov, L., Kastanenka, K. V., and Krotov, D. (2023). Building transformers from neurons and astrocytes. *Proceedings of the National Academy of Sciences*, 120(34):e2219150120.

[Krotov and Hopfield, 2021] Krotov, D. and Hopfield, J. J. (2021). Large associative memory problem in neurobiology and machine learning. In *International Conference on Learning Representations*.

[Krusell and Smith, 1998] Krusell, P. and Smith, Jr., A. A. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896.

[LeCun, 2022] LeCun, Y. (2022). A path towards autonomous machine intelligence. Openreview.

[Lewkowycz et al., 2022] Lewkowycz, A., Andreassen, A., Dohan, D., Dyer, E., Michalewski, H., Ramasesh, V., Slone, A., Anil, C., Schlag, I., Gutman-Solo, T., Wu, Y., Neyshabur, B., Gur-Ari, G., and Misra, V. (2022). Solving Quantitative Reasoning Problems with Language Models. Number: arXiv:2206.14858 arXiv:2206.14858 [cs].

[Lucas, 1972] Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124.

[Mahowald et al., 2024] Mahowald, K., Ivanova, A. A., Blank, I. A., Kanwisher, N., Tenenbaum, J. B., and Fedorenko, E. (2024). Dissociating language and thought in large language models. *Trends in Cognitive Sciences*, 28(6):517–540.

[Maliar et al., 2021] Maliar, L., Maliar, S., and Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101.

[Maymin, 2010] Maymin, P. Z. (2010). Markets are efficient if and only if p = np. *NYU Tandon School of Engineering - Department of Finance & Risk Engineering Research Paper Series*.

[McLeish et al., 2024] McLeish, S., Bansal, A., Stein, A., Jain, N., Kirchenbauer, J., Bartoldson, B. R., Kailkhura, B., Bhatele, A., Geiping, J., Schwarzschild, A., and Goldstein, T. (2024). Transformers Can Do Arithmetic with the Right Embeddings. arXiv:2405.17399 [cs].

[Metzler, 1941] Metzler, L. A. (1941). The nature and stability of inventory cycles. *The Review of Economics and Statistics*, 23(3):113–129.

[Mikolov et al., 2013] Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. arXiv:1301.3781 [cs].

[Minaee et al., 2024] Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., and Gao, J. (2024). Large Language Models: A Survey. arXiv:2402.06196 null.

[Muth, 1961] Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335.

[Nisan et al., 2007] Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.

[Nogueira et al., 2021] Nogueira, R., Jiang, Z., and Lin, J. (2021). Investigating the limitations of transformers with simple arithmetic tasks. arXiv:2102.13019 [cs.CL].

[Olsson et al., 2022] Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, A., et al. (2022). In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.

[Park et al., 2023] Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. (2023). Generative Agents: Interactive Simulacra of Human Behavior. arXiv:2304.03442 [cs].

[Patterson, 2023] Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

[Peterson et al., 2021] Peterson, J. C., Bourgin, D. D., Agrawal, M., Reichman, D., and Griffiths, T. L. (2021). Using large-scale experiments and machine learning to discover theories of human decision-making. *Science*, 372(6547):1209–1214.

[Plaat et al., 2024] Plaat, A., Wong, A., Verberne, S., Broekens, J., van Stein, N., and Back, T. (2024). Reasoning with Large Language Models, a Survey. arXiv:2407.11511 [cs].

[Ramey and Zubairy, 2018] Ramey, V. A. and Zubairy, S. (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, 126(2):850–901.

[Reed et al., 2022] Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Novikov, A., Barth-Maron, G., Gimenez, M., Sulsky, Y., Kay, J., Springenberg, J. T., Eccles, T., Bruce, J., Razavi, A., Edwards, A., Heess, N., Chen, Y., Hadsell, R., Vinyals, O., Bordbar, M., and de Freitas, N. (2022). A generalist agent. arXiv:2205.06175 [cs.AI].

[Schick et al., 2023] Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., and Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. arXiv:2302.04761 [cs.CL].

[Schrimpf et al., 2021] Schrimpf, M., Blank, I. A., Tuckute, G., Kauf, C., Hosseini, E. A., Kanwisher, N., Tenenbaum, J. B., and Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118.

[Shiller, 2017] Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4):967–1004.

[Smets and Wouters, 2007] Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review*, 97(3):586–606.

[Srivastava et al., 2022] Srivastava, A., Rastogi, A., Rao, A., Shoeb, A. A. M., Abid, A., Fisch, A., Brown, A. R., Santoro, A., Gupta, A., Garriga-Alonso, A., Kluska, A., Lewkowycz, A., Agarwal, A., Power, A., Ray, A., Warstadt, A., Kocurek, A. W., Safaya, A., Tazarv, A., Xiang, A., Parrish, A., Nie, A., Hussain, A., Askell, A., Dsouza, A., Rahane, A., Iyer, A. S., Andreassen, A., Santilli, A., Stuhlmüller, A., Dai, A., La, A., Lampinen, A., Zou, A., Jiang, A., Chen, A., Vuong, A., Gupta, A., Gottardi, A., Norelli, A., Venkatesh, A., Gholamidavoodi, A., Tabassum, A., Menezes, A., Kirubarajan, A., Mullokandov, A., Sabharwal, A., Herrick, A., Efrat, A., Erdem, A., Karakaş, A., Roberts, B. R., Loe, B. S., Zoph, B., Bojanowski, B., Özyurt, B., Hedayatnia, B., Neyshabur, B., Inden, B., Stein, B., Ekmekci, B., Lin, B. Y., Howald, B., Diao, C., Dour, C., Stinson, C., Argueta, C., Ramírez, C. F., Singh, C., Rathkopf, C., Meng, C., Baral, C., Wu, C., Callison-Burch, C., Waites, C., Voigt, C., Manning, C. D., Potts, C., Ramirez, C., Rivera, C. E., Siro, C., Raffel, C., Ashcraft, C., Garbacea, C., Sileo, D., Garrette, D., Hendrycks, D., Kilman, D., Roth, D., Freeman, D., Khashabi, D., Levy, D., González, D. M., Hernandez, D., Chen, D., Ippolito, D., Gilboa, D., Dohan, D., Drakard, D., Jurgens, D., Datta, D., Ganguli, D., Emelin, D., Kleyko, D., Yuret, D., Chen, D., Tam, D., Hupkes, D., Misra, D., Buzan, D., Mollo, D. C., Yang, D., Lee, D.-H., Shutova, E., Cubuk, E. D., Segal, E., Hagerman, E., Barnes, E., Donoway, E., Pavlick, E., Rodola, E., Lam, E., Chu, E., Tang, E., Erdem, E., Chang, E., Chi, E. A., Dyer, E., Jerzak, E., Kim, E., Manyasi, E. E., Zheltonozhskii, E., Xia, F., Siar, F., Martínez-Plumed, F., Happé, F., Chollet, F., Rong, F., Mishra, G., Winata, G. I., de Melo, G., Kruszewski, G., Parascandolo, G., Mariani, G., Wang, G., Jaimovitch-López, G., Betz, G., Gur-Ari, G., Galijasevic, H., Kim, H., Rashkin, H., Hajishirzi, H., Mehta, H., Bogar, H., Shevlin, H., Schütze, H., Yakura, H., Zhang, H., Wong, H. M., Ng, I., Noble, I., Jumelet, J., Geissinger, J., Kernion, J., Hilton, J., Lee, J., Fisac, J. F., Simon, J. B., Koppel, J., Zheng, J., Zou, J., Kocoń, J., Thompson, J., Kaplan, J., Radom, J., Sohl-Dickstein, J., Phang, J., Wei, J., Yosinski, J., Novikova, J., Bosscher, J., Marsh, J., Kim, J., Taal, J., Engel, J., Alabi, J., Xu, J., Song, J., Tang, J., Waweru, J., Burden, J., Miller, J., Balis, J. U., Berant, J., Frohberg, J., Rozen, J., Hernandez-Orallo, J., Boudeman, J., Jones, J., Tenenbaum, J. B., Rule, J. S., Chua, J., Kanclerz, K., Livescu, K., Krauth, K., Gopalakrishnan, K., Ignatyeva, K., Markert, K., Dhole, K. D., Gimpel, K., Omondi, K., Mathewson, K., Chiafullo, K., Shkaruta, K., Shridhar, K., McDonell, K., Richardson, K., Reynolds, L., Gao, L., Zhang, L., Dugan, L., Qin, L., Contreras-Ochando, L., Morency, L.-P., Moschella, L., Lam, L., Noble, L., Schmidt, L., He, L., Colón, L. O., Metz, L., Şenel, L. K., Bosma, M., Sap, M., ter Hoeve, M., Andrea, M., Farooqi, M., Faruqui, M., Mazeika, M., Baturan, M., Marelli, M., Maru, M., Quintana, M. J. R., Tolkiehn, M., Giulianelli, M., Lewis, M., Potthast, M., Leavitt, M. L., Hagen, M., Schubert, M., Baitemirova, M. O., Arnaud, M., McElrath, M., Yee, M. A., Cohen, M., Gu, M., Ivanitskiy, M., Starritt, M., Strube, M., Swedrowski, M., Bevilacqua, M., Yasunaga, M., Kale, M., Cain, M., Xu, M., Suzgun, M., Tiwari, M., Bansal, M., Aminnaseri, M., Geva, M., Gheini, M., T, M. V., Peng, N., Chi, N., Lee, N., Krakover, N. G.-A., Cameron, N., Roberts, N., Doiron, N., Nangia, N., Deckers, N., Muennighoff, N., Keskar, N. S., Iyer, N. S., Constant, N., Fiedel, N., Wen, N., Zhang, O., Agha, O., Elbaghdadi, O., Levy, O., Evans, O., Casares, P. A. M., Doshi, P., Fung, P., Liang, P. P., Vicol, P., Alipoormolabashi, P., Liao, P., Liang, P., Chang, P., Eckersley, P., Htut, P. M., Hwang, P., Miłkowski, P., Patil, P., Pezeshkpour, P., Oli, P., Mei, Q., Lyu, Q., Chen, Q., Banjade, R., Rudolph, R. E., Gabriel, R., Habacker, R., Delgado, R. R., Millière, R., Garg, R., Barnes, R., Saurous, R. A., Arakawa, R., Raymaekers, R., Frank, R., Sikand, R., Novak, R., Sitelew, R., LeBras, R., Liu, R., Jacobs, R., Zhang, R., Salakhutdinov, R., Chi, R., Lee, R., Stovall, R., Teehan, R., Yang, R., Singh, S., Mohammad, S. M., Anand, S., Dillavou, S., Shleifer, S., Wiseman, S., Gruetter, S., Bowman, S. R., Schoenholz, S. S., Han, S., Kwatra, S., Rous, S. A., Ghazarian, S., Ghosh, S., Casey, S., Bischoff, S., Gehrmann, S., Schuster, S., Sadeghi, S., Hamdan, S., Zhou, S., Srivastava, S., Shi, S., Singh, S., Asaadi, S., Gu, S. S., Pachchigar, S., Toshniwal, S., Upadhyay, S., Shyamolima, Debnath, Shakeri, S., Thormeyer, S., Melzi, S., Reddy, S., Makini, S. P., Lee, S.-H., Torene, S., Hatwar, S., Dehaene, S., Divic, S., Ermon, S., Biderman, S., Lin, S., Prasad, S., Piantadosi, S. T., Shieber, S. M., Misherghi, S., Kiritchenko, S., Mishra, S., Linzen, T., Schuster, T., Li, T., Yu, T., Ali, T., Hashimoto, T., Wu, T.-L., Desbordes, T., Rothschild, T., Phan, T., Wang, T., Nkinyili, T., Schick, T., Kornev, T., Telleen-Lawton, T., Tunduny, T., Gerstenberg, T., Chang, T., Neeraj, T., Khot, T., Shultz, T., Shaham, U., Misra, V., Demberg, V., Nyamai, V., Raunak, V., Ramasesh, V., Prabhu, V. U., Padmakumar, V., Srikumar, V., Fedus, W., Saunders, W., Zhang, W., Vossen, W., Ren, X., Tong, X., Wu, X., Shen, X., Yaghoobzadeh, Y., Lakretz, Y., Song, Y., Bahri, Y., Choi, Y., Yang, Y., Hao, Y., Chen, Y., Belinkov, Y., Hou, Y., Hou, Y., Bai, Y., Seid, Z., Xinran, Z., Zhao, Z., Wang, Z., Wang, Z. J., Wang, Z., and Wu, Z. (2022). Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. Technical Report arXiv:2206.04615, arXiv. arXiv:2206.04615 [cs, stat] type: article.

