Bounded rationality in Keynesian beauty contests: a lesson for central bankers?

Felix Mauersberger, Rosemarie Nagel, and Christoph Bühren

Abstract

The great recession (2008) triggered an apparent discrepancy between empirical findings and macroeconomic models based on rational expectations alone. This gap led to a series of recent developments of a behavioral microfoundation of macroeconomics combined with the underlying experimental and behavioral Beauty Contest (BC) literature, which the authors review in this paper. They introduce the reader to variations of the Keynesian Beauty Contest (Keynes, *The general theory of employment, interest, and money*, 1936), theoretically and experimentally, demonstrating systematic patterns of out-of-equilibrium behavior. This divergence of (benchmark) solutions and bounded rationality observed in human behavior has been resolved through stepwise reasoning, the so-called level k, or cognitive hierarchy models. Furthermore, the authors show how the generalized BC function with limited parameter specifications encompasses relevant micro and macro models. Therefore, the stepwise reasoning models emerge naturally as building blocks for new behavioral macroeconomic theories to understand puzzles like the lacking rise of inflation after the financial crisis, the efficacy of quantitative easing, the forward guidance puzzle, and the effectiveness of temporary fiscal expansion.

(Published in Special Issue [Bio-psycho-social foundations of macroeconomics](#))

JEL E12 E13 E7 D80 D9 C91

Keywords Beauty Contest game; expectation formation; equilibration; level k reasoning; behavioral macroeconomics; game theory; experimental economics

Authors

Felix Mauersberger, University of Bonn, Germany

Rosemarie Nagel, ✉ ICREA-UPF Barcelona GSE, Spain, rosemarie.nagel@upf.edu

Christoph Bühren, Clausthal University of Technology, Germany

Citation Felix Mauersberger, Rosemarie Nagel, and Christoph Bühren (2020). Bounded rationality in Keynesian beauty contests: a lesson for central bankers? *Economics: The Open-Access, Open-Assessment E-Journal*, 14 (2020-16): 1–38.
<http://dx.doi.org/10.5018/economics-ejournal.ja.2020-16>

Received October 31, 2019 Published as Economics Discussion Paper November 1, 2019

Revised April 1, 2020 Accepted April 16, 2020 Published June 4, 2020

© Author(s) 2020. Licensed under the [Creative Commons License - Attribution 4.0 International \(CC BY 4.0\)](#)

1 Introduction

Markets like the stock market are prone to booms and crashes. Recent crashes have been the financial crisis 2008-2011 and the Corona crisis in early 2020. During the Corona crisis, Wall Street suffered its worst day since the stock market crash in 1987. Such crashes have been explained by the self-fulfillingness of beliefs: if traders are pessimistic regarding future stock prices, they start selling stocks today. The drop in demand for stocks leads to a decrease in stock prices and the traders find their pessimistic beliefs confirmed (see, e.g., Shiller, 2015). Once a downturn is triggered, it might be persistent. Some naive traders observing the crash may anchor their expectation on past realizations of stock market prices. A slightly more sophisticated trader might anticipate such behavior and might sell her stocks today since she is afraid of having to sell the stock at a lower price in the future. Such a transaction might be worthwhile for a trader to prevent losses, although she might not truly believe that the stocks will actually be worth less in the future. This example shows that naive market participants can drive behavior at a large scale.

This insight is not new and has recently been incorporated into macroeconomic models (see the seminal paper by García-Schmidt and Woodford, 2019). The behavioral assumptions have been inspired by a number of past laboratory experiments, in particular by the so-called Beauty Contest (BC) game. Out of these laboratory experiments, a behavioral model of step level reasoning, the so-called *level-k model* has emerged. Our survey reviews the experimental beauty contest literature and shows how the idea of level-k thinking is useful as a behavioral foundation of macroeconomics. We present a four-part structure: the first part is about the level-k model and how it describes out-of-equilibrium behavior in a large set of experiments from over 25 years; the second section is about systematic patterns of out-of-equilibrium behavior in variations of Keynesian Beauty contest experiments which can also be relevant in real-world situations or in behavioral macroeconomic models; third, we show how a generalized BC function with limited parameter specifications encompasses relevant micro and macro models; fourth, we review new papers in macroeconomics applying the level k model to explain different macroeconomic puzzles.

The perhaps most widely used equilibrium concept in macroeconomics is the *Rational Expectations Equilibrium (REE)*. REE presumes that all agents have perfect knowledge about their environment and make on average correct decisions (Muth, 1961). While REE as a concept remains a useful benchmark and retains macroeconomic analyses tractable, many empirical studies challenge this idea. For instance, there is an extensive literature that uses survey data on inflation expectations to challenge the validity of the rational expectation assumption (see Mavroeidis et al., 2014, for a survey). The notion of rational expectations has also been widely challenged in laboratory studies (see Hommes, 2011, for a survey).¹

In response to the increasing use of non-rational expectations in economics, macroeconomic papers such as Woodford (2013) and García-Schmidt and Woodford (2019) make use of the notion of *temporary equilibrium*. By temporary equilibrium they mean that market outcomes at any point in time result from optimizing decisions by households and firms but that their expectations need not be correct. Yet, since the predominant solution concept in macroeconomics is still the rational expectations equilibrium and human behavior (sometimes) converges to it, we take REE as the benchmark for our further exposition. Level k is one specification of non-rational expectations

¹Laboratory studies have become an increasingly popular methodology to address questions related to macroeconomics (for surveys, see Hommes, 2013; Amano et al., 2014; Duffy, 2014; Cornand and Heinemann, 2019; Hommes, 2020).

together with an iterated best reply structure. However, in the microeconomic literature, it has not been referred to as equilibrium choices, as beliefs and resulting best replies of different players are not consistent to each other.

Due to the critiques regarding the rational expectations hypothesis, especially after the 2008 crisis, we argue in this paper that macroeconomics may benefit from *behavioral* micro-foundations. We review the recent behavioral macroeconomics literature, with the seminal paper by Woodford (2013) and García-Schmidt and Woodford (2019), which incorporates into economic models behaviorally founded assumptions that are well-established in behavioral and experimental economics through laboratory studies with human subjects. Laboratory methods have the advantage that they represent controlled environments where human behavior and responses to exogenous interventions can more cleanly be documented than this would be possible in the field.

Specifically, we consider variations of the so-called Keynesian Beauty Contest (Keynes, 1936). In this citation, Keynes links the stock market to a newspaper beauty contest with the following game:

“the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole.”

Any six faces can constitute a Nash equilibrium. Such kind of multiplicity is one of the greatest challenges for theoretic modeling since pure rational reasoning will not solve the coordination problem. However, Keynes’s emphasis is on his behavioral reasoning model. Not all participants will think or form beliefs alike, and therefore actions can be very different.

“It is not a case of choosing those [faces] that, to the best of one’s judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.”

Yet, there is a caveat. The game, as depicted by Keynes, does not give rise to heterogeneity, neither in beliefs nor in actions, if the same average preference is presupposed. Once six faces are assumed to be the average preference, they have to be chosen by any player. A higher order of reasoning will collapse to the same first order belief and thus the same choices of faces.

While Keynes’s game has been used in economics to exemplify the challenges of the multiplicity of equilibria, the following variation has become the paradigmatic game to visualize heterogeneity in human behavior closely related to Keynes’s behavioral reasoning. Imagine every participant in a group of people is asked to choose a number from 0 to 100 (instead of faces). The person whose choice is closest to two-thirds times the average of all chosen numbers wins a fixed prize. Thus, we just change Keynes’s target “average” to the target, say $2/3$ -average, accordingly with a multiplication factor non-equal to one. This change has three advantages: first of all, the equilibrium is unique: all players choose zero. Therefore, any deviation from zero must be due to bounded rationality or beliefs about such behavior. Furthermore, this game has become famous in the behavioral game theory literature, primarily because of the compelling tensions between the theoretical solution and experimental results, with a substantial heterogeneity of behavior. Such discrepancies have been bridged with Keynes’s iterated reasoning model, called level k . Level k

nesting the extreme cases of full rationality, corresponding to Nash equilibrium, as well as complete irrationality, corresponding to playing randomly or choosing focal points. The interplay of theory and behavioral reasoning is the leading challenge we cover in this paper. Thirdly, by changing the parameter from positive (e.g., $2/3$) to negative values (e.g., $-2/3$), we can discuss the difference between positive (e.g., the higher others' choices, the higher my choice) and negative feedback (e.g., the higher other players' choices the lower my choice) situations which have very different convergence patterns.

This kind of heterogeneity in behavior and beliefs of different players has led to the so-called level k (Nagel, 1995), and some variations (Stahl and Wilson, 1995), or the cognitive hierarchy model (Camerer et al., 2004). An extensive survey about these models is presented by Crawford et al. (2013). Deviations from equilibrium have been replicated by many experiments of similar variations of this or other games, both in the laboratory and in the field. Over time there might be gradual convergence to (one of) the equilibrium point(s). Level k models can explain such behavior.

We highlight the importance of the BC game and its generalization for macroeconomics by arguing that there is a close link between the BC game and the micro-foundations of standard workhorse macroeconomic models. The focus of this paper is on New-Keynesian macroeconomics with standard frameworks as they can be found in textbooks like Woodford (2003), Galí (2008), and Walsh (2010). Yet, we also address other areas of macroeconomics such as economic growth.²

As shown by authors such as Woodford (2013) and Angeletos and Lian (2018), agents' utility (or profit) in New-Keynesian models depends on the agent's own action and the aggregate of all other agents' actions in the economy. Specifically, in standard New-Keynesian models, the firms' optimal price depends on the (discounted) average price across firms in the economy. Similarly, the household's optimal consumption depends on the (discounted) average consumption across households in the economy.

Since there have been many critiques of the REE and since the evidence of level k as a heuristic of decision-making in the BC game is overwhelming, we argue that considering the implications of level k at the macroeconomic level is a logical step to open the box for out-of-equilibrium considerations in academic research.

The reason why level k is useful for macroeconomics is that it provides inertia in responses to aggregate shocks. García-Schmidt and Woodford (2019) use level k to explain the sluggishness in the response of inflation to low interest rates after the financial crisis 2007-2011. Furthermore, they find that under level- k behavior there is no forward guidance puzzle. Angeletos and Lian (2018) revisit the effects of forward guidance and also consider the implication of level k for fiscal expansion. Other papers apply level k to analyze the effects of incomplete markets (Farhi and Werning, 2019), and quantitative easing (Iovino and Sergeyev, 2018).

The paper is structured as follows: Section 2 introduces "level k " as an alternative concept to rational expectations. Section 3 provides a review of BC experiments with level k patterns, which we consider relevant for behavioral macroeconomic modeling. Section 4 presents the generalized BC game as a canonical framework closely linked to micro-founded macroeconomic models. Section 5 introduces the recent literature of behavioral micro-foundation of macroeconomic models with level k features. Section 6 concludes.

²Mauersberger and Nagel (2018) show how a generalized BC game (by adding a constant and an idiosyncratic shock to the discounted average) can span a wide variety of aggregative games: Cournot models, forecasting models, etc.. Benhabib et al. (2015) demonstrate that this framework is a simplification of a general equilibrium model with sentiments.

2 The Level k model

This section briefly reviews the basic level k model (using the p-Beauty contest game introduced in the beginning), and variations which provide cognitive processes and a realistic description of expectation formation in experimental games, particularly when a game is played for the first time. The basic idea of “level k” reasoning is that, instead of choosing the Nash equilibrium, agents tend to think about what other possibly boundedly rational subjects may choose. The “level k” model of boundedly rational reasoning starts with a specification of a naive approach to the game (“level 0”), and a (finitely) iterated best reply structure as explained in the following. Thus randomness and equilibrium behavior can be easily incorporated at the same time into the model.

Nagel (1995) introduced the level k model to the experimental literature by describing the (out-of-equilibrium) behavior in the BC-game mentioned in the introduction. Subjects are asked to choose a number between 0 and 100, the one being closest to two-thirds of the average of all chosen numbers receiving a fixed price. All other players receive nothing. In the case of a tie, the prize is split between those who tie.

Consider the naive player who chooses a random number from 0 to 100 with equal probability. The expected value of this player’s choice is 50. Such behavior is referred to as “level 0.” A slightly more sophisticated player anticipates such behavior and best-responds to level 0 by choosing $\frac{2}{3} \cdot 50$. This player type is referred to as “level 1.” Such behavior can also be anticipated, and some players may best respond to level 1 and choose $(\frac{2}{3})^2 \cdot 50$. Such a player type is called “level 2.” Similarly, one can define k such thinking steps by “level k ”. A level k player chooses $(\frac{2}{3})^k \cdot 50$. The higher k is, the closer the behavior corresponds to the Nash-equilibrium of zero.

For first-period behavior, level 0 is typically assumed to be a uniform distribution or the focal point 50.³ In the subsequent periods, level k can also be applied. In the BC game, level 0 then becomes the average choice of the previous periods, and all other levels are adjusted from this anchor or reference point, similar to the first period choice 50 (see also Stahl (1996) who combined a reinforcement model and level- k , using the data from Nagel (1995)). In the next section we show different level 0 specifications, depending on the particular functional form of the BC game or other specifications of the rules of the game.

