

Categorization in Games: A Bias-Variance Perspective*

Philippe Jehiel[†] Erik Mohlin[‡]

March 17, 2023

Abstract

We develop a framework for categorization in games, applicable both to multi-stage games of complete information and static games of incomplete information. Players use categories to form coarse beliefs about their opponents' behavior. Players best-respond given these beliefs, as in analogy-based expectations equilibria. Categories are related to previously used strategies via the requirements that categories contain a sufficient amount of observations and exhibit sufficient within-category similarity, in line with the bias-variance trade-off. When applied to classic games including the chainstore game and adverse selection games our framework yields less unintuitive predictions than those arising with standard solution concepts.

Keywords: Bounded rationality; Categorization; Bias-variance trade-off; Adverse selection; Chainstore paradox.

JEL codes: C70, C73, D82, D83, D91.

*This paper has benefited from comments by Tore Ellingsen, Drew Fudenberg, Topi Miettinen, Alexandros Rigos, and Larry Samuelson. We also thank audiences at Lund University (Arne Ryde Workshop on Attention in Decision Making) and Bar-Ilan University (Learning Evolution Games 2019) for comments. Erik Mohlin is grateful for financial support from the Swedish Research Council (grant no. 2015-01751 and 2019-02612) and the Knut and Alice Wallenberg Foundation (Wallenberg Academy Fellowship 2016-0156). Philippe Jehiel thanks the European Research Council (grant no. 742816) for funding.

[†]Paris School of Economics and University College London. Address: PSE, 48 boulevard Jourdan, 75014 Paris, France. E-mail: jehiel@enpc.fr.

[‡]Lund University and the Institute for Futures Studies (Stockholm). Address: Lund University Department of Economics, Tycho Brahes väg 1, 220 07 Lund, Sweden. E-mail: erik.mohlin@nek.lu.se.

1 Introduction

Human decision-makers need to make simplifications in order to navigate through social reality. We need to divide the complex web of interactions into manageable pieces to evaluate different courses of action. We need to extrapolate from past interactions to be able to predict what others will do. Categories serve these functions (Anderson 1991, Laurence and Margolis 1999, Gärdenfors 2000, Murphy 2002, Xu 2007). A categorization bundles distinct objects or situations into groups or categories, whose members are viewed as sufficiently similar to warrant a similar treatment. As a result, categorical reasoning facilitates prediction: when a situation is classified as belonging to a category then by virtue of its similarity with other members of the category we expect similar behavior.

From the perspective of statistics and machine learning, categorizations should satisfy some properties to address the bias-variance trade-off (e.g. Geman et al. 1992). On the one hand, if categories are too coarse, bundling together situations that are too dissimilar, the resulting estimates are likely to be too biased. On the other hand, if categories are too narrow, bundling together too few data points, the resulting estimates will be unreliable, as they are plagued by high variance. Gigerenzer and Brighton (2009) discuss how simple heuristics typically used by humans can be viewed as devices inducing some bias in order to reduce variance. Mohlin (2014) derives properties of categorizations that solve the bias-variance trade-off optimally for the purpose of making predictions.

In economics, a growing literature has introduced categorical thinking into game theory (Samuelson 2001, Jehiel 2005, Jehiel and Samet 2007, Jehiel and Koessler 2008, Azrieli 2009, Mengel 2012, Arad and Rubinstein 2019).¹ A significant part of this literature has worked with exogenously given categories. While in some interactions it may be reasonable to assume that the categorizations are given exogenously, either by the framing of the game, or by players' culture or previous personal experience, in other interactions, it seems more appropriate to view the categories as being formed within the learning environment.²

Our starting point is the analogy-based expectation equilibrium (Jehiel 2005, Jehiel and Koessler 2008) in which categories (analogy classes) are used to form predictions about opponents' play. We endogenize the analogy partitions relying on principles that are inspired by the bias-variance trade-off. Specifically, different cohorts of players are called to move in different situations (nodes or states depending on the application) in a game. After the play of a given game, new players receive some feedback about it

¹Earlier contributions include Dow (1991) and Rubinstein (1998). Fryer and Jackson (2008) examine a non-strategic context.

²Some of the cited papers endogenize the categories assuming there is a fixed cost to adding a category and considering the best possible categorization in terms of the induced payoff consequence minus the cost associated with the categorization. Such approaches require a large amount of rationality (arguably at least as large as in standard economic models) while our interest lies in situations in which the rationality of subjects is more limited.

that consists in the disclosure of behaviors in a subset of situations. For example, in extensive form games, the feedback will typically consist of the on-path behaviors, and in Bayesian games we will consider cases in which behaviors are disclosed if some event (such as trade) occurs. We assume that in each situation, a player may pick a non-intended action with some exogenous probability ε , as in the trembling-hand formulation (Selten, 1975). This could be due to mistakes or experimentation, and it will be used to generate some observations for all situations.

When a new cohort arrives, the corresponding players categorize the situations in which their opponents have to make a move, using the available data from the games played by the previous cohort. Players are endowed with similarity functions, representing their perception of how similar various situations are to each other. They form categories by bundling together situations perceived to be as similar as possible, while respecting the desiderata that each category should contain enough data points, in line with the bias-variance trade-off. We formalize this by imposing that each analogy class should have a mass of observations above a threshold κ , unless doing so creates too high within-category dissimilarity. While our approach allows for any specification of κ and ε , we focus on the case in which κ and ε are small and vanish at such a rate that ε is asymptotically much smaller than κ , implying that on-path situations can be distinguished perfectly but off-path situations have to be bundled (according to their similarity). Both the feedback structure and the similarity functions are primitives of our model.

In a given cohort, players stick to their initial choices of categorizations (analogy partitions) throughout their active period. Within the cohort, players are randomly matched to play the same game many times, leading us to assume that behaviors converge to an analogy-based expectation equilibrium given their maintained categorizations (analogy partitions). The equilibrium behavior induces new data, which are used by the next cohort of players, thereby generating a dynamic system in which analogy partitions are structured by the data generated by the previous cohort, and new data are generated by the analogy-based expectation equilibrium induced by the analogy partitions chosen by the current cohort. To simplify the analysis of the learning dynamics, we assume that each cohort consists of a continuum of players (thereby avoiding the stochastic nature of a dynamics based on finitely many players).

When the learning process admits steady states, we will refer to them as *categorization equilibria*. We discuss their properties in some standard economic applications, contrasting the predictions so obtained with those of standard equilibrium approaches. In case no categorization equilibrium exists, we describe the learning dynamics explicitly, and note that, in our application to adverse selection games, it may take the form of cycling.

Our main results are as follows. We first consider extensive form games of complete information for which the feedback consists of the on-path behaviors. Nodes which are reached without trembles are treated as singleton analogy classes (because they are ob-

served sufficiently often). In this case, when a categorization equilibrium exists, it must be a self-confirming equilibrium (Fudenberg and Levine 1998, chapter 6) since predictions on-path have to be correct. But, categorization equilibria typically impose extra constraints for beliefs off-path as compared with self-confirming equilibria, and these constraints are shaped by the chosen similarity functions (and how the trembling probability ε compares to the minimum size κ requirement).

In the chainstore game (Selten 1978), we assume (for both the monopolist and the challengers) that histories in which there was some entry that was not immediately followed by a fight are treated as very dissimilar from other histories (in which either the challengers never entered or when there was entry it was always followed by a fight).³ We establish the existence of a categorization equilibrium with no entry except in the last few periods, and when there is entry toward the end of the interaction, the monopolist does not fight.

We next discuss how categorization equilibria look like in several well known alternative games discussed in relation to the finite horizon paradox. For example, in the Centipede game (Rosenthal 1981) all categorization equilibria feature immediate Take, as in the Nash equilibrium. We also briefly consider public good games with or without a punishment stage, and we assume there that histories are distinguished according to whether agents contributed and, in the case with a punishment stage, whether non-contributors were punished in the past. We note that full contributions can be achieved in a categorization equilibrium when agents have the opportunity to punish their peers, but not otherwise. This application is developed in the Online Appendix.

We then apply our approach to an adverse selection game of the Akerlof type, modeled as a Bayesian game between an informed seller and an uninformed buyer who values the good more than the seller. The seller submits an ask price and the buyer submits a bid price. The transaction takes place only if the bid price exceeds the ask price, and when there is a transaction, the price is assumed to be the bid price (so that the seller has a weakly dominant strategy consisting in submitting an ask price equal to her valuation). We assume that feedback always includes the ask and bid prices but that the quality is disclosed to outsiders only when there is a transaction (as in Esponda 2008). Regarding similarity, we make the (natural) assumption that qualities (seller types) are more similar to one another when they are less far apart (in the euclidean distance sense). We show that the learning dynamics leads to cycles with intended bid prices always lying above the Nash equilibrium price.

To get a sense of the dynamics, suppose that in a given cohort, buyers set a bid price p^* most of the time (and tremble, picking other bid prices, with a small probability ε).

³We believe there is some salient difference between these two sets of histories justifying our assumption. Of course, other assumptions on the similarity functions are possible and our general framework allows us to deal with any such assumptions if desired.

In the next cohort, qualities resulting in an ask price below p^* are treated as singleton analogy classes (because they are on-path and the corresponding data are abundant). But, higher qualities must be bundled into coarse analogy classes consisting of intervals of qualities (because the corresponding data are scarce as they are observed only following trembles).

Starting from a price p^* that corresponds to the Bayes–Nash equilibrium, we establish that the (intended, non-trembling) bid price has to be strictly larger in the next generation. The price chosen by non-experimenting buyers may then increase over the next few generations, but at some point when it gets too high, increasing the bid price will not look profitable and at this point buyers will quote the Nash equilibrium bid price instead. The bid prices will cycle from then on, always being weakly above the price arising in the full rationality benchmark. The key to understanding this is to realize that if the bid price becomes very large, then the next generation will have all the needed data to compute the correct best response, which is the Nash equilibrium price. However, when this bid price is low (as in the Nash equilibrium), increasing the price will seem attractive. Specifically, it will seem to the buyers that raising the bid price slightly above p^* has a slight effect on the price to be paid, but increases average quality of traded goods more than slightly. Thus, it will seem as if the utility function of the buyers had a local optimum above p^* .⁴

Our paper can be related to several strands of the literature. First, since we assume that the threshold κ and the trembling probability ε are such that on-path nodes are treated as singleton analogy classes in extensive form games, our categorization equilibria in such games can be viewed as offering a selection device for self-confirming equilibria (Fudenberg and Levine 1993, 1998). Fudenberg and Levine (2006) propose a different selection referred to as subgame-confirmed equilibria in which strategy profiles are Nash equilibria on-path and self-confirming equilibria one step away from the path. They have provided a learning foundation for it in a class of extensive form games (in which players move once) when players are patient and experiment optimally.⁵ Our approach differs from Fudenberg and Levine (2006) in several respects. Most importantly, in our case, when beliefs are incorrect, they are related to the true behaviors via a categorization which is itself structured by the similarity functions, the experimentation probability and the minimum size requirement imposed on categories.

Second, a number of approaches have been proposed to avoid the unintuitive predictions obtained in finite horizon interactions, including Kreps et al. (1982)’s crazy type approach and Neyman (1985)’s finite automaton approach (see also Rubinstein 1998).

⁴Whether this local optimum is perceived to be a global optimum by the buyers depends on the level of the current price p^* . For p^* below or at the the Nash equilibrium price, it will in fact seem to the buyers that the best response is above p^* . For p^* sufficiently above the Nash equilibrium price it will instead seem to the buyers that the best response is at the Nash equilibrium price.

⁵See also Kalai and Neme (1992) who have proposed another notion of equilibrium for extensive form games in which strategy profiles are Nash equilibria up to p steps away from the path.

We note that these approaches avoid the classical predictions in all versions of the finite horizon paradox (this applies also to Jehiel’s (2005) ABEE-approach). By contrast, in our setting, we get around the unintuitive predictions in some versions of the finite horizon paradox games (namely those involving more complex interactions such as the chain store game or the public good game with punishment stage) but not in others (e.g. the centipede game or the public good game without punishment stage). This is in line with the experimental findings of Fehr and Gächter’s (2002) who have shown that contributions in public good games are higher and do not diminish over time when there are punishment opportunities.⁶

Third, a number of approaches have revisited the classic adverse selection games introduced by Akerlof (1970) and studied whether relaxations of the buyer’s rationality could generate more trading activity. These include Eyster and Rabin (2005)’s cursed equilibrium, Jehiel and Koessler (2008)’s analogy-based expectation equilibrium, and Esponda (2008)’s behavioral equilibrium.⁷ Our modelling of such interactions is inspired by Esponda (2008), in particular with respect to the feedback function. But, our derivation of categorization-based expectations based on that feedback is different, leading to more trade than in the rational case (in contrast to Esponda’s finding), as well as cycling (which has no counterpart in the other approaches).⁸ Our predictions for this type of interactions are broadly in line with the experimental findings reported in Fudenberg and Peysakhovich (2016). They observe more trade than predicted by the Nash equilibrium and they suggest comparative statics with respect to the difference of valuation between the seller and the buyer that agree with our predictions.

Fourth, our paper can be related to a growing literature on misspecifications in games, which, in addition to the already mentioned cursed equilibrium (Eyster and Rabin, 2005) and analogy-based expectation equilibrium (Jehiel, 2005), include the Berk-Nash equilibrium (Esponda and Pouzo, 2016) and the Bayesian Network Equilibrium (Spiegler, 2016). Some papers have suggested endogenizing the misspecifications based on evolutionary arguments (in particular He and Libgober 2022, Fudenberg and Lanzani 2020, and Heller and Winter 2020), but to the best of our knowledge, none of these papers have developed an approach based on the bias-variance trade-off to endogenize misspecifications.

⁶There are other plausible explanations for this phenomenon, e.g. preferences for reciprocity.

⁷See also Miettinen (2009) on the relationship between these various approaches.

⁸A few recent papers identify cycles of beliefs in the context of misspecified models. In Esponda, Pouzo and Yamamoto (2021) and Bohren and Hauser (2021) (see also Nyarko 1991), the evidence accumulated while taking a particular action may push beliefs in a direction that makes another action seem optimal, and once this new action is taken the data that are being generated induce a belief that makes the previous action seem optimal again. In Fudenberg Romanyuk, and Strack (2017) cycles may arise from the fact that the learner never ceases to perceive an information value of experimenting with another action. None of these papers feature endogenous categorizations.

2 Framework

We present our approach within a unified setup covering both multi-stage games of complete information and (static) Bayesian games. Specifically, we consider games with two players $i \in I = \{1, 2\}$ such that player $i \in I$ faces various possible situations referred to as $x_i \in X_i$, and in situation x_i player i has to choose an action $a_i \in A_i(x_i)$. Extension to more players is straightforward. In an extensive-form game with complete information, \mathcal{X}_i will represent the nodes at which player i must move. In a Bayesian game, \mathcal{X}_i will represent the set of types of player i . In the former case, the profile of actions chosen by the two players at the various nodes determines which nodes are visited. In the latter case, nature chooses the profile of types according to some probability assumed to be known by both players. For simplicity, we consider the finite case in which the set of situations and the sets of actions are all finite. When we consider the trade application, we will consider straightforward extensions of the definitions to the case of a continuum of actions and situations.

A strategy for player i is defined by $\sigma_i = (\sigma_i(x_i))_{x_i \in \mathcal{X}_i}$ where $\sigma_i(x_i) \in \Delta A_i(x_i)$ describes the probability distribution over possible actions chosen by player i at x_i . A play of the game is described by the set of situations that occurred and the actions taken in those situations, denoted

$$(\hat{a}, \hat{x}) = \{(\hat{a}_i, \hat{x}_i)_{i \in I} : \hat{x}_i \text{ occurred and } i \text{ chose } \hat{a}_i \text{ at } \hat{x}_i\}$$

Regarding the feedback, we assume that after the play of a game only a subset of (\hat{a}, \hat{x}) is disclosed to outsiders (which will be used by new players to form expectations). We refer to such a disclosure as the feedback given the play and denote it by $\phi(\hat{a}, \hat{x})$. In extensive-form games, we assume that only the actions on the equilibrium path are observed (as is commonly assumed, see Fudenberg and Levine, 1998). In Bayesian games, we will use this formulation to accommodate applications like trades in which the actions (bargaining offers) and types (determining the quality of the good) would be disclosed only when the transaction takes place (as in Esponda, 2008).