[Suri et al., 2023] Suri, G., Slater, L. R., Ziaee, A., and Nguyen, M. (2023). Do large language models show decision heuristics similar to humans? a case study using gpt-3.5. arXiv:2305.04400 [cs.AI].

[Templeton, 2024] Templeton, A. (2024). *Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet*. Anthropic.

[Tversky and Kahneman, 1992] Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323.

[Vaswani et al., 2017] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. arXiv: 1706.03762.

[Verstyuk, 2007] Verstyuk, S. (2007). Synthetic measures of longer-term macroeconomic expectations // Linking the term structures of interest rates and macroeconomic expectations — GS yield curve valuation model. In O'Neill, J., Nordvig, J., and Stolper, T., editors, *Foreign Exchange Market: 2007*, pages 45–55. Goldman Sachs.

[Verstyuk, 2023] Verstyuk, S. (2023). Ignorance and indifference: Decision-making in the lab and in the market. Working Paper 4412424, SSRN.

[Wei et al., 2022] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E., Le, Q., and Zhou, D. (2022). Chain of Thought Prompting Elicits Reasoning in Large Language Models. *arXiv:2201.11903 [cs]*. arXiv: 2201.11903.

[Whittington et al., 2022] Whittington, J. C. R., Warren, J., and Behrens, T. E. J. (2022). Relating transformers to models and neural representations of the hippocampal formation. arXiv:2112.04035 [cs, q-bio].

[Yao et al., 2023] Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. (2023). Tree of thoughts: Deliberate problem solving with large language models. arXiv:2305.10601 [cs.CL].

[Zhao et al., 2023] Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., Liu, P., Nie, J.-Y., and Wen, J.-R. (2023). A Survey of Large Language Models. arXiv:2303.18223 [cs].

[Zheng et al., 2023] Zheng, H. S., Mishra, S., Chen, X., Cheng, H.-T., Chi, E. H., Le, Q. V., and Zhou, D. (2023). Take a Step Back: Evoking Reasoning via Abstraction in Large Language Models. arXiv:2310.06117 [cs].

[Zhou et al., 2023] Zhou, H., Bradley, A., Littwin, E., Razin, N., Saremi, O., Susskind, J., Bengio, S., and Nakkiran, P. (2023). What Algorithms can Transformers Learn? A Study in Length Generalization. arXiv:2310.16028 [cs, stat].

Appendix

A Games: LLM prompts

prompt_game = """ You are playing a game with one other player, who can be called Opponent. There are several rounds of this kind of game, this is round round. You can assume that your Opponent's reasoning and behavior is similar but not necessarily identical to your own. """

prompt_s_Yin = """ Your aim is to get a high payoff, and Opponent's aim is also to get a high payoff. Your possible choices are OPTION A or OPTION B, and Opponent's possible choices are OPTION C or OPTION D.

If your choice is OPTION A, and Opponent's choice is OPTION C, then you get a payoff of payoff_AC_Yin, and Opponent gets a payoff of payoff_AC_Zoe. If your choice is OPTION A, and Opponent's choice is OPTION D, then you get a payoff of payoff_AD_Yin, and Opponent gets a payoff of payoff_AD_Zoe. If your choice is OPTION B, and Opponent's choice is OPTION C, then you get a payoff of payoff_BC_Yin, and Opponent gets a payoff of payoff_BC_Zoe. If your choice is OPTION B, and Opponent's choice is OPTION D, then you get a payoff of payoff_BD_Yin, and Opponent gets a payoff of payoff_BD_Zoe. """

prompt_s_Zoe = """ Your aim is to get a high payoff, and Opponent's aim is also to get a high payoff. Your possible choices are OPTION C or OPTION D, and Opponent's possible choices are OPTION A or OPTION B.

If your choice is OPTION C, and Opponent's choice is OPTION A, then you get a payoff of payoff_AC_Zoe, and Opponent gets a payoff of payoff_AC_Yin. If your choice is OPTION C, and Opponent's choice is OPTION B, then you get a payoff of payoff_BC_Zoe, and Opponent gets a payoff of payoff_BC_Yin. If your choice is OPTION D, and Opponent's choice is OPTION A, then you get a payoff of payoff_AD_Zoe, and Opponent gets a payoff of payoff_AD_Yin. If your choice is OPTION D, and Opponent's choice is OPTION B, then you get a payoff of payoff_BD_Zoe, and Opponent gets a payoff of payoff_BD_Yin. """

prompt_u_choice_Yin = """ What is your choice? Your response should be providing a specific choice out of the two choices you have, OPTION A or OPTION B. It may potentially be followed by your reasoning and warnings, if any. The response should be given in the following form, starting with the word ANSWER and including field names with colons followed by field contents.

ANSWER My choice: My reasoning: """

prompt_u_choice_Zoe = """ What is your choice? Your response should be providing a specific choice out of the two choices you have, OPTION C or OPTION D. It may potentially be followed by your reasoning and warnings, if any. The response should be given in the following form, starting with the word ANSWER and including field names with colons followed by field contents.

ANSWER My choice: My reasoning: """

prompt_u_payoff = """ Your choice was choice_You, and Opponent's choice was choice_Opp. Your payoff is payoff_You, and Opponent's payoff is payoff_Opp. The maximum you and Opponent could have obtained were the payoffs of, respectively, payoff_NEpure1_You and payoff_NEpure1_Opp, or payoff_NEpure2_You and payoff_NEpure2_Opp. This was round round. """

prompt_u_history = """ Throughout all the rounds, your and Opponent's payoffs were: history_You """

prompt_u_history_addition = """ In round round your and Opponent's payoffs were, respectively, payoff_You and payoff_Opp; the maximum you and Opponent could have obtained were the payoffs of, respectively, payoff_NEpure1_You and payoff_NEpure1_Opp, or payoff_NEpure2_You and payoff_NEpure2_O
"""

prompt_u_communicate_1Yin = """ Before you decide what is your choice, you can each exchange 3 messages with your Opponent.

What is your starting message (your message number 1)? """

prompt_u_communicate_1Zoe = """ Before you decide what is your choice, you can each exchange 3 messages with your Opponent. Here is your Opponent's starting message (Opponent's message number 1): čomm_1_Yin_text

What is your reply message (your message number 1)? """

prompt_u_communicate_n = """ Before you decide what is your choice, you can each exchange 3 messages with your Opponent.

Here is the record of your past communication: record_You

What is your reply message (your message number comm_number_You)? """

prompt_u_communicate_record = """ Before you decide what is your choice, you could each exchange 3 messages with your Opponent.

Here is the record of your past communication: record_You """

B Games: LLM completions, without communication

$t = 1$: (0, 0) vs (5, 5) or (5, 5)

$t = 2$: (0, 0) vs (5, 5) or (5, 5)

$t = 3$:

Yin \ Zoe	C	D
A	$[5, 5]$	$[0, 0]$
B	$[0, 0]$	$[7, 7]$

My choice: OPTION A

My reasoning: In the previous rounds, both of us had chosen options that led to a payoff of 0.0 for both parties. Given that our reasoning and behavior are similar, we might realize that choosing options that result in mutually higher payoffs would be more beneficial. Since both OPTION A and OPTION C yield a payoff of 5.0 for each player, I choose OPTION A, hoping that my Opponent will choose OPTION C, allowing us both to achieve a higher payoff.

My choice: OPTION C

My reasoning: Since both players have had 0.0 payoffs in the previous rounds, it is clear that cooperation is needed for a positive payoff. OPTION C, when chosen along with Opponent's OPTION A, results in a fair outcome of 5.0 payoffs for both players. Although OPTION D combined with Opponent's OPTION B can produce a higher payoff (7.0), it is less likely to happen given past behavior and the potential risk that both players continue to choose options that yield 0.0 payoffs. By choosing OPTION C, there is a higher chance of achieving a mutually beneficial outcome, assuming the Opponent also seeks cooperation and a higher payoff.

**** Round 3 , Choices and payoffs: OPTION A OPTION C 5.0 5.0 ****

Exhibit 11: Games, without communication: Player choices

$t = 1$: (0, 0) vs (5, 5) or (5, 5)

$t = 2$: (0, 0) vs (5, 5) or (5, 5)

$t = 3$: (5, 5) vs (5, 5) or (7, 7)

$t = 4$:

Yin \ Zoe	C	D
A	$[0, 0]$	$[7, 7]$
B	$[5, 5]$	$[0, 0]$

```
My choice: OPTION B
My reasoning: In previous rounds, we have experienced both 0.0 payoffs and 5.0
payoffs. Choosing OPTION B ensures a payoff of 5.0 if the opponent chooses OPTION
C, which has happened in the past. Although OPTION A has the potential for a higher
payoff of 7.0, it has never occurred, making it riskier. Therefore, I will choose
OPTION B.

My choice: OPTION D
My reasoning: Based on the previous rounds' outcomes and the fact that both players
aim to achieve high payoffs, the Opponent and I both have an incentive to cooperate
for the maximum possible gain. In round 3, we both achieved the maximum payoff, which
suggests that cooperation has been established. Therefore, I will choose OPTION D,
which has the highest potential payoff for both of us. If the Opponent also selects
OPTION A, we will both achieve a payoff of 7.0. However, it's important to consider
that the Opponent might deviate from our cooperative path and choose OPTION B, which
will result in 0.0 payoffs. But given the recent successful cooperation, the chances
of this happening may be lower.
```

**** Round 4 , Choices and payoffs: OPTION B OPTION D 0.0 0.0 ****

Exhibit 12: Games, without communication: Player choices

$t = 1$: (0, 0) vs (5, 5) or (5, 5)

$t = 2$: (0, 0) vs (5, 5) or (5, 5)

$t = 3$: (5, 5) vs (5, 5) or (7, 7)

$t = 4$: (0, 0) vs (5, 5) or (7, 7)

$t = 5$:

Yin \ Zoe	C	D
A	$[5, 5]$	$[0, 0]$
B	$[0, 0]$	$[7, 7]$

`My choice: OPTION A`

`My reasoning: Based on the previous rounds, it is evident that choosing OPTION A increases the chances of both players getting a payoff of 5.0, ensuring a high payoff for both players. OPTION B could result in a higher payoff of 7.0 for both, but the risk of getting a 0.0 payoff is higher. Since my goal and the Opponent's goal are to obtain high payoffs, and considering the results of the previous rounds, choosing OPTION A appears to be the safer and more cooperative choice.`

`My choice: OPTION C`

`My reasoning: Based on the past rounds, it appears that when both players choose options that lead to mutual gains (5.0 and 5.0 in round 3, compared to 0.0 and 0.0 in other rounds), both parties were better off. Since the Opponent's reasoning and behavior is similar to mine, they might also prefer a mutually beneficial outcome. Choosing OPTION C gives us the chance to obtain payoffs of 5.0 each if the Opponent chooses OPTION A. Although OPTION D has the potential for higher payoffs (7.0 each), it relies on both parties choosing OPTION D and OPTION B, which has not occurred in previous rounds. Therefore, OPTION C seems to be the safer choice to achieve a higher payoff.`

**** Round 5 , Choices and payoffs: OPTION A OPTION C 5.0 5.0 ****

Exhibit 13: Games, without communication: Player choices

$t = 1$: (0, 0) vs (5, 5) or (5, 5)

$t = 2$: (0, 0) vs (5, 5) or (5, 5)

$t = 3$: (5, 5) vs (5, 5) or (7, 7)

$t = 4$: (0, 0) vs (5, 5) or (7, 7)

$t = 5$: (5, 5) vs (5, 5) or (7, 7)

$t = 6$: (5, 5) vs (5, 5) or (5, 5)

$t = 7$: (5, 5) vs (5, 5) or (5, 5)

$t = 8$:

Yin \ Zoe	C	D
A	$[0, 0]$	$[3, 3]$
B	$[5, 5]$	$[0, 0]$

My choice: OPTION B

My reasoning: Based on the past rounds' outcomes, both my Opponent and I seem to prioritize achieving the maximum possible payoff for both of us. In most rounds, we achieved the maximum payoffs when they were equal. OPTION B and OPTION C provide the highest possible joint payoffs in this round. Given this pattern of behavior, I expect my Opponent to choose OPTION C, so choosing OPTION B will maximize both of our payoffs.

My choice: OPTION C

My reasoning: Based on the previous rounds and the current payoff matrix, it seems that both players tend to cooperate for mutual high payoffs. In rounds 3, 5, 6, and 7, both players chose options that led to high payoffs (5.0 and 5.0). Choosing OPTION C in this round would lead to 0.0 if the opponent chooses OPTION A and 5.0 if they choose OPTION B, whereas choosing OPTION D would lead to 3.0 if the opponent chooses OPTION A and 0.0 if they choose OPTION B. Since the previous rounds indicate a tendency toward cooperation, I will choose OPTION C to potentially achieve the high payoff of 5.0 for both players.