Thus, this kind of behavioral modeling together with an experimental foundation constitutes several compelling features: It proposes a flexible starting point of level 0, based on naive behavior and taking into account the context of the game. It encompasses random behavior and the equilibrium as extreme benchmarks. Iterated reasoning has been used in theoretical concepts as rationalizability (Bernheim, 1984; Pearce, 1984) or eductive reasoning Guesnerie (1992). However, such papers do not consider the problem of a low level of reasoning among human subjects and how to resolve the indeterminacy of stopping the iterative process. Empirical observations in the laboratory and field have revealed that the majority of human subjects are comprised of levels 0 to 3.

Following Nagel (1995), the literature has extended the level k model and addressed possible drawbacks (see also the survey of Crawford et al. (2013), and we add some critical comments at the end of the next section). The focus of the microeconomics literature has been on first period choices, when subjects have no experience with the particular situation, or on repeated interactions without

³Note that game theory offers an alternative equilibration process, starting from 100. In the case of $p = 2/3$: all numbers above $100 \cdot 2/3$ are (weakly) dominated, thus can be deleted by rational players. An infinite process of iterated deletion of such strategies leads to zero. Most subjects do not derive their choices from such a procedure.

feedback on choices of others or payoffs, to exclude issues of learning.⁴ In contrast, the behavioral macroeconomic literature we will discuss below, applied level-k related models to behavior over time,

One questionable feature of the original level-k model is it implicitly assumes that *all* other players adopt a certain level of reasoning, assuming that all others choose one level lower than oneself. Stahl and Wilson (1994, 1995) construct level 0, 1, and 2 types and Nash equilibrium as above, and add wordly types which assume that all others are distributed over level 0 and level 1 types. Their experimental set-up (using 3x3 games) does not allow for level 3 or higher types.

Camerer et al. (2004) introduce a one-parameter cognitive hierarchy model such that level 2 and higher types assume that all other players are distributed according to a Poisson distribution over lower types, with one free Poisson parameter. Therefore, this model becomes a predictive model, as the econometrician can perform out of sample validation, within the same game or across different experimental situations, given a Poisson parameter estimate based on actual observations. Camerer et al. (2004) justify their model by the widely documented finding in the psychology literature that individuals tend to be overconfident about their own ability over other estimated abilities.

Alaoui and Penta (2016) outline a model in which players' depth of reasoning is endogenous. Their motivation is to use the endogenous depth of reasoning (EDR) model to make inferences and sharp predictions that hold across different games. They propose that individuals act as if they face a trade-off between costly reasoning and the benefit of doing so. The costs are related to the sophistication of the player. The benefit instead is related to the game payoffs. Behavior is determined by the individual's depth of reasoning and his or her belief about the opponents' reasoning processes.

3 Variations of the Keynesian Beauty Contest game

In this section, we introduce some parameter variations of the original Beauty Contest.⁵ We show to what extent such changes imply differences in the theoretical properties of the new games. Furthermore, we link the experiments to relevant issues in macroeconomic theory. As mentioned before, REE is probably the most widely used concept for macroeconomic analyses. Thus, we discuss variations together with resulting behavioral and theoretical changes. More specifically, maintaining the same equilibrium, differences in out-of-equilibrium properties through parameter changes might also induce behavioral changes. When behavior systematically differs from equilibrium choices, making predictions using equilibrium concepts alone is questionable. Nevertheless, we want to emphasize that this does not make equilibrium concepts disposable. Quite on the contrary, if, for example, the equilibrium produces higher welfare than the gains subjects attain, the next step should be to investigate which institutional changes are necessary to reach the desirable theoretical outcomes. We will provide such examples. Nevertheless, out-of-equilibrium behavior might be more likely observed with untrained subjects. Thus, studies allowing for experience, that is repeated interactions, are also discussed. We, therefore, separate initial behavior of experiments from observations over time. For more extensive reviews along the same line, refer to Amano et al. (2014), Arifovic and Duffy (2018), Hommes (2013), and Mauersberger and Nagel (2018).

⁴Exceptions are papers by Dale O. Stahl and coauthors. See, e.g., Haruvy and Stahl (2012) and the references therein.

⁵See also, for a historical account, Nagel et al. (2016).

Table 1: Beauty Contest experiments from different studies

Authors	variation(s)	BC target	choices	equilibrium	subject pool	no. of subjects	mean	std. dev.	level 0	winning number	spikes at
Nagel (1995)	1a	$p=2/3$	[0,100]	0	students	67	36.73	20.21	50	24.49 (k=2)	33.33, 22.22
Bosch et al. (2002)	2	$p=2/3$	[0,100]	0	economists	146	17.15	22.64	50	11.43 (k=4)	0
		$p=2/3$	[0,100]	0	newspaper readers	7900	23.08	20.24	50	15.39 (k=3)	50, 33.33, 22.22, 0
Nagel et al. (2017)	4 (two persons)	$p=2/3$ (fixed prize)	[0,100]	0 (dominance)	students (class)	160	36.98	22.72	50	24.65 (k=2)	50, 0
		$p=2/3$ (distance prize)	[0,100]	0	students (class)	149	37.36	28.04	50	24.91 (k=2)	55, 33.33, 0
Benhabib et al. (2019)	5a, 6, and 7a	$p=2/3$	any number	0	students	40	14.03	16.47	0	9.35 (k unknown)	no clear spikes (15, 10)
		$p=-2/3$	any number	0	students	40	1.87	27.5	e or 0	-1.25 (k=0)	0
		$p=2/3 + e$ from $N(0,10)$	any number	0	students	39	2.37	5.73	e or 0	1.58 (k=0)	0
Bühren, Nagel (2019)	5 and 6	$p=2/3, c=10$	[0,100]	30	students (class)	68	42.11	22.35	50	38.07 (k=2)	50, 43
		$p=2/3, c=10$	any number	30	students (class)	72	14.53	14.53	10	19.69 (k=1 or 2)	10
		$p=2/3, c=10$	[0,100]	30	economists	59	34.98	19.82	50 (or 10)	33.32 (k=4)	30
		$p=2/3, c=10$	any number	30	economists	65	20.84	15.15	10	23.89 (k=3)	30
Heemeijer et al. (2009)	8 (p-parameters unknown)	$p=20/21$	[0,100]	60	students	36	64.31	121.1	50	64.11	no clear spikes
		or $p=-20/21$	[0,100]	60	students	42	59.79	6.15	50	60.20	no clear spikes

Table 1 shows different parameter variations of the basic BC game together with average behavior and its variance in the initial period, game-theoretic predictions, and level k outcomes (level 0, the level k at the winning number, and modal level k s). We also mention the subject pool, as trained subjects behave potentially differently in the first period. To draw a connection between table 1 and macroeconomics, the study of aggregate economic outcomes, we intentionally concentrate on aggregate measures, i.e., averages across agents' choices, in BC games, rather than on the distributions on behavior.

The original Keynesian Beauty Contest ($p = 1$): Multiplicity of equilibria The multiplicity of equilibria poses one of the most challenging obstacles for making predictions and achieving coordination among agents. This challenge is also well-known in macroeconomics. As exemplified by Woodford (2013, p.316): “The existence of a large set of possible equilibrium outcomes, including the possibility of fluctuations in response to sunspots or large fluctuations in response to small changes in fundamentals, is regarded as an undesirable form of instability.” Keynes’s 1936 face competition game which we presented in the introduction, can be re-framed with a game in which faces become numbers from 0 to 100 and a target of “1 times the average” unlike in Nagel’s variation, in which the multiplication factor was $p = 2/3$. Those choosing the average win some fixed prize, and all others earn nothing. In such a game, any number can be an equilibrium choice. That is, if all players choose the same number, then nobody should be able to get a higher payoff by deviating. Coricelli and Nagel (2009) show that most players choose the focal point, 50, thus there is not much dis-coordination, unlike predicted by theory.

However, discoordination comes back in the following game with a slightly different payoff: All winners receive higher payoffs if they chose higher numbers. Again, any number, chosen by all, constitutes an equilibrium. Yet, it is not clear whether all players pick the highest number, which constitutes the so-called Pareto-optimal solution. This means that in any other equilibrium payoffs are lower for at least one player. Such multiplicities in both games are considered as a serious problem in game theory or other fields, such as macroeconomics. Even a group of rational agents may not know how to coordinate. Van Huyck et al. (1990, 1991) study behavior in such experiments with subjects choosing among integers between 1 and 7 and the median or, alternatively, the minimum choice, determines the payoff for all players minus the individual cost of deviating from the target order statistic. The higher the median or minimum choice, the higher the payoffs. Thus, there is a Pareto ranking of equilibria. In the first period, the experiments show high heterogeneity of choices among human subjects, leading to a dispersion of individual payoffs. Over time, there is convergence to one equilibrium, typically the middle choice in the median game and the lowest number in the minimum game.

For macroeconomists, it should be essential to understand the empirical selection of the equilibria. Wording alone (median and minimum) produces huge differences unattained by theoretical considerations. The authors also vary the number of players. While two players easily achieve optimal outcomes, already a group of four players or more fails to achieve high payoffs. Applying level k results in heterogeneity through different level 0 formulations: randomness would result in a best reply to the midpoint of the interval, while focal point, payoff maximum or risk dominance criteria can also lead to the maximum or minimum strategy. Higher-order beliefs all collapse to the same strategy as suggested by level 1 reasoning: best reply to level 1 (say choice 4) will also lead to choice 4.

The economic literature has produced several strands coming out of such theoretical and empir-

ical findings. Theorists have implemented different ways to select one equilibrium out of the many (see, e.g., global game literature, which adds payoff perturbations to obtain a unique equilibrium). Indeed, experiments confirmed that the global game selection provides a good descriptive model of behavior, no matter whether there is complete or incomplete information about payoffs in such coordination games (see Heinemann et al., 2004). Experimenters have asked the question which equilibrium subjects attain over time as mentioned above and have proposed different mechanisms to improve efficiencies. We will present a theoretical change to achieve uniqueness, which also has been relevant to macroeconomics. Yet, the uniqueness of equilibrium does not necessarily restore uniqueness in behavior by real subjects. This is the leading challenge we cover in this paper.

Variation 1a: a unique vs. multiple equilibria and the first application of the level k model In the following experiment, with a target, e.g., $2/3$ -average a unique equilibrium results. Again, we observe heterogeneity of behavior, but most, if not all, choices are off the equilibrium. Yet, we will show that there is a pattern of behavior that can be reconciled with the level k model we presented in the previous section. Indeed, that model was built with the following game and data.

The experimental study Nagel (1995) about the basic p -average games ($p=1/2, 2/3, 4/3$) contains the first specification and visualization of the so-called level k model. Each subject, typically invited from an undergraduate student pool, played within the same group and parameter p for four periods. After each period, all choices, the average, the target, and the winning numbers were written on the blackboard. Every subject was only allowed to participate in only one session.

When $p < 1$, there is a unique Nash equilibrium, choosing zero. For $p > 1$, there are two equilibria: all players choose zero, or all players choose 100. In the latter case, the equilibrium of zero is highly unstable, meaning that if one player slightly deviates choices tend not to return to this equilibrium; an equilibrium of 100 exists because if every other player chooses 100, the optimal choice is also 100, which is the stable equilibrium. The level k model can explain the out-of-equilibrium behavior, for both, the unique equilibrium and the multiple equilibrium case.

We highlight four striking findings for $p = 2/3$. Firstly, in the first period, there is a high heterogeneity of behavior, although the game is the same for all subjects. Yet, there are clear patterns. Zero is rarely chosen, and most choices lie between 20 and 50. Secondly, choices are clustered around the numbers that correspond to the levels of reasoning. For the experiment with a target of two-thirds times the average, the level k choices correspond to 50 for level 0, $33\frac{1}{3}$ for level 1, $22\frac{2}{9}$ for level 2, etc.. Figure 1 shows the behavior of Nagel's lab-experiment in the upper left-hand side with $p = 2/3$. For $p = 1/2$ or $p = 4/3$, there are salient spikes in the distribution of play at $(p)^k \cdot 50$, for $k = 1$ or $k = 2$. Thus, it is clear that the Nash equilibrium is not a good choice when the game is played for the first time. But note that someone who chooses, say 33, can be a person who chooses randomly or believes that most people will choose randomly, and best replies to it.⁶ In variation 2, we show how subjects trained in game theory play against each other.

Variation 1b: behavior over time and level k It is not surprising that first-period behavior is far from equilibrium. Therefore, repeating the game is a crucial design feature in experimental economics. As a third finding in the same paper, choices tend to approach the unique equilibrium,

⁶There is a large experimental literature that tries to recover reasoning procedure distinguishing between the sophistication of players, e.g., through observing team reasoning. For a brief survey of such methods, see Mauersberger and Nagel (2018).

Bosch, Montalvo, Nagel, Satorra, AER 2002

Figure 1: Relative frequencies of first-period choices reported in Bosch-Domenech et al. (2002)

albeit slowly, as the BC game is repeated within the same subject pool. In the version with four-thirds times the average, behavior converges to 100, being the stable equilibrium. The level k model can again explain the reason for slow convergence. A crucial question is the modeling of the level 0 type. Choices above the mean of the previous period are rarely chosen (typically less than 10% see Nagel (1995), table 2). Therefore, Nagel decided to take the mean as the critical level 0-type, which is also supported by written comments. With this in mind, there are only a few observations of level k greater than 3, as observed in the first period. The modal choice is typically at level 2. Thus, the level of reasoning does not increase over time beyond higher levels than 3. Note, any other reasonable assumed lower level 0 behavior would conclude also with no increase of levels over time.

