



Research Paper

Individual evolutionary learning in repeated beauty contest games [☆]Mikhail Anufriev ^{a,b}, John Duffy ^{c,d}, Valentyn Panchenko ^{e,*}^a Department of Economics, University of Technology Sydney, Australia^b Department of Finance, Faculty of Economics, Technical University of Ostrava, Czechia^c Department of Economics, University of California, Irvine, United States of America^d ISER, Osaka University, Japan^e Economics, UNSW Business School, University of New South Wales, Australia

ARTICLE INFO

JEL classification:

C63
C72
C73
D83

Keywords:

Beauty contest game
Learning
Evolutionary dynamics
Testbed
Agent-based model

ABSTRACT

The Individual Evolutionary Learning (IEL) algorithm was proposed as a portable learning model for games with large strategy spaces. In principle, IEL benchmark simulations could substitute or supplement expensive experiments with human subjects. We evaluate the ability of the IEL model to replicate experimental findings observed in repeated Keynesian Beauty Contest (KBC) games, which have a large strategy space. The IEL specification with standard parameter values is able to capture major dynamical features and differences between treatments in both one-dimensional (Nagel, 1995; Duffy and Nagel, 1997) and two-dimensional (Anufriev et al., 2022b) versions of KBC games. We compare IEL with some other simple learning models and find that it performs relatively better across multiple treatments. We also use IEL to predict behavior in repeated KBC games that have not yet been conducted experimentally.

1. Introduction

Modeling how agents learn in large or complex strategy spaces poses a significant challenge. The high dimensionality of such spaces makes it difficult to represent and analyze all possible strategies comprehensively. A promising solution is to use the individual evolutionary learning model pioneered by Jasmina Arifovic and colleagues.¹ This evolutionary-based machine learning algorithm, grounded in genetic algorithms and classifier systems, stands out for its ability to navigate expansive and complex strategy spaces with minimal initial setup. Perhaps more importantly, this algorithm appears to closely emulate the behavior of human subjects across a diverse range of experiments. Indeed, Arifovic thought that computational experiments using IEL could complement or even

[☆] This paper is dedicated to the memory of Jasmina Arifovic, who inspired us to use evolutionary machine learning algorithms in our own research. We thank the editor and two referees for their helpful comments and suggestions. We gratefully acknowledge support for this project from The Australian Research Council Discovery Project funding scheme (project DP200101438), and from the Czech Science Foundation (GACR) under project 22-28882S and VSB-TU Ostrava under the SGS project SP2022/4. The Matlab code and data used in this study are available at <https://osf.io/zf3ct/>.

* Corresponding author.

E-mail addresses: Mikhail.Anufriev@uts.edu.au (M. Anufriev), duffy@uci.edu (J. Duffy), v.panchenko@unsw.edu.au (V. Panchenko).

¹ The IEL is based upon early work by Arifovic (1994, 1995, 1996) using genetic algorithms. The formalization and use of the IEL model can be found in Arifovic and Ledyard (2001, 2004, 2007, 2011, 2018), Arifovic and Karaivanov (2010), Arifovic et al. (2013), Anufriev et al. (2013), Anufriev et al. (2022a), Arifovic et al. (2019), and Arifovic et al. (2023).

<https://doi.org/10.1016/j.jebo.2023.12.010>

Received 27 April 2023; Received in revised form 3 December 2023; Accepted 7 December 2023

Available online 4 January 2024

0167-2681/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

replace expensive laboratory experiments with human subjects in settings with complex and/or large strategy space. In the first paper introducing the IEL algorithm, Jasmina Arifovic and John Ledyard write²:

Absent a unique, compelling model of behavior, some economists have turned to the experimental economics laboratory as a testbed for new mechanisms in much the same way that early aircraft designers turned to the wind-tunnel to test their designs... But this is expensive and time consuming... Existing models of learning would have a problem with an environment where the strategy space contains a continuum of strategies.

In this paper, we revisit this proposed use of IEL to predict the results of experiments with large strategy spaces. Specifically, we assess the IEL model's ability to forecast the behavior of participants in repeated plays of the Keynesian Beauty Contest (KBC) game. This game is a workhorse experimental framework for studying how participants learn to play a dominance-solvable equilibrium with a large strategy space. Here we consider nine different treatments of the repeated KBC, some of which have multiple Nash equilibria. While learning models can be used as a selection device, the IEL model has not yet been tested in the context of selecting between multiple Nash equilibria. Additionally, the different treatments exhibit different dynamical properties, e.g., regarding the speed and monotonicity of convergence to equilibrium. We leverage this behavioral variability to assess whether a standard “off-the-shelf” parameterization of the IEL model can capture these treatment differences. To our knowledge, this is the first study that uses the IEL model to investigate such games.

We further compare the IEL approach with other models that have been used to study the repeated KBC. In particular, we consider the learning direction model, which we complement with an error correction model and we also consider an adaptive learning model that we found useful in explaining behavior in more complex, two-dimensional repeated KBC. We show that the standard, off-the-shelf parameterization of the IEL model provides a better fit to the data in most treatments of the classic one-dimensional KBC games and in all treatments of the two dimensional KBC game when the models are initialized with the choices that human subjects made in the first period of those experiments. An important advantage of the IEL model is that it does not require the estimation of any parameters; while the model parameters could be estimated, we stick with the off-the-shelf parameterization that Arifovic and associates have used in many prior implementations. Further, even when IEL is not initialized using the experimental data, it still captures the differences across the treatments rather well. Finally, we apply the IEL model to experimental treatments of the KBC that have not yet been tested experimentally, and we show how the IEL model generates clear, testable predictions for experimental outcomes. We conclude that the IEL model is well-suited as a testbed device that can be used to evaluate and predict behavior in new experimental treatments either before or in lieu of running them in the laboratory.

The rest of the paper is organized as follows. Section 2 describes the KBC game and the experimental data that we use in this paper. It also discusses the predictions and findings for various KBC treatments. Section 3 formalizes how we apply the IEL model to the repeated KBC game. We aim here to gain some understanding of the IEL dynamics. Section 4 confronts the IEL model and other learning models with the data from repeated KBC experiments and provides evidence that the off-the-shelf IEL model is able to match qualitative if not quantitative features of the period-by-period experimental data. Section 5 evaluates the suitability of using IEL as a testbed and looks at whether the IEL simulations with no experimental input (no initial conditions) can continue to capture the key differences between the experimental treatments. Section 6 provides some final remarks.

2. Keynesian beauty contest game

In the standard, one-dimensional (1D) version of the KBC game, each player i in a group of N players is asked to guess a number, $x_i \in [0, 100]$. Each player's payoff depends on how close his/her guess is to some target number, x^* . To determine this target, all guesses are aggregated using some statistic, $m(x_1, \dots, x_N)$, typically the mean, i.e., $m(x_1, \dots, x_N) = \sum x_i / N$. Then, the target number is defined as a commonly known linear function of this statistic. For example, in the pioneering study of Nagel (1995), the target is

$$x^* = p \cdot m(x_1, \dots, x_N) = p \cdot \frac{1}{N} \sum_{i=1}^N x_i. \quad (1)$$

When $p \in (0, 1)$, the unique dominance solvable Nash equilibrium prediction is for all players to guess 0.

Several generalizations of the simple KBC game have been pursued. Duffy and Nagel (1997) study the robustness of the subjects' behavior to different order statistics, $m(\cdot)$. In their paper, apart from the treatment with the mean-KBC game with the target as in Eq. (1), they also consider treatments where $m(\cdot)$ is the median or the maximum of individual guesses. Güth et al. (2002) adds a constant to the target equation (1), shifting the NE of the game to the interior of the guessing interval. Sutan and Willinger (2009) study the game with a negative value for p . Finally, Anufriev et al. (2022b) study a two-dimensional (2D) version of the KBC game, which we will refer to here as the *planar* KBC game.

Regarding behavior in this game, Nagel (1995) first showed in the original KBC game, that subjects do apply some iterated reasoning but the vast majority of subjects use only a limited number of steps of such iterated reasoning. Given some initial reference point, x_0 , (typically $x_0 = 50$), “level-1” types choose $p \cdot x_0$ as their guess. More sophisticated, “level-2” types, presume that all other players are level-1 types, and best respond to that choice by guessing $p(p x_0) = p^2 x_0$, and so on. The same iterated “level- k ” reasoning

² See the Introduction in Arifovic and Ledyard (2001). The closest published paper is Arifovic and Ledyard (2011).

describes subjects' behavior when the KBC game is repeated and subjects are given feedback on the past round average and winning number. Nagel (1995) puts forward a hypothesis that beyond the first round of play, individuals formulate their guesses as some factor of the past average. That is, adding the subscript t for a round, a guess $x_{i,t}$ of agent i is the previous mean, \bar{x}_{t-1} , times the so-called adjustment factor, $a_{i,t}$. The adjustment factors are updated from period to period following the learning direction model proposed by Selten and Stoecker (1986) which posits that agents find their *ex post* optimal adjustment factor and then change their individual factors in the corresponding direction, increasing them if their guesses were too low and decreasing them if their guesses were too high. Nagel (1995) showed that many subjects in the mean-KBC game behave consistently with the learning direction model. Duffy and Nagel (1997) find further support for this model in the initial rounds of the median- and maximum-based KBC.

The learning direction model does not specify a functional form for the updating of the factors over time. For this reason, it was not included in the large set of models that we compared in (Anufriev et al., 2022b) using the planar KBC data. In that paper, we ran a contest among several different learning models (though not IEL). The set of models included dynamic versions of the homogeneous level- k reasoning model as well as mixed level- k models including the cognitive hierarchy model of Camerer et al. (2004). Also included were models based on past averaging of targets, the EWA model of Camerer and Ho (1999), the structural mixed level- k learning model of Gill and Prowse (2016), and some non-learning models. We found that a simple model that mixed between levels 0 and 1, equivalent to an adaptive learning model, performed very well in the planar KBC game. It performs better, for example, than the EWA model of Camerer and Ho (1999) and is similar to a structural model with increasing level types as proposed by Gill and Prowse (2016). The adaptive model was only slightly outperformed by the winner of our learning model contest, which adds a very small fraction of Nash Equilibrium choices to the adaptive model.

In the next two subsections, we describe the experimental data from the previous literature that we use in this paper and the predictions and findings of the various KBC models that we will try to match with the IEL model and other models.