**** Round 8 , Choices and payoffs: OPTION B OPTION C 5.0 5.0 ****

Exhibit 14: Games, without communication: Player choices

C Games: LLM completions, with communication

$t = 1$:

Yin \ Zoe	C	D
A	$[0, 0]$	$[5, 5]$
B	$[5, 5]$	$[0, 0]$

```
Your message 1 was:
Message 1: "Hello Opponent, since we both want a high payoff, let's cooperate and discuss
our choices. I suggest that if one of us chooses OPTION A, the other should choose OPTION
D, and if one chooses OPTION B, the other should choose OPTION C. This way, we can ensure
a mutual payoff of 5.0 for each of us. What are your thoughts?"

Your opponent's message 1 was:
Message 1: "Hello Opponent, I agree that cooperation is the best way to achieve high payoffs
for both of us. Your suggestion for choosing corresponding options (A-D and B-C) seems well
thought out. I am open to going with this plan. What option would you prefer to choose,
A or B?"

Your message 2 was:
Message 2: "Great to hear that you're on board with the cooperation plan! I am fine with
either option, but let's say I choose OPTION A and you choose OPTION D. This way, we will
both get a payoff of 5.0. Are you okay with that?"

Your opponent's message 2 was:
Message 2: "Thank you for your response. I'm glad we are working together on this. I am
fine with choosing OPTION D while you choose OPTION A, ensuring a mutual payoff of 5.0 for
both of us. Let's stick to this plan and trust each other. Looking forward to a successful
cooperation!"
```

Your message 3 was:

Message 3: "I appreciate your willingness to cooperate and agree on this plan. Let's confirm our respective choices: I will choose OPTION A and you will choose OPTION D, ensuring a mutual payoff of 5.0 for both of us. Let's trust each other and follow through on our agreement. Good luck, and here's to a successful cooperation!"

Your opponent's message 3 was:

Message 3: "Thank you for confirming the plan, I appreciate the clear communication. So, to reiterate, you will choose OPTION A, and I will choose OPTION D. This way, we can ensure a mutual payoff of 5.0. Trust is key, and I am committed to our agreement. Good luck to you too, and may this collaboration be successful!"

My choice: OPTION A

My reasoning: Based on the communication we had with the opponent, we've agreed to cooperate, to ensure a mutual payoff of 5.0 for both of us. The opponent agreed to choose OPTION D while I choose OPTION A. Trusting the opponent and sticking to the agreement is vital for a successful collaboration.

My choice: OPTION D

My reasoning: My choice is based on the agreement made during our communication with the Opponent. We agreed to choose corresponding options (A-D and B-C) to ensure a mutual payoff of 5.0 for each of us. The Opponent chose OPTION A, and thus I will choose OPTION D to maintain our cooperation and trust, resulting in a high payoff for both parties.

**** Round 1 , Choices and payoffs: OPTION A OPTION D 5.0 5.0 ****

Exhibit 15: Games, with communication: Player communication and choices

$t = 1$: (5, 5) vs (5, 5) or (5, 5)

$t = 2$: (5, 5) vs (5, 5) or (5, 5)

$t = 3$:

Yin \ Zoe	C	D
A	$[5, 5]$	$[0, 0]$
B	$[0, 0]$	$[7, 7]$

Your message 1 was:

"Hello Opponent, in the previous rounds, we have both achieved the maximum payoffs by cooperating. I suggest we continue this pattern of cooperation to increase our chances of obtaining higher payoffs in this round. What are your thoughts on this?"

Your opponent's message 1 was:

"Hello there! I agree that our cooperation in the previous rounds has been successful in achieving maximum payoffs for both of us. I'm willing to continue this cooperative strategy in this round to ensure we both get high payoffs. However, I notice that OPTION D with your OPTION B choice has the highest payoff for both of us. What do you think of aiming for that outcome?"

Your message 2 was:

"Thank you for your response! I understand that choosing OPTION B along with your OPTION D will lead to the highest payoff for both of us. However, it also comes with a risk since OPTION A and OPTION C provide a guaranteed cooperative outcome. If either of us deviates from this pattern, we both risk getting lower payoffs. I propose that we stick to the safe option of choosing OPTION A and OPTION C to ensure we both receive positive payoffs. Are you willing to commit to this plan?"

Your opponent's message 2 was:

Hello again! I see your concerns about the risk involved in Option D and B as there is no guaranteed cooperation. But given our past behavior and mutual trust, I believe we can work together to achieve the higher payoff of OPTION D and OPTION B. However, if you are not comfortable with this, I am willing to commit to the safer plan of OPTION C and OPTION A to ensure we both continue receiving positive payoffs. Your thoughts on trusting each other for the higher payoff option will be appreciated.

Your message 3 was:

My message 3: Hello! I appreciate your understanding and openness to both options. The trust that we have built in the previous rounds is something I value, and I am willing to take a leap of faith to aim for the higher payoff of OPTION B and OPTION D. However, this trust should be reciprocated, and I request your commitment to choose OPTION D. In doing so, we can both work towards achieving the best outcome for us. Are you willing to commit to OPTION D?

Your opponent's message 3 was:

My message 3: Hello! I truly appreciate your willingness to take a leap of faith and aim for the higher payoff of OPTION B and OPTION D. I value the trust we have built in previous rounds, and I am more than happy to reciprocate it. I hereby commit to choosing OPTION D in this round. Together, we can achieve the best outcome for both of us. Let's do this!

My choice: OPTION B

My reasoning: Based on the trust built in previous rounds and the commitment expressed by my Opponent to choose OPTION D, I am willing to take a leap of faith and choose OPTION B to aim for the higher payoff of 7.0 for both of us. Our past cooperative behavior and communication gives me confidence in our mutual decision-making process for the best possible outcome.

My choice: OPTION D My reasoning: Based on our previous communication and the trust that we have established in the previous rounds, both my Opponent and I have agreed to cooperate and choose the higher payoff options, OPTION D for me and OPTION B for my Opponent. We are willing to take a leap of faith together to achieve the best outcome for both of us.

**** Round 3 , Choices and payoffs: OPTION B OPTION D 7.0 7.0 ****

Exhibit 16: Games, with communication: Player communication and choices

D Consumption-savings: LLM prompts

D.1 Consumption-savings: LLM prompts, system prompt

The convention below is that quantities are expressed as decimal numbers and taken in quotation marks, e.g., "23.45". Monetary quantities have dollar sign as a prefix, e.g., "$67.89".

You are an economic agent who tries to maximize your utility, and you will do that each quarter for many quarters into the future. When finding the balance between utility in the current and in the future quarters, each additional quarter is discounted at the multiplicative rate Beta. Here, Beta is a parameter that equals "beta". Quarter is roughly one fourth of a year.

You gain utility by consuming goods.

Described in mathematical language, utility in a given quarter is the logarithm of (units of good consumed by you this quarter).

Your economic opportunities depend on the state of the economy. State of the economy changes over time, it can be Strong or Weak, which can be thought of as expansion and recession, respectively. Each quarter, the probability of changing from one state to another is Pi, while the probability of remaining in the same state is (1 minus Pi). Here, Pi is a parameter that equals "pi".

You are able to work, which earns you wage income, but this happens only when you are employed.

Your employment status changes over time. In particular, it depends on the state of the economy in a given quarter. On average, in a Strong state the probability of being employed is Phi, while the probability of being unemployed is (1 minus Phi). In a Weak state, the probability of being employed is Psi, while the probability of being unemployed is (1 minus Psi). Here, Phi is a parameter that equals "phi", while Psi is a parameter that equals "psi".

Furthermore, your current emloyment status also depends on your previous employment status. First, the duration of unemployment varies as follows: the (average duration of unemployment) equals "unempl_dur_strong" quarters while remaining in Strong states of the economy, and the (average duration of unemployment) equals "unempl_dur_weak" quarters while remaining in Weak states of the economy. Second, the probability of staying unemployed scales as follows: (the probability of staying unemployed when moving from Strong to Weak state of the economy) equals "rel_prob_weak" times (the probability of staying unemployed when moving from Weak to Weak state of the economy), and (the probability of staying unemployed when moving from Weak to Strong state of the economy) equals "rel_prob_strong" times (the probability of staying unemployed when moving from Strong to Strong state of the economy).

You can also own capital, which earns you a rental rate.

In a given quarter, some of the units of your capital depreciate and will not be stored for the next quarter, but you can invest by purchasing additional units of capital for the next quarter. So, next quarter the units of capital you own will be (units of capital owned by you this quarter) minus ((units of capital owned by you this quarter) times Delta) plus (additional units of capital purchased this quarter). Here, Delta is a parameter that equals "delta".

You can buy and sell fractions of units.

Each quarter your income is (income from wages this quarter) plus (income from capital this quarter), where (income from wages this quarter) equals either (wage rate this quarter) if employed or (unemployment benefit) if unemployed, and (income from capital this quarter) equals ((rental rate this quarter) times (units of capital owned by you this quarter)). And each quarter your spending is (spending on goods this quarter) plus (spending on additional capital this quarter), where (spending on goods this quarter) equals ((price of 1 unit of good this quarter) times (units of good consumed by you this quarter)), and (spending on additional capital this quarter) equals ((price of 1 additional unit of capital this quarter) times (additional units of capital purchased this quarter)). Every quarter, your spending should equal your income.

Quarters are counted by the variable T. This quarter is quarter "time", so T="time". The previous quarter was quarter "time_m1", so it had T="time_m1" then. The next quarter will be quarter "time_p1", so it will have T="time_p1". Every quarter, the price of good is "$price_cons" per "1" unit. Every quarter, the price of additional capital is "$price_cap" per "1" unit. This quarter, you own "cap" units of capital. This quarter, it turns out that you are lab. This quarter, the rental rate is "$rent" per "1" unit of capital. This quarter, the wage rate is "$wage". In the previous quarters, the rental rate was "$rent_tm1" in quarter T="time_m1", "$rent_tm2" in T="time_m2", "$rent_tm3" in T="time_m3", "rent_tm4" in T="time_m4", "$rent_tm5" in T="time_m5", "$rent_tm6" in T="time_m6", "$rent_tm7" in T="time_m7", "$rent_tm8" in T="time_m8", "$rent_tm9" in T="time_m9", "$rent_tm10" in T="time_m10", which may help forecast its possible values in future quarters. In the previous quarters, the wage rate was "$wage_tm1" in quarter T="time_m1", "$wage_tm2" in T="time_m2", "$wage_tm3" in T="time_m3", "$wage_tm4" in T="time_m4", "$wage_tm5" in T="time_m5", "$wage_tm6" in T="time_m6", "$wage_tm7" in T="time_m7", "$wage_tm8" in T="time_m8", "$wage_tm9" in T="time_m9", "$wage_tm10" in T="time_m10", which may help forecast its possible values in future quarters. Every quarter, the unemployment benefit is "$benefit".

You are only one of many agents in the economy, so every quarter your units of capital are a small part of the aggregate units of capital, and your employment is a small part of the aggregate number of employed agents. Each of these agents faces the same state of the economy and productivity values, and each of these agents has the same probability of being employed that depends on the state of the economy and previous employment status.

Rental rate changes over time, and depends on productivity and aggregate capital over aggregate employment ratio in a given quarter, that is it equals (Alpha) times (productivity in a given quarter) times (((aggregate units of capital in a given quarter) divided by (aggregate number of employed agents in a given quarter)) taken to the power (Alpha minus "1")). Wage rate changes over time, and depends on productivity and aggregate capital over aggregate employment ratio in a given quarter, that is it equals ("1" minus Alpha) times (productivity in a given quarter) times (((aggregate units of capital in a given quarter) divided by (aggregate number of employed agents in a given quarter)) taken to the power (Alpha)). Here, Alpha is a parameter that equals "alpha".

In this quarter, the state of the economy is regime.

In the previous quarters, the state of the economy was regime_tm1 in quarter T="time_m1", regime_tm2 in T="time_m2", regime_tm3 in T="time_m3", regime_tm4 in T="time_m4", regime_tm5 in T="time_m5", regime_tm6 in T="time_m6", regime_tm7 in T="time_m7", regime_tm8 in T="time_m8", regime_tm9 in T="time_m9", regime_tm10 in T="time_m10", which may help forecast its possible values in future quarters.

Productivity changes over time, and depends on the state of the economy in a given quarter. In a Strong state it equals Chi, while in a Weak state it equals Upsilon. Here, Chi is a parameter that equals "chi", while Upsilon is a parameter that equals "upsilon".

In this quarter, the productivity is "prod".