Figure 2 shows the average behavior over time, which drops as if the average level k is just 1 or 2 from one period to the next, with level 0 in period t being the average of period $t - 1$. Given the rare choices above the average of the previous, the interval to choose from shrinks over time. Nagel (1995) gives a behavioral justification that the average level k is typically not increasing over time with a simple adjustment model. She shows that those subjects who iterated too many steps as compared to the optimal level k will adjust less in the next period and the other way around for

those who iterated too few steps. Convergence is faster with smaller p -parameters and also if an outlier has less influence on the order statistic, as seen in figure 2 in the $1/2$ -median game (Duffy and Nagel, 1997) and very much slower in small groups with three players than in larger groups with 15 players.

Figure 2: Mean behavior over time from Nagel (1995), Ho et al. (1998) and Duffy and Nagel (1997)

Such experimentally founded inertia are especially important to take into account in behavioral macroeconomic modeling, in which typically exogenous shocks drive behavior out-of-equilibrium. Here, all out-of-equilibrium behavior is solely induced by the mind of the subjects. A convergence back to old seemingly stationary points typically are not reached quickly as the great recession has shown. Therefore, we can assume that also in reality humans might only slowly converge to equilibrium for not believing that others adjust rapidly. Such kind of sluggishness is especially true in situations with positive feedback as in our $2/3$ example, when beliefs that others play high, induces oneself also best respond with high choices. We will see below that negative feedback situations converge faster.

Variation 2: diverse subject pools We have shown heterogeneity within the same subject pools, which are typically students. For about 25 years, subjects have been invited from different populations. Figure 1 shows the data of over 7,000 participants playing this game with $p = 2/3$,

with data of Nagel (1995), and choices from subjects in field studies, including participants from newspaper experiments (the rules of the game were submitted to Financial Times, Spektrum der Wissenschaft, and Expansión), classrooms, a newsgroup, and conference data (see Bosch-Domenech et al., 2002). As one can see, the (average) behavior is far from equilibrium and varies between the groups. Yet, there is a similar pattern (as described above) across all data sets: the spikes are at 0 (at least in some), 22.22, 33.33, and 50. Only a few spikes can be observed at 66.67 and 100, with the latter being some subjects who say that they want to spoil the gain for sophisticated players. Averages vary from 20 to 35 between experiments because the distributions at the spikes across the populations are quite different. Experts choose the lower number, yet the equilibrium is not chosen by more than 40%. Thus, the winning number is 2 to 3 steps away from 50, or, in other words, close to level 2 or 3 reasoning. Note also that the number of players is very different for each subject pool (see table 1), varying between 15 for the lab experiments and several thousand in the newspaper experiments. Levine et al. (2017) invited subjects from less sophisticated subject pools as the students or the newspaper readers mentioned above, which we do not show here. Winning numbers in that paper are typically near level 1, which means that most people choose random numbers and an average (near) 50.

One of the strong conclusions from 25 years of experimental games is indeed that level 2, and 3 are good predictive guesses for best reply behavior, given that there is a (clear) level 0 anchor. In games which include level 1 behavior near an equilibrium, then this point is a good predictor. Games that allow an iterated best reply structure and require many levels of such reasoning are typically far away from the equilibrium, especially in initial periods (see also, e.g., Crawford et al., 2013). As a conclusion so far, we can emphasize that theoretical analysis of the game and its off equilibrium structure together with experimental general findings related to level- k informs the researcher ex-ante about potential predictive qualitative behavior. Thus, theoretical properties and empirical foundations work hand-in-hand.

Variation 3: A unique Pareto optimal solution with distance payoffs The following variation provides a clear case, why equilibrium analysis and properties are important, even if actual behavior cannot be predicted by such solutions. We introduce a (small) variation about the payoff structure, which has fundamental effects on the theoretical properties but not on behavior. We replace the tournament structure (“winner takes it all”) by a distance function. Everyone is paid according to the deviation of their own choice from the target, e.g. the payoff is $100 - (\text{own decision} - \frac{2}{3} \cdot \text{average})^2$. Here, in the Nash equilibrium, all players still choose zero. However, now the sum of all payoffs is maximized. In other words, the welfare-maximizing state is reached in this equilibrium. Thus, this game has a unique Pareto optimal equilibrium, which means that playing the equilibrium makes all players better off than in any other strategy combination. Non-compliance, therefore, will lead to large losses which should be prevented through well-designed institutions. In variation 1, all strategy combinations are Pareto-optimal, as a shift from one winner to another would make the former worse off.

There is another main difference, being particularly interesting for epistemic game theory: Iterated elimination of strictly dominated strategies leads to the unique equilibrium. Thus, a rational player who assumes that all others are rational and so on, can eventually decide to choose zero. In the prior case, a rational player can only eliminate dominated strategies but not necessarily further domination levels. However, all these theoretical differences do not affect “boundedly rational” agents. The behavior in the two different games is indistinguishable. (Kocher et al.,

2002)

The conclusion from experiments with this particular game shows that the equilibrium concept is important, even when it has no predictive power. When behavior fails to reach the welfare optimizing equilibrium, a call for optimal institutional changes is then desirable. As a side note, alternative institutions are also installed when inefficient equilibria are reached, for example, in public good experiments (see, e.g. Gächter et al., 2004). We will see that variation 5, allowing for open intervals, and variation 6 with exogenous signals constitute such rule or institutional changes for our games.

Variation 4: The two players (dominant strategy equilibrium) versus many players cases This variation deals with two fundamental issues. One will be related to the notion of dominant strategies induced by the new game. The other concern is about the question of whether subjects inhibit the low level of reasoning because they expect others not to play the equilibrium or whether they are boundedly rational themselves. Since fairness concerns do not matter in the BC game, the out-of-the-equilibrium behavior is therefore only induced through bounded rationality of the player him or herself or the beliefs about others. In this variation, we can closely analyze the rationality of a player. Instead of many players, the game is played among two individuals (Grosskopf and Nagel, 2008). The two-person $2/3$ average game with fixed payoffs introduces a surprising complication as in a logical puzzle. The reader should pause to think what he would play against an undergraduate student, trying to be closer than the other player to the target, $2/3$ average of his or her and the other player's choice, to win the fixed prize. It is the simplest of all such games, yet not even most game theorists see what to play optimally against boundedly rational players. The equilibrium is again zero, but with the additional caveat, that it is in (weakly) dominant strategies. This means that it always wins to play zero, but also higher numbers can win against a bounded rational player. The reason is straightforward: With two differently chosen numbers, the lower number always wins, since the midpoint of two numbers multiplied by $2/3$ is always closer to the lower number. However, when the payoff depends on the above-mentioned distance function (Nagel et al., 2016), then again, the iterative elimination procedure leads to zero, yet, with reasoning steps at 100, 50, 25, etc., due to the influence of one's own number.

Finding the dominant strategy is not as easy in this two-person fixed payoff game as with more than two players, as we have seen above. When $n > 2$ rationality only precludes numbers above 66.66. Yet, in this latter case, all those who use level 1 or higher will never pick such numbers, while in the $n = 2$ case, all choices above zero are weakly dominated by zero. Thus, the players make several cognitive mistakes in two-person games. Firstly, they confuse the fixed payoff game with a game to be as close as possible to the target and not just being closer than the other player. Furthermore, they do not see a large impact on the target through their own choice. The level k model, with 50 as level 0, seems to be valid as a descriptive theory for all three variations. Chou et al. (2009) call this problem "game form recognition". This means that players in the two-person and fixed payoff game do not play the game the experimenter suggests to them.

(Not) surprisingly, figure 3 shows that there are no differences in behavior, neither between the two persons with different payoff functions (distance vs. tournament payoffs) nor between two vs. many players. Bosch-Rosa and Meissner (2019) even go a step further and introduce the one-player BC game where a single person picks two numbers between 0 and 100 to be closest to two-thirds times the average. Similarly to Chou et al. (2009), they conclude that most subjects misunderstand the structure of the game, however, there is a substantial number of zero-players in this set-up..

Figure 3: Relative frequencies of choices reported by Nagel et al. (2016) in the two-person games with fixed payoffs (upper left), distance payoffs (upper middle) and with economics professors in various conferences (upper right); lower left for $n > 2$ games with economic professors and students in advanced economic classes; in lower middle newspaper contestants; lower right chess players online. Data source: UR: Grosskopf and Nagel (2008); LL and LM: Bosch-Domenech et al. (2002); LR: Bühren and Frank (2012)

Figure 4 shows the expected payoff for each choice given the actual distribution in the particular treatments. When the weakly dominant equilibrium strategy is played in the fixed payoff treatment, it attains the highest payoff. The optimal outcome in the other games is at or near level 2 or 3, respectively, as discussed in the previous experimental treatments.

The finding in this variation is important for theorists and policy implications: Dominant strategies constitute the easiest recommendation of what to play in a game from a theoretical point of view as long as the equilibrium is not payoff-inferior. In the two-player fixed payoff treatment, one can test whether a player is rational (chooses a weakly dominant strategy) or not. However, not even economists (who played before seminars and talks) necessarily pick zero, yet more do so than in the game with more than two players. Grosskopf and Nagel (2008) show that behavior in a two-person-fixed payoff case converges over time to equilibrium, though a bit faster than in the $n > 2$ case. Both these findings should be a warning for anybody making predictions, blindly following a desirable theoretical feature. Theorists have reacted to such findings with new theoretic properties based on empirical findings. According to Li (2017), the two-person game has non-obviously dominant strategies. Therefore, there is a theoretical underpinning that such strategies are difficult to find which has to be considered in mechanism design, for example.

Figure 4: Expected payoffs of choices by treatment (see Nagel et al., 2016)

Variation 5a: Closed vs. open intervals with $p = \frac{2}{3}$ It can be very important whether there are boundary restrictions or not in a model - this applies to behavior in the BC game and several macroeconomic problems. For example, Benhabib et al. (2001) show the existence of two steady states if the “zero lower bound” on nominal interest rates is binding: one steady state with inflation near the target and one with the nominal interest rate near zero and the possibility of deflation. They show that the deflationary steady state with interest rates at the zero lower bound is stable even under active monetary policy.

Introducing a design without boundaries in the basic BC game does not change the equilibrium as compared to the boundary case. Yet, the behavioral changes should be instructive for macroeconomic policymaking. If there are no boundaries, any real number can be chosen: positive, negative, or zero, and therefore one cannot perform iterated elimination of dominated strategies or iterated best reply. Thus, game theory does not provide an out-of-equilibrium structure (iterated dominance) as before. Bühren and Nagel (2019) observe very different behavior in the $2/3$ -average game for both sets. In the open set, behavior is much closer to zero than in the boundary condition. Level k reasoning is relevant for the closed set. In the open choice set, level 0 can only be zero itself as the “midpoint” of the real line. Yet, most people seem to assume quite some noise in the behavior of others. Thus, the average is typically between 10 and 20, albeit far smaller than in the original case when it is typically above 30 with undergraduate students. For positive p -parameters, nobody chooses negative numbers (see also table 1). Below we will also discuss the contrast with p being negative or positive.

Boundaries are often used in policy making, especially when it comes to inflation target publications and zero lower bound restrictions. The findings of this variations shows that equilibrium properties alone cannot predict the outcome. Yet, out-of-equilibrium properties as taken into ac-

count by iterated reasoning and level k indicate some bounded rational features which variation should be closer to equilibrium. The maybe surprising feature is that boundaries can drive behavior further away from equilibrium than unbounded intervals. Here level 0 on the equilibrium plus noise indicates behavior closer to equilibrium as compared to the case with a large bounded interval 0 to 100, with 50 as a focal point.

Variation 5b: Closed vs. open intervals in BC games with a constant Behavior is even more drastically different when a constant is added to the target used so far. Bühren and Nagel (2019) also compare behavior with and without boundaries and an added constant, $2/3$ -average plus c (similar to Kocher et al., 2002). When ten is added to the $2/3$ -average game, the Nash equilibrium is 30, no matter if the choice set is an open interval or between 0 and 100 (see also variation 6). However, also for the treatments with an added constant, behavior is different with and without boundaries. Bühren and Nagel (2019) show for the no boundaries case that participants typically take 10 as the starting point (anchor), from which some of them proceed with k level reasoning (level 1 = $2/3 \cdot 10 + 10 = 16.67$, level 2 = 21.11, level 3 = 24.07...).

In the boundaries case, participants chose as the starting point mainly 50 (level 1 = $2/3 \cdot 50 + 10 = 43.33$, level 2 = 38.89, level 3 = 35.93...). The two starting points or level 0 reasoning explains the difference of the two treatments: In the boundary case, the average is 42.11, while, in the unbounded case, the average is 14.52, as shown in table 1. Also, conference participants who play among themselves or against undergraduate students show the same kind of discrepancy between the two versions, yet with an average level higher by one or two steps as compared to the undergraduate students (see table 1).

Again, policy makers have to consider such behavioral changes when moving from a bounded case to an unbounded case and the other way round. The constant may be ignored in one case or the other for determining level 0 reasoning.