2.1. Experimental designs and data

While there are many experimental tests of the KBC game, we use only data from experiments where this game was repeatedly played as we are interested in *learning* dynamics. Thus, in all our data, there are $T > 1$ repetitions or “rounds” of the KBC, played by the same participants who get standard feedback between rounds.³ We use subscript $t = 1, \dots, T$ to denote the rounds.

The first data set we use is the experimental data of Nagel (1995), **N95** henceforth. In total, there are 10 experimental sessions from this study. All of them use the functional form (1) (based on the mean of all guesses) to describe the target number. Sessions 1–3 have $p = 1/2$, sessions 4–7 have $p = 2/3$, and sessions 8–10 have $p = 4/3$. The number of subjects, N , is between 15 and 18; the variance is due to different show-up rates. All sessions have $T = 4$ rounds. One important design feature of these (and other early) experiments is the payoff structure. In each round, all N subjects competed for a fixed, known prize, P . This prize was equally split among all winners, i.e., those participants whose guesses were equally close to the target value x^* (most often there is a single winner).

Formally, let $x_{i,t}$ be the guess of individual i at time t , and let \bar{x}_t be the average of all N guesses. Further, let x_t^* be the target number based on these guesses, so that $x_t^* = p\bar{x}_t$, as in (1). The set of winners, given the individual guesses and the target, is defined by

$$W_t = \left\{ i : x_{i,t} \in \arg \min_i |x_{i,t} - x_t^*| \right\}. \quad (2)$$

Let $\#W_t \geq 1$ denote the number of elements in this set. Then, the payoff function for each round of the **N95** experiment is

$$\pi_{i,t}(x_{1,t}, \dots, x_{N,t}) = \begin{cases} \frac{P}{\#W_t} & \text{if } i \in W_t, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

This payoff function is referred to as an all-or-nothing tournament structure in the literature.

The second dataset we use is due to Duffy and Nagel (1997), **DN97** henceforth. They study the robustness of the KBC game to modification of the statistic $m(\cdot)$ used in the target equation (1). Specifically, they fix $p = 1/2$ but consider three ways of combining individual guesses into a statistical measure: the mean (as in Eq. (1)), and also the median, and the maximum. Provided that the target number is computed using one of these three statistics, their incentive structure in all 12 sessions is also a tournament, as described by Eqs. (2) and (3). Sessions 1–4 use the median version of the KBC game so that $m(x_{1,t}, \dots, x_{N,t})$ is the median of N guesses and the target x_t^* is half of this median. All sessions are run for 4 rounds, except for session 4 which is run for $T = 10$ rounds. Sessions 5–8 use the mean version of the KBC game, so that $m(x_{1,t}, \dots, x_{N,t}) = \bar{x}_t$. Thus in these sessions, participants play exactly the same game as in **N95**, with $p = 1/2$. They do it for 4 rounds, except for session 8 which is run for 10 rounds. Finally, sessions

³ This standard feedback after each round t consists of the own guess of participant i , $x_{i,t}$, the relevant statistic used to aggregate guesses, $m(x_{1,t}, \dots, x_{N,t})$, (e.g., the mean), the target number, x_t^* , and the realized payoff of participant i . For KBC games with an “all-or-nothing” tournament structure (as explained below), feedback also includes the winning number or guess.

Table 1

Experimental data used in the paper. **N95** stands for Nagel (1995), **DN97** stands for Duffy and Nagel (1997), and **ADP22** stands for Anufriev et al. (2022b).

Experiment	Ses	N	T	p or \mathbf{P}	m	payoff	Dynamics
N95	1	15	4	1/2	mean	Eq. (3)	converge to 0
	2	16	4	1/2	mean	Eq. (3)	
	3	17	4	1/2	mean	Eq. (3)	
	4	18	4	2/3	mean	Eq. (3)	converge to 0, slower than above
	5	15	4	2/3	mean	Eq. (3)	
	6	16	4	2/3	mean	Eq. (3)	
	7	18	4	2/3	mean	Eq. (3)	
	8	18	4	4/3	mean	Eq. (3)	converge to 100
	9	15	4	4/3	mean	Eq. (3)	
	10	16	4	4/3	mean	Eq. (3)	
DN97	1, 2	15	4	1/2	median	Eq. (3)	converge to 0
	3	13	4	1/2	median	Eq. (3)	
	4	13	10	1/2	median	Eq. (3)	
	5	16	4	1/2	mean	Eq. (3)	converge to 0, slower than above
	6	14	4	1/2	mean	Eq. (3)	
	7	15	4	1/2	mean	Eq. (3)	
	8	14	10	1/2	mean	Eq. (3)	
	9–11	15	4	1/2	max	Eq. (3)	slowly decreases does not converge
	12	15	10	1/2	max	Eq. (3)	
ADP22	1–4	10	15	\mathbf{P}_{Sink}	mean	Eq. (5)	converge to (90, 20)
	5–8	10	15	$\mathbf{P}_{\text{SaddleNeg}}$	mean	Eq. (5)	converge to (90, 100) converge to (0, 38)
	9–11	10	15	$\mathbf{P}_{\text{SaddlePos}}$	mean	Eq. (5)	
	12–16	10	15	$\mathbf{P}_{\text{Source}}$	mean	Eq. (5)	

9–12 use the maximum version of the KBC game. In this case $m(x_{1,t}, \dots, x_{N,t}) = \max(x_{1,t}, \dots, x_{N,t})$, and $x_t^* = \max(x_{1,t}, \dots, x_{N,t})/2$. All sessions are run for 4 rounds, except for session 12 which is run for 10 rounds.

The last data set we use is from Anufriev et al. (2022b), **ADP22** henceforth. In this paper, participants play the *planar* KBC game. They submit a pair of numbers, $(x_{i,t}, y_{i,t})$, with each of these numbers belonging to the interval $[0, 100]$. The mean statistic is used to aggregate all first and all second guesses. The target values, (x_t^*, y_t^*) , are then linear functions of these means. Using matrix notation, the targets can be written as

$$\begin{pmatrix} x_t^* \\ y_t^* \end{pmatrix} = \mathbf{P} \begin{pmatrix} \bar{x}_t \\ \bar{y}_t \end{pmatrix} + \mathbf{b}, \quad \text{where} \quad \bar{x}_t = \frac{1}{N} \sum_{i=1}^N x_{i,t}, \quad \bar{y}_t = \frac{1}{N} \sum_{i=1}^N y_{i,t}. \quad (4)$$

The payoff structure in the Anufriev et al. (2022b) study is different from the all-or-nothing tournament structure. The payoff earned by participant i is

$$\pi_{i,t}(x_{1,t}, \dots, x_{N,t}, y_{1,t}, \dots, y_{N,t}) = \frac{500}{5 + |x_{i,t} - x_t^*| + |y_{i,t} - y_t^*|} \quad (5)$$

points per round.⁴ There are four treatments, called Sink, SaddleNeg (for saddle negative), SaddlePos (for saddle positive), and Source. Each treatment involves a different matrix \mathbf{P} and vector \mathbf{b} . The matrices are

$$\mathbf{P}_{\text{Sink}} = \begin{pmatrix} 2/3 & 0 \\ -1/2 & -1/2 \end{pmatrix}, \quad \mathbf{P}_{\text{SaddleNeg}} = \begin{pmatrix} 2/3 & 0 \\ -1/2 & -3/2 \end{pmatrix},$$

$$\mathbf{P}_{\text{SaddlePos}} = \begin{pmatrix} 2/3 & 0 \\ -1/2 & 3/2 \end{pmatrix}, \quad \mathbf{P}_{\text{Source}} = \begin{pmatrix} 3/2 & 0 \\ -1/2 & -3/2 \end{pmatrix},$$

while the vectors \mathbf{b} in the treatments are chosen in such a way that all treatments have the same Nash Equilibrium, where every participant guesses $x = 90$ and $y = 20$. Note that all these matrices are upper-diagonal, but their diagonal elements are different in different treatments. As they vary, the predicted dynamics under a simple best-response model vary as well, and those dynamics give rise to the names of the four treatments; see **ADP22** for further discussion. This experiment has 16 sessions (4 per treatment), each with 10 participants and 15 rounds.

2.2. Predictions, experimental findings and models

A summary of the experimental designs, predictions, and the findings from the experiments that we will study using the IEL model is provided in Table 1.

⁴ Thus subjects are incentivized to get as close to each target as possible. By contrast with the tournament structure in (2)–(3), every participant in this study earns a positive point payoff in every round.

Theoretically, all KBC games with $p = 1/2$ and $p = 2/3$ in **N95** and all games in **DN97** have a unique Nash Equilibrium (NE) where everyone submits a guess of 0. These games are solvable via iterated deletion of dominated strategies. In the **N95** game with $p = 4/3$, in addition to the same NE, there is another NE, where everyone guesses 100. In all four treatments of the planar KBC game of **ADP22**, there is a unique interior NE where everyone submits the guesses $(x, y) = (90, 20)$. Moreover, the SaddlePos and Source treatments have additional NE on the boundaries of the strategy space, but in those equilibria, guesses are far from the target values, making $(90, 20)$ the unique Pareto-optimal NE. The theoretical predictions given by the NE are clearly not confirmed in the first periods of all experiments (and in other studies that focus on a single-period KBC game). However, in some cases, the NE can be seen as the limiting point of the experimental game dynamics. We now discuss the findings from the KBC experiments we study as well as the models that have been used to describe the dynamic of expectations in repeated versions of the KBC game.

N95 finds that over time, the median experimental guesses are moving towards the NE of 0 in the sessions with $p = 1/2$ and $p = 2/3$ and to the NE of 100 in the sessions with $p = 4/3$. This happens, as the majority of participants decrease their guesses between rounds in sessions with $p < 1$ and increase their guesses between rounds when $p = 4/3$. The rate of decrease in these median guesses is significantly higher (at the 5% level) when $p = 1/2$ than when $p = 2/3$.

DN97, where $p = 1/2$, find that the median guesses of subjects decline in all treatments but that the rate of change in the median guess is significantly larger in the median game than in the mean game and is significantly larger in the mean game than in the maximum game. Comparing the median guess dynamics in the longer session with 10 rounds, subjects converge to the NE of 0 for the median game (changing from 19 in the first round to 0 in the last round), and for the mean game (from 30 in the first round to 0.16 in the last round), but subjects do not converge to 0 in the maximum game by the end of 10 rounds (the change is from 40 in the first round to 20 in the last round).

