Cursed Sequential Equilibrium*

Meng-Jhang Fong† Po-Hsuan Lin‡ Thomas R. Palfrey§

January 2023 (revised April 2023)

Abstract

This paper develops a framework to extend the strategic form analysis of cursed equilibrium (CE) developed by Eyster and Rabin (2005) to multi-stage games. The approach uses behavioral strategies rather than normal form mixed strategies, and imposes sequential rationality. We define cursed sequential equilibrium (CSE) and compare it to sequential equilibrium and CE. We provide a general characterization of CSE and apply it to five applications in economics and political science. These applications illustrate a wide range of differences between CSE and Bayesian Nash equilibrium or CE: in signaling games; games with preplay communication; reputation building; sequential voting; and the dirty faces game where higher order beliefs play a key role. Several of these applications illustrate how and why CSE implies systematically cursed behavior in dynamic games of incomplete information with private values, while CE coincides with Bayesian Nash equilibrium for such games.

JEL Classification Numbers: C72, D83
Keywords: Multi-stage Games, Private Information, Cursed Equilibrium, Learning

---

*Grants from the National Science Foundation (SES-0617820) and the Gordon and Betty Moore Foundation (1158) supported this research. We thank participants of the Caltech Proseminar in October 2022, the Caltech Theory Seminar, and Colin Camerer for comments, and thank Matthew Rabin for earlier discussions on the subject during his visit at Caltech as a Moore Distinguished Scholar. We are grateful to Shengwu Li and Shani Cohen for recent correspondence that helped to clarify the differences between the CSE and SCE approaches to the generalization of cursed equilibrium for dynamic games.

†Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125 USA. mjfong@caltech.edu

‡Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125 USA. plin@caltech.edu

§Corresponding Author: Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, California 91125 USA. trp@hss.caltech.edu. Fax: +16263958967 Phone: +16263954088

1 Introduction

Cursed equilibrium (CE) proposed by Eyster and Rabin (2005) is a leading behavioral equilibrium concept that was developed to explain the "winner's curse" and related anomalies in applied game theory. The basic idea behind CE is that individuals do not fully take account of the dependence of other players' strategic actions on private information. Cursed behavior of this sort has been detected in a variety of contexts. Capen et al. (1971) first noted that in oil-lease auctions, "the winner tends to be the bidder who most overestimates the reserves potential" (Capen et al. (1971), p. 641). Since then, this observation of overbidding relative to the Bayesian equilibrium benchmark, which can result in large losses for the winning bidder, has been widely documented in laboratory auction experiments (Bazerman and Samuelson, 1983; Kagel and Levin, 1986; Kagel et al., 1989; Forsythe et al., 1989; Dyer et al., 1989; Lind and Plott, 1991; Kagel and Levin, 2009; Ivanov et al., 2010; Camerer et al., 2016). In addition, the neglect of the connection between the opponents' actions and private information is also found in non-auction environments, such as bilateral bargaining games (Samuelson and Bazerman, 1985; Holt and Sherman, 1994; Carrillo and Palfrey, 2009, 2011), zero-sum betting games with asymmetric information (Rogers et al., 2009; Søvik, 2009), and voting and jury decisions (Guarnaschelli et al., 2000).

While CE provides a tractable alternative to Bayesian Nash equilibrium and can explain some anomalous behavior in games with a winner's-curse structure, a significant limitation is that it is only developed as a strategic form concept for simultaneous-move Bayesian games. Thus, when applying the standard CE to dynamic games, the CE analysis is carried out on the strategic form representation of the game, implying that CE cannot distinguish behavior across dynamic games that differ in their timing of moves but have the same strategic form. That is, players are assumed to choose type-dependent contingent strategies simultaneously and *not* update their beliefs as the history of play unfolds. A further limitation implied by the strategic form approach is that CE and standard Bayesian Nash equilibrium make identical predictions in games with a private-values information structure (Eyster and Rabin (2005), Proposition 2). In this paper we extend the CE in a simple and natural way to multi-stage games of incomplete information. We call the new equilibrium concept *Cursed Sequential Equilibrium* (CSE).