Variation 6: Signals drawn from normal distributions with (non-)zero means We have mentioned above that experience (through repetitions or training as game theory) can drive behavior to equilibrium. However, the time out-of-equilibrium can be very costly as shown in variation 3 with a Pareto optimal equilibrium. Therefore, here we introduce a parameter change that induces immediate equilibrium behavior at least on average. So far, we have shown that behavior is typically out-of-equilibrium due to cognitive difficulties to find the equilibrium and/or beliefs that others cannot find it. Keynes formulates these difficulties with

“human decisions affecting the future, whether personal, political or economic, cannot depend on strict mathematical expectation, since the basis for making such calculations does not exist ... [In making decisions] our rational selves choose between the alternatives as best as we are able, calculating where we can, but often falling back for our motive on whim or sentiment or chance.” [Keynes (1936), pp. 162-163]

Our main focus has been to determine level 0 behavior as driven through chance by naive players or by an added constant. The following variation introduces a simple way of offering the players “falling back on sentiments”. Every subject i receives an individual idiosyncratic signal added to the target number (Benhabib et al., 2019). This variation is a direct application from a simplified general equilibrium model with sentiments (Benhabib et al., 2015). The target becomes $2/3 \cdot \text{average} + \epsilon_i$, where ϵ_i is drawn from a normal distribution with mean zero and a positive

variance. Choices, therefore, need to be made from the entire real numbers (open interval) as introduced in variation 5a.

Note that the signal is payoff relevant. Subjects know the distribution of the signals but are not informed about the signal realizations of other subjects. The intuitive choice of a boundedly rational subject would be to choose the sentiment ϵ_i , which becomes an “anchor” or a focal point. It happens that each player choosing her ϵ_i will at the same time constitute an equilibrium. While computing the equilibrium is more elaborate than in all other cases, the intuition why choosing one’s idiosyncratic signal is an equilibrium is simple: If all players choose their idiosyncratic signals, the mean will be zero in expectation (since the signals are centered around zero). Thus, $2/3$ -average is zero, and the target becomes each player’s own signal ϵ_i . If every player behaved this way, the average choice would be zero. Indeed, on average, subjects choose such an anchor or close to it. This means that, to the best of our knowledge, for the first time in a long history of BC games, equilibrium is reached on average instantaneously.

Yet, it is too early to be content. From an experimentalist’s perspective, signals have to be chosen wisely. Consider the case, when the mean of the normal distribution is not zero, then ϵ_i is not an equilibrium choice. For example, if the mean of the distribution is 10, then the equilibrium choice of a player i is $20 + \epsilon_i$. Thus, in equilibrium, players should choose 30 on average. Here again, subjects fail to make the right equilibrium choice and more likely choose their signal, and thus the average is close to 10 as in the case of variation 5b, with the target $2/3 \cdot \text{average} + 10$, with a known $c = 10$ (Bühren and Nagel, 2019). When the mean of the signal distribution is not zero, level k again becomes a descriptive model, with the level 0 player choosing the signal, and level k iterations until the equilibrium is reached.

There is, of course, another caveat. What do such signals mean in reality, or how can a policymaker choose or distribute such signals “in a wise way”? Are they “constructed” by the agents under concern themselves? It is easy to imagine that such signals come from past play, which is around some perceived or realized mean. However, if we consider a traffic experiment, a machine could give a player a signal which route to choose, and then even though all are free in their choice, on average the decision are right and traffic jams are avoided.

Variation 7a: Strategic substitutes vs. complements in a one-shot game It is well-known that models with strategic substitutes and complements can have very different stability properties in equilibrium.⁷ This difference can be easily shown with BC experiments. So far, we typically maintained $p = \frac{2}{3}$. In this variation, we just change the sign of the target to $-\frac{2}{3} \cdot \text{average}$ and compare behavior in the two cases, $p = \frac{2}{3}$ or $-\frac{2}{3}$. The equilibrium does not change.

The two games differ by their feedback structure. The feedback structure can be characterized by the two related concepts of *strategic complements* (also called positive feedback) and *strategic substitutes* (negative feedback) (see Bulow et al., 1985). Agents’ decisions are complements if they have an incentive to match other agents’ decisions. Conversely, agents’ decisions are substitutes if agents have an incentive to do the opposite of what others are doing. For instance, if a firm can increase its profit by charging the same price as others, then prices are strategic complements. If firms can make more profit by charging high prices when their competitors charge low prices (or vice versa), then prices are strategic substitutes.

⁷For other surveys covering the topic strategic substitutes vs. complements, see, e.g., Camerer and Fehr (2006); Hommes (2013, 2020).

In simple BC games where agents need to be closest to a target of p times the average, whether the system exhibits strategic substitutes or strategic complements only depends on the coefficient p . If $p > 0$, the system exhibits strategic complements, since if all others increase their choices, an individual also has an incentive to increase her choice. Conversely, if $p < 0$, the system exhibits strategic substitutes, since if all others increase their choices, the individual has an incentive to decrease her choice and vice versa.

Theoretically, the same prediction would be obtained. Yet, out-of-equilibrium changes occur. The reader is invited to see how best replies work in the two different p -situations (hint: start with any arbitrary average and find the best reply to it and continue such procedure several times). One will quickly see that behavior is closer to equilibrium in the case of p with a negative sign as compared to a positive sign. In the former, a negative average incurs a positive best reply and thus already one or two level of reasoning result in positive and negative realizations. As a result, players produce a smaller average and closer equilibrium behavior (zero) as compared to the treatment with $p > 0$ with strategic complements. Under strategic complements, higher average choices provoke higher best replies, driving the aggregated outcome further away from equilibrium zero, see Benhabib et al. (2019) and table 1 which produces an average of 14.03 for complements and 1.87 for substitutes and $p = 2/3$. In the next variation, we discuss some policy implications for such game theoretic properties.

Variation 7b: Strategic substitutes vs. complements with behavior over time The basic BC games are ideal for showing the difference between these theoretical structures, especially over time. The convergence to equilibrium in case of complementarity is especially slow if inexperienced subjects. Heemeijer et al. (2009) conduct a BC game, in which subjects only receive qualitative information about the reaction function. The reaction function are linear so that each individual's optimal choice under complete information would be $c + p \cdot \text{average choice}_t + \epsilon_t$, where c is a constant and $\epsilon_t \sim N(0, 0.25)$ is a small stochastic shock. In one treatment, choices are strategic complements ($p = \frac{20}{21}$), while in the other treatment choices are strategic substitutes ($p = -\frac{20}{21}$). The parameter c is calibrated so that the fundamental equilibria are the same across treatments.

As shown in figure 5, strategic substitutes lead to much closer equilibrium choices than in a situation with strategic complements. An arbitrarily high average belief in the latter will induce a high best reply, while in the former, a low negative choice will result. Thus, if people just differ by one reasoning step, the average of such types will come closer to equilibrium in strategic substitute cases. This is a great example in which there is little information in both cases. Equilibrium can be attained in one and not in the other case. Also, this (non-)convergence does not depend on the sophistication of the subjects, which are here the typical student population.

Cooper et al. (2017) introduce the subjects to a BC game ($2/3$ average plus 10, in case of complements, and $-2/3$ average +10 for substitutes) for some periods. Then they introduce a different kind of shock, related to the parameters: they implement a known change of p and c . After this new information, behavior converges quickly to equilibrium in the treatment with strategic substitutes. In contrast, it only converges slowly in case of strategic complements, although both treatments have trained the subjects in the same way in the first set of parameters.

Assenza et al. (2018) give an example of a policy recommendation where an environment with strategic complements changes to an environment with a mixture of strategic complements and strategic substitutes. They show that disobeying the Taylor principle in a New-Keynesian framework (i.e., $\phi_\pi > 1$; see section 4.3.3) leads to a system which purely exhibits strategic complements.

Figure 5: Experiments by Heemeijer et al. (2009) with strategic substitutes, $p = -\frac{20}{21}$ (left); experiments with strategic complements $p = \frac{20}{21}$ (right)

However, fulfilling the Taylor principle introduces strategic substitutability into the system so that the New-Keynesian framework becomes a mixture of strategic substitutes and strategic complements. Their experimental results show that if the Taylor principle is not fulfilled, outcomes generally do not converge to the RE steady state. However, they also show in their experiments that just obeying the Taylor principle ($\phi_\pi = 1.005$) does not result in convergence. Yet, they experimentally suggest that a stronger reaction coefficient ($\phi_\pi = 1.5$) is sufficient to ensure convergence to the steady state.

Variation 8: Not knowing the underlying parameters In many BC experiments, subjects were informed about the precise functional form of the best-response function. In other words, they knew the target to pick a number to be closest to two-thirds times the average. Yet, there is a large number of experiments in which only qualitative information is provided to subjects. An example would be conducting a BC game in which subjects are only given the information that “the higher the average choice, the higher is the number you need to choose yourself.” This is a common procedure in the experimental macroeconomics and the experimental finance literature since neither market participants nor the policymaker may know the underlying law of motion of a real market.

Sonnemans and Tuinstra (2010) review the literature on BC games with and without knowledge of the underlying parameters. They conclude that the nature of the strategic environment (being captured by the ‘p-parameter’ in standard number guessing games, which indicates whether choices are strategic substitutes or strategic complements) is the vital feature that determines whether choices converge (quickly) to equilibrium. Details of the experimental design, such as the target number, the information given to the participants, and the incentive structure, play less of a role

for convergence. An example for a study where the parameters are systematically varied and at the same time unknown to subjects is Heemeijer et al. (2009). They show that dynamics are more unstable when $p > 0$ than when $p < 0$.

In setups with incomplete information, the level k model cannot be directly applied. Instead, something similar might happen: Over time, an adjustment factor, how much to deviate from the previous result, is something similar to a step level or gradient learning model. If one sees step sizes increasing from one period to the next, then one might increase even more the step size in subsequent periods. Or when step sizes do not change, one should maintain the same level. With fMRI techniques, one can visualize whether subjects consider the reasoning processes of others or not even if the game structure itself does not reveal it through behavior (see Nagel et al. (2018) in market entry games or Hampton et al. (2008) in repeated inspection games). In both papers, the authors visualize brain activity in specific theory of mind areas, e.g., in the mPFC, which correlates higher-order belief reasoning.

If laboratory experiments show clear discrepancies between the normative benchmark and observed play under optimal information conditions, then in the real world, it might be an even more obvious problem to play a game that does not correspond to the actual true game. A systematic test about whether knowledge of the underlying model matters is provided by Mirdamadi and Petersen (2018). In an experimental New-Keynesian economy, where agents need to forecast macroeconomic variables, they vary the information participants receive about the economy's data-generating process (no information, qualitative information, and quantitative information). They show that providing quantitative information about the underlying data-generating process consistently reduces inflation forecast errors as well as the dispersion of inflation expectations.

Summary comment about experimental insights for a behavioral micro-foundation of macroeconomic models In this section, we have presented leading examples of discrepancies between human behavior and theoretical solutions in BC games. We mitigated such gaps, allowing for repetitions, exogenous noisy signals, strategic substitutes and complements through parameter variations, and sophistication of subjects. Introducing dominant strategies was not a sufficient condition to obtain equilibrium behavior. The reason is the lack of cognitive ability to find such strategies. We chose to present behavior in variations of the BC contest since it forms the closest relationship to macroeconomic modeling of expectation formation. This will be shown more formally in the next section. The patterns in out-of-equilibrium behavior were largely explained through the level k model, a cognitive model of strategic reasoning, which has finally been implemented into macroeconomics more than 25 years after the first publications in the experimental literature. Note that the out-of-equilibrium dynamics in the experiments did not depend on exogenous shocks. Instead, such shocks and heterogeneity are created in the heads (mind), thus endogenously, by the subjects themselves, who therefore form (higher-order) beliefs of such shocks of other players.

A critical comment about level k and some alternatives So far, we have shown results of actual behavior in variations of the BC game, which we consider important building blocks for a behavioral macro foundation. Most of the studies were closely tied with the level k model as a good descriptive theory. A critical comment is now arising, how general this model is. First of all, it is one of the few cognitive reasoning models for strategic interaction. This means that the model is mimicking the reasoning procedure of actual subjects. This has been visualized not only through choices but through different kinds of measures as asking subjects directly for comments

(Bosch-Domenech et al., 2002), through brain-activity (Coricelli and Nagel, 2009; Nagel et al., 2018), letting them talk in teams of two subjects per one choice, and other measures, see for such measures (Mauersberger and Nagel, 2018, section 5).

Several critical questions arise with this model of which the formulation of level 0, the anchor or reference point, is the most important one. Most experimental work bases the reference point of the first period according to (uniform) random choices. Yet, if there are clear focal or salient points, perceived even for naive players, those might be better first anchors. There is not yet an established general theory that constitutes a salient or focal point for game theory besides a list of examples of focal points in different specific games, as we also showed in our variations above.

So far, no paper systematically analyses empirical papers with “failures” of level k , as typically done in meta-studies. Within similar games (games that are only different due to parameter changes), subjects maintain similar levels. This is not necessarily true across different games. In the latter case, the average level is typically also between level 0 and level 3, but subjects are fairly inconsistent from one game to another with respect to their levels (Georganas et al., 2015). For macroeconomics, this could be good news if at least the average level within a specific population is fairly robust, as individual behavior may be of less interest.

Another critical shortcoming of this model is that a player can typically form beliefs about level k choices only if the environment is fairly well specified. This critique was discussed in the context of variation 8 (Not knowing the underlying parameters.)