2.1 Analogy-Based Expectations

Player i categorizes \mathcal{X}_j (the set of player j 's situations) into analogy classes $\mathcal{C}_i^1, \dots, \mathcal{C}_i^K$ that constitute a partition $\mathcal{C}_i = \{\mathcal{C}_i^1, \dots, \mathcal{C}_i^K\}$ of \mathcal{X}_j . An analogy class $\mathcal{C}_i^k \in \mathcal{C}_i$ of player i satisfies the requirement that if x_j and x'_j belong to the same analogy class \mathcal{C}_i^k , then the action spaces of player j at x_j and x'_j are the same. We let $\beta_i(\mathcal{C}_i^k)$ denote the analogy-based expectation of player i about the play of player j in \mathcal{C}_i^k . It is a probability distribution over the action space of player j in \mathcal{C}_i^k meant to capture how player i views

player j 's representative behavior in \mathcal{C}_i^k .⁹ For every $x_j \in \mathcal{X}_j$, we let $\mathcal{C}_i(x_j)$ be the unique analogy class \mathcal{C}_i^k to which x_j belongs. We refer to $\beta_i = (\beta_i(\mathcal{C}_i^k))_{k=1}^K$ as the analogy-based expectation of player i .

Given β_i , player i will expect player j to behave according to the strategy defined by $\sigma_j^{\beta_i} = \left(\sigma_j^{\beta_i}(x_j) \right)_{x_j \in \mathcal{X}_j}$, with $\sigma_j^{\beta_i}(x_j) = \beta_i(\mathcal{C}_i(x_j))$. That is, player i expects player j in situation x_j to behave according to the representative behavior in the analogy class $\mathcal{C}_i(x_j)$ to which x_j belongs as defined by $\beta_i(\mathcal{C}_i(x_j))$.¹⁰

We have in mind that most of the time player i will play a best-response to $\sigma_j^{\beta_i}$, and the rest of the time player i will experiment or tremble and try all available actions. We also have in mind that experimentations/trembles occur independently at the various x_i . In other words, our treatment of experimentation is similar to the extensive-form version of the trembling-hand equilibrium (Selten, 1975). Formally,

Definition 1 σ_i is an ε_i -perturbed best-response to β_i if σ_i is a best-response to $\sigma_j^{\beta_i}$ subject to the constraint that at every x_i , $\sigma_i(x_i)$ assigns a probability no less than ε_i to every action at x_i and the probability distributions $\sigma_i(x_i)$ are independent across the various x_i .

Remark 1 (a) In the definition of ε_i -perturbed best-response, we implicitly assume that the probability of experimentation is the same for all actions at x_i , and the same at all x_i . We could obviously extend this to allow for more general experimentation strategies, but this would bring no additional insight. (b) The best-response is implicitly defined at the ex ante stage here, but given that we consider games with perfect recall and all situations are reached with positive probability (due to trembling), the same choice of strategy would arise had we required an interim or sequential notion of best-response at every x_i .

In general, we allow for the possibility that players i and j have different probabilities of experimentation, and we denote the profile of experimentation probabilities by $\varepsilon = (\varepsilon_i, \varepsilon_j)$. This will allow us to accommodate applications in which we believe one player is much less likely to experiment than the other player because say the former but not the latter has a dominant strategy. The situations that are reached with positive probability in the absence of trembles/experimentation ($\varepsilon = 0$) will be referred to as *on-path situations*. The remaining situations, which are reached with positive probability only when there are trembles/experimentation ($\varepsilon_i, \varepsilon_j > 0$) are *off-path situations*. This distinction will play a role when we endogenize the analogy partitions.

⁹In the case of n players \mathcal{C}^i partitions $\times_{j \neq i} X_j$, with the requirement that if x_j and x'_l (possibly belonging to different players) belong to the same analogy class $\mathcal{C}_k^i \in \mathcal{C}^i$, then the action spaces of player j at x_j and player l at x'_l are the same. Furthermore, $\beta_i(\mathcal{C}_k^i)$ denotes the analogy-based expectation of player i about the play of all players acting at the situations belonging to \mathcal{C}_k^i .

¹⁰The rationale for this is that this is the simplest representation of player j 's strategy consistent with β_i (see Jehiel 2022 for elaboration).

The analogy-based expectations are required to be related to the strategy profile and the feedback structure through a consistency requirement. This formalizes the idea that given the analogy partitions and the feedback structure, the play has reached a steady state.

A strategy profile σ together with a feedback structure ϕ and potential chance moves, induces a probability $\mu^\sigma(a_j, x_j)$ that action a_j in situation x_j be disclosed.¹¹ We assume that ϕ is such that for every ε -perturbed strategy profile σ , every analogy class \mathcal{C}_i^k is reached and disclosed with strictly positive probability. That is,¹²

$$\mu^\sigma(\mathcal{C}_i^k) = \sum_{x'_j \in \mathcal{C}_i^k, a'_j \in \mathcal{A}_j(x'_j)} \mu^\sigma(a'_j, x'_j)$$

is strictly positive for every \mathcal{C}_i^k .

Definition 2 *The analogy-based expectation β_i is consistent with the ε -perturbed strategy profile σ and the feedback ϕ if for every \mathcal{C}_i^k , and every action a_j in the action space of player j at \mathcal{C}_i^k ,*

$$\beta_i(\mathcal{C}_i^k)[a_j] = \frac{1}{\mu^\sigma(\mathcal{C}_i^k)} \sum_{x_j \in \mathcal{C}_i^k} \mu^\sigma(a_j, x_j), \quad (1)$$

where $\beta_i(\mathcal{C}_i^k)[a_j]$ refers to the probability assigned to action a_j by $\beta_i(\mathcal{C}_i^k)$.

Combining Definition 1 and Definition 2 we define:

Definition 3 *Given a profile of analogy partitions $\mathcal{C} = (\mathcal{C}_1, \mathcal{C}_2)$, and a feedback structure ϕ , an ε -perturbed analogy-based expectation equilibrium is a strategy profile $\sigma = (\sigma_1, \sigma_2)$ such that there exists a profile of analogy-based expectations $\beta = (\beta_1, \beta_2)$ satisfying for $i = 1, 2$:*

- (a) σ_i is an ε_i -perturbed best-response to β_j ,
- (b) β_i is consistent with (σ, ϕ) as defined in (1).

Remark 2 (a) *When the feedback ϕ is complete (i.e. when it contains information about the entire profile (a, x) for all choices of action profiles) or when it contains information only about the equilibrium path in extensive form games of complete information, the above definition is equivalent to the one provided in Jehiel (2005) for extensive form games or Jehiel and Koessler (2008) for Bayesian games. For more general specifications of the feedback structure ϕ , the above definition can be viewed as a natural generalization of*

¹¹We do not include a reference to ϕ in μ^σ since ϕ will be taken as fixed and exogenous throughout.

¹²Observe that $\mu^\sigma(\mathcal{C}_i^k)$ is not a probability as it could be greater than 1 in some cases. This reflects that in extensive-form games, a single play of the game typically allows one to reach more than one situation. Also note that μ is normalized so that there is a mass 1 of games being played.

the analogy-based expectation equilibrium as previously defined.¹³ (b) In the the definition of equilibrium, we have in mind that players are unaware of the payoff and information structure of their opponent as well as of the feedback correspondence.

2.2 Endogenous Categorizations

The general idea we wish to formalize is that given the dataset generated by the strategy profile σ and the feedback structure ϕ , players categorize situations according to their similarity (as derived by some exogenously given player-specific similarity function) and the amount of corresponding data. Roughly, players would like each analogy class to include situations as similar to one another as possible while consisting of sufficiently many data.

We capture this in a reduced form fashion as follows. Each player i is endowed with a similarity function $\zeta_i : 2^{\mathcal{X}_j} \rightarrow [0, 1]$ defined over subsets of \mathcal{X}_j where for every $\mathcal{C}_i^k \subseteq \mathcal{X}_j$, $\zeta_i(\mathcal{C}_i^k) \in [0, 1]$ is a measure of the similarity among the various x_j in \mathcal{C}_i^k . We assume that a singleton set has maximum similarity, i.e. $\zeta_i(\{x_j\}) = 1$ for all $x_j \in \mathcal{X}_j$. We allow for similarity functions such that for some subset $X \subseteq \mathcal{X}_j$ consisting of two or more situations it holds that $\zeta_i(X) = 0$ in which case the situations in X are considered maximally dissimilar. We also allow for similarity functions such that $\zeta_i(X) = 1$ for some non-singleton X , in which case the situations in X are considered maximally similar.¹⁴

Let $\mathcal{P}(\mathcal{X}_j)$ denote the set of partitions of the set \mathcal{X}_j . Given the data generated by σ and ϕ , player i would ideally want to pick her analogy partition $\mathcal{C}_i = \{\mathcal{C}_i^k\}_{k=1}^K$ so as to solve

$$\max_{\mathcal{C}_i \in \mathcal{P}(\mathcal{X}_j)} W_i(\kappa_i(\zeta_i(\mathcal{C}_i^1), \mu^\sigma(\mathcal{C}_i^1)), \dots, \kappa_i(\zeta_i(\mathcal{C}_i^K), \mu^\sigma(\mathcal{C}_i^K))), \quad (2)$$

for some functions W_i and κ_i assumed to be weakly increasing in their arguments. Intuitively, and in line with the bias variance trade-off, player i would like each analogy class to be associated with a large mass of data (as reflected in the monotonicity of κ_i with respect to its second argument) and to contain situations as similar as possible (as reflected in the monotonicity of κ_i with respect to its first argument). The aggregation of the various $\kappa_i(\zeta_i(\mathcal{C}_i^k), \mu^\sigma(\mathcal{C}_i^k))$ is made through W_i which can take different forms, as long as it is weakly increasing in each argument.

A more structured version of this maximization problem could require that situations are categorized so as to produce the largest overall similarities (for example, measured

¹³In Jehiel (2022), it is explained how one can view the behavioral equilibrium of Esponda (2008) as an analogy-based expectation equilibrium by transforming the static Bayesian game into a well chosen multi-stage game. Whether such a transformation can be done here for general ϕ structure should be the subject of further investigation.

¹⁴If two situations $x_i, x'_i \in \mathcal{X}_i$ have different actions sets, i.e. $\mathcal{A}_i(x_i) \neq \mathcal{A}_i(x'_i)$, we assume that any subset that contains both situations has maximal dissimilarity. This will imply that an adjusted analogy partition will never bundle nodes with different action sets, as in Jehiel (2005). In the present paper this assumption will only matter in our discussion of public goods games with a punishment stage.

as the sum of similarities $\zeta_i(\mathcal{C}_i^k)$ over the different analogy classes \mathcal{C}_i^k subject to the constraint that each analogy class should have total mass no smaller than some threshold κ_i . Even this more structured problem would be hard to solve in general.¹⁵ In light of this difficulty, we impose the following weaker desiderata on analogy partitions.

Definition 4 *Given σ and a threshold $\kappa > 0$, we say that $\mathcal{C} = (\mathcal{C}_i, \mathcal{C}_j)$ is κ -adjusted to σ if for each player i , her analogy partition $\mathcal{C}_i = \{\mathcal{C}_i^1, \dots, \mathcal{C}_i^K\}$ satisfies the following criteria*

1. *For each $x \in \mathcal{X}_j$ with $\mu^\sigma(\{x\}) \geq \kappa$, there exists k such that $\mathcal{C}_i^k = \{x\}$.*
2. *If $X \subseteq \mathcal{X}_j$ and $\zeta_i(X) = 0$, there exists no k such that $\mathcal{C}_i^k = X$.*
3. *Let \mathcal{X}_j^{sing} denote the set of situations put into singleton analogy classes in \mathcal{C}_i . If \mathcal{C}_i^k is such that $\mu^\sigma(\mathcal{C}_i^k) < \kappa$, then for any $X \subseteq \mathcal{X}_j \setminus (\mathcal{C}_i^k \cup \mathcal{X}_j^{sing})$, it holds that $\zeta_i(\mathcal{C}_i^k \cup X) = 0$.*
4. *For any non-singleton pair $\mathcal{C}_i^k, \mathcal{C}_i^{k'}$ in \mathcal{C}_i , there is no disjoint pair X, X' such that $X \cup X' = \mathcal{C}_i^k \cup \mathcal{C}_i^{k'}$, $\mu^\sigma(X) \geq \kappa$, $\mu^\sigma(X') \geq \kappa$, and $\zeta_i(X) + \zeta_i(X') > \zeta_i(\mathcal{C}_i^k) + \zeta_i(\mathcal{C}_i^{k'})$.*

Roughly, the first condition says that if a situation is encountered enough times (as parameterized by κ), it is treated as a singleton analogy class (as there is no need to bundle it with other situations to meet the minimum mass criterion). In applications, we will have in mind that on-path situations satisfy this minimum mass requirement. The second condition requires that when a subset of situations are considered to be maximally dissimilar, they cannot be bundled into an analogy class, which seems like a weak and natural condition to impose. The third condition says that the only reason for an analogy class not to meet the minimum mass condition is that adding other situations to the analogy class would induce maximum dissimilarity. The fourth condition requires a kind of local optimality when restricting attention to pairs of analogy classes and considering as criterion the sum of similarities over the various analogy classes.¹⁶

One could potentially include additional desiderata such as relating the number K of analogy classes to σ . A natural such refinement is to require that K is minimally chosen so as to guarantee that $\mu^\sigma(\mathcal{C}_i^k) \geq \kappa$ for all k . However, in general it may be a hard problem to solve, and we have chosen not to include it in the above desiderata. Still,

¹⁵This problem is similar to (but more complicated than) the knapsack problem studied in computer science that can be described as follows. There is a set of objects, each characterized by a value and a capacity requirement. The task is to select a subset of objects that maximizes the sum of value while not exceeding a capacity constraint. This problem is known to be NP-hard, thereby formalizing how difficult it may be to solve in general. It is simpler than our problem to the extent that the criterion in the knapsack problem (sum of the values of each included item) is additive whereas the similarity criterion in our problem is potentially more complex. (Even imposing that W_i is additively separable in its arguments would not make our criterion any simpler than the one in the knapsack problem.)

¹⁶More generally, one could require that for some given increasing function W it is not the case that $W(\zeta_i(X), \zeta_i(X')) > W(\zeta_i(\mathcal{C}_i^k), \zeta_i(\mathcal{C}_i^{k'}))$.

such a refinement will be considered in the context of the trade application in which case it is simpler for players to check and implement.

2.3 Categorization Equilibrium

For fixed ε and κ , we define:

Definition 5 *A profile (σ, \mathcal{C}) is an (ε, κ) -categorization equilibrium if*

- (a) *σ is an ε -perturbed analogy-based expectation equilibrium given \mathcal{C} and*
- (b) *\mathcal{C} is κ -adjusted to σ .*

An (ε, κ) -categorization equilibrium can be understood as a steady state as follows. Assume the system has stabilized to (σ, \mathcal{C}) . When looking at the data generated by previous matches, players would be led to choose analogy partitions \mathcal{C} that are κ -adjusted to the strategy profile σ used in those matches. When trying next to form analogy-based expectations using such analogy partitions, they would be led to have beliefs as defined in (1) given that the play is governed by σ . They would then play as assumed in σ given that σ is an ε -perturbed analogy-based expectations equilibrium (ABEE) for \mathcal{C} , thereby leading the desired steady state property.¹⁷

In our analysis, we are interested in the limit of (ε, κ) -categorization equilibria as ε and κ go to zero. We also consider the case in which we additionally require that κ/ε grows large in the limit. Formally,

Definition 6 *A profile (σ, \mathcal{C}) is a categorization equilibrium if there are sequences $(\varepsilon^m)_m$ and $(\kappa^m)_m$ converging to zero and a sequence $(\sigma^m)_m$ converging to σ , such that (σ^m, \mathcal{C}) is an $(\varepsilon^m, \kappa^m)$ -categorization equilibrium for all m . If $\lim_{m \rightarrow \infty} \kappa^m/\varepsilon^m = \infty$ then (σ, \mathcal{C}) is referred to as coarse categorization equilibrium.*

The fact that κ^m and ε^m vanish implies that on-path nodes are treated as singleton analogy classes. The extra requirement for a coarse categorization equilibrium, that κ^m is arbitrarily small relative to ε^m , means that off-path nodes are bundled together unless doing so generates analogy classes with maximal dissimilarity. In other words, this captures situations with maximally coarse categorizations of the off-the-path situations.

¹⁷We implicitly describe here the case in which all players assigned to the same role would end up with the same analogy partitions (which would require that every such subject employs the same algorithm to find out the analogy partitions). Extensions to non-unitary versions (see Fudenberg Levine, 1993) are possible, but would bring no additional insights in the applications.

2.4 Preliminary Result

In the special case of extensive form games of complete information (i.e. allowing for simultaneous moves but no asymmetric information), assuming that the feedback consists in disclosing the played path, we have:


Proposition 1 *Consider an extensive-form game of complete information and assume that the feedback consists of observing the path of play.*

- (a) *If (σ, \mathcal{C}) is a categorization equilibrium then σ is a (unitary) self-confirming equilibrium (Fudenberg and Levine 1993, 1998).*
- (b) *If σ is a subgame perfect Nash equilibrium (SPNE) then there is a \mathcal{C} such that (σ, \mathcal{C}) is a categorization equilibrium.*
- (c) *If σ is a subgame perfect Nash equilibrium (SPNE) then there may be no \mathcal{C} such that (σ, \mathcal{C}) is a coarse categorization equilibrium.*

Proof. (a) Since $\kappa^m \rightarrow 0$ and $\varepsilon^m \rightarrow 0$ players must have correct expectations about behaviors on the path, given criterion 1 in definition 4. The result follows.