In Section 2, we present the framework and our extension of cursed equilibrium to dynamic games. We consider the framework of multi-stage games with observed actions, introduced by Fudenberg and Tirole (1991b), where players' private information is represented by types, with the assumption that the set of available actions is independent of their types at each public history. Our new solution concept is in the same spirit of the cursed equilibrium—in our model, at each stage, players will (partially) neglect the dependence of the other players' behavioral strategies on their types, by placing some weight on the incorrect belief that all types adopt the average *behavioral* strategy. Specifically, at each public history, this corresponds to the average distribution of actions *given the current belief about others' types at that stage*. Therefore, as players update their beliefs about others' private information via Bayes' rule, but with incorrect beliefs about the other players' behavioral strategies, in later stages this can lead them to have *incorrect* beliefs about the other players' average distribution of actions.

Following Eyster and Rabin (2005)'s notion of cursedness, we parameterize the model by a single parameter $\chi \in [0, 1]$ which captures the degree of cursedness and define fully cursed ($\chi = 1$) CSE analogously to fully cursed ($\chi = 1$) CE. Recall that in a fully cursed ($\chi = 1$) CE, each type of each player chooses a best reply to expected (cursed) equilibrium distribution of other players' actions, averaged over the type-conditional strategies of the other players, with this average distribution calculated using the prior belief on types. Loosely speaking, a player best responds to the average CE strategy of the others. In a χ-CE, players are only partially cursed, in the sense that each player best responds to a χ-weighted linear combination of the *average* χ-CE strategy of the others and the *true* (type-dependent) χ-CE strategy of the others.

The extension of this definition to multi-stage games with observed actions is different from χ-CE in two essential ways: (1) the game is analyzed with behavioral strategies; and (2) we impose sequential rationality and Bayesian updating. In a fully cursed ($\chi = 1$) CSE, (1) implies at every stage t and each public history at t, each type of each player i chooses a best reply to the expected (cursed) equilibrium distribution of other players' stage-t actions, averaged over the type-conditional stage-t behavioral strategies of other players, with this average distribution calculated using i's current belief about types at stage t. That is, player i best responds to the average stage-t CSE strategy of others. Moreover, (2) requires that each player's belief at each public history is derived by Bayes' rule wherever possible, and best replies are with respect to the continuation values computed by using the fully cursed beliefs about the behavioral strategies of the other players in current and future stages.

A χ-CSE, for $\chi < 1$, is then defined in analogously to χ-CE, except for using a χ-weighted linear combination of the *average* χ-CSE *behavioral strategies* of others and the *true* (type-dependent) χ-CSE *behavioral strategies* of others. Thus, similar to the fully cursed CE, in a fully cursed ($\chi = 1$) CSE, each player believes other players' actions at each history are *independent* of their private information. On the other hand, $\chi = 0$ corresponds to the standard sequential equilibrium where players have correct perceptions about other players' behavioral strategies and are able to make correct Bayesian inferences.[1]

After defining the equilibrium concept, in Section 3 we explore some general properties of the model. We first prove the existence of a cursed sequential equilibrium in Proposition 1. Intuitively speaking, CSE mirrors the standard sequential equilibrium. The only difference is that players have incorrect beliefs about the other players' behavioral strategies at each stage since they fail to fully account for the correlation between others' actions and types at every history. We prove in Proposition 2 that the set of CSE is upper hemi-continuous with respect to χ. Consequently, every limit point of a sequence of χ-CSE points as χ converges to 0 is a sequential equilibrium. This result bridges our behavioral solution concept with the standard equilibrium theory. Finally, we also show in Proposition 4 that χ-CSE is equivalent to χ-CE for one-stage games, demonstrating the connection between the two behavioral solutions.