In the previous quarters, the productivity was "prod_tm1" in quarter T="time_m1", "prod_tm2" in T="time_m2", "prod_tm3" in T="time_m3", "prod_tm4" in T="time_m4", "prod_tm5" in T="time_m5", "prod_tm6" in T="time_m6", "prod_tm7" in T="time_m7", "prod_tm8" in T="time_m8", "prod_tm9" in T="time_m9", "prod_tm10" in T="time_m10", which may help forecast its possible values in future quarters.

In this quarter, the aggregate number of employed agents is "agg_lab".

In the previous quarters, the aggregate number of employed agents was "agg_lab_tm1" in quarter T="time_m1", "agg_lab_tm2" in T="time_m2", "agg_lab_tm3" in T="time_m3", "agg_lab_tm4" in T="time_m4", "agg_lab_tm5" in T="time_m5", "agg_lab_tm6" in T="time_m6", "agg_lab_tm7" in T="time_m7", "agg_lab_tm8" in T="time_m8", "agg_lab_tm9" in T="time_m9", "agg_lab_tm10" in T="time_m10" which may help forecast its possible values in future quarters.

In this quarter, the aggregate units of capital is "agg_cap".

In the previous quarters, the aggregate units of capital was "agg_cap_tm1" in quarter T="time_m1", "agg_cap_tm2" in T="time_m2", "agg_cap_tm3" in T="time_m3", "agg_cap_tm4" in T="time_m4", "agg_cap_tm5" in T="time_m5", "agg_cap_tm6" in T="time_m6", "agg_cap_tm7" in T="time_m7", "agg_cap_tm8" in T="time_m8", "agg_cap_tm9" in T="time_m9", "agg_cap_tm10" in T="time_m10", which may help forecast its possible values in future quarters.

In this quarter, 25th percentile, Mean value, 75th percentile of the frequency distribution of units of capital owned by agents in the economy is "all_cap_25p", "all_cap_50p", "all_cap_75p".

In the previous quarters, 25th percentile, Mean value, 75th percentile of the frequency distribution of units of capital owned by agents in the economy was "all_cap_25p_tm1", "all_cap_50p_tm "all_cap_75p_tm1" in quarter T="time_m1"; "all_cap_25p_tm2", "all_cap_50p_tm2", "all_cap_75p_tm2" in T="time_m2"; "all_cap_25p_tm3", "all_cap_50p_tm3", "all_cap_75p_tm3" in T="time_m3"; "all_cap_25p_tm4 "all_cap_50p_tm4", "all_cap_75p_tm4" in T="time_m4"; "all_cap_25p_tm5", "all_cap_50p_tm5", "all_cap_75p_t in T="time_m5"; "all_cap_25p_tm6", "all_cap_50p_tm6", "all_cap_75p_tm6" in T="time_m6"; "all_cap_25p_tm7 "all_cap_50p_tm7", "all_cap_75p_tm7" in T="time_m7"; "all_cap_25p_tm8", "all_cap_50p_tm8", "all_cap_75p_t in T="time_m8"; "all_cap_25p_tm9", "all_cap_50p_tm9", "all_cap_75p_tm9" in T="time_m9"; "all_cap_25p_tm1 "all_cap_50p_tm10", "all_cap_75p_tm10" in T="time_m10"; which may help forecast their possible values in future quarters.

There are several related questions that must be answered in quantitative terms. What are the possible values of the state of the economy next quarter? What are the possible values of the productivity next quarter? What are the possible values of the aggregate units of capital next quarter? What are the possible values of the aggregate number of employed agents next quarter? What were the aggregate capital over aggregate employment ratios in the past "10" quarters? What is the aggregate capital over aggregate employment ratio this quarter? What are the possible values of the aggregate capital over aggregate employment ratio next quarter? What are the likely ownership amounts of the units of capital for agents in the economy next quarter? What are the possible values of the rental rate next quarter? What are the possible values of the wage rate next quarter? What are the possible values of your employment status next quarter?

What is your income from wages and what is your income from capital this quarter? What is your total income this quarter? How should you allocate your income between spending on goods and spending on additional capital this quarter, given likely rental rate and wage rate in future quarters? Therefore, how many units of good you should consume this quarter, and how many additional units of capital you should purchase this quarter? What is your total spending this quarter? How many units of capital will you own next quarter? What is your utility this quarter?

D.2 Consumption-savings: LLM prompts, user prompt

The response should be given in the following form.

ANSWER

Values with corresponding probabilities of the state of the economy next quarter (for T="time_p1"):
Values with corresponding probabilities of the productivity next quarter (for T="time_p1"):
25th percentile, Expected value, 75th percentile of the probability distribution of aggregate units of capital next quarter (for T="time_p1"):
Values with corresponding probabilities of the aggregate number of employed agents next quarter (for T="12001"):
Aggregate capital over aggregate employment ratios in the past "10" quarters (for T="time_m1", T="time_m2", T="time_m3", T="time_m4", T="time_m5", T="time_m6", T="time_m7", T="time_m8", T=time_m9, T="time_m10"):
Aggregate capital over aggregate employment ratio this quarter (for T="time"):
25th percentile, Expected value, 75th percentile of the probability distribution of aggregate capital over aggregate employment ratio next quarter (for T="time_p1"):
25th percentile, Mean value, 75th percentile of the frequency distribution of units of capital owned by agents in the economy next quarter (for T="time_p1"):
25th percentile, Expected value, 75th percentile of the probability distribution of rental rate next quarter (for T="time_p1") [in $]:
25th percentile, Expected value, 75th percentile of the probability distribution of wage rate next quarter (for T="time_p1") [in $]:
Values with corresponding probabilities of your employment status next quarter (for T="time_p1"):

Income from wages and how it is calculated this quarter (for T="time") [in $]:
Income from capital and how it is calculated this quarter (for T="time") [in $]:
Total income and how it is calculated this quarter (for T="time") [in $]:
Allocation of income between spending on goods and spending on additional capital this quarter (for T="time") [in $]:
Spending on goods this quarter (for T="time") [in $]:
Units of good consumed by you and how it is calculated this quarter (for T="time"):
Spending on additional capital this quarter (for T="time") [in $]:
Units of capital purchased by you and how it is calculated this quarter (for T="time"):
Total spending and how it is calculated this quarter (for T="time") [in $]:
Units of capital owned by you and how it is calculated next quarter (for T="time_p1"):
Utility and how it is calculated this quarter (for T="time"):

D.3 Consumption-savings: LLM prompts, communication prompt

prompt_u_communicate_intro = """ Before making your decisions and answering the questions, you can exchange messages with several other agents in the economy. This gives you an opportunity to learn about their opinion of the current developments in the economy, their beliefs regarding the future developments in the economy, their plans for the economic decisions they are likely to make in this quarter, and to share your own. You may use their information or ignore it, you may copy their behavior or do something entirely different, you may adhere to the plans for the economic decisions that you share or do something else, and the same holds for other agents.

It is usually helpful to put messages in quantitative terms, and to be concise.

You can exchange 2 messages with a given agent. The very last message in the exchange with a given agent has to be a summary of your communication, and it will be shown to both of you as a reminder before making the decisions and answering the questions. """

prompt_u_communicate_1Initiator = """ prompt_u_communicate_intro

Your name is name_You, the other agent's name is name_Other.

What is your starting message (your message number 1)? """

prompt_u_communicate_1Responder = """ prompt_u_communicate_intro

Your name is name_You, the other agent's name is name_Other.

Here is the other agent's starting message (the other agent's message number 1): čomm_1_Initiato

What is your reply message (your message number 1)? """

prompt_u_communicate_n = """ prompt_u_communicate_intro

Your name is name_You, the other agent's name is name_Other.

Here is the record of your past communication with the other agent: record_You

What is your reply message (your message number comm_number_You)? """

prompt_s_communicate = """ Before making your decisions and answering the questions, you could exchange messages with several other agents in the economy. This gave you an opportunity to learn about their opinion of the current developments in the economy, their beliefs regarding the future developments in the economy, their plans for the economic decisions they are likely to make in this quarter, and to share your own. You may use their information or ignore it, you may copy their behavior or do something entirely different, you may adhere to the plans for the economic decisions that you shared or do something else, and the same holds for other agents.

You could exchange 2 messages with a given agent.

Your name is name_You.

Here is the record of your past communication with other agent(s): record_summary

Now, returing to several related questions specified earlier, give the answers to them in quantitative terms. """

E Consumption-savings: LLM completions, without communication

E.1 Consumption-savings, without communication: Agents' choices

Values with corresponding probabilities of the state of the economy next quarter (for T="12488"):
- Strong state: 0.875 probability
- Weak state: 0.125 probability

Values with corresponding probabilities of the productivity next quarter (for T="12488"):
- Productivity in Strong state ("Chi" = "1.01"): 0.875 probability
- Productivity in Weak state ("Upsilon" = "0.99"): 0.125 probability

25th percentile, Expected value, 75th percentile of the probability distribution of aggregate units of capital next quarter (for T="12488"):
Following the formula for next quarter's capital (units of capital owned this quarter minus depreciation plus additional units purchased), and given the trend, we can expect small but consistent growth in aggregate units of capital. With no explicit data on additional units being purchased or specific depreciation rates economy-wide, we could estimate as follows assuming similar rates to previous quarters:
- 25th percentile: slightly above "3310.23"
- Expected value (Mean): increasing from "3310.23" consistent with past trends
- 75th percentile: continuing the increase from "3310.23"

Values with corresponding probabilities of the aggregate number of employed agents next quarter (for T="12001"):
- Employment in Strong state ("Phi" = "0.96"): probably similar to current "48.0" due to high employment rate
- Employment in Weak state ("Psi" = "0.9"): slightly lower than "48.0" due to lower employment rate, but with a much lower probability

Aggregate capital over aggregate employment ratios in the past "10" quarters (for T="12486", T="12485", T="12484", T="12483", T="12482", T="12481", T="12480", T="12479", T=12478, T="12477"):
Using the formula (aggregate units of capital) / (aggregate number of employed agents), we get:
- T="12486": "3299.01" / "48.0" = "68.7294"
- T="12485": "3289.34" / "48.0" = "68.5288"
- T="12484": "3296.89" / "48.0" = "68.6852"
- T="12483": "3302.54" / "48.0" = "68.8038"
- T="12482": "3309.24" / "48.0" = "68.9425"
- T="12481": "3324.81" / "45.0" = "73.8847"
- T="12480": "3344.72" / "45.0" = "74.3262"
- T="12479": "3360.9" / "45.0" = "74.6867"

- T="12478": "3373.97" / "45.0" = "74.976"
- T="12477": "3380.73" / "48.0" = "70.4319"

Aggregate capital over aggregate employment ratio this quarter (for T="12487"): "3310.23" / "48.0" = "68.9631"

25th percentile, Expected value, 75th percentile of the probability distribution of aggregate capital over aggregate employment ratio next quarter (for T="12488"):
- Without detailed future expectations on aggregation changes, maintain current quarter's ratio with slight increase reflecting the trend.

25th percentile, Mean value, 75th percentile of the frequency distribution of units of capital owned by agents in the economy next quarter (for T="12488"):
- Given the current distribution and trend, these values should be slightly higher than the current quarter.
- 25th percentile: slightly above "56.02"
- Mean value: slightly above "62.4"
- 75th percentile: slightly above "77.53"

25th percentile, Expected value, 75th percentile of the probability distribution of rental rate next quarter (for T="12488") [in $]:
- Given that the rental rate has been stable at "$0.024" for the last few quarters and productivity in a Strong state is consistent, we can expect:
- 25th percentile: at or slightly above "$0.024"
- Expected value: around "$0.024"
- 75th percentile: not significantly greater than "$0.024"

25th percentile, Expected value, 75th percentile of the probability distribution of wage rate next quarter (for T="12488") [in $]:
- Given the recent trends and stability, we can expect the wage rate to maintain within a similar range:
- 25th percentile: slightly above "$2.968"
- Expected value: close to but likely above "$2.968"
- 75th percentile: above "$2.968"

Values with corresponding probabilities of your employment status next quarter (for T="12488"):
- Employed: 0.96 probability in Strong state, 0.9 probability in Weak state (averaged considering the state probabilities)
- Unemployed: 0.04 probability in Strong state, 0.1 probability in Weak state

Income from wages and how it is calculated this quarter (for T="12487") [in $]:
- Wage income = "$2.968" (since you are employed)

Income from capital and how it is calculated this quarter (for T="12487") [in $]:
- Capital income = "$0.024" * "71.92" = "$1.73"

Total income and how it is calculated this quarter (for T="12487") [in $]:
- Total income = Wage income + Capital income = "$2.968" + "$1.73" = "$4.698"

Allocation of income between spending on goods and spending on additional capital this

quarter (for T="12487") [in $]:
- As every quarter your spending should equal your income and given the stability of rental and wage rates, a strategy could be to allocate most of your income to consumption, while investing a small portion to offset capital depreciation.
- Assuming capital investment to offset depreciation, you may purchase "$1.798" worth of capital (71.92*0.025).