(b) Let L be the length of the longest path of play. This is the highest number of mistakes needed to reach any terminal node under any strategy profile. By choosing sequences $(\varepsilon^m)_m$ and $(\kappa^m)_m$ such that $\lim_{m \rightarrow \infty} \kappa^m / (\varepsilon^m)^L < 1$ we ensure that there is some M such that for any $m > M$ any (ε, κ) -categorization equilibrium will put all off-path nodes in singleton analogy classes. This implies that all players have correct expectations at all nodes. And since (for any finite m) all nodes are reached with positive probability all players will play ε -best responses at all nodes, converging to exact best responses as $m \rightarrow \infty$.

(c) Consider the following version of the centipede game where players 1 and 2 take turn choosing between Pass and Take. The unique SPNE is TP for Player 1 and PT



for Player 2 (indicated by the fat arrows). Both of Player 2's nodes are off-path and reached by a single mistake (by Player 1 at the first node). If $\lim_{m \rightarrow \infty} \kappa^m / \varepsilon^m = \infty$ then Player 1 will bundle these two nodes together (assuming Player 1 does not perceive them as maximally dissimilar) and form the expectation that Player 2 passes with probability 1/2. Given this belief, Player 1 perceives the expected utility of passing at both of her nodes to be 2.5 making it seem optimal to deviate from the strategy SPNE. ■

Remark 3 *When considering other classes of games such as Bayesian games and other feedback structures ϕ such as the one considered in Section 4, one may wonder how categorization equilibria relate to self-confirming equilibria defined in the sense of Battigalli (1987) or Dekel et al (2004). We note that if players are aware of the tremble structure as well as ϕ , categorization equilibria need not be self-confirming equilibria given ϕ , even assuming that κ is set so that on the path situations are treated as singleton analogy classes.*¹⁸

Remark 4 *In relation to the difference between (b) and (c) of Proposition 1 we note that if $\kappa^m / (\varepsilon^m)^l < 1 < \kappa^m / (\varepsilon^m)^{l+1}$ then any node that is at most l steps off the equilibrium path will be placed in a category of its own under any $(\varepsilon^m, \kappa^m)$ -categorization equilibrium, whereas nodes that are further away from the equilibrium path may be bundled more coarsely.*

Remark 5 *In the Online Appendix we provide two examples in which (σ, C) is a categorization equilibrium but σ is not a Nash equilibrium. Constructing such examples either require that the feedback differs from the path of play (in which case a normal form game with just two players can be used to illustrate the claim) or (if the feedback is the path of play) that one considers games with at least three players and some asymmetric information. In the latter case we adapt an example from Fudenberg and Levine (1993) used to illustrate that a self-confirming equilibrium may differ from a Nash equilibrium.*

2.5 Dynamics

In some cases there will be no categorization equilibrium, hence the need to describe the learning dynamics differently. In an attempt to illustrate the possibility of cycling that could emerge then, we will consider the following dynamics chosen for its simplicity. Let $\mathcal{C}(t) = (\mathcal{C}_i(t), \mathcal{C}_j(t))$ denote a profile of analogy partitions at time t . Given $\mathcal{C}(t)$, a corresponding analogy-based expectation equilibrium σ^t is played at t . At $t + 1$, a new profile of analogy partitions $\mathcal{C}(t + 1) = (\mathcal{C}_i(t + 1), \mathcal{C}_j(t + 1))$ applies so that for each player i , $\mathcal{C}_i(t + 1)$ is adjusted to σ^t . Given $\mathcal{C}(t + 1)$, a corresponding analogy-based expectation equilibrium σ^{t+1} is played at $t + 1$, and so on. The dynamics is parameterized by the initial choice of analogy partitions $\mathcal{C}(1)$, as well as the choice of analogy-based expectation equilibrium in any period t if there are several of them, and the choice of adjusted analogy partitions if there are several of them.

The interpretation of this dynamics is as follows. There are populations of players i and j matched in pairs in every period. Each period consists of a large (potentially

¹⁸This is so because, the conjecture $\sigma_j^{\beta_i}$ may be at odds with the observations $\sigma_j^{\beta_i}$ when the trembling structure as well as ϕ is known (for example, this is the case in the trade application developed below when the trembles are not concentrated on bid prices above the maximum value of the seller).

infinite) number of subperiods with random re-matching in every subperiod. Players are renewed in every period. At the start of a given period, players adjust their choice of analogy partitions to the observations that they see from period $t - 1$. They stick to their analogy partitions throughout the period. We assume that the experience gained during the subperiods lead the players to converge to a steady state, and accordingly to behave as in an analogy-based expectation equilibrium given the profile of analogy partitions prevailing in the period.

The implicit assumption behind this dynamic is that once a given player has settled on an analogy partition, she never reconsiders it in his lifetime and the choice of analogy partition is based on the behavioral data observed in the last period (which are the more recent observations at the start of the period). It is not difficult to establish that if such a dynamic learning model has a steady state it corresponds to a categorization equilibrium. However, we will consider the dynamics to cover cases in which there is no steady state and cycles emerge instead.

3 Chainstore and Other Finite-Horizon Games

In this Section, we apply our approach to the classic chainstore game, and we illustrate how the monopolist may deter entry in most periods in a categorization equilibrium. By contrast, in the centipede game, we note that categorization equilibria coincide with the usual subgame perfect Nash equilibrium (SPNE). We also briefly apply the approach to finitely repeated public good games, and illustrate how positive contributions can be achieved in categorization equilibria if the game includes a punishment stage but not otherwise.

3.1 Chainstore Game

3.1.1 Set-Up

In the finitely repeated chainstore game an *incumbent monopolist* faces a sequence of T *challengers*. Each challenger chooses to Enter (E) or to stay Out (O). If the challenger enters then the monopolist chooses whether to Accommodate (A) or Fight (F). The stage game payoffs of the monopolist and a challenger are denoted u_M and u_C , respectively, with $u_C(E, A) > u_C(O) > u_C(E, F)$ and $u_M(O) > u_M(E, A) > u_M(E, F)$. In words, the challenger prefers entering and facing an accommodating incumbent over not entering, and prefers not entering over entering and facing a fighting incumbent. The monopolist prefers the challenger to stay out over accommodating an entering challenger, and prefers the latter over fighting an entering challenger. Each challenger maximizes her payoff (in the stage she participates) and the monopolist maximizes the sum of stage game payoffs.

In the unique SPNE of this game, challengers choose E in every period and this is followed by A . This is shown by backward induction where these behaviors are established in the last period irrespective of the history and then established in period t assuming they prevail from period $t + 1$ onward.

The prediction of the SPNE in the chainstore game has been considered unintuitive, as the monopolist would seem to be able to deter entry by playing F in case of entry. While this kind of reasoning cannot arise in a SPNE, we will see it can in a categorization equilibrium.

To make the chainstore game fit into our general two-player framework, we assume the challengers at the various time periods t form a single player, the challenger.¹⁹ We also assume that the trembling probability is the same for the monopolist and the challenger.

3.1.2 Similarity

Our key assumptions concern similarity between histories. We will have in mind that nodes with histories in which there was an E action not matched in the same period by an F action are very dissimilar from other nodes with other histories (i.e., histories in which there was either no E action or each E action were matched with an F action). We believe there is an important qualitative difference between these two sets of histories making our assumption plausible. In effect, it will force us to have analogy classes that do not mix these two subsets of histories (according to requirement 2 of Definition 4).

On the challenger's side, an important aspect is whether the similarity of histories depends on calendar time t . In our basic construction, we will assume that the challenger considers histories at two different calendar times to be sufficiently dissimilar so that histories at different periods never belong to the same analogy class. This would fit with the interpretation of the challenger as being an aggregate of different players, one at each calendar time, where the period t challenger would focus exclusively on behaviors previously observed in period t . Below we will also discuss how the analysis is affected when the challenger does not treat histories at different calendar times to be maximally dissimilar, more in line with our treatment of similarity for the monopolist.

Formally, we first consider the nodes at which the challenger must make a decision and refer to the set of these nodes as \mathcal{Q}_C . We consider two subsets of \mathcal{Q}_C :

$$\begin{aligned}\mathcal{Q}_C^{Tough} &= \{q \in \mathcal{Q}_C : \text{No } E \text{ or all } E \text{ immediately followed by } F \text{ in history of } q\}; \\ \mathcal{Q}_C^{Soft} &= \{q \in \mathcal{Q}_C : \text{Some } E \text{ immediately followed by } A \text{ in history of } q\}.\end{aligned}$$

We require that for any $q^{Tough} \in \mathcal{Q}_C^{Tough}$ and $q^{Soft} \in \mathcal{Q}_C^{Soft}$, if q^{Tough} and q^{Soft} belong to X , then $\xi_M(X) = 0$. That is, any set containing a node in \mathcal{Q}_C^{Tough} and a node in \mathcal{Q}_C^{Soft} is

¹⁹This has no effect on the analysis of SPNE.

considered to be maximally dissimilar by the monopolist. Any subset X containing only elements in \mathcal{Q}_C^{Tough} or only elements in \mathcal{Q}_C^{Soft} is supposed to satisfy $\xi_M(X) > 0$.

Regarding the nodes in which the monopolist must make a decision, we denote the set of those corresponding to period t by \mathcal{Q}_M^t and we distinguish in \mathcal{Q}_M^t two subsets:

$$\begin{aligned}\mathcal{Q}_M^{t,Tough} &= \{q \in \mathcal{Q}_M^t : \text{No } E \text{ or all } E \text{ immediately followed by } F \text{ in history of } q\}; \\ \mathcal{Q}_M^{t,Soft} &= \{q \in \mathcal{Q}_M^t : \text{Some } E \text{ immediately followed by } A \text{ in history of } q\}.\end{aligned}$$

We require that if Y contains two nodes q and q' that either, (a) correspond to two different time periods, or (b) do not both belong to $\mathcal{Q}_M^{t,Tough}$, or (c) do not both belong to $\mathcal{Q}_M^{t,Soft}$ for some t , then $\xi_C(Y) = 0$. Any Y not containing two such elements is supposed to satisfy $\xi_C(Y) > 0$.

3.1.3 Equilibrium

Strategy profile.

We define the threshold

$$k^* = \min \{k \in \mathbb{N} \text{ such that } u_M(E, F) + k u_M(O) \geq (k + 1) u_M(E, A)\}, \quad (3)$$

and consider the following strategy profile σ_T :

- Challenger $t \leq T - k^*$ strategy. If E was always matched with F in the past, or if there was no E in the past, play O . Otherwise play E .
- Challenger $t > T - k^*$ strategy. Play E .
- Monopolist strategy. At $t > T - k^*$, play A . At $t \leq T - k^*$; play F if E was always matched with F in the past, or if there was no E in the past; otherwise play A .

On the path of play induced by this strategy profile, the challenger enters only in the last k^* periods, and the monopolist accommodates those entries (while she would fight the challenger if entering in earlier periods).

Categorization profile.

In addition to the strategy profile, a categorization equilibrium consists of a categorization profile \mathcal{C} , illustrated in Figure 1.

- Each *on-path node* is in a separate analogy class.
- The monopolist categorizes *off-path challenger nodes* based on whether there was previously an act of E that was not met by F . The first analogy class bundles all off-path nodes with a history in which E was always met by F , and the second

analogy class bundles all the remaining off-path nodes. Formally, let \mathcal{Q}_C^{off} be the set of monopolist decision nodes that are located off the equilibrium path,

$$\begin{aligned}\mathcal{C}_M^1 &= \left\{ q \in \mathcal{Q}_C^{off} \cap \mathcal{Q}_C^{Tough} \right\}; \\ \mathcal{C}_M^2 &= \left\{ q \in \mathcal{Q}_C^{off} \cap \mathcal{Q}_C^{Soft} \right\}.\end{aligned}$$

- Challengers categorize *off-path monopolist nodes* based on the stage of the game only.²⁰ Formally, let \mathcal{Q}_M^{off} be the set of off-path monopolist decision nodes. For each t let

$$\begin{aligned}\mathcal{C}_{Ct}^1 &= \left\{ q \in \mathcal{Q}_M^{off} \cap \mathcal{Q}_M^{Tough} : q \text{ is in round } t \right\}; \\ \mathcal{C}_{Ct}^2 &= \left\{ q \in \mathcal{Q}_M^{off} \cap \mathcal{Q}_M^{Soft} : q \text{ is in round } t \right\}.\end{aligned}$$

As we show now, the above strategy profile and categorization profile together form a coarse categorization equilibrium when T is large enough.

Proposition 2 *There exists a T^* such that if $T > T^*$, then (σ_T, \mathcal{C}) is a coarse categorization equilibrium of the chainstore game with T periods, implying that in the absence of trembles the challenger enters only in the last k^* periods, and the monopolist fights the challenger in all but the last k^* periods.*

Proof. We need to show that for $T > T^*$ there is a sequence $(\sigma_T^m)_m$ converging to σ_T , such that $(\sigma_T^m, \mathcal{C})$ is an $(\varepsilon^m, \kappa^m)$ -categorization equilibrium for all m . We define σ_T^m as the strategy profile which at each node puts probability ε^m on the action that σ_T puts zero probability on. Since there are only two actions at each node this is enough to specify σ_T^m . Since the starting point of $(\varepsilon^m, \kappa^m)$ is arbitrary it is sufficient to show the following: There exists a T^* such that for any $T > T^*$ there exists an m^* such that if $T > T^*$ and $m > m^*$ then σ_T^m is an $(\varepsilon_T^m, \kappa_T^m)$ -categorization equilibrium of the chainstore game with T periods.

1. First we explain why \mathcal{C} is adjusted to σ_T^m for all $m > m^*$ (and all T).

- (a) For any T , if m is large enough, then $\kappa_T^m < (1 - \varepsilon_T^m)^T$, ensuring that on-path nodes have a mass exceeding the threshold κ_T^m and thus are treated as singleton analogy classes, by point 1 of Definition 4.

²⁰We note that there are other categorizations that could be combined with σ^T to form a CE. For example we could let challengers bundle all monopolist nodes from the same period in a separate category for each time period. They would still have correct expectations.




Figure 1: Illustration of the strategy profile σ_T and the categorization \mathcal{C} for $T = 4$, and $k^* = 2$. Fat arrows indicate the proposed strategy profile σ_T . On-path nodes, in white, are distinguished perfectly by both the monopolist and the challenger. Light grey off-path nodes belong to the set \mathcal{Q}^{Tough} and dark grey off-path nodes belong to the set \mathcal{Q}^{Soft} . The monopolist separates them into the analogy classes \mathcal{C}_M^1 and \mathcal{C}_M^2 respectively. The challenger further separates them by time period, here represented by the nodes inside the dotted rectangles.

- (b) For off-path nodes following histories in which there was some E not matched with F , our similarity assumptions imply that nodes in \mathcal{Q}_C^{Soft} cannot be bundled with nodes that are not in \mathcal{Q}_C^{Soft} , and nodes in $\mathcal{Q}_M^{t,Soft}$ cannot be bundled with nodes that are not in $\mathcal{Q}_M^{t,Soft}$, according to point 2 of Definition 4.²¹
- (c) Furthermore, all off-path nodes in \mathcal{Q}_C^{Soft} have to be bundled together and all off-path nodes in $\mathcal{Q}_M^{t,Soft}$ have to be bundled together (but separately for each t) according to point 3 of Definition 4. This follows from the assumption that $\lim_{m \rightarrow \infty} \kappa_T^m / \varepsilon_T^m = \infty$, which implies that the total mass of the off-path nodes vanishes relative to the threshold κ .
- (d) The situation is analogous for off-path nodes following histories in which there was no E or any E was immediately followed by an F . The off-path nodes of the challenger \mathcal{Q}_C^{off} have to be partitioned into \mathcal{C}_M^1 and \mathcal{C}_M^2 , and the off-path

²¹The total mass of such histories would typically fall short of the κ_T^m threshold, but the dissimilarity with other histories would not allow further bundling (see below for further discussion of this).

nodes of the monopolist have to be partitioned, for each t , into \mathcal{C}_{Ct}^1 and \mathcal{C}_{Ct}^2 .