In multi-stage games, cursed beliefs about behavioral strategies will distort the evolution of a player's beliefs about the other players' types. As shown in Proposition 3, a direct consequence of the distortion is that in χ-CSE players tend to update their beliefs about others' types too passively. That is, there is some persistence in beliefs in the sense that at each stage t, each χ-cursed player's belief about any type profile is

[1]For the off-path histories, similar to the idea of Kreps and Wilson (1982), we impose the χ-*consistency* requirement (see Definition 2) so the assessment is approachable by a sequence of totally mixed behavioral strategies. The only difference is that players' beliefs are incorrectly updated by assuming others play the χ-cursed behavioral strategies. Hence, in our approach if $\chi = 0$, a CSE is a sequential equilibrium.

at least χ times the belief about that type profile at stage $t-1$. Among other things, this implies that if the *prior* belief about the types is full support and $\chi > 0$, the full support property will persist at all histories, and players will (possibly incorrectly) believe every profile of others' types is possible at every history.

This dampened updating property plays an important role in our framework. Not only does it contribute to the difference between CSE and the standard CE through the updating process, but it also implies additional restrictions on off-path beliefs. The effect of dampened updating is starkly illustrated in the pooling equilibria of signaling games where every type of sender behaves the same everywhere. In this case, Proposition 5 shows if an assessment associated with a pooling equilibrium is a χ-CSE, then it also a χ'-CSE for all $\chi' \leq \chi$, but it is not necessarily a pooling equilibrium for all $\chi' > \chi$. This contrasts with one of the main results about CE, that if a pooling equilibrium is a χ-CE for some χ, then it is a χ'-CE for all $\chi' \in [0, 1]$ (Eyster and Rabin (2005), Proposition 3).

This suggests that perhaps the dampened updating property is an equilibrium selection device that eliminates some pooling equilibrium, but actually this is *not* a general property. As we demonstrate later, the χ-CE and χ-CSE sets can be non-overlapping, which we illustrate with a variety of applications. The intuition is that in CSE, players generally do not have correct beliefs about the opponents' average behavioral strategies. The pooling equilibrium is just a special case where players have correct beliefs.

In Section 4 we explore the implications of cursed sequential equilibrium with five applications in economics and political science. Section 4.1 analyzes the χ-CSE of signaling games. Besides studying the theoretical properties of pooling χ-CSE, we also analyze two simple signaling games that were studied in a laboratory experiment (Brandts and Holt, 1993). We show how varying the degree of cursedness can change the set of χ-CSE in these two signaling games in ways that are consistent with the reported experimental findings. Next, we turn to the exploration of how sequentially cursed reasoning can influence strategic communication. To this end, we analyze the χ-CSE for a public goods game with communication (Palfrey and Rosenthal, 1991; Palfrey et al., 2017) in Section 4.2, finding that χ-CSE predicts there will be less effective communication when players are more cursed.

Next, in Section 4.3 we apply χ-CSE to the centipede game studied experimentally by McKelvey and Palfrey (1992) where one of the players believes the other player might be an "altruistic" player who always passes. This is a simple reputation-building game, where selfish types can gain by imitating altruistic types in early stages of the game. The public goods application and the centipede game are both private-values environments, so these two applications clearly demonstrate how CSE departs from CE and the Bayesian Nash equilibrium, and shows the interplay between sequentially cursed reasoning and the learning of types in private-value models.

In strategic voting applications, conditioning on "pivotality"—the event where your vote determines the final outcome—plays a crucial role in understanding equilibrium voting behavior. To illustrate how cursedness distorts the pivotal reasoning, in Section 4.4 we study the three-voter two-stage agenda voting game introduced by Ordeshook and Palfrey (1988). Since this is a private value game, the predictions of the χ-CE and the Bayesian Nash equilibrium coincide for all χ. That is, cursed equilibrium predicts no matter how cursed the voters are, they are able to correctly perform pivotal reasoning. On the contrary, our CSE predicts that cursedness will make the voters less likely to vote strategically. This is consistent with the empirical evidence about the prevalence of sincere voting over sequential agendas when inexperienced voters have incomplete information about other voters' preferences (Levine and Plott, 1977; Plott and Levine, 1978; Eckel and Holt, 1989).