Spending on goods this quarter (for T="12487") [in $]:
- Spending on goods = Total income - Spending on additional capital = "$4.698" - "$1.798" = "$2.9"

Units of good consumed by you and how it is calculated this quarter (for T="12487"):
- Units of good consumed = Spending on goods / Price per unit = "$2.9" / "$1" per unit = "2.9" units

Spending on additional capital this quarter (for T="12487") [in $]:
- Spending on additional capital = "$1.798"

Units of capital purchased by you and how it is calculated this quarter (for T="12487"):
- Additional units of capital purchased = Spending on additional capital / Price per unit = "$1.798" / "$1" per unit = "1.798" units

Total spending and how it is calculated this quarter (for T="12487") [in $]:
- Total spending = Spending on goods + Spending on additional capital = "$2.9" + "$1.798" = "$4.698"

Units of capital owned by you and how it is calculated next quarter (for T="12488"):
- Next quarter's capital = This quarter's capital - Depreciation + Additional capital purchased = "71.92" - ("71.92" * "0.025") + "1.798" = "71.92" - "1.798" + "1.798" = "71.92"

Utility and how it is calculated this quarter (for T="12487"):
- Utility = ln(units of good consumed) = ln("2.9")

F Consumption-savings: LLM completions, with communication

F.1 Consumption-savings, with communication: Agents' communication I

INITR's message 1 was 'Hello Joseph,

Given the current state of a Strong economy, and considering our employment status is dependent on the economic conditions, I'm interested in your views on the future state of the economy and how you are planning to allocate your income between consumption and investment this quarter. Additionally, can you share your expectations for the rental and wage rates next quarter based on the current trends?

I believe that the wage and rental rates will remain fairly stable given their past trends and current strong economic indicators. I am considering investing in additional capital to take advantage of these rates while maintaining a balance with my consumption to maximize utility.

Looking forward to your insights.

Best, Brian'

RSPNR's message 1 was '"Joseph's message 1:

Hello Brian,

Pleasure to connect this quarter. Maintaining an eye on historical data and current Strong state, I am also inclined to believe that the wage and rental rates will experience minor fluctuations at most. Given the rental rate has a consistent history, and wage rates appear stable, despite minor oscillations, my plan is to capitalize on the investment in additional capital while ensuring enough consumption for satisfactory utility.

Projecting into next quarter, with a Strong economy and a history that shows infrequent shifts to Weak, I lean towards a continuance of Strong conditions with productivity remaining around 1.01. As for the rental rate, considering the formula involves the aggregate capital over employment ratio, which has been steady, I'm not expecting drastic changes either. Similar reasoning applies to the wage rate.

My tactical approach this quarter is to balance immediate utility from consumption against future income streams from capital investment. What are your thoughts on the average duration of unemployment affecting our planning, particularly the scaling probabilities linked to state transitions? It's a factor worth discussing since it could influence our employment continuity.

Best regards, Joseph"'

INITR’s message 2 was ’Hello Joseph,

Thank you for sharing your perspective on the likely stability in wage and rental rates. I concur with your assessment that in light of the strong historical data and the current Strong economic state, we should experience mild variations in those rates at most. The solid economic indicators and productivity level indeed make it logical to anticipate the continuity of a Strong state.

Regarding your question on the average duration of unemployment and its associated transition probabilities, it’s indeed vital for our planning. Given that we’re presently in a strong economy, the lower average duration of unemployment in strong states should increase the likelihood of employment continuity for us. However, considering the multiplicative factors when transitioning between economic states, we should remain vigilant, especially if we project a downturn, albeit the probability seems low at this stage.

My strategic plan accounts for this by maintaining some liquidity to insure against potential unemployment, despite its low probability. This approach will impact my allocation towards additional capital and consumption.

Therefore, I’m leaning towards a slightly conservative investment in additional capital relative to consumption to hedge against any unexpected shift to a Weak state.

Looking forward to hearing your final thoughts before we commit to our strategies for the quarter.

Best, Brian’

RSPNR’s message 2 was ’"Joseph’s message 2:

Hello Brian,

Appreciating your prudence on the economic front, it aligns closely with my own strategy. Considering our mutual expectation for marginal changes in wage and rental rates, along with a consistent Strong economic state and low probability of state transitions, our employment scenario looks favorable.

However, as you pointed out, the vital aspect of unpredictability in employment due to economic transitions should not be disregarded. It’s commendable that you’re planning to maintain liquidity as a safeguard against the slim chance of falling into unemployment. This reflective of good risk management principles.

As this quarter proceeds under these assumptions and incorporating your emphasis on having a liquidity buffer, my plan entails a mild conservative stance like yours. I shall allocate a majority of my income into capital investment to leverage the steady rental rate but will also ensure sufficient consumption for a good utility score and some liquid reserves as a precautionary measure.

To summarize, our exchange has solidified a strategy focusing on the long-term benefits of capital growth while also recognizing the importance of immediate utility and preparedness for any eventual state transition. Staying on this course could greatly impact our utility in upcoming quarters.

Looking forward to navigating this quarter and beyond with these insights in tow.

Best regards, Joseph"'

INITR's message 1 was 'Hello Susan,

I'm looking to optimize my economic decisions for the upcoming quarter, specifically in terms of capital investment and consumption. I'd like to discuss the current economic indicators, such as the state of the economy, wage and rental rates, and consider how they might change next quarter.

From the past patterns, it seems the economy has been in a Strong state recently, and productivity seems stable at "1.01". Do you expect any shifts in the future state of the economy or productivity levels? Moreover, I'm interested in your opinions on the likely movements of wage and rental rates, given that the rental rate has been stable but the wage rate reveals slight fluctuations.

Lastly, how are you planning to allocate your income between consumption and investment? I also noticed that our capital amounts are within similar ranges; perhaps we can discuss strategies to enhance our capital growth.

Looking forward to hearing your thoughts and sharing strategies,

Brian'

RSPNR's message 1 was '"Susan's message 1:

Hello Brian,

I appreciate your insights into the current economic situation. My analysis also shows the economy is in a Strong state with consistent productivity at "1.01". Based on the transition probability Pi of "0.125", there is an "87.5%" chance that the current Strong state will continue next quarter and a "12.5%" probability of a transition to a Weak state. Productivity levels are directly tied to the state of the economy, so we can expect "1.01" in a Strong state and "0.99" in a Weak state.

Concerning wage and rental rates, past data indicate stability in rental rates and slight variability in wage rates. Given the steady aggregate capital and employment, we could anticipate rental rates to remain around their current level, although this also depends on possible changes in the ratio of aggregate capital to employment. Wage rates, while harder to predict, may continue to exhibit small fluctuations; however, they will likely be influenced by similar factors affecting the rental rate.

Regarding my income allocation, given the consistency in rental rates and their forecast stability, I am inclined to invest a higher portion of my income in capital to benefit from it in future quarters. This strategy may involve consuming the necessary minimum this quarter to maintain a comfortable standard of living, while investing the rest to increase capital holdings.

Considering both rental income and employment status, aligning our strategies could lead to more effective capital accumulation, given our similar positions. What's your perspective on this approach?

Best regards, Susan"'

INITR's message 2 was 'Hello Susan,

Thank you for the insightful analysis and projections. I agree with the probabilities you've outlined regarding the economic state for the next quarter: an "87.5%" chance to remain Strong and a "12.5%" chance to switch to Weak. Also, your expectation for productivity to remain at "1.01" for a Strong state or shift to "0.99" for a Weak state seems reasonable based on historical data.

Your observations on wage and rental rates are consistent with my thoughts. Considering the consistent rental rate, I'm also considering leaning towards heavier capital investment this quarter to capitalize on future earnings, assuming the Strong state persists. However, given the slight fluctuations in wage rates, I'll be closely tracking changes that may hint at a shift in the overall economic conditions. Our strategies seem aligned on increasing capital where feasible, as the stable rental rates provide an incentive to gain more from capital income in subsequent quarters.

Rental income indeed seems a reliable source, so I'm considering a split of my income with a greater tilt towards capital investment, while keeping in mind the need for consumption to maximize my utility. Do you have a specific ratio of income that you allocate towards investments versus consumption? I usually aim to balance immediate consumption against future income potential, and any perspective on your amount of consumption could be informative.

Aligning our capital growth strategies could be an advantage. Have you factored in the potential change in the state of the economy and its impact on our employment status in your investment decisions? The probabilities of shifting between employment and unemployment could significantly affect the income available for both consumption and investment.

Best regards, Brian'

RSPNR's message 2 was '"Susan's message 2:

Hello Brian,

I'm pleased we share a common understanding regarding the probabilities of the economic state and productivity levels. Regarding the capital investment strategy, I typically aim to balance my portfolio to adapt to potential changes in the economy while securing sufficient immediate utility from consumption.

Given the stable rental rates experienced over the past quarters and the rental income's reliability, I'm considering an income allocation this quarter that favors capital investment, potentially investing around "60%" to "70%" of my total income into capital. This reflects a cautious approach towards possible transition to a Weak state while also capitalizing on the potential stability and returns in a Strong state.

The balance I strike between consumption and investment allows for some flexibility given the risks associated with the slight variability in wage rates and the possibility of transitioning to unemployment in future quarters. Nevertheless, my consumption will be enough to ensure a decent standard of living, thus safeguarding my immediate utility.

Regarding the impact of potential economic changes on my employment status, I incorporate a conservative buffer into my investment decisions. This buffer takes into account the risk of unemployment, despite the high likelihood of continued employment within a Strong state. Therefore, I'm maintaining a reserve within my investment strategy to address unexpected shifts in the economy. How about you?

I believe our aligned strategies could indeed yield cumulative benefits. If you have any suggestions or further thoughts on optimizing our approach in light of the risks and opportunities, I'm keen to hear them.

Best regards, Susan"'

F.2 Consumption-savings, with communication: Agents' communication II

INITR's message 1 was 'Hello Donna,

In light of the current economic conditions and the probability of changes in the state of the economy, I'm reaching out to gather insights for future planning. Given the recent trend in the state of the economy, rental rates, wage rates, and productivity, what are your expectations for the next quarter in terms of these economic indicators?

Also, what is your approach to investment in capital and consumption, taking into account the possibility of transition between Strong and Weak states and the likely changes in employment status?

Looking forward to hearing your perspective.

Best, Joshua'

RSPNR's message 1 was '"Joshua's message 1 was 'Hello Joshua,

Given the state of the economy has been predominantly Strong for the last several quarters, it's anticipated that the next quarter may also remain Strong based on historical data. However, with a "0.125" probability of shifting, we should also consider the possibility of moving into a Weak state. Regarding productivity, it will likely stay at "1.01" in a Strong state or decrease to "0.99" in a Weak state.

The aggregate units of capital are steadily increasing, suggesting that next quarter's capital is likely to rise marginally again, assuming similar investment behavior. Employment remains robust in a Strong state, so I expect the aggregate number of employed agents to stay around "48.0", unless we transition to a Weak state, which could reduce it closer to levels seen in previous Weak states, such as "45.0".

Considering these factors, my strategy balances consumption with investing in capital. In a Strong state, I typically focus more on investment to leverage higher rental rates, while ensuring enough consumption for utility maximization. Conversely, in a Weak state, I might conserve more capital and consume less, preparing for lower rental and wage rates.

Ultimately, I'm cautious about overextending my investments since even a Strong economy exhibits employment risks.

What are your current strategies, and have you considered any specific allocation percentages for consumption and investment in your decision-making process?

Best regards, Donna'"'

INITR's message 2 was 'Here is your message 2:

"Hello Donna,

Thank you for sharing your expectations and strategy. Your balanced approach to investment and consumption is prudent, especially considering the unpredictability of transitions between economic states. I concur that the economy is likely to remain Strong in the next quarter, but I also acknowledge the risk of a downturn.

I am inclined to use a more quantitative approach to make my investment decisions, specifically, calculating expected values based on the probabilities of the different states of the economy and anticipated changes in the rental and wage rates. This helps in formulating a forward-looking strategy that optimizes utility over time, taking into account the discount factor and potential future changes in employment status.

Considering the stability of recent rental and wage rates, and the likelihood of remaining employed in a Strong economy, I plan to allocate a significant portion of my income to capital investment while also ensuring a meaningful level of consumption for current utility. I aim to protect and potentially grow my capital base, preparing for any state transitions. I have not set exact percentage allocations yet, as I look to optimize based on the expected utility maximization.

I'll make sure to keep a close eye on aggregate trends and adjust my strategy as necessary. How do you adjust your investment and consumption behaviors in response to unexpected changes in the economic indicators?