2. Second we examine the analogy-based expectations

- (a) Players have correct expectations at on-path nodes.
- (b) Players also have correct expectations at nodes following off-path histories in which there was some E not matched with F , i.e. at off-path nodes in \mathcal{Q}_C^{Soft} and $\mathcal{Q}_M^{t,Soft}$. This is so because after such histories, the challenger consistently chooses E and the monopolist consistently chooses A after E .
- (c) Next consider off-path monopolist nodes following histories in which there was no E or any E was immediately followed by an F , i.e. off-path nodes in $\mathcal{Q}_M^{t,Tough}$ for some t . (Such a node is only reached when the challenger plays E before $t \leq T - k^*$.) Challengers have correct expectations since they do not bundle together nodes from different time periods. (Indeed this would be true even if challengers did not distinguish between $\mathcal{Q}_M^{t,Tough}$ and $\mathcal{Q}_M^{t,Soft}$.)
- (d) It only remains to check the monopolist's expectations at off-path nodes in \mathcal{Q}_C^{Tough} . As $\varepsilon^m \rightarrow 0$ the expectations here are determined by behavior at nodes with histories containing a single experimentation. The fraction of such nodes at which the challenger chooses E vanishes as $T \rightarrow \infty$. It follows that as T gets large, the monopolist will expect that O is chosen with a probability close to 1.

3. Third and finally we verify that σ_T^m induces a ε_T^m -best-responses given the analogy-based expectations. We have found that the challengers have correct expectations and it is easy to see that they best-responds to the monopolist's strategy, so we focus on the monopolist.

- (a) Monopolist in period $t \leq T$ at an off-path node in \mathcal{Q}_M^{Tough} . By playing F , the monopolist expects that with a probability close to 1, a string of O occur from then on until the end of the game. By playing A , the monopolist correctly expects a string of (E, A) until the end of the game. The former is at least as good as the latter if

$$u_M(E, F) + (T - t) u_M(O) \geq (T - 1 + 1) u_M(E, A),$$

For $t \leq T - k^*$ this is satisfied, but for $t > T - k$ it is not satisfied, by the definition of k^* .

- (b) Monopolist at the on-path node in period $t = T - k^* + 1$. This node is in \mathcal{Q}_M^{Tough} , immediately preceded by the first instance of E . By deviating from σ^T and playing F , the monopolist expects that with a probability close to

1, a string of O occur from the next period until the end of the game. By complying with σ^T and playing A , the monopolist correctly expects a string of (E, A) until the end of the game. Deviation is then perceived unprofitable if by the same condition as before.

- (c) Monopolist at an off-path node in Q_M^{Soft} . Regardless of what happens in the current period, the monopolist (correctly) expects the challengers to play E in all subsequent periods. The best response is to play A from now until the end of the game.
- (d) Monopolist at an on-path node in period $t > T - k^* + 1$. In the history of such a node there has been at least one instance of E that was not immediately followed by A , i.e. the node is in Q_M^{Tough} . The monopolist (correctly) expects the challengers to play E in all subsequent periods. The best response is to play A from now until the end of the game.

■

To emphasize the logic of the proposed categorization equilibrium, observe that the only mistaken expectations are those of the monopolist regarding off-path nodes in Q_C^{Tough} . In particular, if E occurs in period $t = T - k^*$ (i.e. the last period in which the challenger is supposed to stay out) then the monopolist mistakenly expects that by playing F , the challengers will be induced to stay out from then on, whereas in reality, no matter what the monopolist does there will be entry in all the remaining periods. This mistake is caused by the fact that there isn't enough mass of data on behavior at the subsequent challenger nodes (due to our assumption that $\lim_{m \rightarrow \infty} \kappa_T^m / \varepsilon_T^m = \infty$) so that they have to be bundled with many other nodes in Q_C^{Tough} at which indeed fighting after entry deters future entry in most periods.

3.1.4 Discussion

Why do we need $\lim_{m \rightarrow \infty} \kappa_T^m / \varepsilon_T^m = \infty$?

As explained above, this allows us to obtain that histories off-the path in Q_C^{Tough} are bundled into coarse analogy classes consisting of a large number of histories, which in turn leads the monopolist to expect after every such history that the challenger is very unlikely to enter next. If instead we had assumed that $\lim_{m \rightarrow \infty} \kappa_T^m / \varepsilon_T^m = 0$, then every such history would have been a singleton analogy class. Thus, following an entry at $t = T - k^* - 1$, the monopolist would have expected that a string of (E, A) would come next whether he plays A or F , thereby leading him to prefer A , in violation of the recommendation of the assumed strategy. When $\lim_{m \rightarrow \infty} \kappa_T^m / \varepsilon_T^m = 0$, only the usual SPNE could arise as a categorization equilibrium. This is similar to (b) in Proposition 1.

What if histories in Q_C^{Soft} are bundled together whether on or off the path?

In our above construction, off-path histories in Q_C^{Soft} were bundled together. Since reaching such histories requires two trembles, they do not together reach the threshold mass κ_T no matter how large T is.²² In light of this, one may find it reasonable to bundle off-path histories in Q_C^{Soft} with (at least) some on-path histories in Q_C^{Soft} . We note that this would not change our results. The reason is that the behaviors are the same at all histories in Q_C^{Soft} , hence the same strategy profile would constitute a categorization equilibrium with this alternative categorization.²³

What about other categorization equilibria?

It is clear that one cannot support categorization equilibria with fewer periods of entry when the challenger is behaving optimally, as the challenger would always enter in the last k^* periods anticipating that the monopolist would find it optimal to play A (as implied by the definition of k^*). But, one can easily support equilibria with more periods of entry. In fact, take any $k^{**} > k^*$. It is readily verified that replacing k^* by k^{**} in the above strategy profile would be a categorization equilibrium for T large enough.

What if the challenger does not distinguish histories according to time? This is a small modification of our similarity assumptions. Compared to the above setting, the only difference is that for the challenger we now consider $Q_M^{Tough} = \cup_t Q_M^{t,Tough}$ and $Q_M^{soft} = \cup_t Q_M^{t,Soft}$ and we require that if Y contains two nodes q and q' that do not both belong to Q_M^{Tough} nor both belong to Q_M^{soft} , then $\tilde{\xi}_C(Y) = 0$ (while any set Y not having this property satisfies $\tilde{\xi}_C(Y) > 0$). We define a corresponding categorization profile \tilde{C} which only differs from C in that the challengers' categorizations do not differentiate periods, i.e. $\tilde{C}_C^1 \cup_t C_{Ct}^1$ and $\tilde{C}_C^2 \cup_t C_{Ct}^2$.

It can be checked that the above strategy profile σ_T would still be part of a categorization equilibrium (together with \tilde{C}) under the conditions of Proposition 2.²⁴ However, we now observe that there are other categorization equilibria, this time relying on erroneous expectations of the challengers. Still defining k^* as above, we consider the following strategy profile $\tilde{\sigma}_T$:

- Challenger strategy. If E was always matched with F in the past, or if there was

²²This is so because $\lim_{T \rightarrow \infty} \frac{\kappa_T}{\varepsilon_T} = \infty$ and $\lim_{T \rightarrow \infty} (1 - \varepsilon_T)^T = 1$ would imply that $\lim_{T \rightarrow \infty} \frac{\kappa_T}{T(\varepsilon_T)^2} = \infty$.

²³This would correspond to the idea that histories in Q_C^{Soft} are viewed as perfectly similar. It would violate our above axiomatic approach to categorization that imposes that if a singleton node has a mass above the threshold it is treated as a singleton analogy class, but arguably the motivation for such a property was that bundling such a node with other nodes would strictly decrease the similarity in the corresponding set.

²⁴Intuitively, σ_T remains a categorization equilibrium because in period $t = T - k^* + 1$, the challenger treats the history with E at t and no E before as a singleton analogy class, so that he finds it optimal to choose E then. And it can be checked that the bundling of histories in Q_M^{Tough} before $T - k^* + 1$ does not induce mistakes as the monopolist consistently chooses F after any such history.

no E in the past, play O . Otherwise play E .

- Monopolist strategy. At $t > T - k^*$, play A . At $t \leq T - k^*$; play F if E was always matched with F in the past, or if there was no E in the past; otherwise play A .

On the path of play induced by this strategy profile the challenger never enters. In case there is entry the monopolist fights the challenger in all but the last k^* periods.

Proposition 3 *There exists a T^* such that if $T > T^*$, then $(\tilde{\sigma}_T, \tilde{\mathcal{C}})$ is a coarse categorization equilibrium of the chainstore game with T periods, implying that in the absence of mistakes there is not entry, and the monopolist fights the challenger in all but the last k^* periods.*

In this construction, the monopolist plays a best-response to the challenger's strategy and the mistaken belief concerns the challenger who refrains from entering in all periods. She stays out at histories with no earlier (E, A) because she fears the monopolist would fight with a large probability in case of entry. This expectation arises due the bundling of many histories in \mathcal{Q}_M^{Tough} and the observation that according to $\tilde{\sigma}_T$ the monopolist would play F at such histories in all but the last k^* periods. The detailed derivation of the proposition appears in the Appendix.

It should be noted that $\tilde{\sigma}_T$ cannot part of a categorization equilibrium when using the similarity relation assumed in Proposition 2, i.e. when the challenger is induced to categorize different time periods separately. This is so because the challenger would then have to expect that in the last k^* period histories in \mathcal{Q}_M^{Tough} the monopolist would play A after entry, thereby leading challengers to choose E in those events in contrast to the prescription of $\tilde{\sigma}_T$. We see here the effect of the similarity functions in shaping the categorization equilibria.

3.2 Other finite horizon games

In this Section, we briefly discuss the implications of our approach in some alternative classic games considered in relation to the finite horizon paradox.

3.2.1 Centipede game

In the centipede game two players take turn choosing between the actions *pass* and *take*. The players have equally many nodes and the total number of nodes is T . The payoffs are such that a player acting at a node prefers taking at that node over the other player taking at the next node, but prefers both to pass at the these two nodes over either one taking at them.²⁵ A version of the standard backward induction argument can be employed. Note

²⁵Formally, consider player i acting at node $t \leq T - 2$. If i takes she earns $u_i(a_{i,t} = take)$. If i passes and then player $j \neq i$ takes at the next node i earns $u_i(a_{j,t+1} = take)$. If i passes at t , and $j \neq i$ passes

that taking is dominant at node T so for any categorization equilibrium there is a first round t in which a player takes. This means that nodes $1, \dots, t$ are on the equilibrium path. Due to the assumption that κ is asymptotically small in a categorization equilibrium, each node on the equilibrium path is a singleton analogy class. Thus, if player j takes at node t then player i correctly understands that taking is the optimal choice for her at node $t - 1$. We arrive at the following result:

Proposition 4 *Every categorization equilibrium of the centipede game prescribes taking at the first node.*

The comparison between the analysis of categorization equilibria in the chain store game and in the centipede game is instructive. While no departure from SPNE arises in the centipede game (as just established), we found that behaviors very different from those in the SPNE can arise in the chainstore game. It is the more complex extensive form structure in the chainstore game as compared to the centipede game that is responsible for the difference. In the centipede game, all self-confirming equilibria (in which the feedback is the path of play) coincide with the SPNE, thereby explaining why categorization equilibria coincide with SPNE in this case (see Proposition 1a). By contrast, in the chainstore game, our proposed similarity functions together with the assumed trembling behavior have allowed us to construct categorization equilibria that imply plausible behaviors that are not consistent with SPNE.

3.2.2 Public Good Game

We now turn to the finitely repeated linear public goods game, where in each period agents privately decide on how much to contribute to the public good. Since the marginal cost of contribution is assumed to lie in between the private and the social benefit efficiency demands that everyone contributes in all periods, but in the unique subgame perfect equilibrium no one contributes in any period. In some variants, agents at the end of a period can also decide whether or not to punish other agents (after they have observed the profile of contributions in the current period). The subgame perfect Nash equilibrium still predicts no contribution, as well as no punishment in any period. This is in sharp contrast with behaviors experimentally observed in such games: Fehr and Gächter (2002) have documented significant levels of contribution, especially when agents have the possibility of punishing their peers, noting that contributions do not decrease over time in the presence of the punishment option.

In the Online Appendix S.2, we apply our framework to such games, assuming that histories in which agents have failed to contribute (in case there is no punishment stage),

at $t + 1$, and i takes at $t + 2$ then i earns $u_i(a_{i,t+2} = take)$. These payoffs satisfy $u_i(a_{j,t+1} = take) < u_i(a_{i,t} = take) < u_i(a_{i,t+2} = take)$. For player i acting at node $T - 1$ we have $u_i(a_{j,T} = take) < u_i(a_{i,T-1} = take) < u_i(a_{i,T} = pass)$, and for player i acting at node T we have $u_i(a_{i,T} = pass) < u_i(a_{j,T} = take)$.

and histories in which agents failed to contribute and were not punished, or contributors were punished (in case there is a punishment stage) are very dissimilar from other histories where there was always contribution (in case there is no punishment stage), or non-contributors were always punished and contributors never punished (in case there is a punishment stage). We establish that Categorization equilibria with significant levels of contribution can be supported when agents have the opportunity to punish but not otherwise. Such a finding is in agreement with our above observation that when the interaction is sufficiently complex (with more nodes off the path) behaviors that depart from the SPNE predictions can be supported but not otherwise. It is also consistent with the experimental findings reported in Fehr and Gächter (2002).

4 On Cycling in Adverse Selection Games

In this Section, we apply our approach to the classic lemon's market game (in a simple double auction format as in Esponda 2008). We make the point that our approach may account for the emergence of cycling with periods of little trade followed by periods of more trade. Roughly, cycling emerges because the little trade regime triggers more optimism about trade on the buyer side due to the scarcity of the data relating sellers' behaviors and the quality of goods, thereby leading to a high trade regime. And, as data are less scarce in the high trade regime, buyers are led back to play rationally as in the low trade regime when making their expectations based on the data from the high trade regime. We make the framework and the analysis precise in the rest of the Section.

4.1 Set-Up

4.1.1 Market

Consider a market for trade of indivisible objects with random quality ω distributed on $\Omega = [0, 1]$ according to a continuous density function g , with cumulative G . Sellers know the quality ω of their good. But buyers do not observe qualities; they only know the distribution of ω . The valuation of a given seller coincides with the quality ω of his good. The corresponding valuation of a buyer is $v = \omega + b$, where $b \in (0, 1)$ represents gains from trade. We posit a one-to-one trading mechanisms between pairs consisting of one seller and one buyer drawn at random from their respective pools. In each pair, the seller and the buyer act simultaneously. The seller quotes an ask price $a(\omega)$ that depends on the quality ω . The buyer quotes a bid price $p \in [0, 1]$. The market mechanism is such that if $p < a$ there is no trade, and if $p \geq a$ trade occurs at price p . Hence, if there is trade the buyer obtains utility $u(p) = v - p$, and the seller obtains utility p . If there is no trade, the seller gets ω and the buyer gets 0. This can be viewed as a Bayesian game

between one seller informed of the state ω and one buyer not observing ω with action profiles and payoffs as just shown.

In this modeling of the trading mechanism, setting the ask price equal to the quality $a(\omega) = \omega$ is a weakly dominant strategy for the seller (just as bidding one's own valuation is a weakly dominant strategy in the second-price auction). In what follows, it will be assumed that the seller always employs this weakly dominant strategy and never plays any other strategy.

To make the analysis simpler, we assume that $b < (g(1))^{-1}$ and that G has the *monotone reversed hazard rate property* (or equivalently that G is strictly log-concave). That is, for all p ,

$$\frac{\partial}{\partial p} \left(\frac{g(p)}{G(p)} \right) < 0.$$

Moreover, we assume the following smoothness condition: $|g'(p)| < g(p)$ for all p .²⁶

We model the set of qualities and prices as a continuum, for tractability reasons. In reality quality and prices may be discrete and our continuum model would describe the limiting case with arbitrarily many discrete values.

4.1.2 The Categorization Setup

To apply the general framework introduced above we identify Ω with \mathcal{X} , and we adopt the (straightforward) extensions of our definitions to deal with the case of a continuum of states and a continuum of actions.

Feedback. Since coarse reasoning will only concern the buyer, it is enough to specify which profiles (ω, a) of quality ω and ask prices a are disclosed to new buyers. In line with Esponda (2008), we posit that (ω, a) appears in the feedback only when there is trade, i.e. when $a < p$. This defines the ϕ function for the application.

Trembles. We will assume that only buyers tremble.²⁷ Specifically, with probability $1 - \varepsilon$ a buyer picks a best response to her expectations and with probability ε she experiments. When trembling, we assume that buyers choose experimental bids according to a pdf f and cdf F with full support on $[0, 1]$.²⁸

Similarity. When categorizing Ω , the buyer employs a similarity function related to the average euclidean distance in the set so that if \mathcal{C} is a subset of Ω , $\xi(\mathcal{C})$ will be thought

²⁶While not essential for our main conclusion regarding the presence of price cycles, these extra assumptions will simplify the analysis by ensuring that there is a unique interior Nash equilibrium. The assumptions will also allow us to more completely analyse the cycling phenomenon.