Finally, in Section 4.5 we study the relationship between cursedness and epistemic reasoning by considering the two-person dirty faces game previously studied by Weber (2001) and Bayer and Chan (2007). In this game, χ-CSE predicts cursed players are, to some extent, playing a "coordination" game where they coordinate on a specific learning speed about their face types. Therefore, from the perspective of CSE, the non-equilibrium behavior observed in experiments can be interpreted as possibly due to a coordination failure resulting from cognitive limitations.

The cursed sequential equilibrium extends the concept of cursed equilibrium from static Bayesian games to multi-stage games with observed actions. This generalization preserves the spirit of the original cursed equilibrium in a simple and tractable way, and provides additional insights about the effect of cursedness in dynamic games. A contemporaneous working paper by Cohen and Li (2023) is closely related to our paper. That paper adopts an approach based on the coarsening of information sets to define sequential cursed equilibrium (SCE) for extensive form games with perfect recall. The SCE model captures a different kind of cursedness[2] that arises if a player neglects the dependence of other players' unobserved (i.e., either future or simultaneous) actions on the history of play in the game, which is different from the dependence of other players' actions on their type (as in CE and CSE). In the terminology of Eyster and Rabin (2005) (p. 1665), the cursedness is with respect to endogenous information, i.e., what players observe about the path of play. The idea is to treat the unobserved actions of other players in response to different histories (endogenous information) similarly to how cursed equilibrium treats players' types. A two-parameter model of partial cursedness is developed, and a series of examples demonstrate that for plausible parameter values, the model is consistent with some experimental findings related to the failure of subjects to fully take account of unobserved hypothetical events, whereas behavior is "more rational" if subjects make decisions after directly observing such events. At a more conceptual level, our paper is related to several other behavioral solution concepts developed for dynamic games, such as agent quantal response equilibrium (AQRE) (McKelvey and Palfrey, 1998), dynamic cognitive hierarchy theory (DCH) (Lin and Palfrey, 2022; Lin, 2022), and the analogy-based expectation equilibrium (ABEE) (Jehiel, 2005; Jehiel and Koessler, 2008), all of which modify the requirements of sequential equilibrium in different ways than cursed sequential equilibrium.

2 The Model

Since CSE is a solution concept for dynamic games of incomplete information, in this paper we will focus on the framework of multi-stage games with observed actions (Fudenberg and Tirole, 1991b). Section 2.1 defines the formal structure of multi-stage games with observed actions, followed by Section 2.2, where the χ-cursed sequential equilibrium is formally developed.

[2] We illustrate some implications of these differences in the application to signaling games in Section 4.1. For a more detailed discussion of the differences between CSE and SCE, see Fong et al. (2023).

2.1 Multi-Stage Games with Observed Actions

Let $N = \{1, \ldots, n\}$ be a finite set of players. Each player $i \in N$ has a *type* θ_i drawn from a finite set Θ_i. Let $\theta \in \Theta \equiv \times_{i=1}^{n} \Theta_i$ be the type profile and $\theta_{-i} \in \Theta_{-i} \equiv \times_{j \neq i} \Theta_j$ be the type profile without player i. All players share a common (full support) prior distribution $\mathcal{F}(\cdot) : \Theta \to (0, 1)$. Therefore, for every player i, the belief of other players' types conditional on his own type is

$$\mathcal{F}(\theta_{-i}|\theta_i) = \frac{\mathcal{F}(\theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mathcal{F}(\theta'_{-i}, \theta_i)}.$$

At the beginning of the game, players observe their own types, but not the other players' types. That is, each player's type is his own private information.