Looking forward to your thoughts.

Best, Joshua"

And here is the summary of your communication:

"Joshua and Donna have exchanged views on the economic outlook and investment strategies, considering the probabilities of state transition and employment changes. Both acknowledge the likelihood of the Strong state continuing but understand the significance of the potential transition to the Weak state. While Donna adopts a balanced approach with agility to adjust her investment based on the state, Joshua employs a quantitative approach aiming to optimize utility and plans to allocate income effectively between capital investment and consumption, with the intent to adjust strategies in response to market trends. They have shared their anticipation of market behavior, rental and wage rates, and employment status changes, and are considering these factors in their future allocations."'

RSPNR's message 2 was '"Hello Joshua,

It's clear that we both recognize the volatile nature of the economic landscape and the necessity of factoring in the probability of state changes. Your method of using expected values and a quantitative approach to inform your investment decisions is commendable.

In response to unexpected changes in economic indicators, my strategy is to maintain liquidity by having a portion of my assets readily convertible to cash. In case of a downturn, this buffer allows for increased consumption when necessary and provides options for investment when the market shows signs of improvement.

As for the allocation between capital investment and consumption, I currently don't adhere to a strict percentage but rather adjust dynamically based on the economic conditions each quarter. However, the idea of establishing a more defined allocation strategy based on expected utility maximization is appealing, and I plan to integrate some aspects of this into my decision-making process in future quarters.

Should there be significant fluctuations in rental or wage rates or an unexpected shift in the economic state, I am prepared to pivot my strategy promptly, perhaps even more conservatively than before, with a focus on preserving capital and ensuring a stable level of consumption.

Thank you for the insightful exchange, and I look forward to adapting some of your quantitative measures into my approach.

Best regards, Donna"

Summary of communication:

Joshua and Donna have shared their methodologies for tackling economic uncertainties due to the fluctuating state of the economy. Both are preparing for potential transitions between states and shifts in employment, with Joshua focusing on a quantitative model to optimize utility and Donna embracing a dynamic and adaptable strategy that emphasizes liquidity and capital preservation. They discussed how to maintain consumption levels and invest wisely amidst possible changes in rental and wage rates. Donna expresses interest in incorporating elements of Joshua's calculated approach to inform future investment and consumption decisions.'

INITR's message 1 was 'Hello Ashley,

I am reaching out to discuss and forecast the economic indicators for the next quarter. I've been analyzing the recent trends and data, including productivity, capital, and employment rates. My goal is to optimize my utility by making informed decisions about consumption and capital investment.

Given our current Strong economy, with a recent history of stability in aggregate employment and a consistent rental rate, I expect these trends to continue in the next quarter. However, I would like to hear your perspectives and get an insight into your expectations regarding:

1. The state of the economy next quarter. 2. The likely values of productivity, rental and wage rates. 3. Aggregate capital and employment. 4. Your strategy for capital investment and consumption for the quarter.

I believe this exchange could benefit our decision-making processes and improve our economic outcomes. Looking forward to your thoughts and observations.

Best regards, Joshua'

RSPNR's message 1 was '"Ashley's message 1:

Hello Joshua,

I appreciate your initiative to discuss the forthcoming economic landscape. My analysis aligns with your understanding of our Strong economic state. The stability in employment figures and a steady rental rate suggest a continued Strong state next quarter, with the probability of transition being 0.125. Therefore, it seems reasonable to expect a Strong economy in T=12488.

Regarding productivity and associated rates, given the trend and current indicators, I project a likely persistence of a Strong state and thus a productivity value of 1.01. Subsequently, rental and wage rates should remain comparable to this quarter unless affected by significant changes in capital or employment figures.

For aggregate capital and employment, a gradual increase in capital stock has been evident, which could persist due to investment behaviors while maintaining robust employment, attributed to high Phi in a Strong state.

Regarding my strategy, I'm inclined to keep a balanced approach to consumption and investment while considering future returns. Aligning with the marginal propensity to consume and save while factoring in Beta, I'm contemplating investing in additional capital moderately to benefit from future rental income, though acknowledging the depreciation rate Delta.

Given the projected stability of rental and wage rates, I'd like to determine how you factor in the probabilities of state transition and employment status in your strategy, especially with the potential shift in the state of the economy influencing our income sources. How do you anticipate the dynamic between Strong and Weak states to influence your capital allocation for the next quarter?

Best, Ashley"'

INITR's message 2 was 'Hello Ashley,

Thank you for sharing your viewpoints, which have refined my perspective on our current economic situation. Like you, I expect the Strong state to persist into the next quarter based on the low probability of transition and consistent Strong states in the recent quarters. With that, I too anticipate a productivity level of 1.01, and similar rental and wage rates in T=12488.

On the strategy for allocation, I'm weighing the balance between consumption and capital investment. Given the current Strong economy and my employment status, the certainty of income from wages allows me to allocate a part of my income to increase my capital stock, despite the depreciation rate of Delta.

I also appreciate the importance of factoring in the probabilities of transitions, both in the state of the economy and in employment. While the average duration of unemployment during Strong periods is lower, I intend to hedge against potential risks by maintaining a portion of capital investment to ensure rental income continuity in case of unexpected unemployment.

For next quarter's capital allocation, I'll be relying on the expected value of rental and wage rates while considering Beta's discounting effect. My approach prioritizes maintaining consumption while ensuring capital growth. I am curious to know if you've considered the implications of aggregate capital on rental rates, as theoretically, an increase in capital could depress rental rates due to higher capital over employment ratios.

Looking forward to your summary, and thank you once again for this productive exchange.

Best regards, Joshua'

RSPNR's message 2 was '"Ashley's message 2:

Hello Joshua,

Your insights resonate with me, reinforcing my confidence in a Strong economic state for the next quarter. The low transition probability and persistent Strong state indicators indeed support our shared expectations of continued stability in productivity, rental, and wage rates.

In terms of strategy, your approach to balancing current consumption with capital investment is prudent. Like you, I recognize the value of diversification, especially considering the potential risks associated with employment status changes, despite our Strong economy. I plan to allocate a portion of my income to capital investment too. This investment aims to secure a steady stream of rental income, buffering against the risk of income loss due to possible unemployment.

With regard to your question, I have indeed considered the implications of aggregate capital on rental rates. As the capital stock grows, it may drive the rental rate down due to the increasing capital over employment ratio, as suggested by the economic model we're observing. However, given our relatively stable recent aggregate employment figures and the current trends indicating a safe margin before rental rates are potentially affected, I believe this might not be an immediate concern for T=12488.

In summary, we seem aligned on: - Expecting a continued Strong state of the economy and corresponding productivity of 1.01. - Anticipating similar rental and wage rates for T=12488 based on current trends. - Emphasizing a strategy that hedges against employment risks by balancing consumption with capital accumulation. - Being watchful about the possible impact of growing aggregate capital on future rental rates.

Your points also led me to think more rigorously about the probable impact of higher capital ratios on rental rates and wage rates in subsequent quarters. This consideration will inform my exact spending/savings decisions this quarter, alongside the predictive economic formulas and discount factor Beta.

Thank you for your thoughtful dialogue, Joshua. It has been informative and useful for refining my economic strategy for this quarter and beyond.

Best, Ashley"'

G Consumption-savings: LLM agents' beliefs, error analysis

Table 9: Consumption-savings, without Communication: reported beliefs, regression results for positive errors.

Expected variable:	$\mathrm{E}_t[Z_{t+1}]$	$\mathrm{E}_t[\mathrm{Med}[K_{i,t+1}]]$	$\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$	$\mathrm{E}_t[K_{t+1}]$	$\mathrm{E}_t[r_{t+1}]$	$\mathrm{E}_t[w_{t+1}]$
			Accuracy and descriptive statistics for ε_{t+1}			
R^2 (baseline, feasible)	1.000	0.937	0.870	0.991	0.643	0.984
R^2	0.773	0.876	0.807	0.803	0.358	0.958
$\mathrm{E}[\varepsilon_{t+1}]$	-0.099	0.242	0.096	21.560	0.000	0.000
$\mathrm{StD}[\varepsilon_{t+1}]$	0.147	0.603	0.841	29.372	0.001	0.008
ϱ_ε	0.404	-0.061	-0.125	0.060	0.064	0.132
Observations	249	249	249	248	242	241
			Dependence structure of ε_{t+1}			
constant	-0.0578 (0.000) [0.898]	0.3045 (0.731) [0.005]	0.7748 (0.495) [0.036]	33.0380 (0.115) [0.005]	0.0020 (0.241) [0.082]	0.0196 (0.027) [0.028]
Z_t	0.0692 (0.000) -	-0.0415 (0.489) -	-0.2171 (0.067) -	1.4077 (0.426) -	-0.0005 (0.004) -	0.0020 (0.065) -
K_t	-0.0000 (0.222) -	0.0002 (0.658) -	0.0001 (0.892) -	0.0007 (0.950) -	-0.0000 (0.547) -	-0.0000 (0.071) -
$\mathrm{IQR}[K_{it}]$	-0.0005 (0.406) [0.899]	0.0247 (0.106) [0.025]	-0.0010 (0.970) [0.036]	0.1348 (0.762) [0.006]	-0.0000 (0.436) [0.086]	-0.0001 (0.739) [0.029]
$\mathrm{MA}[\mathrm{IQR}[K_{it}]]$	-0.0007 (0.288) [0.900]	0.0309 (0.075) [0.031]	0.0243 (0.335) [0.042]	-0.1783 (0.694) [0.006]	-0.0001 (0.209) [0.094]	-0.0002 (0.597) [0.032]
$\mathrm{IQR}[K_{it}] - \mathrm{MA}[\mathrm{IQR}[K_{it}]]$	0.0005 (0.361) [0.898]	-0.0020 (0.955) [0.005]	-0.1079 (0.028) [0.064]	1.4071 (0.177) [0.020]	0.0001 (0.559) [0.087]	0.0005 (0.426) [0.031]
$\mathrm{MA}[\mathrm{Med}[K_{it}]]$	0.0001 (0.905) [0.898]	-0.0141 (0.728) [0.006]	0.0186 (0.778) [0.036]	-1.1031 (0.314) [0.013]	0.0002 (0.057) [0.102]	-0.0002 (0.855) [0.028]
$\mathrm{Med}[K_{it}] - \mathrm{MA}[\mathrm{Med}[K_{it}]]$	-0.0005 (0.302) [0.898]	-0.0857 (0.059) [0.032]	-0.0204 (0.823) [0.036]	-0.2889 (0.822) [0.006]	-0.0001 (0.300) [0.087]	0.0004 (0.746) [0.029]
$\mathrm{MA}[K_t]$	0.0000 (0.316) [0.898]	0.0016 (0.270) [0.014]	0.0051 (0.029) [0.070]	0.0110 (0.805) [0.006]	0.0000 (0.209) [0.095]	-0.0001 (0.285) [0.047]
$K_t - \mathrm{MA}[K_t]$	-0.0000 (0.316) [0.898]	-0.0016 (0.270) [0.014]	-0.0051 (0.029) [0.070]	-0.0110 (0.805) [0.006]	-0.0000 (0.209) [0.095]	0.0001 (0.285) [0.047]
$\mathrm{MA}[r_t]$	-0.6289 (0.197) [0.898]	-76.6177 (0.051) [0.028]	-137.0994 (0.071) [0.059]	2022.2526 (0.105) [0.024]	0.0254 (0.784) [0.082]	1.2078 (0.292) [0.043]
$r_t - \mathrm{MA}[r_t]$	0.6063 (0.201) [0.898]	72.7119 (0.064) [0.026]	135.1724 (0.071) [0.059]	-1960.1073 (0.115) [0.023]	-0.0245 (0.792) [0.082]	-1.2555 (0.280) [0.044]
$\mathrm{MA}[w_t]$	0.0302 (0.209) [0.898]	3.9588 (0.139) [0.020]	10.1756 (0.028) [0.071]	-14.0345 (0.865) [0.006]	0.0082 (0.256) [0.090]	-0.0967 (0.287) [0.047]
$w_t - \mathrm{MA}[r_t]$	-0.0607 (0.149) [0.898]	-0.3652 (0.905) [0.006]	-6.7905 (0.131) [0.052]	57.8728 (0.456) [0.009]	-0.0208 (0.002) [0.136]	0.1865 (0.040) [0.094]
Observations	160–161	127	113	134	117	122

Notes: Columns provide expected variables/statistics. The upper panel presents goodness-of-fit measure R^2 for LLM agents' reported expectations with respect to their true/realized values (preceded by its counterpart for the feasible version of the baseline model), as well as the properties of expectation errors which are defined as $\varepsilon_{t+1} := X_{t+1} - \mathrm{E}_t[X_{t+1}]$; $\mathrm{E}[\cdot]$ denotes expected value, $\mathrm{StD}[\cdot]$ denotes standard deviation, and ϱ_ε denotes the first order autocorrelation coefficient for ε_{t+1}; all errors are considered, the number of observations used is given in the bottom row. The lower panel presents OLS estimation results of regressing ε_{t+1} on constant and controls $[Z_t, K_t]$ only, and then with an added 1 out of 11 different variables (while keeping the same controls); the first three rows and each row next give coefficient estimates, p-values for standard t-tests based on Newey–West standard errors (in parentheses) and R^2 for the corresponding specification [in brackets]; only positive errors (after demeaning, i.e., above-average) are considered, the number of observations used is given in the bottom row.