This also leads to a concern about making predictions with such a behavioral model. First of all, the level k model is a descriptive model, specifying types of reasoning. It does not make predictions about frequencies of single types besides that high level types (higher than 2) are rarely observed in the usual undergraduate or general public populations. However, experts’ behavior, i.e., behavior of those who are familiar with a particular game through experience, corresponds to higher levels of reasoning. Yet, some predictions can be made if a low level of reasoning (naive behavior or level 1) coincides or is close to an equilibrium point. Then average behavior is typically also (close to) this point. Also, the contrary is true. If high level of reasoning steps are required to reach equilibrium, the model predicts that behavior is far away from equilibrium. Whether behavior converges to equilibrium and whether naive play is dying out is based on the structure of the game (see, e.g., variation 7 strategic substitutes vs. complements).

As a last point, Camerer et al. (2004) have introduced a one-parameter model using a Poisson distribution to specify a belief system of higher than level 1 types over lower level types. Estimating the Poisson parameter (which is typically near 1.5), data for other experiments can be predicted within the same out-of-sample specifications or out-of-game specification. For theoretical modeling and policy implications, such predictions together with known out-of-equilibrium behavior can be fruitful inputs.

4 Microfoundations of macroeconomics

Mauersberger and Nagel (2018) introduce a generalized BC game with a specific best-reply function which we describe in this section. Despite its clear functional restrictions many games can be encompassed under the same specifications: among them are Cournot competition, Bertrand competition, auctions, asset markets, Cobweb models, New-Keynesian models, the 11-20 game, global games, ultimatum games and stag-hunt games, etc.. This section explains in some detail why standard workhorse models in macroeconomics can be considered as a BC game.

Section 4.1 introduces the generalized BC game as discussed by Mauerberger and Nagel (2018).⁸ Section 4.2 discusses neoclassical models of the business cycle in this context. Section 4.3 establishes the link between the standard New-Keynesian model and aggregative games. We first consider the standard New-Keynesian model in section 4.3.1, in which the relationships between the macroeconomic variables have been derived under the assumption of a rational, representative agent. Under certain restrictions on the expectations operator, expectations in those first-order conditions can be replaced by boundedly rational forecasts. We do not explain the details of these restrictions in this chapter and refer the interested reader to Branch and McGough (2009). In section 4.3.2, we consider a more microfounded version of the New-Keynesian model that allows for quite arbitrary non-rational individual forecasts. Section 4.4 establishes the link between growth models and BC games. To the best of our knowledge, this latter link has previously not been established yet.

4.1 Generalization of the Beauty Contest game

Now let us consider a generalization of the BC game in order to present a clear relation between experimental BC games and the microfoundations of macroeconomics. The BC game in which players have to reach a target that involves the average falls in the class of *aggregative games*. Aggregative games have been introduced by Selten (1970). Aggregative games are defined as games in which the payoffs of every player are a function of their own strategy and the sum of all players' strategies. Since the mean is the sum divided by the number of players, BC games with a target containing the average fall into the class of aggregative games.

Aggregative games have been generalized to so-called *fully aggregative games* (Cornes and Hartley, 2012). Cornes and Hartley (2012) propose a more general aggregation function. They require that the aggregation function $g(\cdot)$ needs to be additively separable, i.e. for the strategies of all players $y_1, \dots, y_N$, there exist increasing functions $h_0, h_1, \dots, h_N : \mathbb{R} \rightarrow \mathbb{R}$ such that $g(y_1, \dots, y_N) = h_0(\sum_{i=1}^N h_i(y_i))$. It is easy to see that any aggregative game is also a *fully aggregative game*. Acemoglu and Jensen (2013) analyze how a change in the exogenous parameters affects the Nash equilibrium, also considering the cases of strategic substitutes.

We define the notion of *generalized Beauty Contest game* following Mauerberger and Nagel (2018). While we consider the one-dimensional case where players choose among real numbers, we consider the multidimensional case (and an example thereof) in subsection 4.3. We refer to a game as a *generalized Beauty Contest game* if the best response of a player i at time t can be written as

$$y_t^i = \hat{E}_t^i \{ c + b \cdot f(y_{t-1}^1, y_t^2, \dots, y_t^N) + d \cdot g(y_{t+1}^1, y_{t+1}^2, \dots, y_{t+1}^N) + \epsilon_t^i \} \quad (1)$$

where i is the individual-superscript, N is the number of individuals, c, b, d are coefficients, and the functions $f(\cdot)$ $g(\cdot)$ are the aggregation across all individuals. Now future values of y enter, which are the individual i 's subjective expectations of future realizations of y . We allow for several different aggregations. $\hat{E}_t^i$ denotes the subjective, possibly non-rational expectation of individual i held at time t . Following Woodford (2013), we assume that this expectations operator represents a well-behaved probability measure over possible future evolutions of the variables in the curly brackets $\{\cdot\}$. $f(\cdot)$ and $g(\cdot)$ can be the average, the median, the sum, the minimum, the maximum, the mode, the least chosen action, or the action chosen by at least $h \leq N$ individuals. ϵ_t^i is an idiosyncratic shock.

⁸The origin of this model can be considered to be the paper by Benhabib et al. (2015), who introduce a neoclassical general equilibrium model with sentiments.

Thus, if the function $f(\cdot)$ or $g(\cdot)$ correspond to the sum or the average, the game falls into the class of *aggregative games* and in the class of *fully aggregative games*. If the function $f(\cdot)$ or $g(\cdot)$ corresponds to other order statistics (minimum, maximum, mode, the least chosen action, or the action chosen by at least $h \leq N$ individuals), then the game does not fall into the class of *fully aggregative games* since these are not additively separable functions.

4.2 Neoclassical models of the business cycle

A natural starting point of macroeconomics is asking whether aggregate fluctuations can be understood by a Walrasian model, i.e., a general equilibrium model with competitive markets without any externalities, asymmetric information, frictions or other market imperfections (see, e.g., Romer, 2012, for a detailed textbook exposition).

In this type of framework, Benhabib et al. (2015) assume that production and employment decisions by firms, and consumption and labor supply decisions by households are made before goods are produced and exchanged and before the realization of market-clearing prices. They show that a simple, abstract, log-linearized version of their model can be written by the equation

$$y_t^i = \hat{E}_t^i \{ b y_t + \epsilon_t^i \}$$

where y_t^i is the optimal output of firm i , and ϵ_t^i is an idiosyncratic demand shock. y_t represents the average output across all firms in the economy. This optimality condition can be nested into the canonical framework given by equation (1), since it represents the special case where $f(\cdot)$ is the average and $c = d = 0$. As shown by Benhabib et al. (2015), under the standard Dixit–Stiglitz specification, output decisions are strategic substitutes so that $b < 0$. Under rational expectations, they show that there is a unique fundamental equilibrium where $y_t^i = y_t, \forall i$.⁹

4.3 New-Keynesian models

An alternative view of the business cycle in macroeconomics is one in which market prices are not fully flexible due to, for example, menu costs or other frictions. This class of models is often referred to as “New-Keynesian models” (see, e.g., Galí, 2008). We begin by mapping the standard textbook New-Keynesian model to equation (1). Subsequently, we briefly describe a more microfounded version of that model in which one can use a wide range of non-rational expectations.

4.3.1 Standard textbook New-Keynesian model

A heterogeneous expectations version of the textbook New-Keynesian model as encountered in Woodford (2003), Galí (2008) or Walsh (2010) can, under some restrictions of the expectations operator (see Branch and McGough (2009)), be written as

$$y_t = \bar{y}_{t+1}^e - \sigma(i_t - \bar{\pi}_{t+1}^e - \rho) \quad (2)$$

$$\pi_t = \kappa y_t + \beta \bar{\pi}_{t+1}^e \quad (3)$$

$$i_t = \rho + \phi_\pi(\pi_t - \pi) \quad (4)$$

⁹Moreover, Benhabib et al. (2015) show that there is a sentiment-driven equilibrium when signals are combinations of private and public signals which require signal extractions.

where y_t denotes the output gap, i.e. the difference between output in the New-Keynesian economy and output under flexible prices. π_t denotes inflation, t the time-subscript and $\bar{\pi}_{t+1}^e$ and $\bar{y}_{t+1}^e$ the mean expected future values of output gap and inflation, being the average forecast of all subjects. The model is closed under an inflation targeting rule for the nominal interest rate i_t with a constant inflation target π . It is easy to see that π is the long-run rational expectations steady state (not necessarily the unique rational expectations equilibrium though).

We consider a learning-to-forecast game, in which subjects are only paid for forecasting inflation and output gap. After everyone submitted their forecasts, actual output gap and inflation are generated by equations (2) and (3). Note that inflation and output gap of period t depend on $\bar{\pi}_{t+1}^e$ and $\bar{y}_{t+1}^e$, respectively, which are the means across all agents' forecasts for period $t+1$. An individual subject i is then paid according to a distance function such as $U_{t+1}^i = A - Q(\hat{E}_t^i \pi_{t+1} - \pi_{t+1})^2$. A and Q are positive constants. This distance function measures how close the agent's inflation forecast for period $t+1$ (given in period t), $\hat{E}_t^i \pi_{t+1}$, is to the actual inflation in $t+1$, π_{t+1} . This is analogous for the output gap forecast.

For the sake of simplicity, we describe the New-Keynesian model as a univariate forecasting game of inflation at time t , π_t .¹⁰

To reduce the dimensionality of this forecasting game to a single dimension, we use the ad-hoc assumption that expectations of the output gap are equal to its long-run steady state value, obtained by using equation (3):

$$\bar{y}_{t+1}^e = \kappa^{-1}(1 - \beta)\pi \quad (5)$$

Substitute (5) into (2) to obtain

$$y_t = \kappa^{-1}(1 - \beta)\pi + \sigma\phi_\pi\pi - \sigma\phi_\pi\pi_t - \sigma\bar{\pi}_{t+1}^e \quad (6)$$

$$\pi_t = \kappa y_t + \beta\bar{\pi}_{t+1}^e \quad (7)$$

Insert (6) into (7) one yields the value of inflation that subjects needed to forecast at time $t - 1$:

$$\pi_t = c + d\bar{\pi}_{t+1}^e \quad (8)$$

with $c \equiv \frac{(1-\beta)+\kappa\sigma\phi_\pi}{1+\kappa\sigma\phi_\pi}\pi$ and $d \equiv \frac{\beta-\kappa\sigma}{1+\kappa\sigma\phi_\pi}$. It becomes apparent that the New-Keynesian model can be considered a BC game with $c > 0, d > 0, b = 0, f(.) = \bar{\pi}_{t+1}^e$ being the average. A learning-to-forecast experiment, in which subjects' only task is forecasting inflation, based on the model equation of the New-Keynesian framework (8) has first been introduced by Pfajfar and Žakelj (2014). These authors vary the Taylor rule for different treatments, and find that the variance in inflation expectations decreases starkly as the computerized central bank adopts a more aggressive response to deviations of inflation from the target.

4.3.2 A more microfounded New-Keynesian model

Woodford (2013) and García-Schmidt and Woodford (2019) consider a version of the New-Keynesian model that is populated by heterogeneous households and firms that do not necessarily have rational expectations. Specifically, these papers assume that forecasts are not necessarily model

¹⁰One can also consider this model as a BC game, in which the variable of interest to be a vector. Anufriev et al. (2019) conduct an experiment on this type of game.

consistent, not necessarily the same across agents and expectations are not necessarily those that agents previously expected to hold.

In their framework, it is shown that the optimal choice of consumption and price not only depend on the macroeconomic outcomes in the next period but instead on the outcomes far into the future:

$$c_t^i = \sum_{k=0}^{\infty} \beta^k \hat{E}_t^i \{ \lambda c_{t+k} - \sigma(i_t - \pi_{t+k}) \} \quad (9)$$

$$\pi_t^j = \sum_{k=0}^{\infty} (\alpha\beta)^k \hat{E}_t^j \{ \phi c_{t+k} + \xi \pi_{t+k} \} \quad (10)$$

where t is the time-period subscript, i is the household superscript, j is the firm superscript, and $\hat{E}_t^i$ denotes the boundedly rational expectation held at time t . β is the discount factor that discounts future payoffs and α is an exogenous probability that the price remains the same as in the last period $t - 1$. $\lambda, \sigma, \phi, \xi$ are positive constants.

Equation (9) presents the consumption rule for the households: c_t^i is a measure of individual total consumption expenditure at time t , c_t denotes average consumption expenditure at time t . i_t is the nominal interest rate which is set by the central bank and can thus be arbitrarily specified. Thus, the (rational expectations) equilibrium and whether it is unique totally depends on the monetary policy rule. Equation (10) presents the pricing rule of the firms: π_t^j denotes firm j 's individual inflation at time t (compared to average prices at time $t - 1$) and π_{t+k} denotes average inflation at time t .

Equations (9) and (10) have a recursive form so that they can be rewritten as

$$c_t^i = \lambda \hat{E}_t^i c_t - \sigma(i_t - \hat{E}_t^i \pi_t) + \beta \hat{E}_t^i c_{t+1}^i \quad (11)$$

$$\pi_t^j = \phi \hat{E}_t^j c_t + \xi \hat{E}_t^j \pi_t + \alpha \beta \hat{E}_t^j \pi_{t+1}^j \quad (12)$$

Thus, individually optimal consumption and prices depend on agents' beliefs about the average consumption and price in the economy and the forecast of agents' own optimal choice in the next period. This equation can be written in matrix form in the following way:

$$\begin{bmatrix} c_t^i \\ \pi_t^j \end{bmatrix} = \begin{bmatrix} -\sigma i_t \\ 0 \end{bmatrix} + \begin{bmatrix} \lambda & \sigma \\ \phi & \xi \end{bmatrix} \begin{bmatrix} \hat{E}_t^i c_t \\ \hat{E}_t^j \pi_t \end{bmatrix} + \begin{bmatrix} \beta & 0 \\ 0 & \alpha\beta \end{bmatrix} \begin{bmatrix} \hat{E}_t^i c_{t+1}^i \\ \hat{E}_t^j \pi_{t+1}^j \end{bmatrix} \quad (13)$$

Strictly speaking, equation (13) does not fit into our scheme of a BC game in equation 1 because of the terms $\hat{E}_t^i c_{t+1}^i$ and $\hat{E}_t^j \pi_{t+1}^j$. At the same time, (13) provides an interesting game structure that has been analyzed as a laboratory experiment by Mauersberger (2016).