The game is played in stages $t = 1, 2, \ldots, T$. In each stage, players simultaneously choose actions, which will be revealed at the end of the stage. The feasible set of actions can vary with histories, so games with alternating moves are also included. Let $\mathcal{H}^{t-1}$ be the set of all possible histories at stage t, where $\mathcal{H}^0 = \{h_\emptyset\}$ and $\mathcal{H}^T$ is the set of terminal histories. Let $\mathcal{H} = \cup_{t=0}^{T} \mathcal{H}^t$ be the set of all possible histories of the game, and $\mathcal{H} \backslash \mathcal{H}^T$ be the set of non-terminal histories.

For every player i, the available information at stage t is in $\mathcal{H}^{t-1} \times \Theta_i$. Therefore, player i's information sets can be specified as $\mathcal{I}_i \in \mathcal{Q}_i = \{(h, \theta) : h \in \mathcal{H} \backslash \mathcal{H}^T, \theta_i \in \Theta_i\}$. That is, a type θ_i player i's information set at the public history h^t can be defined as $\bigcup_{\theta_{-i} \in \Theta_{-i}} (h^t, \theta_i, \theta_{-i})$. With a slight abuse of notation, it will be denoted as (h^t, θ_i). For the sake of simplicity, we assume that, at each history, the feasible set of actions for every player is independent of their type and use $A_i(h^{t-1})$ to denote the feasible set of actions for player i at history h^{t-1}. Let $A_i = \times_{h \in \mathcal{H} \backslash \mathcal{H}^T} A_i(h)$ denote player i's feasible actions in all histories of the game and $A = A_1 \times \cdots \times A_n$. In addition, we assume A_i is finite for all $i \in N$ and $|A_i(h)| \geq 1$ for all $i \in N$ and any $h \in \mathcal{H} \backslash \mathcal{H}^T$.

A behavioral strategy for player i is a function $\sigma_i : \mathcal{Q}_i \to \Delta(A_i)$ satisfying $\sigma_i(h^{t-1}, \theta_i) \in \Delta(A_i(h^{t-1}))$. Furthermore, we use $\sigma_i(a_i^t|h^{t-1}, \theta_i)$ to denote the probability player i chooses $a_i^t \in A_i(h^{t-1})$. We use $a^t = (a_1^t, \ldots, a_n^t) \in \times_{i=1}^{n} A_i(h^{t-1}) \equiv A(h^{t-1})$ to denote the action profile at stage t and a_{-i}^t to denote the action profile at stage t without player i. If a^t is the action profile realized at stage t, then $h^t = (h^{t-1}, a^t)$. Finally, each player i has a payoff function $u_i : \mathcal{H}^T \times \Theta \to \mathbb{R}$, and we let $u = (u_1, \ldots, u_n)$

be the profile of payoff functions. A multi-stage game with observed actions, Γ, is defined by the tuple $\Gamma = \langle T, A, N, \mathcal{H}, \Theta, \mathcal{F}, u \rangle$.

2.2 Cursed Sequential Equilibrium

In a multi-stage game with observed actions, a solution is defined by an "assessment," which consists of a (behavioral) strategy profile σ, and a belief system μ. Since action profiles will be revealed to all players at the end of each stage, the belief system specifies, for each player, a conditional distribution over the set of type profiles conditional on each history. Consider an assessment (μ, σ). Following the spirit of the cursed equilibrium, for player i at stage t, we define the *average behavioral strategy profile of the other players* as:

$$\bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \mu_i(\theta_{-i} | h^{t-1}, \theta_i) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i})$$

for any $i \in N$, $\theta_i \in \Theta_i$ and $h^{t-1} \in \mathcal{H}^{t-1}$.