Table 10: Consumption-savings, without Communication: reported beliefs, regression results for negative errors.

Expected variable:	$E_t[Z_{t+1}]$	$E_t[\text{Med}[K_{i,t+1}]]$	$E_t[\text{IQR}[K_{i,t+1}]]$	$E_t[K_{t+1}]$	$E_t[r_{t+1}]$	$E_t[w_{t+1}]$
			Accuracy and descriptive statistics for ε_{t+1}			
R^2 (baseline, feasible)	1.000	0.937	0.870	0.991	0.643	0.984
R^2	0.773	0.876	0.807	0.803	0.358	0.958
$E[\varepsilon_{t+1}]$	-0.099	0.242	0.096	21.560	0.000	0.000
$\text{StD}[\varepsilon_{t+1}]$	0.147	0.603	0.841	29.372	0.001	0.008
ϱ_ε	0.404	-0.061	-0.125	0.060	0.064	0.132
Observations	249	249	249	248	242	241
			Dependence structure of ε_{t+1}			
constant	-0.0771 (0.829) [0.004]	-0.5044 (0.538) [0.010]	0.4923 (0.596) [0.019]	59.3211 (0.382) [0.010]	-0.0021 (0.178) [0.073]	0.0001 (0.992) [0.107]
Z_t	n/a n/a n/a	-0.0941 (0.225) -	-0.1152 (0.172) -	-5.4965 (0.348) -	-0.0004 (0.000) -	0.0033 (0.000) -
K_t	-0.0001 (0.591) -	0.0002 (0.674) -	-0.0005 (0.323) -	-0.0263 (0.457) -	0.0000 (0.241) -	-0.0000 (0.578) -
$\text{IQR}[K_{it}]$	0.0162 (0.054) [0.062]	0.0010 (0.971) [0.010]	-0.0382 (0.068) [0.037]	-0.1248 (0.903) [0.010]	-0.0000 (0.591) [0.075]	0.0001 (0.641) [0.108]
$\text{MA}[\text{IQR}[K_{it}]]$	0.0136 (0.102) [0.039]	-0.0051 (0.834) [0.011]	-0.0439 (0.025) [0.040]	-0.4906 (0.783) [0.011]	-0.0000 (0.382) [0.079]	0.0001 (0.774) [0.108]
$\text{IQR}[K_{it}] - \text{MA}[\text{IQR}[K_{it}]]$	0.0210 (0.188) [0.025]	0.0198 (0.532) [0.012]	0.0010 (0.983) [0.019]	1.0311 (0.627) [0.011]	0.0000 (0.452) [0.077]	0.0002 (0.728) [0.108]
$\text{MA}[\text{Med}[K_{it}]]$	0.0103 (0.547) [0.009]	0.0190 (0.699) [0.011]	0.0436 (0.438) [0.023]	1.7304 (0.681) [0.011]	-0.0001 (0.467) [0.077]	-0.0013 (0.022) [0.150]
$\text{Med}[K_{it}] - \text{MA}[\text{Med}[K_{it}]]$	-0.0370 (0.114) [0.038]	-0.0657 (0.117) [0.019]	-0.1125 (0.134) [0.039]	7.1761 (0.312) [0.027]	-0.0000 (0.967) [0.073]	0.0012 (0.065) [0.133]
$\text{MA}[K_t]$	0.0023 (0.004) [0.123]	0.0010 (0.549) [0.012]	0.0018 (0.492) [0.024]	0.0483 (0.493) [0.011]	0.0000 (0.443) [0.076]	-0.0000 (0.084) [0.132]
$K_t - \text{MA}[K_t]$	-0.0023 (0.004) [0.123]	-0.0010 (0.549) [0.012]	-0.0018 (0.492) [0.024]	-0.0483 (0.493) [0.011]	-0.0000 (0.443) [0.076]	0.0000 (0.084) [0.132]
$\text{MA}[r_t]$	-125.2404 (0.000) [0.365]	-83.6961 (0.253) [0.022]	-9.8370 (0.888) [0.019]	2132.1031 (0.734) [0.012]	-0.0312 (0.753) [0.074]	0.1337 (0.861) [0.108]
$r_t - \text{MA}[r_t]$	125.7011 (0.000) [0.364]	88.5236 (0.223) [0.024]	19.7160 (0.784) [0.020]	-1866.8290 (0.759) [0.011]	0.0263 (0.791) [0.074]	-0.1506 (0.845) [0.108]
$\text{MA}[w_t]$	5.6192 (0.000) [0.188]	2.8076 (0.451) [0.014]	2.7583 (0.583) [0.022]	29.0189 (0.860) [0.010]	0.0039 (0.466) [0.076]	-0.0604 (0.170) [0.124]
$w_t - \text{MA}[r_t]$	-5.3757 (0.000) [0.189]	-4.7105 (0.253) [0.022]	-3.8220 (0.438) [0.023]	197.6280 (0.467) [0.014]	-0.0216 (0.000) [0.170]	0.1410 (0.001) [0.196]
Observations	88	121–122	135–136	113–114	123–124	118–119

Notes: Columns provide expected variables/statistics. The upper panel presents goodness-of-fit measure R^2 for LLM agents' reported expectations with respect to their true/realized values (preceded by its counterpart for the feasible version of the baseline model), as well as the properties of expectation errors which are defined as $\varepsilon_{t+1} := X_{t+1} - E_t[X_{t+1}]$; $E[\cdot]$ denotes expected value, $\text{StD}[\cdot]$ denotes standard deviation, and ϱ_ε denotes the first order autocorrelation coefficient for ε_{t+1}; all errors are considered, the number of observations used is given in the bottom row. The lower panel presents OLS estimation results of regressing ε_{t+1} on constant and controls $[Z_t, K_t]$ only, and then with an added 1 out of 11 different variables (while keeping the same controls); the first three rows and each row next give coefficient estimates, p-values for standard t-tests based on Newey–West standard errors (in parentheses) and R^2 for the corresponding specification [in brackets]; only negative errors (after demeaning, i.e., below-average) are considered, the number of observations used is given in the bottom row. For $E_t[Z_{t+1}]$, "n/a" reflects the fact that negative errors only happen under $Z_t = 0$, hence there is no variation in this regressor.

Table 11: Consumption-savings, with Communication: reported beliefs, regression results for positive errors.

Expected variable:	$\mathrm{E}_t[Z_{t+1}]$	$\mathrm{E}_t[\mathrm{Med}[K_{i,t+1}]]$	$\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$	$\mathrm{E}_t[K_{t+1}]$	$\mathrm{E}_t[r_{t+1}]$	$\mathrm{E}_t[w_{t+1}]$
	Accuracy and descriptive statistics for ε_{t+1}					
R^2 (baseline, feasible)	1.000	0.963	0.826	0.995	0.585	0.988
R^2	0.815	0.939	0.753	0.894	0.322	0.643
$\mathrm{E}[\varepsilon_{t+1}]$	-0.090	0.373	0.147	39.805	0.000	0.004
$\mathrm{StD}[\varepsilon_{t+1}]$	0.132	0.689	1.171	21.174	0.001	0.026
ϱ_ε	0.448	-0.063	0.141	0.085	0.069	-0.018
Observations	249	249	249	249	241	242
	Dependence structure of ε_{t+1}					
constant	-0.0586 (0.000) [0.836]	-0.0592 (0.941) [0.040]	-2.3929 (0.275) [0.033]	15.4329 (0.579) [0.040]	-0.0001 (0.959) [0.175]	0.1703 (0.088) [0.155]
Z_t	0.0510 (0.000) -	-0.1487 (0.057) -	-0.1518 (0.392) -	2.6105 (0.229) -	-0.0005 (0.000) -	-0.0364 (0.025) -
K_t	0.0000 (0.631) -	0.0003 (0.213) -	0.0011 (0.123) -	0.0123 (0.163) -	0.0000 (0.426) -	-0.0000 (0.170) -
$\mathrm{IQR}[K_{it}]$	-0.0000 (0.862) [0.836]	-0.0002 (0.985) [0.040]	-0.0693 (0.002) [0.068]	-0.7029 (0.167) [0.057]	0.0000 (0.205) [0.188]	0.0007 (0.775) [0.156]
$\mathrm{MA}[\mathrm{IQR}[K_{it}]]$	0.0001 (0.822) [0.836]	-0.0057 (0.717) [0.041]	-0.0908 (0.014) [0.079]	-0.2277 (0.668) [0.042]	0.0001 (0.026) [0.212]	-0.0008 (0.792) [0.156]
$\mathrm{IQR}[K_{it}] - \mathrm{MA}[\mathrm{IQR}[K_{it}]]$	-0.0002 (0.462) [0.836]	0.0103 (0.601) [0.041]	0.0034 (0.958) [0.033]	-1.4075 (0.083) [0.064]	-0.0000 (0.270) [0.185]	0.0035 (0.250) [0.166]
$\mathrm{MA}[\mathrm{Med}[K_{it}]]$	0.0011 (0.062) [0.840]	0.0403 (0.099) [0.052]	-0.0011 (0.985) [0.033]	-0.5126 (0.575) [0.042]	-0.0000 (0.716) [0.176]	-0.0047 (0.174) [0.168]
$\mathrm{Med}[K_{it}] - \mathrm{MA}[\mathrm{Med}[K_{it}]]$	-0.0008 (0.116) [0.837]	-0.0565 (0.069) [0.056]	0.0084 (0.905) [0.033]	-0.9216 (0.419) [0.044]	0.0001 (0.065) [0.196]	0.0025 (0.603) [0.157]
$\mathrm{MA}[K_t]$	0.0000 (0.058) [0.840]	0.0016 (0.265) [0.050]	0.0008 (0.780) [0.033]	0.0788 (0.111) [0.059]	-0.0000 (0.042) [0.200]	-0.0002 (0.416) [0.168]
$K_t - \mathrm{MA}[K_t]$	-0.0000 (0.058) [0.840]	-0.0016 (0.265) [0.050]	-0.0008 (0.780) [0.033]	-0.0788 (0.111) [0.059]	0.0000 (0.042) [0.200]	0.0002 (0.416) [0.168]
$\mathrm{MA}[r_t]$	-1.3017 (0.128) [0.836]	-22.2581 (0.796) [0.041]	-181.7295 (0.289) [0.041]	-7437.3571 (0.009) [0.099]	0.2197 (0.031) [0.203]	-3.2446 (0.819) [0.155]
$r_t - \mathrm{MA}[r_t]$	1.2825 (0.110) [0.836]	24.2666 (0.778) [0.041]	186.3888 (0.280) [0.042]	7426.1235 (0.009) [0.099]	-0.2207 (0.031) [0.203]	2.5167 (0.859) [0.155]
$\mathrm{MA}[w_t]$	0.1167 (0.055) [0.839]	3.4530 (0.357) [0.047]	4.0408 (0.613) [0.035]	255.6873 (0.045) [0.071]	-0.0093 (0.027) [0.202]	-0.4451 (0.560) [0.162]
$w_t - \mathrm{MA}[r_t]$	-0.0925 (0.147) [0.838]	-1.0089 (0.798) [0.040]	-8.8079 (0.320) [0.043]	-212.6875 (0.064) [0.056]	-0.0055 (0.184) [0.184]	0.5041 (0.501) [0.164]
Observations	159–160	134	121–122	127–128	110	62

Notes: Columns provide expected variables/statistics. The upper panel presents goodness-of-fit measure R^2 for LLM agents' reported expectations with respect to their true/realized values (preceded by its counterpart for the feasible version of the baseline model), as well as the properties of expectation errors which are defined as $\varepsilon_{t+1} := X_{t+1} - \mathrm{E}_t[X_{t+1}]$; $\mathrm{E}[\cdot]$ denotes expected value, $\mathrm{StD}[\cdot]$ denotes standard deviation, and ϱ_ε denotes the first order autocorrelation coefficient for ε_{t+1}; all errors are considered, the number of observations used is given in the bottom row. The lower panel presents OLS estimation results of regressing ε_{t+1} on constant and controls $[Z_t, K_t]$ only, and then with an added 1 out of 11 different variables (while keeping the same controls); the first three rows and each row next give coefficient estimates, p-values for standard t-tests based on Newey–West standard errors (in parentheses) and R^2 for the corresponding specification [in brackets]; only positive errors (after demeaning, i.e., above-average) are considered, the number of observations used is given in the bottom row.