However, by adopting a simplified view, one could consider the New-Keynesian model in (13) as a BC game by crudely considering $\begin{bmatrix} \beta & 0 \\ 0 & \alpha\beta \end{bmatrix} \begin{bmatrix} \hat{E}_t^i c_{t+1}^i \\ \hat{E}_t^j \pi_{t+1}^j \end{bmatrix}$ as an idiosyncratic term ϵ_t^i - at least from the analysts perspective. Thus, a paper that can be considered as a (simplified) experimental implementation of equation (13) is Benhabib et al. (2019). The differences between Benhabib et al. (2019) and Mauersberger (2016) are: (a) that Benhabib et al. is a static BC game; (b) most importantly, that Benhabib et al. specify $\hat{E}_t^i c_{t+1}^i$ (or $\hat{E}_t^j \pi_{t+1}^j$) as an exogenous shock while in Mauersberger (2018), $\hat{E}_t^i c_{t+1}^i$ and $\hat{E}_t^j \pi_{t+1}^j$ are forecasts and thus choices of the subjects.

4.3.3 Strategic substitutes and strategic complements

In the experimental section, we discussed the differences in behavior when the system exhibits strategic substitutes or strategic complements which only depends on the coefficient p . If $p > 0$, the system exhibits strategic complements, and if $p < 0$, the system exhibits strategic substitutes.

The New-Keynesian model is a more complicated environment than simple BC games. In the New-Keynesian model, there are two endogenous variables, inflation and consumption, and the interest rate, a policy variable. Depending on the policy reaction to inflation volatility, the system may exhibit purely strategic complements or a mixture of strategic substitutes and complements. See Assenza et al. (2018) for a more detailed exposition.

Equations (11) and (12) show that the only source introducing strategic substitution into the system is the policy interest rate of the central bank. A commonly used interest rate rule is

$$i_t = \bar{i} + \phi_\pi(\pi_t - \pi) \quad (14)$$

where $\bar{i}$ is a constant corresponding to the long-run steady state of the interest rate and π is the central bank's inflation target. The constant ϕ_π captures the strength of the response of the central bank to inflation fluctuations.¹¹ Using (14) in (11) shows that the monetary policy authority introduces some strategic substitutability into the system if $\phi_\pi > 1$. Then the New-Keynesian model becomes a mix between strategic substitutes and strategic complements.

The condition $\phi_\pi > 1$ is also the sufficient conditions that guarantees a unique rational expectations equilibrium. In the absence of any shock, the rational expectations equilibrium becomes π for inflation and the long-run steady state $\frac{\sigma(\bar{i}-\pi)}{\lambda+\beta-1}$ for consumption.

4.4 Growth

This section is a digression to long-run macroeconomics. By *long-run*, we mean that the capital stock is not fixed, being a key assumption in neoclassical growth models. Standard growth models like the Solow-Swan model (Solow, 1956) and the Ramsey-Cass-Koopman models (Ramsey, 1928; Cass, 1965; Koopmans, 1965) describe the growth of the capital stock in the following way:

$$K_{t+1} - K_t = (f(K_t) - c_t) - \delta K_t \quad (15)$$

K_t is the aggregate capital stock of the economy in any period t , f is the neoclassical production function, c_t the consumption and delta the depreciation rate.

Lei and Noussair (2002) implement this model as an experimental coordination game with a decentralized market for capital. In this version, the economy is populated by a finite number of individuals N , which are indexed by j . The capital stock of the economy is then the sum of individual capital holdings k_t^j : $K_t = \sum_{j=1}^N k_t^j$. With e.g. a quadratic utility function $U(c_t) = 310C_t - 5C_t^2$, there exist a unique pair (C^*, K^*) of optimal consumption and capital that determine the so-called golden rule steady state. Social welfare is thus maximized if

$$K_t = \sum_{j=1}^N k_t^j = K^* \quad (16)$$

¹¹Following Galí (2008), we use the subscript π to highlight that this constant captures the response of the central bank to inflation as opposed to any other economic variable.

which can, for a specific individual i , be rewritten as $k_t = K^* - \sum_{j=1, j \neq i}^N k_t^j$. This is a BC game with $c = K^*$ and $b = -1$ and f being the sum of the investment choices of all individuals.

5 Behavioral microfoundations of macroeconomics

5.1 Level k

There is a recent growing literature, exploring the implications of finite depths of reasoning in macroeconomics.¹² One technical difference is that the microeconomic literature on Beauty Contest games rather deals with *static* BC games, closely related to Keynes newspaper contests, i.e., setups in which agents need to form beliefs about the current decisions of others. Contrary to that, papers in macroeconomics also study *dynamic* BC games, in which agents must form beliefs about future actions of others.

Early approaches of level k in macroeconomics include Evans and Ramey (1992), who introduce calculation costs for using the (correct) model equations. Another early approach introduced by Guesnerie (1992, 2009) "eductive stability" of the REE. The question of eductive stability is whether, for any initial (naive) belief, outcomes converge to the REE using infinitely many iterations of the model equations.

More recently, García-Schmidt and Woodford (2019) introduce level k into a New-Keynesian model that allows for heterogeneous boundedly rational forecasts in the same way as Woodford (2013). The motivation for their paper is the empirical observation that a prolonged period of low nominal interest rates during the financial crisis has not resulted in high inflation. This insight has led to increased interest in the "Neo-Fisherian" hypothesis, according to which low nominal interest rates may themselves cause lower inflation. García-Schmidt and Woodford (2019) challenge the Neo-Fisherian paradox by proposing that agents in their model use level k to form expectations. Under level k, a commitment to maintaining a low nominal interest rate for longer should always result in both higher output and higher inflation rather than deflation. However, in the case of a long-horizon commitment to low interest rates, the expansionary or inflationary effects are less pronounced than under rational expectations without any uncertainty. The explanation is as follows: They assume that the expectations of the naive level 0 type correspond to a state of the economy in which no shock has occurred. Since other types respond to such behavior, there will be no deflationary spiral. However, the response to the naive type and other boundedly rational types causes inertia in the adjustment. In other words, there will not be an immediate adjustment to high output and high inflation, as in the rational case. García-Schmidt and Woodford (2019), therefore, present a solution to the forward guidance puzzle. The forward guidance puzzle is the finding that, under common knowledge, a credible announcement regarding the monetary policy of e.g. 1,000,000 years in the future has the same or even greater effects today as an announcement regarding the monetary policy of tomorrow. García-Schmidt and Woodford (2019) argue that there is no puzzle. Their point is that the naive type drives aggregate behavior, so that these authors conclude that the 'central bank's intentions regarding future policy should have no effects on equilibrium outcomes.'

Angeletos and Lian (2018) study the effects of uncertainty about others' actions in a large class of games that nests but is not limited to the New-Keynesian model. Based on the insight that the

¹² Alternative approaches to the level k literature to introduce bounded rationality into macroeconomics have been adaptive learning (see e.g., Evans and Honkapohja (2001)), diagnostic expectations (Bordalo et al., 2018) and limited attention (see e.g., Sims, 2003, Maćkowiak and Wiederholt (2009) and Gabaix (2014)).

New-Keynesian model can be considered as a BC game, Angeletos and Lian obtain the following results in the absence of common knowledge and the presence of higher-order beliefs (with a level of reasoning greater than one): first, any general-equilibrium effect is mitigated for a coefficient smaller than one. Provided that the coefficient is smaller than one, all things equal, agents with higher-order beliefs are less responsive than agents with lower-order beliefs. This is because after the shock, the level 0 type adjusts his behavior immediately and expects to optimize choosing the previous period average. The higher the level of reasoning the further away is an agent from the level 0 choice and the closer she is to the rational expectations equilibrium. Second, the further in the future an event occurs, the smaller the effect on the present - a phenomenon that they call “horizon effect.” This can be explained by the fact that longer horizons involve more iterations of the forward-looking, Euler-type equations, which have a dampening effect for beliefs of higher order. Third, under quite general conditions, the effect size goes to zero as the horizon $T \rightarrow \infty$. This is due to the fact that infinite levels of reasoning are anchored to the common prior and are hence unresponsive to the news even in the presence of small idiosyncratic shocks.

Angeletos and Lian (2018) show that this has two important implications for New-Keynesian macroeconomics: first, it solves the forward-guidance puzzle. Second, under common knowledge, it has been shown that the current effects of a fiscal expansion of a given magnitude increases as the stimulus is announced to take place further in the future. This is no longer the case under imperfect knowledge, which provides a rationale for “front-loading.”

Farhi and Werning (2019) present another application of level k in macroeconomics. They show their results both qualitatively and quantitatively. They demonstrate that future interest rate changes have smaller effects than present interest rate changes under level k thinking, which they call “mitigation effect.” Like Angeletos and Lian (2018), they also find that introducing level k thinking leads to a “horizon effect”, i.e., interest rate changes further out in the future have smaller effects in the presence. However, their calculations show that both the mitigation and the horizon effect under level k thinking are relatively modest in size. They show that a sizable effect is only yielded by combining level k with incomplete markets, introducing uninsurable idiosyncratic risk.

Iovino and Sergeyev (2018) investigate the implications of standard central bank balance sheet policies, i.e., quantitative easing and foreign exchange interventions, under the assumption that agents engage in level k thinking in an overlapping generations model. They find different implications from the rational expectations benchmark: First, in contrast to rational expectations, where central bank implications are neutral, central bank interventions are effective under level k thinking under mild conditions. The reason is that since agents do not hold rational expectations about future endogenous variables, they underestimate the tax risk resulting from policy interventions and incorrectly forecast future asset prices. As a result, they demand lower risk premia, which increases asset prices and thus renders balance sheet policies effective. Second, they show that individual and cross-sectional average forecast errors about future endogenous variables are predictable after balance sheet interventions. They validate their predictions, using data on the mortgage purchases by US enterprises as a proxy for quantitative easing.

Recently, the idea of higher-order beliefs in forecasting has also raised interest in the empirical literature in macroeconomics. Coibion et al. (2018) introduce a survey to New Zealand firm managers, not only eliciting their expectations about inflation but also asking them about their beliefs of other managers and measuring their depth of reasoning in an incentivized p-Beauty Contest task. They also investigate whether managers’ beliefs influence their actions, finding that firms obtaining any type of information made significant reductions in employment and investment but not in prices

or wages. Furthermore, the authors demonstrate differences in standard implementations of k level modeling. First, most managers that exhibit a level of reasoning lower than $k = 4$ believe others will submit an answer in the same range as theirs. Thus, they do not select a number equal to two-thirds times the estimated average. Second, managers generally believe that some of the other managers exhibit higher levels of reasoning than they do. In contrast to that, level k and other models assume that agents act as if all other agents are lower level thinkers than them.¹³ Third, managers highly underestimate the dispersion of the answers. Altogether, they find no significant evidence that agents' degree of k level thinking is related to either how they update their beliefs based on new information or how changes in information affect their decisions.

5.2 Myopia and discounting

Level k means that agents exhibit bounded rationality when looking at market competitors and other agents. An alternative, recently quite popular way of introducing bounded rationality into macroeconomics is the conjecture that people are inattentive to variables that are further away in the future. This literature is briefly reviewed here because those models relate to our general equation (4.1).

Papers like Gabaix (2019) and Angeletos and Huo (2019) investigate the general equilibrium effects of myopia and discounting. Both papers use representative-agent frameworks. Under a representative-agent framework, the (possibly multidimensional version of the) general BC equation (1) can be simplified into the form

$$y_t = \epsilon_t + dE_t y_{t+1} \quad (17)$$

The two studies focus on Euler equations so that $b = 0$. The individual subscript i is dropped because of the representative-agent framework.

Gabaix (2019) uses the theory of sparse maximization (Gabaix, 2014) under which equation (17) becomes

$$y_t = m_\epsilon \epsilon_t + dm_y E_t y_{t+1} \quad (18)$$

where m_x is a vector of attention weight on the respective variable x , where x is either ϵ or y . Each element of m_x is between 0 and 1. This way, the attention to future variables is dampened. Consider for simplicity the unidimensional case. $m_x = 0$ means that the variable x is entirely disregarded, while $m_x = 1$ corresponds to the fully rational model. Gabaix (2019) shows that such a behavioral New-Keynesian model can explain several macroeconomic puzzles. (i) Fiscal stimuli or “helicopter drops of money” are effective and successfully drag the economy out of the zero lower bound. Specifically, the model allows for the joint analysis of optimal monetary and fiscal policy under behavioral inattention. (ii) The Taylor principle no longer applies in a classical way: even with passive monetary policy, the equilibrium is determinate, while the traditional rational model yields multiple equilibria. Multiple equilibria are a concern because of reduced predictive power, and indeterminacy at the zero lower bound. (iii) The welfare costs at the ZLB are much lower than in the traditional New-Keynesian model. (iv) The behavioral New-Keynesian framework also solves

¹³In a theoretical model, Stahl (1993) introduces a large number of thinking types, from those who think that others are below them and others are above them, with different informational assumptions about other players distributions.

the “forward guidance puzzle”, being the conundrum that in the rational model, shocks to very distant interest rates may have a very large effect on today’s macroeconomic outcomes. Since agents in the behavioral New-Keynesian model are partially myopic, this effect no longer exists. (v) There are qualitatively different implications of optimal policy: the optimal commitment policy with rational agents requires “nominal GDP targeting.” This is not the case with “behavioral” firms, because the benefits of commitment are lower under myopia. (vi) The model is “neo-Fisherian” in the long run, but Keynesian in the short run. This is because a permanent rise in the interest rate lowers inflation in the short run but increases it in the long run.