For their planar, 2D KBC game, **ADP22** find that median guesses converge to the NE of $(90, 20)$ in the Sink and SaddleNeg treatments. However, in the SaddlePos and Source treatments, guesses converge to the Pareto-dominated Nash Equilibria on the boundary. In the SaddlePos, the first guess (x) converges to 90, while the second guess (y) stays near 100. In the Source treatment, the first guess (x) goes to the boundary of 0 and the second guess (y) stays in the vicinity of 38.

We will use the qualitative findings from these experiments, succinctly summarized in the last column of Table 1, as a minimal criterion for judging the success or failure of the IEL and other learning models as applied to the same KBC games played in these experimental studies. If a learning model generates dynamics similar to the experimental data, then such a model can be used for prediction purposes or for further experimental design. In addition to qualitative dynamics, we will also consider the quantitative fit of learning models to the experimental data.

3. Individual evolutionary learning

In this section we introduce the Individual Evolutionary Learning (IEL) model. IEL is a simple evolutionary mechanism of *individual learning*, which reinforces successful and discourages unsuccessful strategies. Variants of this model were used in Arifovic and Ledyard (2001, 2007, 2012); Arifovic et al. (2019) and other papers by Jasmina Arifovic and associates.

The choice of the strategy space is important as it involves strategic considerations that the game imposes. In the context of KBC games, a natural assumption is that the strategy of every agent is a guess (or a pair of guesses for the planar game).⁵ The IEL algorithm models the evolution of guesses through the following key elements:

- specification of a space of strategies;
- endowment of each participant with a finite pool of J strategies;
- selection of one strategy from the pool in a given period, which is then played;
- evolution of the strategy pool between periods, involving *experimentation* (with a probability ρ for each strategy) and *replication* based on relative fitness, determined by *foregone utility*.

As the IEL was built as a portable algorithm, we want to be as close as possible to its standard implementations and will set the parameters of the size of the pool and the probability of experimentation to their off-the-shelf values, $J = 100$ and $\rho = 0.033$, respectively – see, e.g., Arifovic and Ledyard (2001, 2011). Furthermore, we will use the same probability function for the strategy selection as is done in these IEL studies.

Specifically, individual i 's pool of J strategies used at time t , denoted as $G_{i,t}$, consists of potential guesses. A generic element from the pool is represented by g , where $g \in [0, 100]$ for the 1D KBC games, as in **N95** and **DN97** experiments. For the 2D KBC games, such as those in **ADP22**, $g = (g_x, g_y) \in [0, 100] \times [0, 100]$. Let $U_i(g)$ be a non-negative foregone utility value associated with $g \in G_{i,t}$. The *proportional* selection rule is given by⁶

⁵ With hindsight analysis of experimental data, one could argue that participants think not in terms of guesses but in terms of levels of iterated deletion of dominated strategies as in the level- k analysis. However, as our aim is to evaluate the ability of IEL to *predict* the experimental data, we choose a specification that does not rely on the knowledge of previously run experiments.

⁶ Arifovic and Ledyard (2001, 2011) model IEL with the same selection rule. Another popular specification for selection probabilities is the logit probability model. We made a corresponding robustness check and found that the results are not affected if this model is used instead. Intuitively, the exact choice is not as important, because the individual pools under the IEL are often quite homogeneous, dominated by the better-performing strategies. As all our utilities are non-negative, we prefer the parameter-free specification in (6).

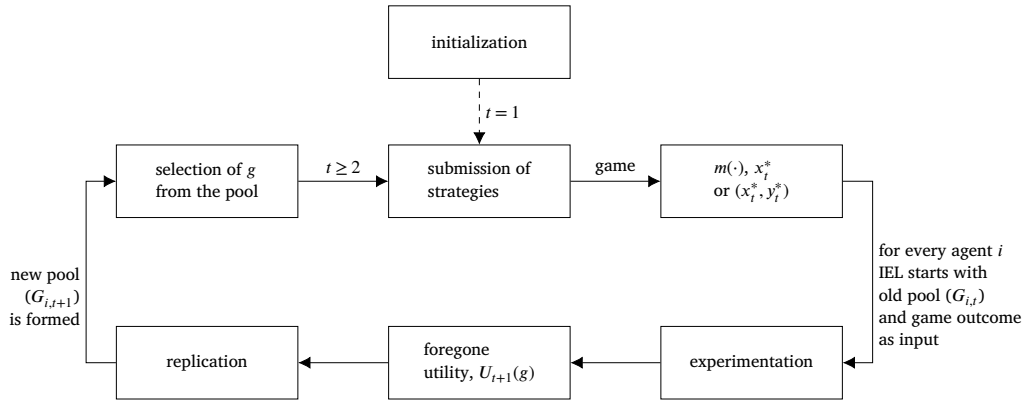


Fig. 1. Flowchart of simulations. Period t starts with a selection of strategies from pools and submission of them as guesses. The IEL algorithm that updates individual pools between periods t and $t + 1$ is shown in the lower part of the flowchart.

$$P(g) = \begin{cases} \frac{U_t(g)}{\sum_{g' \in G_{i,t}} U_t(g')} & \text{if denominator is positive} \\ \frac{1}{J} & \text{otherwise.} \end{cases} \quad (6)$$

At each period $t \geq 2$, artificial agent i draws one strategy from pool $G_{i,t}$ according to distribution (6). We denote the chosen element by $x_{i,t}$ or $(x_{i,t}, y_{i,t})$, depending on whether it is a 1D or 2D KBC game. Subsequently, the rules of the game determine the target x_t^* or targets (x_t^*, y_t^*) .

Fig. 1 illustrates the flowchart of the IEL algorithm. Rule (6) is used for selection when $t \geq 2$. The model is *initialized* in period $t = 1$. Each player's pool of J strategies is populated by uniformly drawing elements from the strategy space. Then one strategy from each player's pool is drawn uniformly at random and submitted as their guess.⁷ Once all strategies (guesses) are submitted, the target is determined and the IEL is ready to update individual pools. For later periods, strategies are drawn by IEL from endogenous pools, as we will now explain.

The evolution of the individual pools between periods t and $t + 1$ and the subsequent selection of strategies from the new pools are based on the foregone utility. This is not the same as the payoff earned in the experiment, as it involves “what if” thinking on the part of players. As such, it makes foregone utility dependent on the information and feedback provided to players between rounds. For the 1D experiments in **N95** and **DN97**, the all-or-nothing tournament structure with the payoff (3) was implemented. After each round participants could see the target and the winning guess as well as all other guesses. “What if” thinking then suggests that guesses that are closer to the target have a higher chance of winning. We formalize it with the following foregone utility specification,

$$U_{t+1}^{\text{ST}}(g) = \frac{1}{1 + (g - x_t^*)^2}. \quad (7)$$

This function is flat at its maximum of 1 reflecting the fact that guesses that are close to the target are likely to win; it drops towards 0 for guesses further away from the target. We call function (7) the *smooth-tournament* foregone utility, to distinguish it from the *tournament* foregone utility, defined as follows: it is equal to 1 if the guess g is no further from the target than the winning number in the most recent round and 0 otherwise. The details of the latter specification are discussed in Section 4.5.

For the 2D experiments in the planar KBC game conducted by **ADP22**, where participants earn points in rounds according to the payoff function (5), we utilize the *hyperbolic* foregone utility specification that mimics the actual payoff function. It is defined for each element $g = (g_x, g_y)$ in the pool as

$$U_{t+1}^{\text{H}}(g) = \frac{500}{5 + |g_x - x_t^*| + |g_y - y_t^*|}. \quad (8)$$

Agents' pools are updated through two procedures: experimentation and replication. At the *experimentation* stage, any element from the previous pool, $G_{i,t}$, of agent i is replaced by a new element with the probability of experimentation $\rho = 0.033$. With probability $1 - \rho$, the element remains unchanged. The new element is uniformly randomly drawn from the strategy space. For example, in the planar KBC, both g_x and g_y are drawn independently from $[0, 100]$.⁸ Denoting the pool obtained after the experimentation stage as G' , the subsequent *replication* stage, aimed at reinforcing successful strategies, proceeds as follows. To fill a position in the updated pool, two elements are randomly chosen from G' with replacement. These two elements are compared using the foregone

⁷ This is a natural, “agnostic” way to initialize IEL in order to use it as a testbed, see footnote 5. In Section 4, where we match IEL to the experimental data, we set first period strategies to the actual guesses made by participants in the first round of the experiments.

⁸ This experimentation is “global,” allowing IEL agents to try elements from any part of the strategy space at any time.

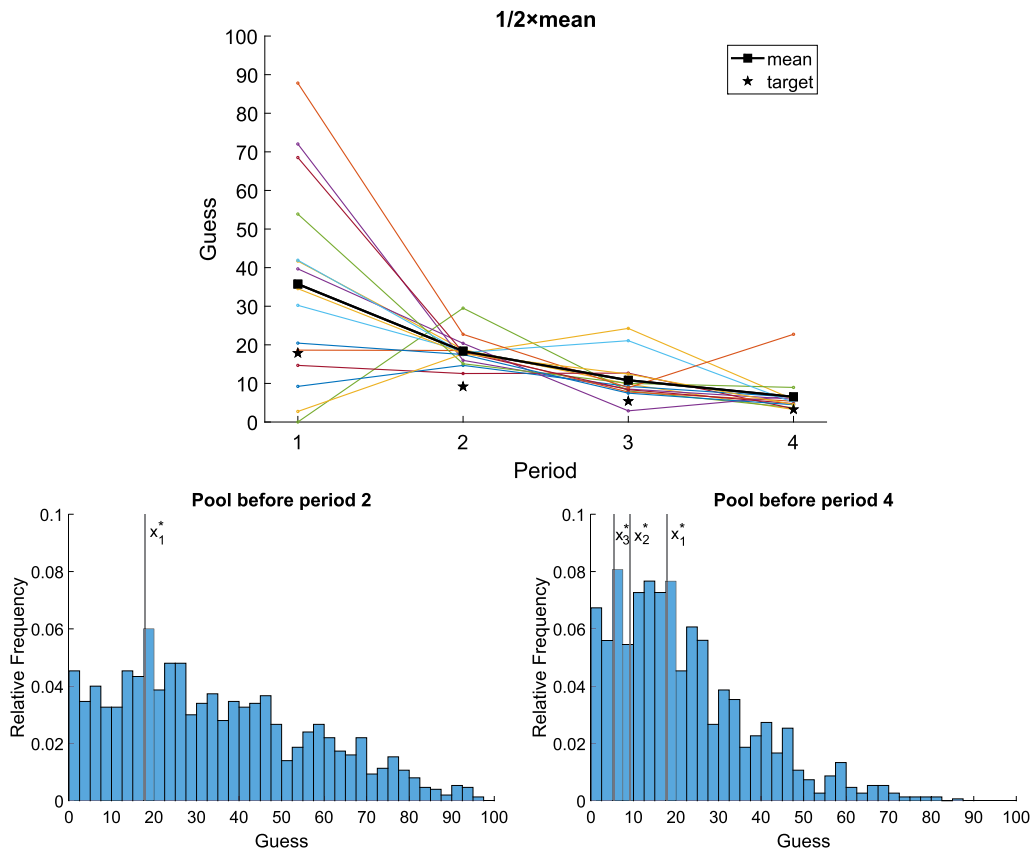


Fig. 2. Dynamics of individual guesses, the averages and the targets over 4 periods (upper panel) and histograms of the merged 15 individual pools in period 2 (left lower panel) and in period 4 (right lower panel).

utility U_{i+1} and the winner of the binary tournament is assigned a place in a new pool, $G_{i,t+1}$. In the case of a draw, each element is assigned to the new pool with a probability of $1/2$.