²⁷This is motivated on the ground that sellers have a weakly dominant strategy, making the discovery of the best strategy probably simpler for sellers than for buyers.

²⁸In line with our trembling formulation described in Section 2, we could impose that $f \equiv 1$ but our results apply to any f , hence our formulation.

of as being the average distance between the various elements in C .²⁹ This notion of similarity will (in line with requirement 4 of Definition 4) give rise to interval analogy partitions in which the set Ω is partitioned into various consecutive intervals.

Threshold Mass. In line with our general assumptions, we have in mind that for on-path qualities ω , i.e., such that (ω, a) is disclosed when the buyer does not tremble, there are enough data for ω to be treated as singleton analogy classes. Since we are considering a setup with a continuum of ω , a strict application of requirement 1 of Definition 4 would not allow this, regardless of how small κ is. We make this assumption in order to simplify the exposition (and note that it could be justified by our view of the continuum as an approximation of the discrete case).³⁰

Formally, denote by p^* the bid price chosen by non-trembling buyers in generation $t - 1$. In generation t , all $\omega \leq p^*$, will be treated as singleton analogy classes so that buyers will understand that the ask price is $a = \omega$ for $\omega < p^*$. However, for $\omega > p^*$, buyers will be using a coarse analogy partition of $(p^*, 1]$ into $K \geq 1$ analogy classes $\mathcal{C}^1, \mathcal{C}^2, \dots, \mathcal{C}^K$ defined by $\mathcal{C}^k = (c_{k-1}, c_k]$ where

$$p^* = c_0 < c_1 < c_2 < \dots < c_{K-1} < c_K = 1.$$

We will require that any \mathcal{C}^k corresponds to a mass no less than κ . As in our definition of categorization equilibrium, we will be interested in the shape of categorizations and bidding strategies in the limiting case of $\kappa \rightarrow 0$ and $\varepsilon \rightarrow 0$.

4.1.3 Preliminary Analysis

Mass of Observations. The density of transactions conditional on trembling is $\tilde{g}_{p^*}(\omega) := g(\omega)(1 - F(\omega))$, and the density of transacted qualities in the dataset is thus given by

$$\mu_{p^*}^{\sigma, \varepsilon}(\omega) = \begin{cases} (1 - \varepsilon)g(\omega) + \varepsilon\tilde{g}_{p^*}(\omega) & \text{if } \omega \leq p^*; \\ \varepsilon\tilde{g}_{p^*}(\omega) & \text{if } p^* < \omega. \end{cases}$$

In what follows we suppress the subscript reference to p^* and write \tilde{g} and μ instead of \tilde{g}_{p^*} and μ_{p^*} , relying on the context to indicate the relevant p^* .

Adjustment of Categorization to Observations. As already mentioned, we view our continuum model as a tractable approximation of a discrete model with finitely many prices and types. In a discrete model, as $\kappa \rightarrow 0$ and $\varepsilon \rightarrow 0$ each type $\omega \leq p^*$ is put in a singleton analogy class. This motivates our focus on the limiting case of an arbitrarily

²⁹We could, with no impact on the analysis, define $\xi(C)$ as the difference between the supremum and infimum ω among the elements of C .

³⁰We briefly discuss later the case in which such situations are pooled as well.

fine-grained categorization below p^* in the continuum model. For types above p^* the number of categories depend on κ and ε in a more complex way. Each analogy class above p^* should satisfy $\kappa \leq \int_{c_{k-1}}^{c_k} \mu_{p^*}^{\sigma, \varepsilon}(\omega)(s) ds$. Consequently, the number of categories above p^* (any $p^* < 1$) is

$$K = \max \left\{ 1, \left\lfloor \frac{1}{\kappa} \int_{p^*}^1 \mu(s) ds \right\rfloor \right\} = \max \left\{ 1, \left\lfloor \left(\tilde{G}(1) - \tilde{G}(p^*) \right) \frac{\varepsilon}{\kappa} \right\rfloor \right\} \leq \max \left\{ 1, \frac{\varepsilon}{\kappa} \right\}.$$

If $\kappa/\varepsilon \rightarrow a$ for some constant $a > 0$, then in the limit an adjusted categorization will have K analogy classes where K is bounded from above by $\max \left\{ 1, \frac{1}{a} \right\}$, which is finite, but possibly larger than one. If we impose $\kappa/\varepsilon \rightarrow 0$, as in the definition of coarse categorization equilibrium, then there is a single analogy class above p^* .

Analogy-Based Expectations. Buyers predict the distribution of ask price a of a type ω seller, knowing that trade occurs if $a \leq p$. For a quality $\omega \leq p^*$ the buyers understand that $a(\omega) = \omega$. Consequently, for a quality $\omega \leq p^*$ the buyer understands that the probability of trade is zero conditional on $\omega \leq p$ and one conditional on $p < \omega$,

$$\Pr(\widehat{a \leq p} | \omega) = \Pr(a < p | \omega) = \Pr(\omega < p | \omega) = \mathbb{I}_{\{\omega \leq p\}}.$$

For a quality $\omega > p^*$ the buyer forms a prediction of the ask price distribution associated with qualities in analogy class \mathcal{C}^k using the data generated under experimentation. Using the fact that $a(\omega) = \omega$ we can write the probability density function (pdf) of ask prices conditional on a quality in \mathcal{C}^k as

$$\tilde{g}(a | \omega \in \mathcal{C}^k) = \frac{\tilde{g}(a)}{\int_{\omega \in \mathcal{C}^k} \tilde{g}(\omega) d\omega} = \frac{\tilde{g}(a)}{\tilde{G}(c_k) - \tilde{G}(c_{k-1})}.$$

Thus, the buyer believes that the pdf of ask prices due to sellers with quality in \mathcal{C}^k is

$$h_{\mathcal{C}^k}(a) = \begin{cases} \frac{\tilde{g}(a)}{\tilde{G}(c_k) - \tilde{G}(c_{k-1})} & \text{if } a \in \mathcal{C}^k; \\ 0 & \text{otherwise.} \end{cases}$$

This implies that, for a quality $\omega > p^*$ with $\omega \in \mathcal{C}^k$, the buyer perceives the probability of trade at price p to be

$$\Pr(\widehat{a \leq p} | \omega \in \mathcal{C}^k) = \int_{a=0}^p h_{\mathcal{C}^k}(a) da = \begin{cases} 1 & \text{if } c_k < p; \\ \frac{\tilde{G}(p) - \tilde{G}(c_{k-1})}{\tilde{G}(c_k) - \tilde{G}(c_{k-1})} & \text{if } c_{k-1} < p \leq c_k; \\ 0 & \text{if } p < c_{k-1}. \end{cases}$$

Using the perceived probability of trade as a function of price p , and letting $k(p)$ be such that $p \in (c_{k(p)-1}, c_{k(p)})$ for $p > p^*$, the following lemma derives the perceived expected




Figure 2: Expected utility for different values of p^* .

payoff as a function of p .

Lemma 1 *Let $v(\mathcal{C}_j) := \mathbb{E}[\omega | \omega \in \mathcal{C}_j] + b$. The perceived expected payoff is*

$$\pi^{CE}(p|p^*) = \begin{cases} G(p) (\mathbb{E}[\omega | \omega \leq p] + b - p) & \text{if } p \leq p^* \\ G(p^*) (\mathbb{E}[\omega | \omega \leq p^*] + b - p) \\ \quad + \sum_{k=1}^{k(p)-1} (G(c_k) - G(c_{k-1})) (v(\mathcal{C}^k) - p) \\ \quad + (\tilde{G}(p) - \tilde{G}(c_{k(p)-1})) \frac{G(c_{k(p)}) - G(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} (v(\mathcal{C}^{k(p)}) - p) & \text{if } p > p^*. \end{cases}$$

As an illustration consider the case of a uniform quality distribution g , i.e., $G(p) = p$ for all $p \in [0, 1]$, a uniform experimentation distribution f_{p^*} above p^* , i.e., $F_{p^*}(p) = p/(1 - p^*)$ for all $p \in [p^*, 1]$. We assume κ high enough to induce a single analogy class above p^* . Figure 1 illustrates the payoff function for $b = 0.3$ and four different values of p^* . Note that when $p^* = 0.3$ the buyer will perceive utility to be maximized by $p \approx 0.38$. However, when $p^* = 0.38$ it will seem to the buyer that utility is maximized at $p = 0.3$. This suggests a price cycle, which we will explore in-depth in what follows.

Dynamics. Letting p_t^* denote the price quoted by buyers of generation t when not trembling, our dynamic system (as described in subsection 2.5) is completely characterized by the initial value of this price p^0 and the recursive condition

$$p_{t+1}^* = \arg \max_p \pi^{CE}(p; p_t^*).$$

We are interested in understanding the sequence p_t^* , and our main result will establish that this sequence must be cyclical no matter what p^0 is.

4.2 Results

4.2.1 Nash Benchmark

In a Nash equilibrium buyers have correct expectations about the mapping between quality and ask price. They maximize³¹

$$\pi^{NE}(p) = \int_{\omega=0}^p (\omega + b - p) g(\omega) d\omega = G(p) (\mathbb{E}[\omega | \omega \leq p] + b - p).$$

We have the following result, whose proof is in the Appendix.

Proposition 5 *There exists a unique Nash equilibrium in which the bid price p^{NE} of uninformed buyers is uniquely defined by*

$$\frac{g(p^{NE})}{G(p^{NE})} = \frac{1}{b}.$$

4.2.2 Learning and Cycling with Categories

Our main result is that the sequence of p_t^* in the categorical learning model has no rest point and must cycle over finitely many values $p^{(1)}, \dots, p^{(m)}$, one of them being p^{NE} as previously characterized, and the others being above p^{NE} . In order to establish this, we first derive three properties related to how p_{t+1}^* varies with p_t^* depending on whether p_t^* is below, above, or equal to p^{NE} . These properties are referred to as lemmata and are proven in the Appendix.

The first property demonstrates that if at time t it is the case that $p_t^* = p^{NE}$, then at time $t + 1$ uninformed buyers bid $p_{t+1}^* > p^{NE}$.

Lemma 2 *If $p_t^* = p^{NE}$ then*

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p | p^{NE}) > p^{NE}.$$

The next property demonstrates that if at time t it is the case that $p_t^* > p^{NE}$, then at time $t + 1$ buyers either bid $p_{t+1}^* = p^{NE}$ or $p_{t+1}^* > p_t^*$.

Lemma 3 *If $p_t^* > p^{NE}$, then either*

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p | p_t^*) = p^{NE}$$

or

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p | p_t^*) > p_t^*.$$

³¹In the case of a uniform quality distribution g this becomes $\pi^{NE}(p) = p \cdot (\frac{p}{2} + b - p) = p(b - \frac{p}{2})$, and so the solution for buyers is $p^{NE} = b$.

The third property demonstrates that if at time t it is the case that $p_t^* < p^{NE}$, then at time $t + 1$ buyers bid $p_{t+1}^* > p_t^*$.

Lemma 4 *If $p_t^* < p^{NE}$, then*

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p|p_t^*) > p_t^*.$$

Roughly, these three properties can be understood as follows. As already mentioned, categorical reasoning induces uninformed buyers to correctly infer that the quality corresponding to an ask price a below p^* is a . On the other hand, the coarse bundling for ask prices above p^* leads uninformed buyers to incorrectly infer that ask prices slightly above p^* are associated with an average quality that lies strictly above p^* . Thus, a buyer would choose a bid price strictly above p^* whenever $p^* \leq p^{NE}$ as she would incorrectly perceive a jump in quality when increasing slightly the bid price above p^* (and any bid price below p^* would rightly be perceived to be suboptimal). This is in essence the content of lemmata 4 and 2. By contrast, when $p^* > p^{NE}$, the best bid price below p^* is rightly perceived to be p^{NE} and the same logic leads the uninformed buyer to either choose p^{NE} or a bid price strictly above p^* with the aim of taking advantage of the jump in the perceived quality when the ask price lies above p^* .

The above properties immediately imply that the price dynamics has no rest point, i.e., there is no p_t^* such that

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p|p_t^*) = p_t^*.$$

To see this, assume by contradiction that p^* is a rest point. By Lemma 4, it cannot be that $p^* < p^{NE}$ since $p_t^* = p^* < p^{NE}$ would imply that $p_{t+1}^* > p_t^* = p^*$. By Lemma 2, it cannot be that $p^* = p^{NE}$ since $p_t^* = p^*$ would imply that $p_{t+1}^* > p^{NE}$. Finally, by Lemma 3, it cannot be that $p^* > p^{NE}$ since $p_t^* = p^*$ would imply either that $p_{t+1}^* > p_t^*$ or that $p_{t+1}^* = p^{NE}$ and thus $p_{t+1}^* \neq p_t^*$ (given that $p_t^* = p^* \neq p^{NE}$). Even though there is no rest point, we can establish that there is a price cycle (making use of Lemmas 2-4), that consists of the Nash price and one or more prices above the Nash price.

Proposition 6 *There exists an increasing sequence $(p^{(1)}, \dots, p^{(\tau)})$ with $\tau \geq 2$ and $p^{(1)} = p^{NE}$ such that if $p_t^* = p^{(i)}$ for $i \in \{1, \dots, \tau - 1\}$ then $p_{t+1}^* = p^{(i+1)}$, and if $p_t^* = p^{(\tau)}$ then $p_{t+1}^* = p^{(1)}$. Moreover, the dynamic converges to the set $\{(p^{(1)}, \dots, p^{(\tau)})\}$ from any initial price $p_0 \in [0, 1]$.*

Proof. Assume, to derive a contradiction, that the sequence p_t^* is monotonic. Lemmata 2–4 imply that

$$p_{t+1}^* = \arg \max_{p \in [0,1]} \pi^{CE}(p|p_t^*) > p_t^*$$

for all t . Since $p_t^* \leq 1$ for all t , it follows that $p_t^* \rightarrow \bar{p}$ for some $\bar{p} > p^{NE}$ as $t \rightarrow \infty$. (To see that there is a $\bar{p} > p^{NE}$ note that if $p_1^* \geq p^{NE}$ then $p_t^* \geq p^{NE}$ for all t .) This implies $|p_{t+1}^* - p_t^*| \rightarrow 0$, which, by continuity of $\pi^{CE}(p|p_t^*)$, implies $|\pi^{CE}(p_{t+1}^*|p_t^*) - \pi^{CE}(p_t^*|p_t^*)| \rightarrow 0$. Since $\pi^{CE}(p|p_t^*) = \pi^{NE}(p)$ for $p \in [0, p_t^*]$, we have $|\pi^{CE}(p_{t+1}^*|p_t^*) - \pi^{NE}(p_t^*)| \rightarrow 0$, and consequently $\pi^{CE}(p_{t+1}^*|p_t^*) \rightarrow \pi^{NE}(\bar{p})$. Since the Nash equilibrium p^{NE} is unique it holds that $\pi^{NE}(p^{NE}) > \pi^{NE}(\bar{p})$, and since $\pi^{CE}(p|p_t^*) = \pi^{NE}(p)$ for $p \in [0, p_t^*]$ we get

$$\pi^{CE}(p_{t+1}^*|p_t^*) \rightarrow \pi^{NE}(\bar{p}) < \pi^{NE}(p^{NE}) = \pi^{CE}(p^{NE}|p_t^*).$$

This is in contradiction to $p_{t+1}^* = \arg \max_{p \in [0, 1]} \pi^{CE}(p|p_t^*)$. We conclude that the sequence p_t^* is not monotonic. Lemmata 2–4 imply that it must be cyclical, consisting of cycles with p^{NE} and one or more price above p^{NE} .

Note that the preceding argument can be used to show, that starting at $p_1^* \geq p^{NE}$ there is convergence to the cycle, from which there is no escape. To see this, suppose (to obtain a contradiction) that there is some $p_1^* > p^{NE}$ that does not belong to the cycle (i.e., $p_i^* \neq p^{(1)}$ for all $i \in \{1, \dots, \tau\}$), from which there is no convergence to the cycle. This means that $p_{t+1}^* > p_t^*$ for all t and $p_t^* \rightarrow \bar{p}$ for some $\bar{p} \in [p^{NE}, p^{(\tau)}]$ as $t \rightarrow \infty$. It remains to show that starting at $p_1^* < p^{NE}$ there is convergence to the set $[p^{NE}, 1]$, which is established in Lemma A2 in the Appendix. ■

4.3 Discussion

4.3.1 Coarse categories below p^*

We have assumed above that the buyers perfectly understand how seller types below p^* behave but have a coarse understanding of how seller types above p^* behave. We can instead consider the case where a categorization is used also to bundle types below p^* . The previous assumption that buyers perfectly understand behavior of types below p^* may then be viewed as a limiting case of an infinitely fine-grained categorization below p^* . In a precursor of this paper, devoted solely to the adverse selection game (Jehiel and Mohlin 2021), we show that cycling is a robust phenomenon with a finite number of categories below p^* .³² We also discuss robustness of our results when (a) gains from trade are multiplicative rather than additive (b) some buyers are rational (c) different buyers may use different categorizations, (d) there is buyer competition. Additionally, we compare our results with those obtained in the case of a centralized trading mechanism.