Angeletos and Huo (2019) focus on an abstract but general framework (spanning all representative models that are nested by equation (17)) and show that, under appropriate assumptions, the incomplete-information economy is observationally equivalent to a complete-information, representative-agent economy in which condition (17) is modified to

$$y_t = \epsilon_t + d\omega_f E_t y_{t+1} + \omega_p y_{t-1} \quad (19)$$

where $\omega_f < 1$ and $\omega_p > 0$. The first alteration ($\omega_f < 1$) captures myopia towards the future, the second one ($\omega_p > 0$) represents an anchoring effect of the current outcome to the past outcome. The authors thus explain the hybrid New-Keynesian Phillips curve, postulating that agents are not only forward-looking but also backward-looking.

Woodford (2018) takes a different approach and outlines a New-Keynesian model in which agents look ahead only a finite distance into the future, so that their planning horizons are explicitly truncated. In this model, agents use a value function learned from experience to evaluate situations that may be reached after a finite sequence of actions by themselves and others. Woodford shows that incorporating this kind of behavior into a macroeconomic model raises doubts about the “Neo-Fisherian” approach.

6 Conclusion

We explained the link between macroeconomic models and experimental BC games through a generalized BC function with limited parameter specifications, encompassing numerous macroeconomic models as for example the New-Keynesian models. By reviewing the experimental literature on BC games, we hope to have convinced the critical reader that using level k as a behavioral micro-foundation for macroeconomic modeling is a promising approach as it is also analytically tractable.

Level k allows for iterated best response structures, typically different for naive versus more sophisticated players, similarly as introduced by Keynes (1936) with his stock market reasoning procedure. This implies, that level k can be considered as an endogenous formulation of shocks produced in the heads of real subjects which explains initial and also over time out-of-equilibrium behavior. This model does not rely on exogenous shocks, as typically implemented in empirical and theoretical macroeconomic models. However, with exogenous shocks, sluggishness in unraveling might occur because people are exposed to new anchors (formulated as level 0) and through higher order beliefs.

As known in the theoretical models and visualized in BC experiments, such out-of-equilibrium behavior is especially detrimental in situations with positive feedback structure, as in the stock or housing market. Dynamics under positive feedback exhibit slower convergence to the equilibrium

(or sometimes even no convergence at all) and greater volatility than in negative feedback situations, as in Cobweb markets representing production markets without speculations (see, e.g., Camerer and Fehr, 2006; Heemeijer et al., 2009; Hommes, 2013).

Possible directions of future research regarding bounded rationality in macroeconomics could be the following: First, there is potential for more empirical work regarding level k and higher-order beliefs in real world macroeconomic settings, using data from surveys, field experiments, and professional forecasters. Second, level k is certainly not the only alternative to the growing field of behavioral macroeconomics. While learning models, for example, have long ago been built into macroeconomic theories (see Evans and Honkapohja, 2001), there have been other experimentally founded models such as the heuristic-switching model with heterogeneous expectations by Brock and Hommes (1997) and Anufriev and Hommes (2012). Hommes and Lustenhouwer (2019) use this model to analyze inflation targeting and central bank credibility. Yet, there is certainly more potential for this model to be used in macroeconomic theory. Third, given the initial success of using level k - a model from the experimental literature - in macroeconomics, more collaboration between macroeconomic theorists and experimentalists could be fruitful. Events such as the “BESLab Experimental Economics Summer School in Macroeconomics” were set up to educate young macroeconomic researchers in experimental methods.

We also showed several virtues of the concept of the rational expectation equilibrium and theoretic off-equilibrium-structures as an iterated best reply or iterated elimination of dominated strategies. First of all, the behavior may converge in the long run to the (static) equilibrium, even when players do not know the structure of the model (see, e.g. Assenza et al., 2018). Second, and equally important, even if the equilibrium is unique, Pareto optimal and welfare-maximizing, human choices may not correspond to equilibrium initially or even not over time, due to their cognitive difficulties. In such cases, there is a call for developing ideas on how to induce convergence instantaneously, as demonstrated in variation 6 of section 3 with idiosyncratic shocks. Those signals can serve as anchors or sentiments and thus can even induce naive players to play the equilibrium, instead of “falling back to chance.” John Maynard Keynes has been the “Padrino” of all our endeavors, establishing the original Beauty Contest game and a conjecture about people’s behavior with the iterated reasoning approach.

Acknowledgments

We would like to thank Michael Woodford, four anonymous referees, and David Tuckett, the editor of this special issue for their valuable comments, and the participants of the Rebuilding Macroeconomics conference at the London Treasury in October 2018. Felix Mauersberger gratefully acknowledges the receipt of an Argelander Grant from the University of Bonn. Rosemarie Nagel thanks for the hospitality during her stay at Columbia University within the Cognition and Decision Lab, where parts of this paper were written. Nagel also acknowledges the grants PGC2018-098949-B-I00 from the Ministry of Education Madrid, and the travel grant Salvador de Madariaga 2019.

References

- Acemoglu, D., and Jensen, M. K. (2013). Aggregate comparative statics. *Games and Economic Behavior*, 81(1): 27–49. ISSN 08998256. DOI 10.1016/j.geb.2013.03.009. URL <http://dx.doi.org/10.1016/j.geb.2013.03.009>.
- Alaoui, L., and Penta, A. (2016). Endogenous Depth of Reasoning. *Review of Economic Studies*, 83(4): 1297–1333.
- Amano, R., Kryvtsov, O., and Petersen, L. (2014). Recent Developments in Experimental Macroeconomics. *Bank of Canada Review*, pages 1–11. URL <https://www.bankofcanada.ca/wp-content/uploads/2014/11/boc-review-autumn14-amano.pdf>.
- Angeletos, G.-M., and Huo, Z. (2019). Myopia and Anchoring. *NBER Working Paper*, 24545. URL <https://www.nber.org/papers/w24545>.
- Angeletos, G. M., and Lian, C. (2018). Forward guidance without common knowledge. *American Economic Review*, 108(9): 2477–2512. ISSN 00028282. DOI 10.1257/aer.20161996. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20161996>.
- Anufriev, M., Duffy, J., and Panchenko, V. (2019). Planar Beauty Contests. *SSRN Working Paper*. ISSN 1556-5068. DOI 10.2139/ssrn.3404334. URL <https://www.ssrn.com/abstract=3404334>.
- Anufriev, M., and Hommes, C. (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(225408): 35–64.
- Arifovic, J., and Duffy, J. (2018). Heterogeneous Agent Modeling: Experimental Evidence. In C. H. Hommes, and B. LeBaron (Eds.), *Handbook of Computational Economics, Vol. 4*, chapter 9, pages 491–540. Amsterdam and Oxford: North-Holland. DOI 10.1016/bs.hescom.2018.04.002. URL <https://linkinghub.elsevier.com/retrieve/pii/S1574002118300121>.
- Assenza, T., Heemeijer, P., Hommes, C. H., and Massaro, D. (2018). Managing Self-organization of Expectations through Monetary Policy : a Macro Experiment. *TSE Working Paper*.
- Benhabib, J., Duffy, J., and Nagel, R. (2019). De-framing rules to De-anchor Beliefs in Beauty Contest Experiments: Keynesian level-k vs. Keynesian Sentiments. *mimeo*.
- Benhabib, J., Schmitt-Grohé, S., and Uribe, M. (2001). The Perils of Taylor Rules. *Journal of Economic Theory*, 96: 40–69. ISSN 00220531. DOI 10.1006/jeth.1999.2585. URL <http://linkinghub.elsevier.com/retrieve/pii/S0022053199925851>.
- Benhabib, J., Wang, P., and Wen, Y. (2015). Sentiments and Aggregate Demand Fluctuations. *Econometrica*, 83(2): 549–585. ISSN 0012-9682. DOI 10.3982/ecta11085. URL <https://doi.org/10.3982/ECTA11085>.
- Bernheim, B. D. (1984). Rationalizable Strategic Behavior. *Econometrica*, 52(4): 1007. ISSN 00129682. DOI 10.2307/1911196. URL <https://www.jstor.org/stable/1911197?origin=crossrefhttps://www.jstor.org/stable/1911196?origin=crossref>.

- Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic Expectations and Credit Cycles. *The Journal of Finance*, 73(1): 199–227. ISSN 00221082. DOI 10.1111/jofi.12586. URL <http://doi.wiley.com/10.1111/jofi.12586>.
- Bosch-Domenech, A., Montalvo, J. G., Rosemarie Nagel, and Albert Satorra (2002). One, Two(Three), Infinity,... : Newspaper and Lab Beauty-Contest Experiments. *The American Economic Review*, 92(5): 1687–1701. URL <https://www.jstor.org/stable/3083273>.
- Bosch-Rosa, C., and Meissner, T. (2019). The One Player Guessing Game: A diagnosis on the relationship between equilibrium play, beliefs, and best responses. *Working Paper*.
- Branch, W. A., and McGough, B. (2009). A New Keynesian model with heterogeneous expectations. *Journal of Economic Dynamics and Control*, 33(5): 1036–1051. URL <https://doi.org/10.1016/j.jedc.2008.11.007>.
- Brock, W. A., and Hommes, C. H. (1997). A Rational Route to Randomness. *Econometrica*, 65(5): 1059–1095. ISSN 00129682. DOI 10.2307/2171879. URL <https://www.jstor.org/stable/2171879>.
- Bühren, C., and Frank, B. (2012). Chess Players Performance Beyond 64 Squares: A Case Study on the Limitations of Cognitive Abilities Transfer. *Talent Development and Excellence*, 4(2): 157–169.
- Bühren, C., and Nagel, R. (2019). Strategic Nudging in the Beauty Contest Game. *mimeo2*.
- Bulow, J. I., Geanakoplos, J. D., and Klemperer, P. D. (1985). Multimarket Oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy*, 93(3): 488–511. URL <https://www.jstor.org/stable/1832005>.
- Camerer, C. F., and Fehr, E. (2006). When Does "Economic Man" Dominate Social Behavior? *Science*, 311(5757): 47–52. ISSN 0036-8075. DOI 10.1126/science.1110600. URL <http://www.sciencemag.org/cgi/content/abstract/sci;311/5757/47http://www.sciencemag.org/cgi/doi/10.1126/science.1110600>.
- Camerer, C. F., Ho, T.-h., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119(3): 861–898. URL <https://doi.org/10.1162/0033553041502225>.
- Cass, D. (1965). Optimum Growth in an Aggregative Model of Capital Accumulation. *The Review of Economic Studies*, 32(3): 233. ISSN 00346527. DOI 10.2307/2295827. URL <https://academic.oup.com/restud/article-lookup/doi/10.2307/2295827>.
- Chou, E., McConnell, M., Nagel, R., and Plott, C. R. (2009). The control of game form recognition in experiments: Understanding dominant strategy failures in a simple two person "guessing" game. *Experimental Economics*, 12(2): 159–179. ISSN 13864157. DOI 10.1007/s10683-008-9206-4. URL <https://doi.org/10.1007/s10683-008-9206-4>.
- Coibion, O., Gorodnichenko, Y., Kumar, S., and Ryngaert, J. (2018). Do You Know That I Know That You Know...? Higher-Order Beliefs in Survey Data. *NBER Working Paper Series*. URL <http://www.nber.org/papers/w24987>.