As the new pool of agent i has been formed, stochastic selection of the strategy (guess) from the pool takes place according to the probability rule (6). This IEL algorithm is run independently for every agent, after which the rules of the game determine the new target based on the selected guesses. Then IEL updates the pools again, utilizing the updated foregone utilities. And so on.

Following the above specification of the IEL, we illustrate its application to the KBC via a typical simulation. For this, we take the game with $p = 1/2$ and $m(\cdot)$ as the mean. That is, we consider the $1/2 \times \text{mean}$ game played in four sessions of N95 and four sessions of DN97. Fig. 2 illustrates one simulation for this game with 15 IEL agents.

The upper panel shows the evolution of the 15 IEL agents' individual guesses (thin colored lines), the mean guess (thick black line with squares), and the target value (indicated by stars) over the first 4 periods. In period 1, as the initial strategy pools were uniformly randomly determined, the 15 guesses are randomly distributed over the interval $[0, 100]$. The average of these guesses is 35.7 and the target is 17.9. After the first period, agents' strategy pools go through experimentation and replication, based on the foregone utility, $U_{i,2}(g)$, computed for all agents i and their guesses g according to (7). The lower left panel shows the histogram of the elements in all participants' pools, $G_{i,2}$, that are used to determine guesses in period 2. The vertical line indicates x_1^* , the previous target, where the foregone utility $U_{i,2}$ reaches its maximum. While aggregation over all agents may mask some differences in individual pools, the histogram illustrates the tendency of the IEL model to keep the heterogeneity of the pool and respond to past performance. In comparison with the uniform distribution used in the first period, the pools are shifted to the left with the highest mass around the first-period target. Recall that on top of this, different elements of the pools have different probabilities to be selected in period 2. Based on the pools and probabilities, the period 2 guesses are then generated. These guesses are, in expectation, close to the previous target and consequently lead to a smaller average ($\bar{x}_2 = 18.4$) and target ($x_2^* = 9.2$) than in the previous period.

The tendency of IEL to move guesses, the mean guess, and the target towards 0 over time will continue, even if the stochastic nature of IEL does not imply monotonicity in the dynamics (a spike from one of the IEL agents in period 4, as shown in Fig. 2, illustrates this point). The lower right panel shows the aggregate distribution of pools $G_{i,4}$. In comparison with the pools from period 2, the mass has moved further to the left. The vertical lines indicate the targets in periods 1, 2, and 3. Note that the mode of distribution is around x_3^* but the median is to the right of this. This indicates a certain inertia that IEL possesses and which comes from the fact that each new pool is largely based on the previous pool. That said, the guesses from the pool that are closer to the past target are more likely to be selected in every period given our foregone and IEL selection procedure.

4. Confronting the experimental data with IEL and other learning models

In this section, we evaluate the ability of IEL to reproduce the dynamic paths observed in repeated KBC experiments. To put IEL in perspective, we also compare it with other dynamic models that have been successful in describing such data. In particular, Nagel (1995) and Duffy and Nagel (1997) used the learning direction model and Anufriev et al. (2022b) focused on adaptive learning, which can be viewed as a mixture of level-0 and level-1 behaviors, thus we consider all three classes of models.

Simulations in this section will use the first period of the experimental data to initialize all models. We simulate the models for each available session from **N95**, **DN97**, and **ADP22**. As the initial conditions affect the trajectory strongly, in order to make a meaningful comparison across models, it is best to condition their forecasts on the same initial conditions. Thus, for each experimental session, we take the first-period guesses actually submitted, their aggregate statistic, and the target as the initial data for the simulated models. For the remaining dynamics there is *no further conditioning* on any experimental data.

4.1. IEL model

To generate the IEL dynamics, we use the algorithm as described in Section 3 with the modification discussed above for the initialization stage. To do this, we populate the initial pools, $G_{i,1}$, of all agents uniformly randomly, but do not use the elements of these pools to generate the first-period guesses and targets. Instead, we take the experimental targets, $x_1^{e,*}$, and evaluate individual pools based on these targets to form the pools for the period 2. The IEL dynamics follow from the usual algorithm starting from beginning of period 2.

4.2. Learning direction model

Based on the experimental evidence, **N95** puts forward a hypothesis that individuals make their guesses as some factor of the past averages. That is, a guess $x_{i,t}$ of agent i is the previous mean, \bar{x}_{t-1} , times the so-called adjustment factor, $a_{i,t}$. The learning direction (LD) model concerns how these individual adjustment factors change between different rounds of play. The model posits that participants compare their individual adjustment factor against the *ex post* optimal adjustment factor and then increase their factors if they were too low and decrease them if they were too high. **DN97** shows that this model works well also for the median and maximum treatments in the KBC game.

To write this model formally, we use \mathbf{x}_t as the vector of all guesses at time t . Recall that $m(\mathbf{x}_t)$ is the statistic (such as average, median, or maximum) of the guesses that is relevant for the treatment. The guess of agent i at time $t \geq 2$ is then

$$x_{i,t} = m(\mathbf{x}_{t-1}) \cdot a_{i,t}. \quad (9)$$

The LD model states that individual adjustment factors change between rounds as follows:

$$a_{i,t-1} \leq a_{t-1}^{\text{opt}} \Rightarrow a_{i,t} \geq a_{i,t-1}, \quad (10)$$

where, consistently with (9), the *optimal* adjustment factor for $t \geq 2$ is defined based on the realized target as

$$a_t^{\text{opt}} = \frac{x_t^*}{m(\mathbf{x}_{t-1})}. \quad (11)$$

Thus, if a_t^{opt} was used in place of $a_{i,t}$ in (9), the agent's guess would coincide with the target and the agent would be the winner (or among the winners). **N95** and **DN97** tested the LD model by counting the instances when (10) holds, with individual adjustment factors found from (9) using the actual guesses. They found that the majority of data supports (10).

To generate the dynamics of guesses made by the LD model, some functional form is required for the evolution of the adjustment factors. We use the simplest linear model (closely related to an error correction model) that is consistent with (10) and assume that

$$a_{i,t} = a_{i,t-1} + \gamma(a_{t-1}^{\text{opt}} - a_{i,t-1}), \quad (12)$$

where the parameter $\gamma > 0$.

Eqs. (9) and (11) are defined for $t \geq 2$ only, because the statistic, $m(\mathbf{x}_{t-1})$, is available only after the first session. **N95** defines the individual and optimal adjustment factors for $t = 1$ by using the focal point 50 as an anchor to which adjustment is applied. We follow this route to initialize the LD model dynamics. Using the experimental data (indicated by the superscript e) of the first round, we derive $a_{i,1} = x_{i,1}^e / 50$ for all i , and $a_1^{\text{opt}} = x_1^{e,*} / 50$. Then, starting from period $t = 2$, the following recursive procedure is applied. First, Eq. (12) is used to define the individual adjustment factors for the period. Then Eq. (9) is used to define the individual guess $x_{i,t}$. Using all the guesses, the game rules determine the new target. Finally, Eq. (11) is applied to find the optimal adjustment factor for this period. Then the process repeats.

Note that the idea of the LD model is closely related to the recursive levels of rational reasoning. In the introduction, we discussed that in a game with the target given by Eq. (1), level- k participants guess p^k times the previous round average, which is a special case of (9). The LD model allows for more flexible guesses that coincide with level- k guesses whenever $k = \log a_{i,t} / \log p$ is an integer number. However, the LD model works in this way only when the target equation does not include a constant term. The planar KBC game studied in Anufriev et al. (2022b) included a constant term. This is why the LD model was not among the models considered by **ADP22** and, for the same reason, we do not simulate the LD model in this paper for the **ADP22** data.

Table 2

Estimated coefficients for both the Learning Direction (LD) model and the Adaptive model (ADA), for each experimental session. For the 1D experiments (N95 and DN97), estimates for x -guesses are presented, while for the 2D experiments (ADP22), estimates for both x and y -guesses are shown.

Treatment			Sessions			
LD model, estimates of γ						
N95	1/2 \times mean		−1.32	0.48	0.73	
N95	2/3 \times mean		−0.18	0.87	0.85	0.74
N95	4/3 \times mean		0.13	0.28	0.21	
DN97	1/2 \times mean		0.20	0.95	0.05	1.01
DN97	1/2 \times median		0.42	0.83	0.61	0.67
DN97	1/2 \times max		0.39	0.75	0.41	0.19
ADA model, estimates of α						
N95	1/2 \times mean		0.66	0.91	0.82	
N95	2/3 \times mean		0.32	0.81	0.77	0.82
N95	4/3 \times mean		0.31	0.38	0.41	
DN97	1/2 \times mean		0.62	0.85	0.70	0.83
DN97	1/2 \times median		0.55	0.76	0.63	0.81
DN97	1/2 \times max		0.67	0.82	0.47	0.67
ADP22	Sink	x	0.43	0.78	0.73	0.89
		y	0.67	0.73	0.61	0.86
ADP22	SaddleNeg	x	0.48	0.56	0.71	0.64
		y	0.78	0.65	0.79	0.49
ADP22	SaddlePos	x	0.84	0.53	0.76	0.68
		y	0.24	0.25	0.18	0.19
ADP22	Source	x	0.22	0.34	0.25	0.19
		y	0.46	0.72	0.56	0.55

4.3. Adaptive model

The LD model uses the evolution of the *adjustment factors* as in (10) or (12). Alternatively, we can model the evolution of the *individual guesses* directly as the simple adaptive process:

$$x_{i,t+1} = x_{i,t} + \alpha(x_t^* - x_{i,t}). \quad (13)$$

Here, the parameter $\alpha \geq 0$ measures how strongly an agent adjusts the previous guess in response to the prior guess error. Note that model (13) is equivalent to the adaptive expectations model, where $0 < \alpha < 1$, and so we call it the adaptive (ADA) model.