³²The set up of Jehiel & Mohlin (2021) differs slightly from the current set-up in two ways: (a) buyers categorize ask prices to predict the associated quality, instead of categorizing types to predict ask prices; (b) there is no experimentation, instead a small fraction of buyers are assumed to be informed about the quality of their matched seller.

4.3.2 Mixed strategies

In the above analysis, we have considered pure strategies on the buyer side. Could it be that by allowing mixing on the buyer side, we restore the existence of steady states? We suspect that if we stick to our assumption that (at least) every ω weakly below the support of the bid price strategy of the buyers would be treated as a singleton analogy class, then there is no such steady state. Suppose that the lower bound \underline{p} of the support of buyer' strategy is strictly lower than p^{NE} . By the logic of lemma 4 the best response would be strictly above \underline{p} , so that \underline{p} could not be part of the support in the steady state. Suppose instead that $\underline{p} > p^{NE}$. (a) If the support is coarsely categorized in a neighborhood of \underline{p} then the best response is either strictly above \underline{p} , or equal to p^{NE} . (b) If the support is finely categorized in a neighborhood of \underline{p} each $\omega \leq \underline{p} + \delta$ (for some $\delta > 0$) is perfectly distinguished, then the best response is either strictly above $\underline{p} + \delta$, or equal to p^{NE} . In either case (a) or (b) \underline{p} cannot be the lower bound of a steady state support. Altogether, this is suggestive that it would not be possible to support a steady state even allowing for mixing on the bidding price. We leave for future research whether by allowing for mixing on the analogy partitions, one could support a steady state.

5 Conclusion

In this paper, we have taken inspiration from the bias-variance trade-off to put structure on the analogy partitions used by players in games with multiple situations. Roughly, analogy classes should contain relatively similar situations and they should be visited frequently enough. We have applied the approach to several classic interactions including the chainstore game (Selten, 1978) and adverse selection games (Akerlof 1970), and analyzed how we could obtain predictions closer to what intuition (or experimental evidence) would suggest. We hope that the approach can be applied more broadly in future works.

6 References

- Akerlof, G. A. (1970). The market for lemons: Quality and the market mechanism. *Quarterly. Journal Economics* 84, 488–500.
- Anderson, J. R. (1991). The adaptive nature of human categorization, *Psychological Review* 98(3), 409-429.
- Arad, A., & Rubinstein, A. (2019). Multidimensional reasoning in games: framework, equilibrium, and applications. *American Economic Journal: Microeconomics*, 11(3), 285-318.

- Azrieli, Y. (2009). Categorizing others in a large game, *Games and Economic Behavior* 67(2), 351-362.
- Battigalli, P. (1987). Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali. unpublished undergraduate dissertation, Bocconi University, Milano.
- Berk, Robert (1966). Limiting behavior of posterior distributions when the model is incorrect, *Annals of Mathematical Statistics* 37, 51–58.
- Bohren, J. A., and Hauser, D. N. (2021). Learning with heterogeneous misspecified models: Characterization and robustness. *Econometrica*, 89(6), 3025-3077.
- Chater, N. (1996). Reconciling simplicity and likelihood principles in perceptual organization, *Psychological Review* 103, 566-581.
- Cho, I. K., and Kasa, K. (2015). Learning and model validation. *The Review of Economic Studies* 82(1), 45-82.
- Dekel, E., Fudenberg, D., & Levine, D. K. (2004). Learning to play Bayesian games. *Games and economic behavior*, 46(2), 282-303.
- Dow, J. (1991). Search decisions with limited memory. *The Review of Economic Studies* 58(1), 1-14.
- Erev, I., and Haruvy, E. (2014). Learning and the economics of small decisions. In *The Handbook of Experimental Economics*, Vol 2, eds. Roth, A. E., Kagel, J. Princeton NJ: University Press.
- Esponda, I. (2008). Behavioral equilibrium in economies with adverse selection, *American Economic Review* 98, 1269–91.
- Esponda, I., and Pouzo, D. (2016). Berk–Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* 84(3), 1093–1130.
- Esponda, I., Pouzo, D., and Yamamoto, Y. (2021). Asymptotic behavior of Bayesian learners with misspecified models. *Journal of Economic Theory* 195, 1052–1060.
- Eyster, E., and Rabin, M. (2005). Cursed Equilibrium, *Econometrica* 73(5), 1623–72.
- Fehr, E and Gächter, S. (2000). Cooperation and punishment in public goods experiments, *American Economic Review* 90(4), 980-994.
- Fehr, E and Gächter, S. (2002). Altruistic punishment in humans, *Nature* 415(6868), 137-140.
- Fudenberg, D. and Levine, D.K. (1993). Self-Confirming Equilibrium, *Econometrica* 61, 523-546.

- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*, MIT Press, Cambridge, MA.
- Fudenberg, D., and Levine, D. K. (2006). Superstition and rational learning, *American Economic Review* 96(3), 630-651.
- Fudenberg, D., and Peysakhovich, A. (2016). Recency, records, and recaps: Learning and nonequilibrium behavior in a simple decision problem, *ACM Transactions on Economics and Computation*, 4(4), 1-18.
- Fudenberg, D., Romanyuk, G., and Strack, P. (2017). Active learning with a misspecified prior, *Theoretical Economics* 12(3), 1155–1189.
- Fudenberg, D., and Lanzani, G. (2020). Which misperceptions persist?. Working Paper. Available at SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3709932
- Fryer, R. and Jackson, M. O. (2008). A categorical model of cognition and biased decision making”, *The B.E. Journal of Theoretical Economics (Contributions)* 8(1), 1-42.
- Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*, MIT Press, Cambridge, MA.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural computation*, 4(1), 1-58.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in cognitive science*, 1(1), 107-143.
- Heller, Y., & Winter, E. (2020). Biased-belief equilibrium. *American Economic Journal: Microeconomics*, 12(2), 1-40.
- He, K., and Libgober, J. (2020). Evolutionarily stable (mis)specifications: Theory and applications. arXiv preprint arXiv:2012.15007.
- Jehiel, P. (2005). Analogy-based expectation equilibrium, *Journal of Economic Theory* 123(2), 81–104.
- Jehiel, P. (2022). Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond. Manuscript prepared for the World Congress of the Econometric Society 2020
- Jehiel, P., and Koessler, F. (2008). Revisiting games of incomplete information with analogy-based expectations equilibrium, *Games and Economic Behavior* 62(2), 533–57.

- Jehiel, P., and Mohlin, E. (2021). Cycling and Categorical Learning in Decentralized Adverse Selection Economies. Lund University Working Paper in Economics No 2021:11).
- Jehiel, P. and Samet, D. (2007). Valuation equilibrium, *Theoretical Economics* 2, 163-185.
- Kalai, E. and Neme, A. (1989). The Strength of a Little Perfection, Discussion paper 858, Northwestern University.
- Kreps, D. M., Milgrom, P., Roberts, J., & Wilson, R. (1982). Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2), 245-252.
- Laurence, S. and Margolis, E. (1999). Concepts and cognitive science, in E. Margolis and S. Laurence, eds., *Concepts: Core Readings*, MIT Press, Cambridge, MA, pp. 3-81.
- Mengel, F. (2012). Learning Across Games, *Games and Economic Behavior* 74(2), 601-619.
- Miettinen, T. (2009). The partially cursed and the analogy-based expectation equilibrium, *Economics Letters* 105(2), 162-164.
- Mohlin, E. (2014). Optimal categorization, *Journal of Economic Theory* 152, 356-381.
- Murphy, G. L. (2002). *The Big Book of Concepts*, MIT Press, Cambridge, MA.
- Neyman, A. (1985). Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Economics Letters*, 19(3), 227-229.
- Nyarko, Y. (1991). Learning in mis-specified models and the possibility of cycles. *Journal of Economic Theory* 55, 416-427.
- Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1), 92-100.
- Rubinstein, A. (1998). *Modeling Bounded Rationality*, MIT Press, Cambridge, MA.
- Samuelson, L. (2001). Analogies, adaptation, and anomalies. *Journal of Economic Theory*, 97(2), 320-366.
- Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. MIT press, Cambridge, MA.

- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4, 25-55.
- Selten, R. (1978). The chainstore paradox. *Theory and Decision*, 9(2), 127-159.
- Spiegler, R. (2016). Bayesian networks and boundedly rational expectations, *Quarterly Journal of Economics* 131(3), 1243–1290.
- Xu (2007). Sortal concepts, object individuation, and language, *TREND in Cognitive Science* 11(9), pp. 400-406.

Appendix

A.1 Chainstore Application

Proof of Proposition 3. The proof is similar to that of proposition 3. We focus on the differences.

1. Why \tilde{C} is adjusted to $\tilde{\sigma}_T^m$ for all $m > m^*$ (and all T). Our revised similarity assumptions imply that nodes in $\mathcal{Q}_M^{t,Soft}$ should be bundled with nodes in $\mathcal{Q}_M^{t',Soft}$, and nodes in $\mathcal{Q}_M^{t,Tough}$ should be bundled with nodes in $\mathcal{Q}_M^{t',Tough}$ for $t \neq t'$.
2. Analogy-based expectations.
 - (a) Players have correct expectations at on-path nodes, as in the proof of proposition 3.
 - (b) Players also have correct expectations at off-path nodes in \mathcal{Q}_C^{Soft} and \mathcal{Q}_M^{Soft} , as in the proof of proposition 3.
 - (c) Next consider off-path monopolist nodes in \mathcal{Q}_M^{Tough} . Challengers have erroneous expectations since they bundle together nodes from different time periods. As $\varepsilon^m \rightarrow 0$ the expectations here are determined by behavior at nodes with histories containing a single experimentation with E . The fraction of such nodes at which the monopolist chooses A vanishes as $T \rightarrow \infty$. It follows that as T gets large, the challenger will expect that F is chosen with a probability close to 1.
 - (d) It remains to check the monopolist's expectations at off-path nodes in \mathcal{Q}_C^{Tough} . At all such nodes the challenger plays O unless trembling. Hence the monopolist has correct expectations.
3. Verify that $\tilde{\sigma}_T^m$ induces a ε_T^m -best-response given the analogy-based expectations. We have found that the challengers have correct expectations and it is easy to see that they best-responds to the monopolist's strategy, so we focus on the monopolist.
 - (a) Monopolist at an off-path node in \mathcal{Q}_M^{Tough} . By playing F , the monopolist correctly expects that with a probability close to 1, a string of O occur from then on until the end of the game. (Same belief as in the proof of proposition 3 but now it is a correct belief.) By playing A , the monopolist correctly expects a string of (E, A) until the end of the game (as in the proof of proposition 3). The time period $t \leq T - k^*$ where the incentive to take F is weakest is $t = T - k^*$. Taking F not unprofitable if

$$u_M(E, F) + k^* u_M(O) \geq (k^* + 1) u_M(E, A),$$

which is satisfied by the definition of k^* . At later time periods taking A is strictly profitable.

- (b) Monopolist at an off-path node in Q_M^{Soft} . The monopolist (correctly) expects the challengers to play E in all subsequent periods and best-responds by playing A from now until the end of the game, as in the proof of proposition 3.
- (c) Challenger at an off-path node in Q_M^{Tough} . Here, the challenger will expect that E is met by F with a probability close to 1 (as T gets large), hence plays O .
- (d) Challenger at an off-path node in Q_M^{Soft} . Here the challenger has correct expectations, hence plays E .

■

A.2 Adverse Selection Application

A.2.1 Nash equilibrium

Proof of Proposition 5. Note that

$$\begin{aligned} \frac{\partial}{\partial p} (\mathbb{E} [\omega | \omega \leq p]) &= \frac{1}{G(p)} pg(p) - \left(\int_{\omega=0}^p \omega g(\omega) d\omega \right) \frac{g(p)}{G(p)^2} \\ &= \frac{g(p)}{G(p)} \left(p - \int_{\omega=0}^p \omega \frac{g(\omega)}{G(p)} d\omega \right) \\ &= \frac{g(p)}{G(p)} (p - \mathbb{E} [\omega | \omega \leq p]). \end{aligned}$$

Thus

$$\begin{aligned} \frac{\partial}{\partial p} \pi^{NE}(p) &= g(p) (\mathbb{E} [\omega | \omega \leq p] + b - p) + G(p) \left(\frac{\partial}{\partial p} (\mathbb{E} [\omega | \omega \leq p]) - 1 \right) \\ &= g(p) (\mathbb{E} [\omega | \omega \leq p] + b - p) + G(p) \left(\frac{g(p)}{G(p)} (p - \mathbb{E} [\omega | \omega \leq p]) - 1 \right) \\ &= g(p) (\mathbb{E} [\omega | \omega \leq p] + b - p) + g(p) (p - \mathbb{E} [\omega | \omega \leq p]) - G(p) \\ &= g(p) b - G(p), \end{aligned}$$

and so the first-order condition of $\max_p \pi^{NE}(p)$ is

$$\frac{g(p)}{G(p)} = \frac{1}{b},$$

and the second-order condition is satisfied in virtue of the assumption that $|g'(p)| < g(p)$. Notice that $\lim_{p \rightarrow 0} \frac{g(p)}{G(p)} = \infty$ and $\frac{g(1)}{G(1)} = g(1)$. Hence, by the assumption that $g(1) < 1/b$ and $\frac{\partial}{\partial p} \left(\frac{g(p)}{G(p)} \right) < 0$, the first-order condition has a unique solution that is interior. ■

A.2.2 Deriving Perceived Expected Payoff

Proof of Lemma 1. The perceived expected payoff is

$$\begin{aligned}\pi^{CE}(p|p^*) &= \int_0^{p^*} \Pr(\widehat{a \leq p}|\omega) (\omega + b - p) g(\omega) d\omega \\ &\quad + \int_{p^*}^1 \Pr(\widehat{a \leq p}|\omega \in \mathcal{C}^k) (\omega + b - p) g(\omega) d\omega,\end{aligned}$$

where

$$\begin{aligned}&\int_0^{p^*} \Pr(\widehat{a \leq p}|\omega) (\omega + b - p) g(\omega) d\omega \\ &= \int_0^{p^*} \mathbb{I}_{\{\omega \leq p\}} (\omega + b - p) g(\omega) d\omega \\ &= \begin{cases} \int_0^p (\omega + b - p) g(\omega) d\omega & \text{if } p < p^* \\ \int_0^{p^*} (\omega + b - p) g(\omega) d\omega & \text{if } p \geq p^* \end{cases} \\ &= \begin{cases} G(p) \int_0^p \omega g(\omega|\omega \leq p) d\omega + (b-p) G(p) & \text{if } p < p^* \\ G(p^*) \int_0^{p^*} \omega g(\omega|\omega \leq p^*) d\omega + (b-p) G(p) & \text{if } p \geq p^* \end{cases} \\ &= \begin{cases} G(p) (\mathbb{E}[\omega|\omega \leq p] + b - p) & \text{if } p < p^* \\ G(p) (\mathbb{E}[\omega|\omega \leq p^*] + b - p) & \text{if } p \geq p^* \end{cases}.\end{aligned}$$

and, writing $i(p)$ for the analogy class that contains $\omega = p$,

$$\begin{aligned}&\int_{p^*}^1 \Pr(\widehat{a \leq p}|\omega \in \mathcal{C}^k) (\omega + b - p) g(\omega) d\omega \\ &= \sum_{k=1}^K \left(\int_{c_{k-1}}^{c_k} \left(\int_{a=c_{k-1}}^p h_{\mathcal{C}_j}(a) da \right) (\omega + b - p) g(\omega) d\omega \right) \\ &= \sum_{k=1}^{k(p)-1} \left(\int_{c_{k-1}}^{c_k} (\omega + b - p) g(\omega) d\omega \right) \\ &\quad + \int_{c_{k(p)-1}}^{c_{k(p)}} \left(\frac{\tilde{G}(p) - \tilde{G}(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} \right) (\omega + b - p) g(\omega) d\omega \\ &= \sum_{k=1}^{k(p)-1} \left((G(c_k) - G(c_{k-1})) \int_{c_{j-1}}^{c_j} (\omega + b - p) g(\omega|\omega \in \mathcal{C}_j) d\omega \right) \\ &\quad + \left(\frac{\tilde{G}(p) - \tilde{G}(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} \right) (G(c_{k(p)}) - G(c_{k(p)-1})) \int_{c_{k(p)-1}}^{c_{k(p)}} (\omega + b - p) g(\omega|\omega \in \mathcal{C}^{k(p)}) d\omega \\ &= \sum_{k=1}^{k(p)-1} ((G(c_k) - G(c_{k-1})) (\mathbb{E}[\omega|\omega \in \mathcal{C}^k] + b - p)) \\ &\quad + \left(\tilde{G}(p) - \tilde{G}(c_{k(p)-1}) \right) \frac{G(c_{k(p)}) - G(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} (\mathbb{E}[\omega|\omega \in \mathcal{C}^{k(p)}] + b - p)\end{aligned}$$