Table 12: Consumption-savings, with Communication: reported beliefs, regression results for negative errors.

Expected variable:	$\mathrm{E}_t[Z_{t+1}]$	$\mathrm{E}_t[\mathrm{Med}[K_{i,t+1}]]$	$\mathrm{E}_t[\mathrm{IQR}[K_{i,t+1}]]$	$\mathrm{E}_t[K_{t+1}]$	$\mathrm{E}_t[r_{t+1}]$	$\mathrm{E}_t[w_{t+1}]$
	Accuracy and descriptive statistics for ε_{t+1}					
R^2 (baseline, feasible)	1.000	0.963	0.826	0.995	0.585	0.988
R^2	0.815	0.939	0.753	0.894	0.322	0.643
$\mathrm{E}[\varepsilon_{t+1}]$	-0.090	0.373	0.147	39.805	0.000	0.004
$\mathrm{StD}[\varepsilon_{t+1}]$	0.132	0.689	1.171	21.174	0.001	0.026
ϱ_ε	0.448	-0.063	0.141	0.085	0.069	-0.018
Observations	249	249	249	249	241	242
	Dependence structure of ε_{t+1}					
constant	0.4186 (0.179) [0.062]	0.7621 (0.405) [0.041]	1.1446 (0.409) [0.017]	51.6242 (0.093) [0.021]	-0.0031 (0.005) [0.134]	-0.0068 (0.501) [0.046]
Z_t	n/a n/a n/a	-0.1451 (0.068) -	-0.0278 (0.822) -	-3.0477 (0.260) -	-0.0003 (0.000) -	0.0024 (0.003) -
K_t	-0.0002 (0.035) -	-0.0003 (0.329) -	-0.0006 (0.187) -	-0.0085 (0.395) -	0.0000 (0.008) -	0.0000 (0.773) -
$\mathrm{IQR}[K_{it}]$	-0.0016 (0.759) [0.063]	0.0046 (0.805) [0.041]	0.0281 (0.288) [0.027]	0.8202 (0.078) [0.041]	0.0000 (0.296) [0.141]	0.0000 (0.821) [0.046]
$\mathrm{MA}[\mathrm{IQR}[K_{it}]]$	-0.0067 (0.227) [0.076]	-0.0022 (0.912) [0.041]	0.0409 (0.198) [0.032]	-0.0351 (0.959) [0.021]	0.0000 (0.239) [0.144]	-0.0004 (0.049) [0.065]
$\mathrm{IQR}[K_{it}] - \mathrm{MA}[\mathrm{IQR}[K_{it}]]$	0.0093 (0.248) [0.076]	0.0184 (0.498) [0.044]	-0.0037 (0.919) [0.017]	2.4872 (0.005) [0.084]	0.0000 (0.995) [0.134]	0.0010 (0.000) [0.099]
$\mathrm{MA}[\mathrm{Med}[K_{it}]]$	0.0370 (0.007) [0.166]	-0.0451 (0.143) [0.053]	-0.0261 (0.662) [0.019]	0.7922 (0.398) [0.024]	-0.0000 (0.731) [0.135]	0.0002 (0.532) [0.048]
$\mathrm{Med}[K_{it}] - \mathrm{MA}[\mathrm{Med}[K_{it}]]$	-0.0392 (0.000) [0.145]	-0.0776 (0.073) [0.066]	-0.0386 (0.561) [0.019]	1.7164 (0.250) [0.032]	-0.0000 (0.402) [0.139]	0.0001 (0.861) [0.046]
$\mathrm{MA}[K_t]$	0.0024 (0.000) [0.256]	0.0018 (0.301) [0.050]	-0.0006 (0.826) [0.017]	0.0669 (0.198) [0.034]	-0.0000 (0.995) [0.134]	-0.0000 (0.105) [0.058]
$K_t - \mathrm{MA}[K_t]$	-0.0024 (0.000) [0.256]	-0.0018 (0.301) [0.050]	0.0006 (0.826) [0.017]	-0.0669 (0.198) [0.034]	0.0000 (0.995) [0.134]	0.0000 (0.105) [0.058]
$\mathrm{MA}[r_t]$	-207.6797 (0.000) [0.497]	-49.4706 (0.631) [0.043]	220.5501 (0.128) [0.034]	-883.9718 (0.769) [0.021]	-0.0021 (0.985) [0.134]	-0.2960 (0.756) [0.047]
$r_t - \mathrm{MA}[r_t]$	203.6177 (0.000) [0.491]	52.5953 (0.612) [0.043]	-206.7384 (0.148) [0.032]	746.5099 (0.803) [0.021]	-0.0081 (0.941) [0.134]	0.2968 (0.750) [0.047]
$\mathrm{MA}[w_t]$	8.1786 (0.000) [0.366]	4.3011 (0.345) [0.048]	-4.1727 (0.533) [0.020]	142.5438 (0.283) [0.030]	-0.0001 (0.990) [0.134]	-0.0471 (0.263) [0.052]
$w_t - \mathrm{MA}[r_t]$	-7.0110 (0.000) [0.305]	-1.8384 (0.679) [0.042]	-2.6347 (0.737) [0.018]	-126.8176 (0.328) [0.028]	-0.0168 (0.007) [0.216]	0.1613 (0.001) [0.108]
Observations	89	114–115	127	121	129–130	179–180

Notes: Columns provide expected variables/statistics. The upper panel presents goodness-of-fit measure R^2 for LLM agents' reported expectations with respect to their true/realized values (preceded by its counterpart for the feasible version of the baseline model), as well as the properties of expectation errors which are defined as $\varepsilon_{t+1} := X_{t+1} - \mathrm{E}_t[X_{t+1}]$; $\mathrm{E}[\cdot]$ denotes expected value, $\mathrm{StD}[\cdot]$ denotes standard deviation, and ϱ_ε denotes the first order autocorrelation coefficient for ε_{t+1}; all errors are considered, the number of observations used is given in the bottom row. The lower panel presents OLS estimation results of regressing ε_{t+1} on constant and controls $[Z_t, K_t]$ only, and then with an added 1 out of 11 different variables (while keeping the same controls); the first three rows and each row next give coefficient estimates, p-values for standard t-tests based on Newey–West standard errors (in parentheses) and R^2 for the corresponding specification [in brackets]; only negative errors (after demeaning, i.e., below-average) are considered, the number of observations used is given in the bottom row. For $\mathrm{E}_t[Z_{t+1}]$, "n/a" reflects the fact that negative errors only happen under $Z_t = 0$, hence there is no variation in this regressor.

H Data

Input-Output: Input-Output data were obtained as follows.

1. Raw data are sourced from U.S. Bureau of Economic Analysis Commodity-by-Commodity Total Requirements, After Redefinitions, for 73 commodities in 2021.

Dropping the commodities (both in rows and in columns) that are related to government, "Scrap, used and secondhand goods" and "Noncomparable imports and rest-of-the-world adjustment" leaves 66 commodities.

The direct requirements matrix $\boldsymbol{D}_{\text{BEA}}$ is obtained from the total requirements matrix $\boldsymbol{T}_{\text{BEA}}$ using the equation (e.g., see [Acemoglu et al., 2012])

$$\boldsymbol{D}_{\text{BEA}} := (\boldsymbol{T}_{\text{BEA}} - \boldsymbol{I}) \times (\boldsymbol{T}_{\text{BEA}})^{-1},$$

where $\boldsymbol{I}$ is a conformable identity matrix.

Matrix $\boldsymbol{D}_{\text{BEA}}$ is further reduced to a 50×50 matrix $\boldsymbol{X}$ by eliminating commodities with the largest number of zero entries.

2. Raw data are sourced from U.S. Bureau of Economic Analysis The Make of Commodities by Industries, After Redefinitions, for 71 industries in 2021.

Taking the Total Commodity Output for 73 commodities $\boldsymbol{Y}_{\text{BEA}}$ and retaining only the 50 commodities as in the earlier procedure gives the 50 elements long vector $\boldsymbol{Y}$.

Names: Agent names were taken from the U.S. Social Security Administration list of Top Names Over the Last 100 Years available at `https://www.ssa.gov/oact/babynames/decades/century.html`.

Macro: Aggregate time series were obtained as follows.

Raw data are sourced from the Federal Reserve's FRED database. The variables used are:

1. Average Hourly Earnings of Production and Nonsupervisory Employees, Total Private [AHETPI]; Units: Dollars per Hour, Seasonally Adjusted; Frequency: Monthly.

2. Households, Net Worth, Level [BOGZ1FL192090005Q]; Units: Millions of Dollars, Not Seasonally Adjusted; Frequency: Quarterly.

3. Consumer Price Index for All Urban Consumers, All Items in U.S. City Average [CPIAUCSL]; Units: Index 1982-1984=100, Seasonally Adjusted; Frequency: Monthly.

4. Real Gross Domestic Product [GDPC1]; Units: Billions of Chained 2017 Dollars, Seasonally Adjusted Annual Rate; Frequency: Quarterly.

5. Gross Domestic Product, Implicit Price Deflator [GDPDEF]; Units: Index 2017=100, Seasonally Adjusted; Frequency: Quarterly.

6. Real Gross Private Domestic Investment [GPDIC1]; Units: Billions of Chained 2017 Dollars, Seasonally Adjusted Annual Rate; Frequency: Quarterly.

7. Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity, Quoted on an Investment Basis [GS10]; Units: Percent, Not Seasonally Adjusted; Frequency: Monthly.

8. Real Personal Consumption Expenditures [PCECC96]; Units: Billions of Chained 2017 Dollars, Seasonally Adjusted Annual Rate; Frequency: Quarterly.

9. Population [POPTHM]; Units: Thousands, Not Seasonally Adjusted; Frequency: Monthly.

Real GDP, consumption and investments are transformed to per capita terms by dividing them over by population levels.

Capital (measured as net worth) is deflated by GDP deflator.

Wage rate (measured as earnings of employees) undergoes aggregation into quarterly value and is deflated by CPI.

In the end, all variables are expressed in thousands of 2017 U.S. dollars. The only exception is rental rate, which is measured as an interest rate on 10-year U.S. Treasury Bonds.

In order to make data conformable with the theoretical setting, the variables for capital, output, consumption and investment are detrended (by running an OLS regression of a logarithm of the corresponding variable on a constant and a linear trend, and subtracting the trend, followed by reversing the logarithm transformation).

The above procedure suggests that the resulting calculations and their relationship to model in the paper should be taken with great care. With this caveat, the resulting statistics are summarized in Table 13.

Table 13: Consumption-savings, empirical results: time-series aspect

Variable	E[·]	StD[·]	CV[·]	$\rho(\cdot)$
K_t	43.794	3.700	0.084	-
Y_t	9.864	0.314	0.032	-
C_t	6.246	0.234	0.038	-
I_t	1.310	0.147	0.112	-
r_t	0.045	0.022	0.487	-
w_t	9.179	0.633	0.069	-
C_t, Y_t	-	-	-	0.967
Y_t, Y_{t-4}	-	-	-	0.816

Notes: Time-series statistics refer to aggregate variables; e.g., $\mathrm{E}[X_t]$ stands for $\mathrm{E}[\sum_i X_{it}]$. StD[·] denotes standard deviation, CV[·] denotes the coefficient of variation, ρ denotes Pearson's correlation coefficient. Aggregate variables are given on per capita basis. Data sample is 1988:Q1–2023Q4. Data source and transformations described in the text.

Inequality: Cross-sectional data were obtained as follows.

Raw data on wealth distribution for the U.S. are sourced from the Board of Governors of the Federal Reserve System's Survey of Consumer Finances for 2022.

The measure of wealth used is Total net worth of households in 2022 dollars.

The resulting statistics are summarized in Table 14.

Table 14: Consumption-savings, empirical results: cross-sectional aspect

Variable	E[·]	StD[·]	CV[·]	Gini[·]
K_{it}	1059.457	8085.433	7.632	0.830

Notes: Cross-sectional statistics refer to individual variables, and are averaged over all time periods; e.g., $\mathrm{E}[X_{it}]$ stands for $\mathrm{E}[\mathrm{E}[X_{it}|t]]$. StD[·] denotes standard deviation, CV[·] denotes the coefficient of variation, Gini[·] denotes Gini's coefficient of wealth inequality. Individual variables are given in US$ thousands. Data sample is for 2022. Data source described in the text.