The advantage of this model is that it can be applied without modification to the experiments where the target equation includes a constant term. ADP22 found that the mixed-level 0 and 1 type model, which is related to the adaptive model above, is among the best performing models for characterizing the dynamics of the planar KBC experimental data.

The simulation of the ADA model is straightforward. The experimental data from the first round, $x_{i,1}^e$, and the experimental target, $x_1^{e,*}$, are plugged into (13) to generate the guesses for the period $t = 2$. These guesses define the new target according to the rules of the game, and the process then repeats.

We also consider the special case of the ADA model, where $\alpha = 1$. In this case, the agent submits the previous target, i.e., plays the best response to the last period average. This model is known as the *Cournot Best Response* (CBR). Its analogy in the macro-literature is called *naïve* expectations.

4.4. Model comparisons in terms of simulated paths

We condition simulations of all learning models on the experimental data from the first round and generate the simulated path for the same number of periods as in the corresponding experiment. Further, we set the number of IEL agents to be equal to the number of participants in the experimental session that we seek to reproduce. IEL is a stochastic model and the generated guesses depend on specific realizations of random numbers.⁹ Hence, we produce 1,000 IEL simulations using different random seeds. The parameters of the LD and ADA models (γ and α , respectively) are estimated for each experimental session and are assumed to be the same for all participants in a given session. We report the estimated parameter values in Table 2.

Notice that the parameter estimates of γ for the LD model and for α for the ADA model are quite heterogeneous, even within the same treatment. For example, the 2/3 \times mean treatment of N95 has estimates for γ ranging from −0.18 to 0.87, and for α the

⁹ Learning direction and adaptive models are also stochastic if an error term is added to the model's equation, i.e., Eq. (12) and (13), respectively. We use the conditional forecasts of these stochastic models, and those forecasts coincide with the medians.

Table 3

RMSEs based on the median guesses for the one-dimensional KBC games for periods 2–4 in four learning models (CBR: Cournot best response, ADA: adaptive, LD: learning direction, IEL: individual evolutionary learning). The smallest RMSE is in bold.

		CBR	ADA	LD	IEL
N95	1/2 × mean	4.60	6.27	5.10	4.83
N95	2/3 × mean	5.15	9.55	5.03	5.35
N95	4/3 × mean	3.63	17.68	3.87	6.25
DN97	1/2 × mean	3.59	6.36	4.16	3.57
DN97	1/2 × median	2.35	5.38	41.57	2.45
DN97	1/2 × max	15.60	14.69	40.49	14.49
Overall		7.56	10.79	25.01	7.44

Table 4

RMSEs based on the median guesses for the planar KBC games for periods 2–15 in three learning models (CBR: Cournot best response, ADA: adaptive, IEL: individual evolutionary learning). The smallest RMSE is in bold.

		CBR	ADA	IEL
ADP22	Sink	5.89	5.46	4.14
ADP22	SaddleNeg	17.18	4.72	4.89
ADP22	SaddlePos	9.01	11.72	6.45
ADP22	Source	33.79	8.24	4.91
Overall		19.70	8.02	5.16

range is 0.32 to 0.82. We will use the estimated coefficients from Table 2, when comparing the performance of the models with the experimental data later in this section, see Tables 3 and 4.

Fig. 3 compares the simulated path dynamics of the considered models and the actual experimental data for one session of each of the one-dimensional KBC games in **N95** (left panel) and **DN97** (right panel). The *median* guess among all participants in each period is used as the measure of comparison; those guesses are shown as black lines with circles (labeled data). The shadows of the red right arrow symbol show the quantiles of the distribution of the IEL median guesses (over all agents). The median of this distribution is shown in dark red. The guesses generated by the LD model are shown in magenta, and the guesses generated by the ADA model are shown in blue. For the LD and ADA models, we use the estimated parameters for γ or α reported in Table 2 (for the corresponding session). Finally, the guesses generated by the CBR (Cournot Best Response) model are shown in green.

All considered models track the actual experimental data well. IEL tends to be slightly slower in convergence than the actual experimental data in all games except for the 1/2 × max game. The IEL and CBR models show a closer fit to the data.

In Fig. 4 we compare the dynamics for the planar KBC game of Anufriev et al. (2022b) with x -guesses shown in the left panels and y -guesses shown in the right panels. We compare the same models and use the same color scheme as in Fig. 3. The only difference is that we do not include the learning direction model in these simulations, as it is not suitable for the games with a constant term in the target equations. For these planar KBC data, the CBR model does not perform well, while the IEL and adaptive models are much closer to the data.

To formally assess the fit of the models to the data, we compare the median of individual guesses generated by each model m , \tilde{x}_t^m , against the median of individual guesses in the experimental data, \tilde{x}_t^e . We use the Root Mean Squared Error (RMSE) criterion and, as we condition on period 1 data in all models, we start calculating errors from period $t = 2$. For a given treatment the RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{S} \sum_s \frac{1}{T_s - 1} \sum_{t=2}^{T_s} d_{t,s}},$$

where the first sum is over all S sessions in the game, T_s is the number of periods, and $d_{t,s}$ measures how close the model is to the data in session s at time t . This measure is defined as

$$d_{t,s} = \begin{cases} \left(\tilde{x}_{t,s}^e - \tilde{x}_{t,s}^m \right)^2 & \text{for N95 and DN97} \\ \frac{1}{2} \left(\left(\tilde{x}_{t,s}^e - \tilde{x}_{t,s}^m \right)^2 + \left(\tilde{y}_{t,s}^e - \tilde{y}_{t,s}^m \right)^2 \right) & \text{for ADP22.} \end{cases}$$

We use the median of guesses because for all our treatments guesses always stay within the $[0, 100]$ interval, whereas the range of targets varies substantially depending on the treatment. For the IEL model, the median (of the median of individual guesses) is

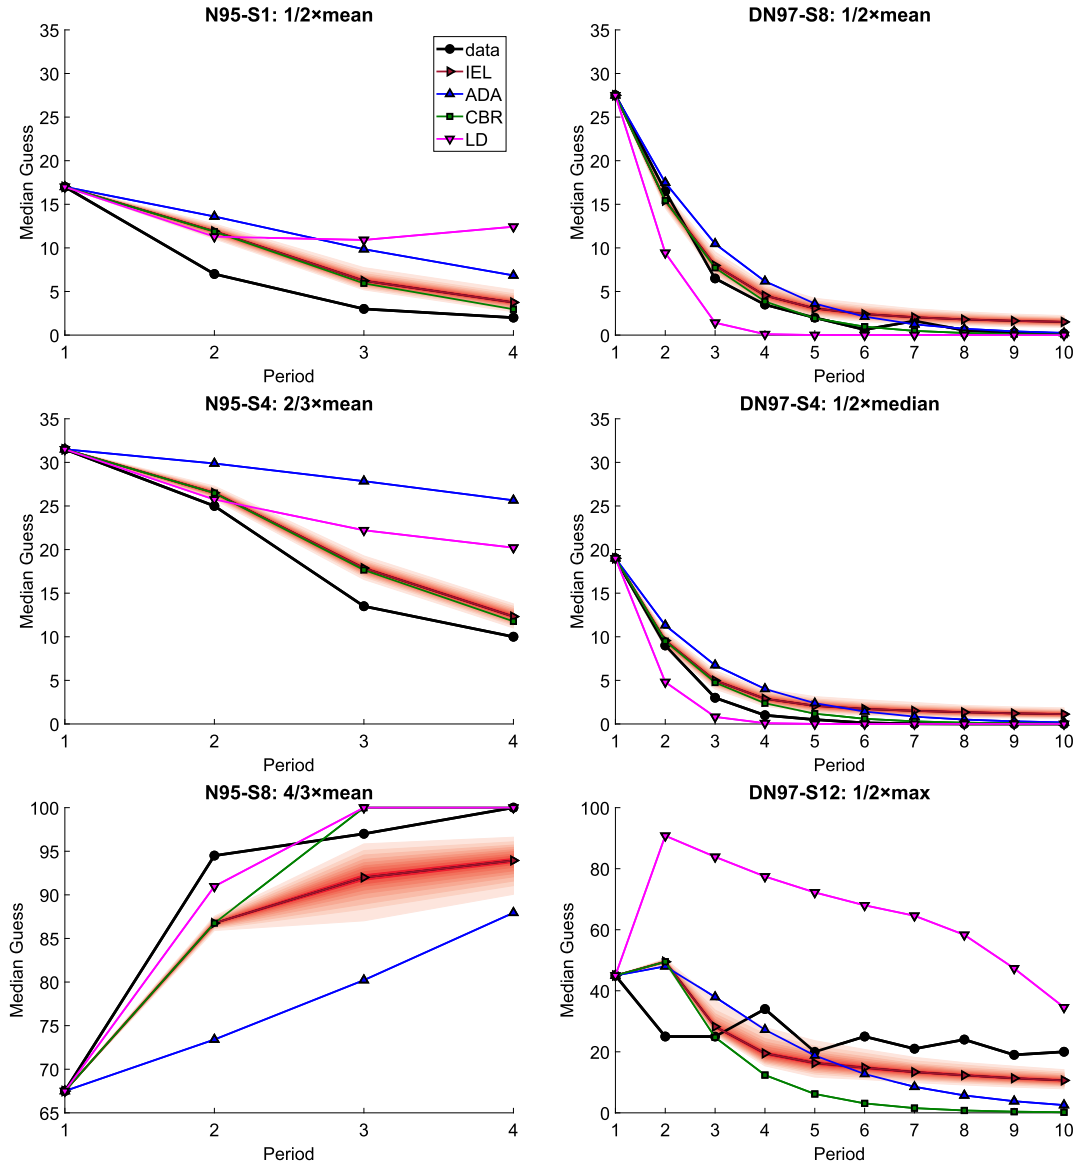


Fig. 3. Performance of learning models for the representative sessions of the KBC experiments. *Left panels:* N95 session 1 for $1/2 \times \text{mean}$ game (top), session 4 for $2/3 \times \text{mean}$ game (middle) and session 8 for $4/3 \times \text{mean}$ game (bottom). *Right panels:* DN97 session 8 for $1/2 \times \text{mean}$ game (top), session 4 for $1/2 \times \text{median}$ game (middle) and session 12 for $1/2 \times \text{max}$ game (bottom). All panels show the experimental data (black), distribution of 1,000 IEL simulations (shades of red) with the median of the simulations (dark red), adaptive (ADA) model (blue), CBR model (green), and LD model (magenta). All simulations are conditional on the experimental guesses in round 1. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

computed over 1,000 simulation runs to account for the randomness of the IEL. For the games of N95 and DN97 we use periods 2 to 4,¹⁰ for the planar KBC of ADP22 we use periods 2 to 15 as in the experiment.