■

A.2.3 Preliminary Observations

Note that $\lim_{p \uparrow c_k} \pi^{CE}(p|p^*) = \lim_{p \downarrow c_k} \pi^{CE}(p|p^*)$, for all $i \in \{1, \dots, K-1\}$, implying that $\pi^{CE}(p|p^*)$ is continuous everywhere. Moreover, $\pi^{CE}(p|p^*)$ is piecewise differentiable with points of non-differentiability only at category boundaries. The first derivative at $p \in (c_{k(p)-1}, c_{k(p)})$ is

$$\begin{aligned} \frac{\partial \pi^{CE}(p|p^*)}{\partial p} &= -G(c_{k(p)-1}) - \left(\tilde{G}(p) - \tilde{G}(c_{k(p)-1}) \right) \frac{G(c_{k(p)}) - G(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} \\ &\quad + \tilde{g}(p) \frac{G(c_{k(p)}) - G(c_{k(p)-1})}{\tilde{G}(c_{k(p)}) - \tilde{G}(c_{k(p)-1})} (\mathbb{E}[\omega | \omega \in \mathcal{C}^{k(p)}] + b - p). \end{aligned} \quad (\text{A1})$$

Letting $p \downarrow p^* = c_{k(p)-1}$ we obtain

$$\left. \frac{\partial \pi^{CE}(p|p^*)}{\partial p} \right|_{p \downarrow p^*} = \tilde{g}(p^*) \frac{G(c_1) - G(p^*)}{\tilde{G}(c_1) - \tilde{G}(p^*)} (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p^*) - G(p^*)$$

Note that

$$\begin{aligned} \tilde{g}(p^*) \frac{G(c_1) - G(p^*)}{\tilde{G}(c_1) - \tilde{G}(p^*)} &= \frac{(1 - F_{p^*}(p^*)) (G(c_1) - G(p^*))}{\int_{p^*}^{c_1} g(\omega) (1 - F_{p^*}(\omega)) d\omega} g(p^*) \\ &= \frac{(1 - F_{p^*}(p^*)) \int_{p^*}^{c_1} g(\omega) d\omega}{\int_{p^*}^{c_1} (1 - F_{p^*}(\omega)) g(\omega) d\omega} g(p^*) > g(p^*) \end{aligned}$$

It follows that

$$\left. \frac{\partial \pi^{CE}(p|p^*)}{\partial p} \right|_{p \downarrow p^*} > g(p^*) (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p^*) - G(p^*) \quad (\text{A2})$$

Without taking the limit $p \downarrow p^*$, for $p \in (p_t^*, c_1)$ we can rewrite (A1) as follows

$$\begin{aligned} \frac{\partial \pi^{CE}(p|p^*)}{\partial p} &= \frac{\tilde{G}(p_t^*) G(p_t^*)}{\tilde{G}(c_1) - \tilde{G}(p_t^*)} \left(\frac{G(c_1) - G(p_t^*)}{G(p_t^*)} - \frac{(\tilde{G}(c_1) - \tilde{G}(p_t^*))}{\tilde{G}(p_t^*)} \right) \\ &\quad + \frac{G(c_1) - G(p_t^*)}{\tilde{G}(c_1) - \tilde{G}(p_t^*)} \tilde{g}(p) \left((\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p) - \frac{\tilde{G}(p)}{\tilde{g}(p)} \right) \end{aligned}$$

We note that

$$\begin{aligned}
\frac{(\tilde{G}(c_1) - \tilde{G}(p_t^*))}{\tilde{G}(p_t^*)} &= \frac{\int_{p_t^*}^{c_1} (1 - F_{p^*}(\omega)) g(\omega) d\omega}{\left(\int_0^{p_t^*} (1 - F_{p^*}(\omega)) g(\omega) d\omega\right)} \\
&< \frac{(1 - F_{p^*}(p_t^*)) \int_{p_t^*}^{c_1} g(\omega) d\omega}{(1 - F_{p^*}(p_t^*)) \int_0^{p_t^*} g(\omega) d\omega} \\
&= \frac{\int_{p_t^*}^{c_1} g(\omega) d\omega}{\int_0^{p_t^*} g(\omega) d\omega} = \frac{G(c_1) - G(p_t^*)}{G(p_t^*)}.
\end{aligned}$$

Moreover, we note that

$$\frac{G(c_1) - G(p_t^*)}{\tilde{G}(c_1) - \tilde{G}(p_t^*)} = \frac{\int_{p_t^*}^{c_1} g(\omega) d\omega}{\int_{p_t^*}^{c_1} (1 - F_{p^*}(\omega)) g(\omega) d\omega} \geq 1, \quad (\text{A3})$$

Thus

$$\frac{\partial \pi^{CE}(p|p^*)}{\partial p} \geq \tilde{g}(p) \left((\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p) - \frac{\tilde{G}(p)}{\tilde{g}(p)} \right) \quad (\text{A4})$$

Finally, we can find a lower bound on the second derivative of $\pi^{CE}(p|p^*)$ with respect to p . For $p \in (p_t^*, c_1)$ we have, using (A3),

$$\begin{aligned}
\frac{\partial^2 \pi^{CE}(p|p^*)}{\partial p^2} &= \tilde{g}'(p) \left((\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p) - 2 \frac{\tilde{g}(p)}{\tilde{g}'(p)} \right) \frac{G(c_1) - G(p_t^*)}{\tilde{G}(c_1) - \tilde{G}(p_t^*)} \\
&\geq \tilde{g}'(p) (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p) - 2\tilde{g}(p). \quad (\text{A5})
\end{aligned}$$

A.2.4 Proof of Lemmata 2-4

Proof of Lemma 2. Since $\pi^{CE}(p|p^{NE})$ coincides with $\pi^{NE}(p)$ on $[0, p^*] = [0, p^{NE}]$, the constrained optimal $p \in [0, p^*]$ is at $p = p^* = p^{NE}$. Differentiating π^{CE} at $p \in \mathcal{C}^1 = (p^{NE}, c_1]$, and letting p go to p^{NE} , we obtain, using (A2),

$$\begin{aligned}
\left. \frac{\partial \pi^{CE}(p|p^{NE})}{\partial p} \right|_{p \downarrow p^{NE}} &> g(p^{NE}) (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p^{NE}) - G(p^{NE}) \\
&= G(p^{NE}) \left(\frac{g(p^{NE})}{G(p^{NE})} (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p^{NE}) - 1 \right) \\
&= G(p^{NE}) \left(\frac{1}{b} (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p^{NE}) - 1 \right) \\
&= \frac{G(p^{NE})}{b} (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] - p^{NE}) \\
&= g(p^{NE}) (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] - p^{NE}) > 0.
\end{aligned}$$

Here, the third and fifth equalities use the fact that $g(p^{NE})/G(p^{NE}) = 1/b$. Since $\pi^{NE}(p)$ is continuous, the desired result is implied. ■

Proof of Lemma 3. Since $\pi^{CE}(p|p_t^*)$ coincides with $\pi^{NE}(p)$ on $[0, p_t^*]$, the constrained optimal $p \in [0, p_t^*]$ is at $p = p^{NE} < p_t^*$. Suppose that $\arg \max_{p \in [p_t^*, 1]} \pi^{CE}(p|p_t^*) = p_t^*$ (requiring $\left. \frac{\partial \pi^{CE}(p|p_t^*)}{\partial p} \right|_{p \downarrow p_t^*} \leq 0$). Then by continuity of $\pi^{CE}(p|p_t^*)$, $\arg \max_{p \in [0, 1]} \pi^{CE}(p|p_t^*) = p^{NE} < p_t^*$. ■

Proof of Lemma 4. Suppose, $p_t^* < p^{NE}$. Then the constrained optimal $p \in [0, p_t^*]$ is at p_t^* . Differentiating π^{CE} at $p \in \mathcal{C}^1 = (p_t^*, c_1]$, and letting p go to p_t^* , we obtain, using (A2),

$$\begin{aligned} \left. \frac{\partial \pi^{CE}(p|p_t^*)}{\partial p} \right|_{p \downarrow p_t^*} &> g(p_t^*) (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p_t^*) - G(p_t^*) \\ &= G(p_t^*) \left(\frac{g(p_t^*)}{G(p_t^*)} (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p_t^*) - 1 \right) \\ &\geq G(p_t^*) \left(\frac{g(p^{NE})}{G(p^{NE})} (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p_t^*) - 1 \right) \\ &= G(p_t^*) \left(\frac{1}{b} (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] + b - p_t^*) - 1 \right) \\ &= g(p_t^*) (\mathbb{E}[\omega | \omega \in \mathcal{C}^1] - p_t^*) > 0. \end{aligned}$$

Hence $\left. \frac{\partial \pi^{CE}(p|p_t^*)}{\partial p} \right|_{p \downarrow p_t^*} > 0$. By continuity of π^{CE} we have $\arg \max_{p \in [0, 1]} \pi^{CE}(p|p_t^*) > p_t^*$. ■

A.2.5 Proof of Convergence to Cycle in Proposition 6

In order to establish convergence to the price cycle from initial prices below p^{NE} first we need one more lemma.

Lemma A1 *There is some $\delta > 0$ such that if $p^* \leq p^{NE}$ then $\mathbb{E}[\omega | \omega \in C_1] > p^* + \delta$.*

Proof of Lemma A1. Assume $p^* \leq p^{NE}$. The mass in each analogy class (above p^*) is at least κ . We establish a lower bound on the width of analogy class \mathcal{C}^1 . Let $g^{\min} = \min_{\omega \in [0, 1]} g(\omega)$ and $g^{\max} = \max_{\omega \in [0, 1]} g(\omega)$. By the full-support assumption we have $g^{\min} > 0$. Note that

$$\begin{aligned} \int_{\omega \in \mathcal{C}^1} \mu(\omega) d\omega &= \varepsilon \int_{\omega \in \mathcal{C}^1} \tilde{g}(\omega) d\omega \leq \varepsilon \int_{\omega \in \mathcal{C}^1} g^{\max} d\omega = \varepsilon (c_1 - p^*) g^{\max} \\ &\Rightarrow c_1 - p^* \geq \frac{\kappa}{\varepsilon g^{\max}}. \end{aligned}$$

Using this we can establish a lower bound on the expected quality in analogy class \mathcal{C}^1 .

Define

$$c_1^*(p^*) = \min \left\{ p^* + \frac{\kappa}{\varepsilon g^{\max}}, \frac{1}{2} (p^{NE} + 1) \right\} \leq c_1,$$

implying that

$$c_1^*(p^*) - p^* \geq \min \left\{ \frac{\min \left\{ \kappa, \varepsilon \left(\tilde{G}(1) - \tilde{G}(p^{NE}) \right) \right\}}{\varepsilon g^{\max}}, \frac{1 - p^{NE}}{2} \right\} := M_1.$$

Note that

$$\begin{aligned} \mathbb{E} [\omega | \omega \in \mathcal{C}^1] &\geq \left(1 - \frac{1}{\mu(\mathcal{C}^1)} \int_{\omega=p^*}^{c_1^*(p^*)} g^{\min}(1 - F(c_1^*(p^*))) d\omega \right) \cdot p^* \\ &\quad + \frac{1}{\mu(\mathcal{C}^1)} \int_{\omega=p^*}^{c_1^*(p^*)} g^{\min}(1 - F(c_1^*(p^*))) d\omega \cdot \left(p^* + \frac{c_1^*(p^*) - p^*}{2} \right). \end{aligned}$$

Moreover,

$$\begin{aligned} \int_{\omega=p^*}^{c_1^*(p^*)} g^{\min}(1 - F(c_1^*(p^*))) d\omega &\geq (c_1^*(p^*) - p^*) g^{\min} \left(1 - F \left(\frac{1}{2} (p^{NE} + 1) \right) \right) \\ &\geq M_1 \cdot g^{\min} \left(1 - F \left(\frac{1}{2} (p^{NE} + 1) \right) \right) := M_2. \end{aligned}$$

Thus we have

$$\begin{aligned} \mathbb{E} [\omega | \omega \in \mathcal{C}^1] &\geq \left(1 - \frac{M_2}{\mu(\mathcal{C}^1)} \right) p^* + \frac{M_2}{\mu(\mathcal{C}^1)} \left(p^* + \frac{c_1^*(p^*) - p^*}{2} \right) \\ &= p^* + \frac{M_2}{\mu(\mathcal{C}^1)} \left(\frac{c_1^*(p^*) - p^*}{2} \right) \geq p^* + \frac{M_2}{2} M_1, \end{aligned}$$

or

$$\mathbb{E} [\omega | \omega \in \mathcal{C}^1] \geq p^* + \frac{1}{2} (c_1^*(p^*) - p^*)^2 g^{\min} \left(1 - F \left(\frac{1}{2} (p^{NE} + 1) \right) \right)$$

■

Lemma A2 *Starting at $p_1^* < p^{NE}$ there is convergence to the set $[p^{NE}, 1]$.*

Proof of Lemma A2. Consider $p_t^* < p^{NE}$. By Lemma 4 we know that $p_{t+1}^* = \max_p \pi^{CE}(p | p_t^*) > p_t^*$. Using Lemma A1 in the proof of Lemma 4 we find that the first derivative of $\pi^{CE}(p | p_t^*)$ wrt to p , is bounded above zero as p goes to p_t^* (from above)

$$\left. \frac{\partial \pi^{CE}(p | p_t^*)}{\partial p} \right|_{p \downarrow p_t^*} > g(p_t^*) (\mathbb{E} [\omega | \omega \in \mathcal{C}^1] - p_t^*) > g(p_t^*) \delta > \delta g^{\min} > 0. \quad (\text{A6})$$

Here $g^{\min} = \min_{p \in [0,1]} g(p) > 0$ by the full support assumption. We can also find a lower bound for the second derivative of $\pi^{CE}(p|p_t^*)$ wrt to p . From A5

$$\begin{aligned} \frac{\partial^2 \pi^{CE}(p|p_t^*)}{\partial p^2} &\geq \tilde{g}'(p) (\mathbb{E}[\omega|\omega \in \mathcal{C}^1] + b - p) - 2\tilde{g}(p) \\ &\geq \left(\min_{p \in [0,1]} \tilde{g}'(p) \right) (p_t^* + \delta + b - p) - 2 \left(\min_{p \in [0,1]} \tilde{g}(p) \right). \end{aligned} \quad (\text{A7})$$

Note that

$$p_{t+1}^* \geq \min \left\{ p \in [p_t^*, 1] : \frac{\partial \pi^{CE}(p|p_t^*)}{\partial p} \leq 0 \right\} \quad (\text{A8})$$

The bounds in (A6) and (A7) together imply that the left hand side of (A8) is bounded above p_t^* . ■

ONLINE APPENDIX

Categorization in Games: A Bias-Variance Perspective

Philippe Jehiel and Erik Mohlin

S.1 Categorization Equilibrium and Nash Equilibrium

Here we present examples demonstrating that, CE may not be not outcome equivalent to any NE, for the reason that this would require inconsistent beliefs, as mentioned in Remark 5.

S.1.1 Example where Feedback Differs from the Path of Play

Consider the following game. Player 1 (row) and Player 2 (column) simultaneously choose between actions A and B, with the following outcomes.


	<i>A</i>	<i>B</i>
<i>A</i>	0, 1	1, 0
<i>B</i>	1, 1	0, 0

The unique Nash equilibrium is (B, A) . Note that B is dominated for Player 2 so we can ignore her belief formation. Suppose that the feedback is such that the outcome of the game is reported if and only if it is (B, B) . This means that an entering cohort will see a record consisting entirely of (B, B) outcomes, and those acting as Player 1 will form the belief that Player 2 plays action B with probability 1. The best response is action A . Thus the unique Categorization equilibrium outcome is (A, A) .

S.1.2 Example where Feedback Coincides with the Path of Play

We now turn to an example where the feedback is the path of play. We need to assume that there are three players so that two of them can disagree about what the remaining player does off the path. Consider the following game.

There is a categorization equilibrium involving the strategy profile (C, E, FG) , according to which Player 1 plays C , Player 2 plays E , and Player 3 plays F at the node following A and plays G at the information set following B and D . Only the root node and the node following C are on the path of play. Suppose that Player 1 deems all of Player 3's nodes sufficiently similar to be bundled together in a single analogy class, whereas Player 1 perceives them sufficiently dissimilar to put each of them in a separate category.