- Cooper, K. B., Schneider, H. S., and Waldman, M. (2017). Limited rationality and the strategic environment: Further theory and experimental evidence. *Games and Economic Behavior*, 106: 188–208. ISSN 08998256. DOI 10.1016/j.geb.2017.09.014. URL <https://linkinghub.elsevier.com/retrieve/pii/S0899825617301690>.
- Coricelli, G., and Nagel, R. (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 106(23): 9163–9168. ISSN 0027-8424. DOI 10.1073/pnas.0807721106. URL <http://www.pnas.org/cgi/doi/10.1073/pnas.0807721106>.
- Cornand, C., and Heinemann, F. (2019). Experiments in macroeconomics: methods and applications. In A. Schram, and A. Ule (Eds.), *Handbook of Research Methods and Applications in Experimental Economics*, chapter 14, pages 269–294. Cheltenham, UK; Northampton, MA; USA: Edward Elgar Publishing. DOI 10.4337/9781788110563.00024. URL <https://www.elgaronline.com/view/edcoll/9781788110556/9781788110556.00024.xml>.
- Cornes, R., and Hartley, R. (2012). Fully aggregative games. *Economics Letters*, 116(3): 631–633. ISSN 01651765. DOI 10.1016/j.econlet.2012.06.024. URL <http://dx.doi.org/10.1016/j.econlet.2012.06.024>.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1): 5–62. ISSN 00220515. DOI 10.1257/jel.51.1.5. URL <https://www.aeaweb.org/articles?id=10.1257/jel.51.1.5>.
- Duffy, J. (2014). Macroeconomics: a survey of laboratory research. *Handbook of experimental economics 2*, pages 1–97. DOI 10.2514/1.36125. URL <http://econ2.econ.iastate.edu/tesfatsi/LaboratoryMacroeconomics.JDuffy2008.pdf>.
- Duffy, J., and Nagel, R. (1997). On the Robustness of Behaviour in Experimental 'Beauty Contest' Games. *The Economic Journal*, 107(445): 1684–1700. ISSN 0013-0133. DOI 10.1111/j.1468-0297.1997.tb00075.x. URL <http://www.jstor.org/stable/2957901>.
- Evans, G. W., and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press. ISBN 9781400824267. DOI 10.1515/9781400824267.
- Evans, G. W., and Ramey, G. (1992). Expectation Calculation and Macroeconomic Dynamics. *American Economic Review*, 82(1): 207–224. URL <https://www.jstor.org/stable/2117611>.
- Farhi, E., and Werning, I. (2019). Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review*, 109(forthcoming): 3887–3928. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20171400>.
- Gabaix, X. (2014). A Sparsity-Based Model of Bounded Rationality. *The Quarterly Journal of Economics*, 129(4): 1661–1710. ISSN 0033-5533. DOI 10.1093/qje/qju024. URL <https://doi.org/10.1093/qje/qju024https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qju024>.
- Gabaix, X. (2019). A Behavioral New Keynesian Model. *Working Paper*, pages 1–58. DOI 10.3386/w22954. URL <http://www.nber.org/papers/w22954.pdf>.

- Gächter, S., Herrmann, B., and Thöni, C. (2004). Trust, voluntary cooperation, and socio-economic background: Survey and experimental evidence. *Journal of Economic Behavior and Organization*, 55(4 SPEC.ISS.): 505–531. ISSN 01672681. DOI 10.1016/j.jebo.2003.11.006. URL <https://doi.org/10.1016/j.jebo.2003.11.006>.
- Galí, J. (2008). *Monetary Policy, Inflation and the Business Cycle*. Princeton, NJ: Princeton University Press. ISBN 978-0-691-13316-4.
- 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. ISSN 0002-8282. DOI 10.1257/aer.20170110. URL <https://pubs.aeaweb.org/doi/10.1257/aer.20170110>.
- Georganas, S., Healy, P. J., and Weber, R. A. (2015). On the persistence of strategic sophistication. *Journal of Economic Theory*, 159: 369–400. ISSN 00220531. DOI 10.1016/j.jet.2015.07.012. URL <https://linkinghub.elsevier.com/retrieve/pii/S0022053115001350>.
- Grosskopf, B., and Nagel, R. (2008). The two-person beauty contest. *Games and Economic Behavior*, 62(1): 93–99. ISSN 08998256. DOI 10.1016/j.geb.2007.03.004. URL <https://doi.org/10.1016/j.geb.2007.03.004>.
- Guesnerie, R. (1992). An exploration of the eductive justifications of the rational-expectations hypothesis. *American Economic Review*, 82(5): 1254–1278. URL <http://www.jstor.org/stable/2117477>.
- Guesnerie, R. (2009). Macroeconomic and Monetary Policies from the Eductive Viewpoint. In K. Schmidt-Hebbel, and C. E. Walsh (Eds.), *Monetary Policy Under Uncertainty and Learning*, pages 171–202. Santiago de Chile: Central Bank of Chile.
- Hampton, A. N., Bossaerts, P., and O’Doherty, J. P. (2008). Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 105(18): 6741–6746. ISSN 00278424. DOI 10.1073/pnas.0711099105. URL <https://pubmed.ncbi.nlm.nih.gov/18427116/>.
- Haruvy, E., and Stahl, D. O. (2012). Between-game rule learning in dissimilar symmetric normal-form games. *Games and Economic Behavior*, 74(1): 208–221. ISSN 08998256. DOI 10.1016/j.geb.2011.06.001. URL <http://dx.doi.org/10.1016/j.geb.2011.06.001https://linkinghub.elsevier.com/retrieve/pii/S0899825611000959>.
- Heemeijer, P., Hommes, C., Sonnemans, J., and Tuinstra, J. (2009). Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33: 1052–1072. ISSN 01651889. DOI 10.1016/j.jedc.2008.09.009. URL <https://doi.org/10.1016/j.jedc.2008.09.009>.
- Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, 72(5): 1583–1599. ISSN 00129682. DOI 10.1111/j.1468-0262.2004.00544.x. URL <https://www.jstor.org/stable/3598899>.

- Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests". *American Economic Review*, 88(4): 947–969. URL <http://www.jstor.org/stable/1170131>.
- Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control*, 35(1): 1–24. ISSN 01651889. DOI 10.1016/j.jedc.2010.10.003. URL <http://dx.doi.org/10.1016/j.jedc.2010.10.003>.
- Hommes, C. (2013). *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge: Cambridge University Press.
- Hommes, C. (2020). Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach. *Journal of Economic Literature*. DOI 10.1257/jel.20191434. URL <https://www.aeaweb.org/articles?id=10.1257/jel.20191434&&from=f>.
- Hommes, C., and Lustenhouwer, J. (2019). Inflation targeting and liquidity traps under endogenous credibility. *Journal of Monetary Economics*, 107: 48–62. ISSN 03043932. DOI 10.1016/j.jmoneco.2019.01.027. URL <https://linkinghub.elsevier.com/retrieve/pii/S0304393219300170>.
- Iovino, L., and Sergeyev, D. (2018). Central bank balance sheet policies without rational expectations. *Working Paper*, pages 1–63.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*. London: Macmillan.
- Kocher, M., Sutter, M., and Gueth, W. (2002). Experimental 'beauty contests' with homogeneous and heterogeneous players and with interior and boundary equilibria. *Economics Letters*, 74: 219–228. URL [https://doi.org/10.1016/S0165-1765\(01\)00544-4](https://doi.org/10.1016/S0165-1765(01)00544-4).
- Koopmans, T. (1965). On the Concept of Optimal Economic Growth. In *The Econometric Approach to Development Planning*. Amsterdam: North-Holland.
- Lei, V., and Noussair, C. N. (2002). An Experimental Test of an Optimal Growth Model. *American Economic Review*, 92(3): 411–433. ISSN 0002-8282. DOI 10.1257/00028280260136246. URL <http://www.ingentaconnect.com/content/aea/aer/2002/00000092/00000003/art00006><http://pubs.aeaweb.org/doi/abs/10.1257/00028280260136246>.
- Levine, S. S., Bernard, M., and Nagel, R. (2017). Strategic Intelligence: The Cognitive Capability to Anticipate Competitor Behavior. *Strategic Management Journal*, 38(12): 2390–2423. ISSN 01432095. DOI 10.1002/smj.2660. URL <http://doi.wiley.com/10.1002/smj.2660>.
- Li, S. (2017). Obviously Strategy-Proof Mechanisms. *American Economic Review*, 107(11): 3257–3287. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20160425>.
- Maćkowiak, B., and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3): 769–803. ISSN 00028282. DOI 10.1257/aer.99.3.769.
- Mauersberger, F. (2016). Monetary Policy Rules in a Non-Rational World : A Macroeconomic Experiment. *Working Paper, Columbia Economics Discussion Papers*.

- Mauersberger, F., and Nagel, R. (2018). Levels of Reasoning in Keynesian Beauty Contests: A Generative Framework. In C. H. Hommes, and B. LeBaron (Eds.), *Handbook of Computational Economics 2: Volume 4*, pages 541–634. Amsterdam and Oxford: North-Holland. DOI 10.1016/bs.hescom.2018.05.002. URL <https://linkinghub.elsevier.com/retrieve/pii/S1574002118300145>.
- Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature*, 52(1): 124–188. ISSN 0022-0515. DOI 10.1257/jel.52.1.124. URL <https://www.aeaweb.org/articles?id=10.1257/jel.52.1.124>.
- Mirdamadi, M., and Petersen, L. (2018). Macroeconomic Literacy and Expectations. *Working Paper*.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, 29(3): 315–335. URL <http://www.jstor.org/stable/1909635>.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5): 1313–1326. URL <https://www.jstor.org/stable/2950991>.
- Nagel, R., Brovelli, A., Heinemann, F., and Coricelli, G. (2018). Neural mechanisms mediating degrees of strategic uncertainty. *Social Cognitive and Affective Neuroscience*, 13(1): 52–62. ISSN 17495024. DOI 10.1093/scan/nsx131. URL <https://pubmed.ncbi.nlm.nih.gov/29228378/>.
- Nagel, R., Bühren, C., and Frank, B. (2016). Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game. *Mathematical Social Sciences*, pages 1–17. ISSN 01654896. DOI 10.1016/j.mathsocsci.2016.09.001. URL <http://dx.doi.org/10.1016/j.mathsocsci.2016.09.001>.
- Pearce, D. G. (1984). Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica*, 52(4): 1029. ISSN 00129682. DOI 10.2307/1911197. URL <https://www.jstor.org/stable/1911197?origin=crossref>.
- Pfajfar, D., and Žakelj, B. (2014). Experimental evidence on inflation expectation formation. *Journal of Economic Dynamics and Control*, 44: 147–168. ISSN 01651889. DOI 10.1016/j.jedc.2014.04.012. URL <http://linkinghub.elsevier.com/retrieve/pii/S0165188914001079>.
- Ramsey, F. P. (1928). A Mathematical Theory of Saving. *The Economic Journal*, 38(152): 543. ISSN 00130133. DOI 10.2307/2224098. URL <https://www.jstor.org/stable/2224098?origin=crossref>.
- Romer, D. (2012). *Advanced Macroeconomics*. New York: McGraw-Hill, 4th edition. ISBN 0073511374.
- Selten, R. (1970). *Preispolitik der Mehrproduktenunternehmung in der statischen Theorie*, volume 16 of *Ökonometrie und Unternehmensforschung / Econometrics and Operations Research*. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-48889-4. DOI 10.1007/978-3-642-48888-7. URL <http://link.springer.com/10.1007/978-3-642-48888-7>.

- Shiller, R. J. (2015). *Irrational Exuberance*, volume 53. Princeton, NJ: Princeton University Press, 3rd edition. ISBN 9781400865536. DOI 10.2307/j.ctt1287kz5. URL <http://arxiv.org/abs/1011.1669><http://dx.doi.org/10.1088/1751-8113/44/8/085201><http://www.jstor.org/stable/10.2307/j.ctt1287kz5>.
- Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3): 665–690. ISSN 03043932. DOI 10.1016/S0304-3932(03)00029-1. URL [https://doi.org/10.1016/S0304-3932\(03\)00029-1](https://doi.org/10.1016/S0304-3932(03)00029-1).
- Solow, R. M. (1956). A Contribution to the Theory of Economic Growth. *The Quarterly Journal of Economics*, 70(1): 65. ISSN 00335533. DOI 10.2307/1884513. URL <https://academic.oup.com/qje/article-lookup/doi/10.2307/1884513>.
- Sonnemans, J., and Tuinstra, J. (2010). Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology*, 31(6): 964–984. ISSN 0167-4870. DOI 10.1016/j.joep.2010.08.005. URL <http://dx.doi.org/10.1016/j.joep.2010.08.005>.
- Stahl, D. O. (1993). Evolution of Smartn Players. *Games and Economic Behavior*, 5(4): 604–617. ISSN 08998256. DOI 10.1006/game.1993.1033. URL <http://www.sciencedirect.com/science/article/pii/S089982568371033X><http://linkinghub.elsevier.com/retrieve/pii/S089982568371033X>.
- Stahl, D. O. (1996). Boundedly Rational Rule Learning in a Guessing Game. *Games and Economic Behavior*, 16: 303–330. ISSN 08998256. DOI 10.1006/game.1996.0088. URL <https://doi.org/10.1006/game.1996.0088>.
- Stahl, D. O., and Wilson, P. W. (1994). Experimental evidence on players’ models of other players. *Journal of Economic Behavior and Organization*, 25(3): 309–327. ISSN 01672681. DOI 10.1016/0167-2681(94)90103-1. URL [https://doi.org/10.1016/0167-2681\(94\)90103-1](https://doi.org/10.1016/0167-2681(94)90103-1).
- Stahl, D. O., and Wilson, P. W. (1995). On Players’ Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1): 218–254. ISSN 08998256. DOI 10.1006/game.1995.1031. URL <https://linkinghub.elsevier.com/retrieve/pii/S0899825685710317>.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure. *American Economic Review*, 80(1): 234–248. ISSN 00028282. DOI <http://www.aeaweb.org/aer/>. URL <http://www.jstor.org/stable/2006745>.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games. *The Quarterly Journal of Economics*, 106(3): 885–910. ISSN 0033-5533. DOI 10.2307/2937932. URL <https://academic.oup.com/qje/article-lookup/doi/10.2307/2937932>.
- Walsh, C. E. (2010). *Monetary Theory and Policy*. Cambridge: MIT Press, 3rd edition.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

Woodford, M. (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Review of Economics*, 5(1): 303–346. ISSN 1941-1383. DOI 10.1146/annurev-economics-080511-110857. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev-economics-080511-110857> <http://www.nber.org/papers/w19368>.

Woodford, M. (2018). Monetary Policy Analysis When Planning Horizons Are Finite. *Working Paper*.