Tables 3 and 4 report the RMSEs for the one-dimensional games of N95 and DN97, and the planar KBC games of ADP22, respectively. We observe that the IEL model produces the smallest RMSE in most of the one-dimensional and in all planar KBC games. IEL has the smallest overall RMSE.

4.5. Robustness checks and alternative designs

We have emphasized that we are using an off-the-shelf parameterization of the IEL model, with an experimentation probability $\rho = 0.033$ and a pool size $J = 100$. We now explore the robustness of our findings concerning alternative choices for these two

¹⁰ In DN97 there were 3 sessions, one for each game, with 10 rounds each. For comparability with all the other sessions that only had 4 rounds we compute the RMSE only for periods 2 to 4 of these sessions.

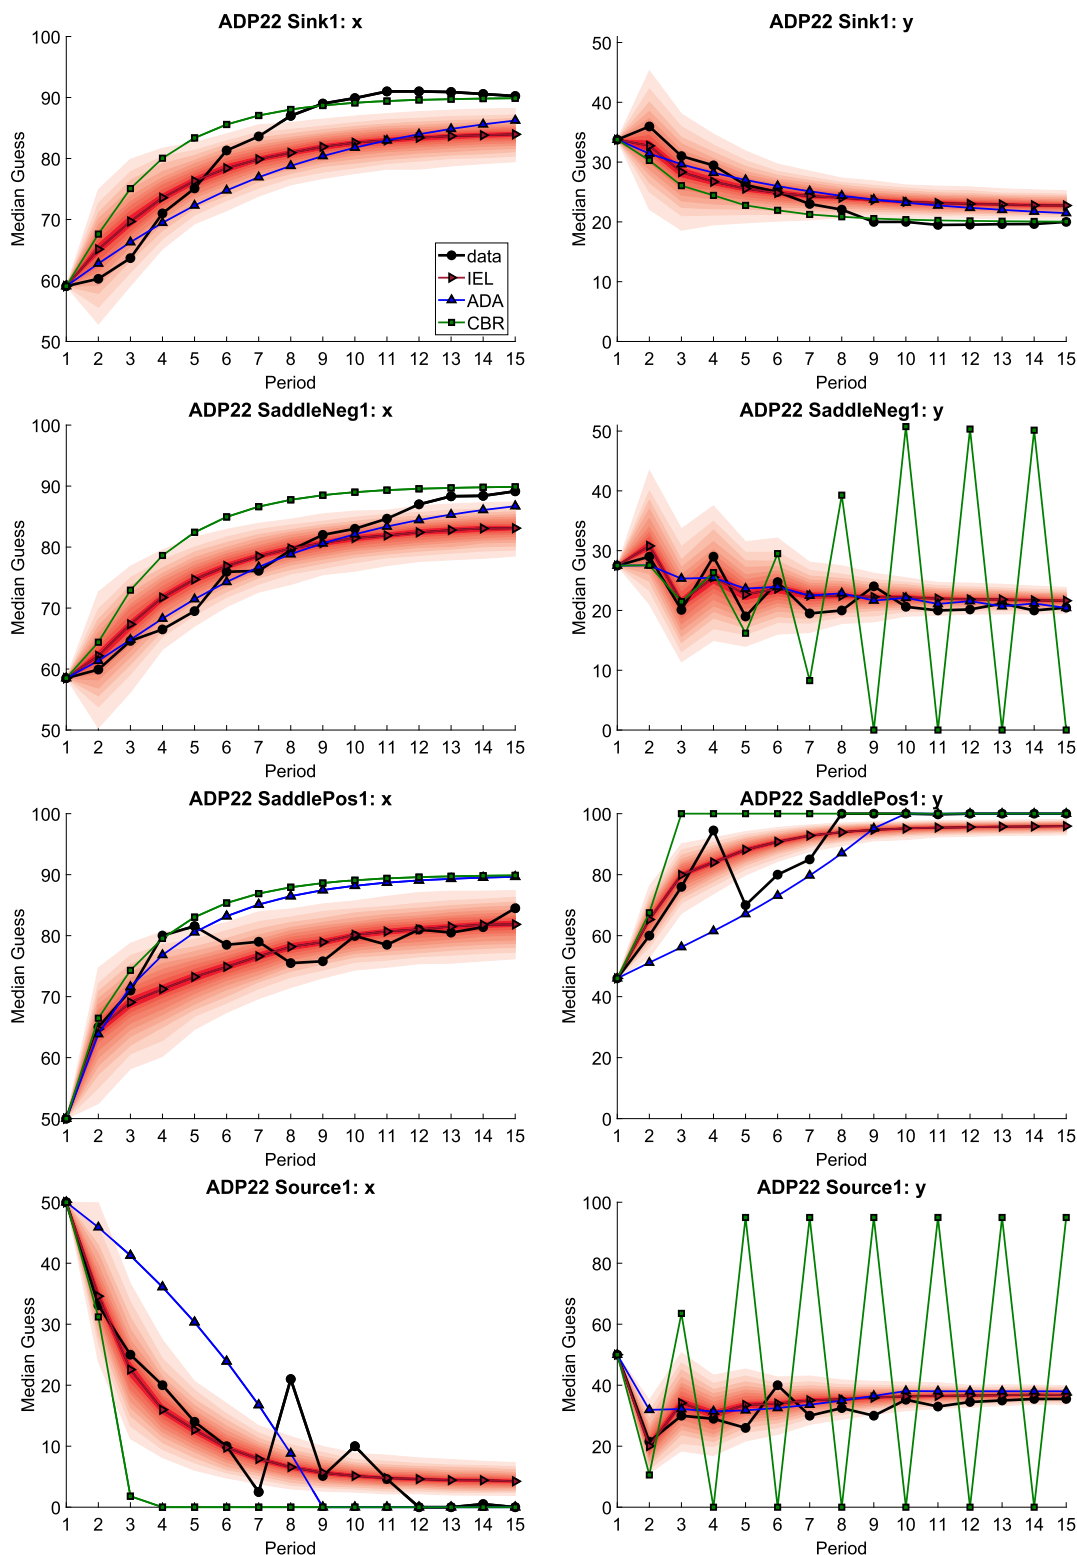


Fig. 4. Performance of learning models in four treatments of the planar KBC game in ADP22. *Left*: x-guesses. *Right*: y-guesses. All panels show the experimental data (black), distribution of 1,000 IEL simulations (shades of red) with the median of the simulations (dark red), adaptive model (blue), and CBR model (green). All simulations are conditional on the experimental guesses in round 1.

Table 5

Sensitivity analysis of the overall RMSE concerning variations in the probability of experimentation, ρ , and pool size, J , for the 1D experiments (N95 and DN97). The RMSEs are reported for the medians from 1,000 runs of the IEL model for each of the 1D treatments.

$\rho \backslash J$	10	30	50	100	150	200
0.01	9.56	7.68	7.50	7.44	7.43	7.43
0.02	9.56	7.68	7.49	7.44	7.43	7.43
0.033	9.51	7.66	7.48	7.44	7.43	7.43
0.05	9.49	7.67	7.51	7.44	7.45	7.43
0.1	9.31	7.65	7.51	7.44	7.44	7.45
0.2	9.16	7.66	7.49	7.47	7.46	7.46
0.33	8.98	7.63	7.52	7.48	7.49	7.49
0.5	8.87	7.63	7.55	7.52	7.52	7.52

Table 6

Sensitivity analysis of the overall RMSE concerning variations in the probability of experimentation, ρ , and pool size, J , for the 2D experiments (ADP22). The RMSEs are reported for the medians from 1,000 runs of the IEL model for each of the 2D treatments.

$\rho \backslash J$	10	30	50	100	150	200
0.01	9.36	6.28	5.60	5.16	5.04	4.99
0.02	9.16	6.19	5.55	5.15	5.03	4.96
0.033	8.93	6.09	5.53	5.16	5.02	4.97
0.05	8.49	5.93	5.48	5.14	5.05	5.00
0.1	7.96	5.88	5.42	5.17	5.11	5.02
0.2	7.40	5.90	5.51	5.32	5.24	5.21
0.33	7.68	6.30	6.04	5.85	5.75	5.73
0.5	9.31	7.99	7.77	7.56	7.48	7.46

parameters of the IEL model. Additionally, we examine the robustness of the method we employ to approximate the hypothetical utility of players in sessions with an all-or-nothing tournament payoff structure, which was frequently used in the initial KBC game experiments in one dimension.

Regarding the parameterization of IEL, Tables 5 and 6 show the overall RMSE of IEL simulations relative to the median guesses from 1,000 runs each of all 1D and 2D treatments, as functions of different combinations of ρ and J . These overall RMSE numbers can be compared with those reported in Tables 3 and 4, specifically with the overall RMSE numbers in those tables. We observe that the overall RMSE is low for the off-the-shelf IEL parameterization of $(\rho, J) = (0.033, 100)$ that we use, and there is little value to changing ρ or J in the 1D treatments. In the 2D treatments, there may be some improvement in terms of RMSE from a larger strategy space (J) in combination with a lower experimentation rate (ρ), but one also needs to consider the computational speed cost of using a larger set of strategies. Overall, Tables 5 and 6 give the impression that the off-the-shelf parameterization of the IEL model is indeed a good, uninformed choice.

We further consider the robustness of our results in the 1D treatments to the manner in which the forgone utility is specified and how experimentation is performed. The smooth-tournament specification U^{ST} defined in Eq. (7) is our baseline choice for foregone utility in the 1D treatments. As an alternative, we use foregone utility directly based on the all-or-nothing tournament structure of the 1D treatments. Using the distance between the target and the winning number, $d_i^W = \min_i |x_{i,t} - x_i^*|$, we define the *winning interval* within the set of permissible guesses as

$$I_t^W = [x_i^* - d_i^W, x_i^* + d_i^W] \cap [0, 100]. \quad (14)$$

Then, the *tournament* foregone utility is defined as

$$U_{t+1}^T(g) = \begin{cases} 1 & \text{if } g \in I_t^W, \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

Using either the smooth-tournament (7) or tournament (15) specification for foregone utility, under global experimentation new values are drawn from the entire allowable range $[0, 100]$. As a further robustness exercise, we also explore the case of the so-called *targeted* experimentation, where new values are uniformly drawn not from the entire allowable range but instead from the winning interval I_t^W defined in (14). This interval depends on the most recent winning number. The targeted experimentation is also parameter free (as the global experimentation) but it directs the experimentation towards the range where previous round behavior could bring success.