To see that this constitutes a categorization equilibrium note that F is dominant for Player 3 at the node following A , and G is dominant for Player 3 at the information set following B and D . Since Player 2 has correct beliefs about the behavior of Player 3 it follows that E is optimal for Player 2. All of Player 3's nodes are reached by a single mistake. Hence Player 1 believes that Player 3 plays F with probability $1/3$ at all of Player 3's nodes (since Player 1 bundles them all together). Player 1 has a correct belief about Player 2's behavior at the on-path node following C . Under these beliefs Player 1 optimally plays C .

In order for Player 2 to take action E she needs to believe that player 3 plays F with at least probability $1/4$ at the information set following B and D . Hence, in a Nash equilibrium implementing the outcome (C, E) Player 3 must follow a strategy that puts at least probability $1/4$ on F at the information set following B and D . In order for Player 1 to take action C rather than action B she needs to believe that player 3 plays F with at most probability $1/5$ at the node following B . Hence in a Nash equilibrium implementing the outcome (C, E) Player 3 must follow a strategy that puts at most probability $1/5$ on F at the information set following B and D . Thus the beliefs required for Players 1 and 2 are inconsistent.

S.2 Public Goods Game

S.2.1 The Game With or Without Punishment

We now apply our approach to public good games. The game has more than two players, but it is straightforward extensions of our basic definitions to the multi-player case. We consider a finitely repeated n -player linear public good game with punishment. The game is repeated T times and players maximize the sum of payoffs. Each round consists of a contribution stage and a punishment stage. Each player holds an endowment of e units. We focus on the simplified case where i can either contribute her entire endowment to

the public good or not contribute at all, $g_i \in G = \{0, e\}$. The payoff of player i is from the contribution stage is

$$u_i^{Cont}(g) = \alpha \sum_{j=1}^n g_j + (e - g_i),$$

where α , with $\frac{1}{n} < \alpha < 1$, captures the marginal per capita return from contributing to the public good. The contribution stage is followed by a punishment stage: each player i can decide whether to punish another player or not. In particular, each player i can subtract punishment points $p_{ij} \in P = \{0, p\}$ from each other player j . For each punishment point a cost of $\beta > 0$ is incurred. This gives rise to the following payoff function,

$$u_i^{Pun}(g, p) = \alpha \sum_{j=1}^n g_j + (e - g_i) - \sum_{j \neq i}^n p_{ji} - \beta \sum_{j \neq i}^n p_{ij}.$$

In the unique SPNE of this game no player contributes, and no player punishes, yielding payoffs of e to everybody. Total payoff is maximized when everyone contributes e , resulting in payoffs of αne .

S.2.2 Zero Contributions without Punishment Stage

We first examine the game without the punishment stage. In this game the stage game payoff of player i is given by u_i^{Cont} . All categorization equilibria are based on the same strategy profile, which coincides with the SPNE, implying that no one contributes. To see why note that in the last round no player contributes, since there is no punishment stage. Suppose there is an equilibrium with full contribution in the second to last round. In this case players on the equilibrium path in the second to last round have a correct belief that no one will contribute in the next round, despite everyone contributing in the second to last round. Thus not contributing in the second to last round is perceived to give a higher payoff, no matter what the off-path expectations about the last round are. Extending this reasoning, we get:

Proposition S1 *Every categorization equilibrium prescribes non-contribution by all players in all rounds.*

Proof. We prove this by induction using the following base case and induction step.

Base case: All categorization equilibria prescribe no-contribution by all players at all information sets in round T .

Induction step: If a categorization equilibrium prescribes no-contribution by all players on the equilibrium path in rounds $\{t + 1, \dots, T\}$ then the categorization equilibrium also prescribes non-contribution by all players on the equilibrium path in round t .

To establish the base case, consider a player i in period T at an information set at which the her strategy prescribes contribution. Regardless of what she expects the other

players to do, no-contribution yields a higher payoff.

To establish the induction step, consider a categorization equilibrium that prescribes no-contribution by all players on the equilibrium path in rounds $\{t + 1, \dots, T\}$. Consider player i in period t at an information set H_t on the equilibrium path (there is only one unless non-degenerate mixed strategies are used). Suppose the strategy prescribes contribution by player i . All on-path nodes are singleton categories. Hence, player i has a correct belief that compliance, i.e. contribution in the current round and no-contribution in the following round yields $\alpha \left(e + \sum_{j \neq i} g_j(H_t) \right) + e(T - t)$. Deviation is expected to yield at least $\alpha \left(\sum_{j \neq i} g_j(H_t) \right) + e + e(T - t)$. The latter is larger than the former. ■

S.2.3 Positive Contributions with Punishment Stage

Our assumption regarding similarity is that players distinguish sharply between two kinds of histories: (i) histories in which all acts of non-contributions were punished (by all those who contributed) and no act of contribution was punished, and (ii) all other histories. A history of either kind is never bundled with a history of the other kind. We also assume that $(n - 1)p \geq e(1 - \alpha)$, meaning that the cost of being punished is high enough relative to the benefit of not contributing. Under these assumptions we can show, that for sufficiently long games (sufficiently large T) there is a categorization equilibrium with contribution in every round, and (off-path) punishment in a no-contribution even except in the last few periods. The construction is similar to the one underlying Proposition 3 for the chainstore game. In the first kind of histories (i) the strategy prescribes contribution and punishment of non-contributors (and only non-contributors), except in the last few rounds in which non-punishment is prescribed. In the second kind of history (ii) the strategy prescribes non-contribution and no punishment. The threat of punishment off-path would not be credible in a standard SPNE. The reason players contribute throughout the interaction in our categorization equilibrium is that the bundling of all off off-path histories of the first kind induce players to believe that they will be punished with probability approaching one (as $T \rightarrow \infty$) if they fail to contribute, even towards the end of the game where in reality they would not be punished. In what follows we provide a detailed description of our construction

Similarity In general it is natural to assume that if two situations $x_i, x'_i \in X_i$ have different actions sets, i.e. $A_i(x_i) \neq A_i(x'_i)$, then any analogy class that contains both situations has maximal dissimilarity. This implies that an adjusted analogy partition will never bundle nodes with different action sets, as in Jehiel (2005). Since contribution decision information sets and punishment information sets have different actions sets any analogy class that contains both kinds of information sets have maximal dissimilarity. Let H^{Con} denote the sets of contribution decision information sets, and let H^{Pun} denote

the set of punishment decision information sets. Since the action sets are different any set that bundles information sets from \mathcal{H}^{Con} and \mathcal{H}^{Pun} have zero similarity. For both \mathcal{H}^{Con} and \mathcal{H}^{Pun} we assume that similarity is mainly determined by whether non-contributors, but not contributors, were punished. Let \mathcal{H}^{Fair} denote the set of information sets with a history such that in each previous round all non-contributors were punished by all contributors, and no contributors were punished.

$$\mathcal{H}^{Fair} = \left\{ \begin{array}{l} \text{In each previous round in the history of } H, \text{ for all } j: \\ g_j = 0 \Rightarrow p_{lj} = p \text{ for all } l \text{ with } g_l = e, \text{ and} \\ g_j = 1 \Rightarrow p_{lj} = 0 \text{ for all } l. \end{array} \right\}$$

Let \mathcal{H}^{Unfair} denote the complement, i.e. information sets with a history such that in at least one previous round there was a non-contributor who was not punished by all contributors, or there was a contributor who was punished. We assume if H and H' belong to X but $H \in \mathcal{H}^{Fair}$ and $H' \in \mathcal{H}^{Unfair}$, then $\xi(X) = 0$. Let

$$\begin{aligned} \mathcal{H}^{Con-Fair} &= \mathcal{H}^{Con} \cap \mathcal{H}^{Fair} \\ \mathcal{H}^{Con-Unfair} &= \mathcal{H}^{Con} \cap \mathcal{H}^{Unfair} \\ \mathcal{H}^{Pun-Fair} &= \mathcal{H}^{Pun} \cap \mathcal{H}^{Fair} \\ \mathcal{H}^{Pun-Unfair} &= \mathcal{H}^{Pun} \cap \mathcal{H}^{Unfair} \end{aligned}$$

Any subset X containing only elements in $\mathcal{H}^{Con-Fair}$ or only elements in $\mathcal{H}^{Con-Unfair}$ satisfies $\xi(X) > 0$. Likewise, any subset X containing only elements in $\mathcal{H}^{Pun-Fair}$ or only elements in $\mathcal{H}^{Pun-Unfair}$ satisfies $\xi(X) > 0$.

Strategy profile We assume

$$(n-1)p \geq e(1-\alpha). \quad (S1)$$

For each $\bar{n} \in \{1, \dots, n-1\}$ let

$$k_{\bar{n}}^* = \min \{k \in \mathbb{N} \text{ such that } (\alpha n + 1)ek^* \geq \beta p \bar{n}\}. \quad (S2)$$

Consider the strategy profile $\hat{\sigma}$, where each individual i plays the following strategy:

- At $H \in \mathcal{H}^{Con-Fair}$, contribute e .
- At $H \in \mathcal{H}^{Con-Unfair}$, do not contribute.
- At $H \in \mathcal{H}^{Pun-Fair}$, where in the immediately preceding contribution stage,
 - i contributed and $\bar{n} \in \{1, \dots, n-1\}$ other players did not contribute: punish if $t \leq T - k_{\bar{n}}^*$, otherwise do not punish.

- i contributed and all other players contributed: do not punish.
 - i did not contribute: do not punish.
- At $H \in \mathcal{H}^{Pun-Unfair}$, do not punish.

On the path of play induced by this strategy profile everyone contributes in all rounds. In case there is non-contribution all contributors punish, except the last period.

Categorization profile Under the categorization profile $\widehat{\mathcal{C}}$, each on-path information set is in a separate analogy class, as usual. Off-path information sets are categorised based on the type of decision (contribution or punishment) and on whether the history was in \mathcal{H}^{Fair} or \mathcal{H}^{Unfair} . Formally, let \mathcal{H}_{-i}^{off} denote the off-path information sets at which players other than i move, define

$$\begin{aligned} \mathcal{C}_{-i}^{Con-Fair} &= \left\{ H \in \mathcal{H}_{-i}^{off} : H \in \mathcal{H}^{Con-Fair} \right\}; \\ \mathcal{C}_{-i}^{Con-Unfair} &= \left\{ H \in \mathcal{H}_{-i}^{off} : H \in \mathcal{H}^{Con-Unfair} \right\}; \\ \mathcal{C}_{-i}^{Pun-Fair} &= \left\{ H \in \mathcal{H}_{-i}^{off} : H \in \mathcal{H}^{Pun-Fair} \right\}; \\ \mathcal{C}_{-i}^{Pun-Unfair} &= \left\{ H \in \mathcal{H}_{-i}^{off} : H \in \mathcal{H}^{Pun-Unfair} \right\}. \end{aligned}$$

Proposition S2 *If (S1) then there exists a T^* such that if $T > T^*$, then $(\hat{\sigma}_T, \widehat{\mathcal{C}})$ is a coarse categorization equilibrium of the chainstore game with T periods, implying that in the absence of mistakes everyone contributes in all rounds.*

Remark S1 *Condition S1 requires that the cost of being punished is high enough relative to the benefit of not contributing, and the definition of k_n^* in (S2) implies that in period $t \leq T - k^*$ the cost of punishing is lower than the loss from others not contributing (in response to non-punishment), whereas in in period $t \leq T - k^*$ the cost of punishing is higher than the loss from others not contributing.*

Proof of Proposition S2. We need to show that for $T > T^*$ there is a sequence $(\hat{\sigma}_T^m)_m$ converging to $\hat{\sigma}_T$, such that $(\hat{\sigma}_T^m, \widehat{\mathcal{C}})$ is an $(\varepsilon^m, \kappa^m)$ -categorization equilibrium for all m . We define $\hat{\sigma}_T^m$ as the strategy profile which at each node puts probability ε^m on the action that $\hat{\sigma}_T$ puts zero probability on. Since there are only two actions at each node this is enough to specify $\hat{\sigma}_T^m$. Since the starting point of $(\varepsilon^m, \kappa^m)$ is arbitrary it is sufficient to show the following: There exists a T^* such that for any $T > T^*$ there is exists an m^* such that if $T > T^*$ and $m > m^*$ then $\hat{\sigma}_T^m$ is an $(\varepsilon_T^m, \kappa_T^m)$ -categorization equilibrium of the chainstore game with T periods.

1. Why $\widehat{\mathcal{C}}$ is adjusted to $\hat{\sigma}_T^m$ for all $m > m^*$ (and all T).

- (a) For any T , if m is large enough, then $\kappa_T^m < (1 - \varepsilon_T^m)^{2nT}$, ensuring that on-path nodes have a mass exceeding the threshold κ_T^m and thus are treated as singleton analogy classes.
- (b) For off-path nodes our similarity assumptions imply that information sets in \mathcal{H}^{Fair} and \mathcal{H}^{Unfair} have to be separated. Likewise, information sets in \mathcal{H}^{Con} and \mathcal{H}^{Pun} have to be separated. No further refinement is allowed (for m large enough).

2. Analogy-based expectations¹

- (a) Players have correct expectations at on-path information sets.
- (b) Players also have correct expectations at off-path information sets in \mathcal{H}^{Unfair} , since after the corresponding histories no one contributes at any information set and no one punishes at any information set.
- (c) Next consider off-path information sets in $\mathcal{H}^{Con-Fair}$. At all such nodes everyone contributes, resulting in correct expectations.
- (d) Finally consider off-path information sets in $\mathcal{H}^{Pun-Fair}$. As $\varepsilon^m \rightarrow 0$ the expectations here are determined by behavior at information sets with histories containing a single act of non-contribution (due to experimentation) in the present round. The fraction of such nodes at which not everyone punishes vanishes as $T \rightarrow \infty$. It follows that as T gets large, expects everyone except the non-contributor to punish with a probability close to 1.

3. Verify that $\hat{\sigma}_T^m$ induces a ε_T^m -best-response given the analogy-based expectations.

- (a) First consider player i at an information set $H \in \mathcal{H}^{Con-Fair}$ (on-path or off-path) in round $t \leq T$. Complying with the proposed strategy profile yields for the continuation

$$EU_i(g_i = e|t) = \alpha ne(T - t + 1).$$

The player believes that if she makes a one-shot deviation then with probability approaching 1 (as $T \rightarrow \infty$) everyone else punishes her, and play remains in $\mathcal{H}^{Con-Fair}$. Hence, a one-shot deviation yields

$$EU_i(g_i = 0|t) = \alpha ne + e(1 - \alpha) + (-(n - 1)p + \alpha ne(T - t))$$

¹In a game with more than two players there are at least two options for how to specify analogy based expectations at off-path information sets. Players may ignore correlation across the other players' actions and form expectations about individual actions (here contributions), or they may form expectations about the distribution of actions (contributions). Here we present results derived for expectations about individual contributions. We can confirm that the results are essentially the same under expectations about the distribution of contributions.

The difference is

$$EU_i(g_i = e|t) - EU_i(g_i = 0|t) = (n-1)p - e(1-\alpha).$$

If (S1) holds then deviation is not profitable.

- (b) Second, consider player i at information set $H \in \mathcal{H}^{Con-Unfair}$ in round $t \leq T$. Complying with the proposed strategy profile yields $EU_i(g_i = 0|t) = (T-t+1)e$. A one-shot deviation yields $EU_i(g_i = 0|t) = \alpha e + (T-t)e$. The former is larger than the latter since $\alpha < 1$.
- (c) Third, consider player i at an information set $H \in \mathcal{H}^{Pun-Fair}$ (on-path or off-path) in round $t \leq T$.
- i. If everyone complied in the contribution stage then (clearly) not punishing is perceived to be optimal.
 - ii. If player i was the only one not to contribute, then (clearly) not punishing is perceived to be optimal.
 - iii. If i contributed and \bar{n} other players did not contribute then i believes that with probability approaching 1 (as $T \rightarrow \infty$) all other contributors will punish the non-contributors, so that punishing yields

$$EU_i(p_{il} = p|t) = -\beta p \bar{n} + \alpha n e (T-t).$$

not punishing leads to \mathcal{H}^{Unfair} , hence yields $EU_i(p_{il} = 0|t) = e(T-t)$. The difference is

$$EU_i(p_{il} = p|t) - EU_i(p_{il} = 0|t) = -\beta p \bar{n} + (\alpha n + 1)e(T-t).$$

This is decreasing in t . For $t = T - k_{\bar{n}}^*$ the difference is

$$EU_i(p_{il} = p|t) - EU_i(p_{il} = 0|t) = -\beta p \bar{n} + (\alpha n + 1)ek_{\bar{n}}^*.$$

By the definition of $k_{\bar{n}}^*$ this non-negative, hence punishing is profitable for $t \leq T - k_{\bar{n}}^*$. For $t > T - k_{\bar{n}}^*$ it is strictly negative so punishing is not profitable.

- (d) Fourth, consider player i at information set $H \in \mathcal{H}^{Pun-Unfair}$ in round $t \leq T$. Clearly, punishing is not perceived as optimal.

■