Table 7

Root mean square errors (RMSE) under alternative specifications for foregone utility and experimentation for the IEL model relative to the experimental data. The RMSEs are reported for the medians from 1,000 runs of the IEL model for each of the 1D treatments.

		Smooth-tour.		Tournament	
		Global	Targeted	Global	Targeted
N95	$1/2 \times \text{mean}$	4.83	4.67	20.43	4.77
N95	$2/3 \times \text{mean}$	5.35	5.23	16.23	5.27
N95	$4/3 \times \text{mean}$	6.25	4.55	28.43	25.89
DN97	$1/2 \times \text{mean}$	3.57	3.57	15.05	3.60
DN97	$1/2 \times \text{median}$	2.45	2.39	19.39	2.35
DN97	$1/2 \times \text{max}$	14.49	14.83	22.02	18.13
Overall		7.44	7.36	20.31	12.75

In Table 7 we report the impacts of various foregone utility schemes and experimentation designs on the RMSE of the median guesses in IEL simulations compared to experimental data. The off-the-shelf specification that we have been using is the smooth-tournament, global experimentation specification. Thus, the third column of Table 7 repeats the RMSE results for IEL from Table 3.

Relative to this baseline, the addition of targeted experimentation to the smooth-tournament specification for foregone utility, U^{ST} , results in a slight improvement in the fit to the experimental data. Targeted experimentation plays a more important role in reducing the RMSE if we use the tournament foregone utility specification, U^{T} . For this utility, under global experimentation, there is a high chance that none of the guesses g in the new pool following experimentation falls within the winning interval I_t^W , especially when this interval is small. In that case, all such guesses receive the foregone utility of 0 and a new guess is selected from the pool uniformly at random. This may repeat for several periods leading to non-convergence and substantial differences from the experimental data, as indicated by the high overall RMSE for this specification in Table 7. The targeted experimentation avoids this issue as new values are drawn from the winning interval I_t^W . While the targeted experimentation greatly helps in the case of tournament foregone utility, as Table 7 makes clear, its improvement over off-the-shelf global experimentation is small under the smooth-tournament structure that we used in our baseline simulations.

In summary, the off-the-shelf values of the IEL parameters and the global experimentation that we used result in RMSEs that are already very close to the smallest RMSEs observed in our robustness exercise.

5. IEL as a testbed

In the previous section, we demonstrated the good fit of IEL when it is initialized using the actual experimental first-period data. However, IEL was initially proposed to simulate experiments *before* they are run, as discussed in the introduction, and we now consider the use of IEL as a testbed.¹¹ In subsection 5.1 we investigate the ability of IEL to reproduce typical features of the KBC game experiments with *random initializations*, i.e., when the IEL agents' pools are formed based on uniform distributions and flat (uninformed) selection functions. Later, in subsection 5.2 we show how IEL can be used to make crisp predictions for some interesting experimental treatments of the repeated KBC that have yet to be conducted.

5.1. Random initialization

In Section 3, we illustrated the consequences of a single random initialization of guesses in the first 4 periods of the $1/2 \times \text{mean}$ game. We now run 1,000 simulations (corresponding to different stochastic initializations and IEL realizations) for a larger number of periods.

Fig. 5 shows the evolution of the IEL-generated guesses in the $1/2 \times \text{mean}$ game. The left panel shows the evolution of the distribution of the median guess over 1000 simulations and its median for 10 periods. The right panel shows the cumulative distribution of individual guesses for periods 1, 5, and 10 from 1000 runs. We observe a relatively fast convergence of the median guess to the equilibrium. The median guess starts at 50 but it is very close to 0 by the end of 10 periods. The variance of the distribution of individual guesses reduces with periods. This is a common feature in the IEL simulations as well as in the actual experimental data across all games. By contrast in the $1/2 \times \text{max}$ game, the reduction in variance is much slower both in the IEL simulations and the real experimental data.

Fig. 6 compares the IEL dynamics across treatments of the **N95** and **DN97** that we can treat as different games. Note that we use the log scale in this figure. The limit point, where the dynamics converge, and the ranking in the speed of convergence are exactly the same as in the experimental data (see the last column of Table 1).¹² For example, the $2/3 \times \text{mean}$ game converges to 0 slower than the $1/2 \times \text{mean}$ game, as observed by Nagel (1995), and the $1/2 \times \text{median}$ game converges to 0 faster than the $1/2 \times \text{mean}$ game, as found

¹¹ A testbed is a controlled and reproducible environment that allows researchers to manipulate variables in order to observe how a system or technology responds under different conditions.

¹² We do not aim to do the exact quantitative comparison here, as there were a relatively small number of periods in the experiment.




Fig. 5. Results from 1,000 IEL simulations of the $1/2 \times \text{mean}$ game with 15 agents. Left: Evolution of median guesses over 10 periods. The median guess is shown by the thick dark line. Right: Empirical CDF of guesses in periods 1, 5 and 10.




Fig. 6. The dynamics of the medians over 1000 IEL simulations in three treatments of N95 (left) and DN97 (right) experiments. The y-axes are shown on the log-scale.

in Duffy and Nagel (1997). Interestingly, when running the $1/2 \times \text{max}$ game for only 10 periods (as DN97 do in their session 10), it is hard to see whether the dynamics converge to equilibrium. IEL suggests that it does eventually converge, but the convergence takes much longer than in the other two treatments.

Fig. 7 displays the time evolution of the IEL simulations for the planar KBC game from ADP22, for both x and y guesses. In all four treatments, both variables converge to the corresponding limit values observed in the experiment, which are shown and labeled in the figure. In particular, the IEL model predicts correctly the two treatments (Sink and SaddleNeg) that exhibit convergence to the Pareto optimal (internal) Nash Equilibrium.

5.2. Using IEL for prediction/design

The IEL model can also be used for simulating and predicting scenarios for other parameterizations, including those where experiments have not yet been run. For example, Sutan and Willinger (2009) (SW09) conduct a *one-shot* KBC game experiment where $N = 8$ participants make guesses in treatments with positive and negative feedback and are paid according to the all-or-nothing tournament structure. The targets in these treatments are determined as

$$x_t^* = \frac{2}{3}(\bar{x}_t + 30) \quad \text{and} \quad x_t^* = 100 - \frac{2}{3}\bar{x}_t,$$

respectively. In both treatments, the unique NE is when everyone guesses 60, which is strictly inside the range of possible guesses $[0, 100]$, different from the experiments in N95 and DN97. Notice further that the first treatment involves positive feedback of the mean guess on the target value while the second treatment involves negative feedback. SW09 find that the subjects' guesses are closer to the NE under negative rather than positive feedback. However, since SW09 did not repeat their game beyond a single round, we do not know how the guesses would evolve further in a repeated game setting. We can use IEL to predict such behavior over time. For that purpose, we again ran 1,000 simulations with 100 periods each and found that for both treatments of SW09, the IEL model always converges to the NE. However, there are differences in the convergent path that we illustrate in Fig. 8 (upper panels) where we show the first 15 periods from our 1,000 IEL simulations for each treatment (negative feedback on the left and positive feedback on the right).

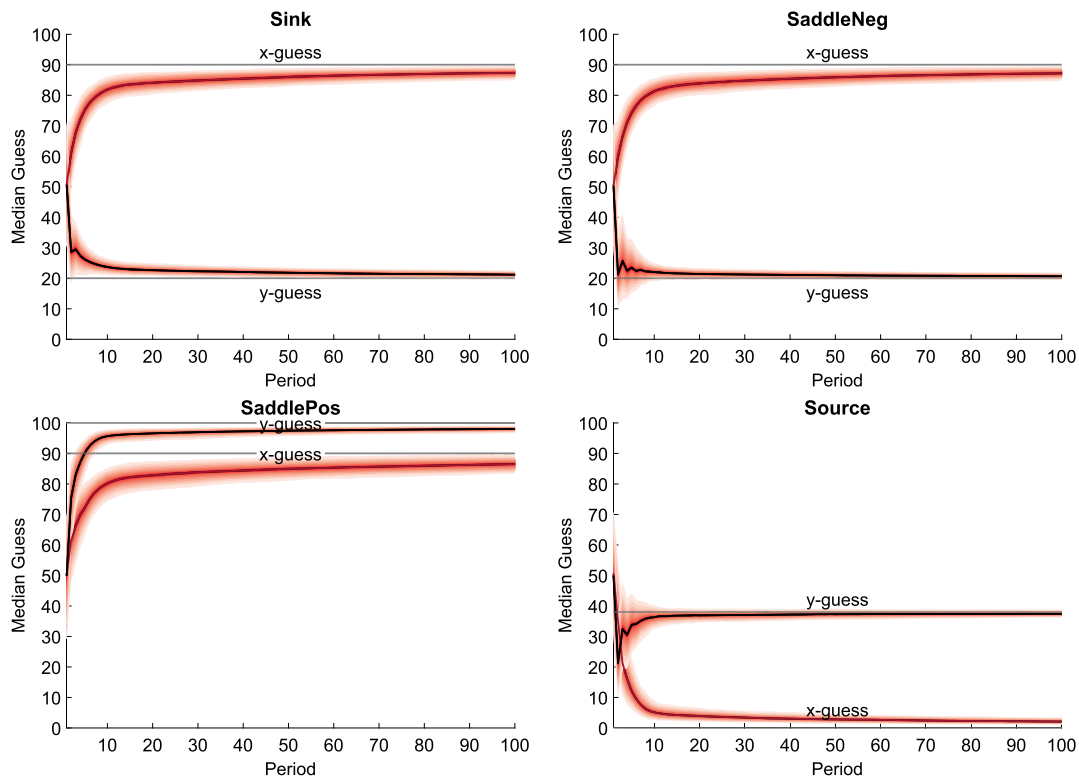


Fig. 7. The dynamics of the median guesses in the four treatments of the ADP22 experiment. The shades of red show the quantities of the distribution over 1000 IEL simulations. The dark red line shows the median of the simulations for the x guess, the black line shows the median for the y guess. The labeled horizontal line shows the limit values, where the corresponding guesses converged in the experiment.

We see that the differences between positive and negative feedback learning observed by SW09 persist beyond the first round. While both systems start at a guess centered around 50 (due to the random initialization) the negative feedback system displays more rapid but oscillatory convergence around the NE value while the positive feedback system converges more slowly and from below, on average.

This difference between negative and positive feedback has been also studied in related literature on “Learning-to-Forecast” experiments, see e.g., Hommes (2011) for a review. These experiments typically employ near-unit-root specifications in the data-generating process, that is, they have a relatively high absolute value for the slope coefficient, p , in the linear map. For example, Heemeijer et al. (2009) (HHST09) report results from negative and positive feedback treatments where the slopes are $p = -20/21$ and $p = 20/21$, respectively. By adding a certain constant term to each equation, they also obtain 60 as the same NE in both of their treatments. While those experiments differ from the KBC game in a number of dimensions (e.g., in the information provided to participants, the incentive structure and the number of participants), the dynamical system is essentially the same and it is of interest to check how the KBC game would be played with this much higher (in absolute value) coefficient for p . Thus, we run 1,000 IEL simulations, all with random initializations and using the same specification as for SW09 (namely $N = 8$, and an all-or-nothing tournament payoff structure) but with the targets determined as in HHST09. The middle panels of Fig. 8 show results from simulations of these negative and positive feedback treatments. Simulations in both treatments converge to the NE of 60 but this process takes much longer than in the SW09 specification with $p = 2/3$. Simulations under positive feedback do not converge within the first 15 periods, in contrast to the positive feedback case of $p = 2/3$, or the case of negative feedback (either $p = 2/3$ or $p = 20/21$).

As a final experiment, we consider how behavior would change with an even lower value for p . The lower panel of Fig. 8 shows the results of 1,000 randomly initialized simulations for the case where the NE remains at 60 but where the target equation has a slope of $1/3$ in absolute value. Compared with the simulations for the negative and positive SW09 treatment specifications, we see that convergence in this case is much faster. Intuitively, as little weight is being placed on the expectations feedback term, the constant term plays more of a stabilizing role and so the equilibrium is reached more quickly.

The bottom line of this analysis is that the absolute value of p , that is, the slope-coefficient of the map, matters for the speed of convergence. The greater the absolute value, the slower the speed of convergence and this is true for both positive and negative feedback. Convergence is usually faster under negative feedback, but for low enough p that difference with respect to the positive feedback case seems to disappear.

Taken together, these simulations demonstrate that an off-the-shelf IEL model, without conditioning on experimental data, can deliver interesting predictions that could be verified in experimental tests, or that could even replace the costly process of collecting such data.




Fig. 8. IEL simulation for **SW09** ($|p| = 2/3$), **HHST09** ($|p| = 20/21$) and the new ($|p| = 1/3$) specifications. Negative feedback treatments are in the left column and positive feedback treatments are in the right column. Simulation results are from 1,000 rounds for different values of the initial seeds. The shades of red show the quantities of the distribution over the 1,000 simulations. The dark black line shows the median of the simulations. The dashed horizontal line at 60 indicates the limiting NE value which is the same in all treatments.

6. Conclusion

The Individual Evolutionary Learning (IEL) model was designed to simulate individual play in environments with large strategy spaces. As with any learning model, IEL has parameters (the size of individual pools, J , and the probability of experimentation of any element in the pool, ρ) and requires making other design choices, such as specifying the strategy space, the type of experimentation (e.g., global or targeted), its distribution (uniform or, say, normal), and the foregone utility. But given these choices, can the IEL model actually be used as a *testbed*, that is as, an alternative to expensive experiments, or as a complementary tool for experimental design?

To address this question we applied the IEL model to three sets of data from repeated KBC games. These are games with a large strategy space that are often played in the lab with 10 – 20 participants. Our approach to the learning model choices was simple: to be as close as possible to the rules of the games. This dictated our specification of the strategy space (as guesses), the type of experimentation (global and targeted), and the foregone utility (identical to the experimental payoff in the planar game and a smoother version of the experimental payoff in the games with tournament payoffs). Further, for the IEL model, we used the off-the-shelf parameter values for J and ρ .

We found that this ‘un-informed’ version of IEL generates dynamics that are close to those observed in the experimental data. We demonstrated this in two ways. First, we illustrated how the IEL simulations match the experimental data for the first few periods of the KBC games. In particular, a comparison of IEL with those learning models that have been used in the original papers to describe the data showed that the IEL is comparable to those other models in terms of RMSE and slightly better overall. Second, we investigated if un-informed-by-the-data simulations of IEL can reproduce the differences between treatments that those original papers have found and we find support for this possibility as well. We have further considered the robustness of our findings to different parameterizations for the IEL model, or for different specifications of foregone utility or experimentation.

Finally, we have shown how we can use the IEL model to predict behavior in experiments that have yet to be conducted. For instance, we provide evidence that the absolute value of the slope coefficient, p , matters for the speed of convergence to an interior Nash equilibrium in the 1D games, which could be verified in a careful experimental test.

Testbed models offer a cost-effective, scalable, and secure approach to conducting experiments and obtaining data that would be challenging or unfeasible to obtain otherwise. By using testbed models, researchers can refine and improve their experimental protocols before moving on to more expensive and intricate testing procedures. Our analysis of the repeated KBC games using the IEL model supports the idea that IEL is a valuable testbed for experimental researchers as it reliably mimics the behavior of human subjects. It can therefore enable a better understanding of likely behavior in complex systems, optimize experimental procedures, and save on resources. Thus, IEL should be considered a critical tool for research in experimental social science.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The Matlab code and data used in this study are available at <https://osf.io/zf3ct/>.

References

- Anufriev, Mikhail, Arifovic, Jasmina, Ledyard, John, Panchenko, Valentyn, 2013. Efficiency of continuous double auctions under individual evolutionary learning with full or limited information. *J. Evol. Econ.* 23, 539–573.
- Anufriev, Mikhail, Arifovic, Jasmina, Ledyard, John, Panchenko, Valentyn, 2022a. The role of information in a continuous double auction: an experiment and learning model. *J. Econ. Dyn. Control* 104387.
- Anufriev, Mikhail, Duffy, John, Panchenko, Valentyn, 2022b. Learning in two-dimensional beauty contest games: theory and experimental evidence. *J. Econ. Theory* 201, 105417.
- Arifovic, Jasmina, 1994. Genetic algorithm learning and the cobweb model. *J. Econ. Dyn. Control* 18 (1), 3–28.
- Arifovic, Jasmina, 1995. Genetic algorithms and inflationary economies. *J. Monet. Econ.* 36 (1), 219–243.
- Arifovic, Jasmina, 1996. The behavior of the exchange rate in the genetic algorithm and experimental economies. *J. Polit. Econ.* 104 (3), 510–541.
- Arifovic, Jasmina, Karaivanov, Alexander, 2010. Learning by doing vs. learning from others in a principal-agent model. *J. Econ. Dyn. Control* 34 (10), 1967–1992.
- Arifovic, Jasmina, Ledyard, John, 2001. Computer testbeds and mechanism design: application to the class of Groves-Ledyard mechanisms for provision of public goods. Working paper. Available at <https://www.researchgate.net/publication/2381600>.
- Arifovic, Jasmina, Ledyard, John, 2004. Scaling up learning models in public good games. *J. Public Econ. Theory* 6 (2), 203–238.
- Arifovic, Jasmina, Ledyard, John, 2007. Call market book information and efficiency. *J. Econ. Dyn. Control* 31 (6), 1971–2000.
- Arifovic, Jasmina, Ledyard, John, 2011. A behavioral model for mechanism design: individual evolutionary learning. *J. Econ. Behav. Organ.* 78 (3), 374–395.
- Arifovic, Jasmina, Ledyard, John, 2012. Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism. *J. Public Econ.* 96 (9–10), 808–823.
- Arifovic, Jasmina, Ledyard, John, 2018. Learning to alternate. *Exp. Econ.* 21 (3), 692–721.
- Arifovic, Jasmina, Bullard, James, Kostyshyna, Olena, 2013. Social learning and monetary policy rules. *Econ. J.* 123 (567), 38–76.
- Arifovic, Jasmina, Boitnott, Joshua F., Duffy, John, 2019. Learning correlated equilibria: an evolutionary approach. *J. Econ. Behav. Organ.* 157, 171–190.
- Arifovic, Jasmina, Duffy, John, Jiang, Janet Hua, 2023. Adoption of a new payment method: experimental evidence. *Eur. Econ. Rev.*, 104410.
- Camerer, Colin, Ho, Hua Teck, 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67 (4), 827–874.
- Camerer, Colin F., Ho, Teck-Hua, Chong, Juin-Kuan, 2004. A cognitive hierarchy model of games. *Q. J. Econ.* 119 (3), 861–898.
- Duffy, John, Nagel, Rosemarie, 1997. On the robustness of behaviour in experimental ‘beauty contest’ games. *Econ. J.* 107 (445), 1684–1700.
- Gill, David, Prowse, Victoria, 2016. Cognitive ability, character skills, and learning to play equilibrium: a level-k analysis. *J. Polit. Econ.* 124 (6), 1619–1676.
- Güth, Werner, Kocher, Martin, Sutter, Matthias, 2002. Experimental ‘beauty contests’ with homogeneous and heterogeneous players and with interior and boundary equilibria. *Econ. Lett.* 74 (2), 219–228.
- Heemeijer, Peter, Hommes, Cars, Sonnemans, Joep, Tuinstra, Jan, 2009. Price stability and volatility in markets with positive and negative expectations feedback: an experimental investigation. *J. Econ. Dyn. Control* 33 (5), 1052–1072.
- Hommes, Cars, 2011. The heterogeneous expectations hypothesis: some evidence from the lab. *J. Econ. Dyn. Control* 35 (1), 1–24. <https://doi.org/10.1016/j.jedc.2010.10.003>.
- Nagel, Rosemarie, 1995. Unraveling in guessing games: an experimental study. *Am. Econ. Rev.* 85 (5), 1313–1326.
- Selten, Reinhard, Stoecker, Rolf, 1986. End behavior in sequences of finite Prisoner’s Dilemma supergames A learning theory approach. *J. Econ. Behav. Organ.* 7 (1), 47–70.
- Sutan, Angela, Willinger, Marc, 2009. Guessing with negative feedback: an experiment. *J. Econ. Dyn. Control* 33 (5), 1123–1133.