In CSE, players have incorrect perceptions about other players' behavioral strategies. Instead of thinking they are using σ_{-i}, a χ-cursed[3] type θ_i player i would believe the other players are using a χ-weighted average of the average behavioral strategy and the true behavioral strategy:[4]

$$\sigma_{-i}^{\chi}(a_{-i}^t | h^{t-1}, \theta_{-i}, \theta_i) = \chi \bar{\sigma}_{-i}(a_{-i}^t | h^{t-1}, \theta_i) + (1 - \chi) \sigma_{-i}(a_{-i}^t | h^{t-1}, \theta_{-i}).$$

The beliefs of player i about θ_{-i} are updated in the χ-CSE via Bayes' rule, whenever possible, assuming other players are using the χ-cursed behavioral strategy rather than the true behavioral strategy. We call this updating rule the *χ-cursed Bayes' rule.* Specifically, an assessment satisfies the χ-cursed Bayes' rule if the belief system is derived from the Bayes' rule while perceiving others are using σ_{-i}^{χ} rather than σ_{-i}.

Definition 1. (μ, σ) *satisfies χ-cursed Bayes' rule if the following is applied to update*

[3]We assume throughout the paper that all players are equally cursed, so there is no i subscript on χ. The framework is easily extended to allow for heterogeneous degrees of cursedness.

[4]If $\chi = 0$, players have correct beliefs about other players' behavioral strategies at every stage.

the posterior beliefs whenever $\sum_{\theta'_{-i}\in\Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1},\theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1},\theta'_{-i},\theta_i) > 0$*:*

$$\mu_i(\theta_{-i}|h^t,\theta_i) = \frac{\mu_i(\theta_{-i}|h^{t-1},\theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1},\theta_{-i},\theta_i)}{\sum_{\theta'_{-i}\in\Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1},\theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1},\theta'_{-i},\theta_i)}.$$

Let Σ^0 be the set of totally mixed behavioral strategy profiles, and let Ψ^{χ} be the set of assessments (μ,σ) such that $\sigma \in \Sigma^0$ and μ is derived from σ using χ-cursed Bayes' rule.[5] Lemma 1 below shows that another interpretation of the χ-cursed Bayes' rule is that players have correct perceptions about σ_{-i} but are unable to make perfect Bayesian inference when updating beliefs. From this perspective, player i's cursed belief is simply a linear combination of player i's cursed belief at the beginning of that stage (with χ weight) and the Bayesian posterior belief (with $1-\chi$ weight). Because σ is totally mixed, there are no off-path histories.

Lemma 1. *For any* $(\mu,\sigma)\in\Psi^{\chi}$*,* $i\in N$*,* $h^t=(h^{t-1},a^t)\in\mathcal{H}\backslash\mathcal{H}^T$ *and* $\theta\in\Theta$*,*

$$\mu_i(\theta_{-i}|h^t,\theta_i) = \chi\mu_i(\theta_{-i}|h^{t-1},\theta_i) + (1-\chi)\left[\frac{\mu_i(\theta_{-i}|h^{t-1},\theta_i)\sigma_{-i}(a^t_{-i}|h^{t-1},\theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1},\theta_i)\sigma_{-i}(a^t_{-i}|h^{t-1},\theta'_{-i})}\right]$$

Proof. See Appendix A. □

This is analogous to Lemma 1 of Eyster and Rabin (2005). Another insight provided by Lemma 1 is that even if player types are independently drawn, i.e., $\mathcal{F}(\theta) = \Pi_{i=1}^n \mathcal{F}_i(\theta_i)$, players' cursed beliefs about other players' types are generally *not* independent across players. That is, in general, $\mu_i(\theta_{-i}|h^t,\theta_i) \neq \Pi_{j\neq i}\mu_{ij}(\theta_j|h^t,\theta_i)$. The belief system will preserve the independence only when the players are either fully rational ($\chi=0$) or fully cursed ($\chi=1$).

Finally, we place a consistency restriction, analogous to consistent assessments in sequential equilibrium, on how χ-cursed beliefs are updated off the equilibrium path, i.e., when $\sum_{\theta'_{-i}\in\Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1},\theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1},\theta'_{-i},\theta_i) = 0$.

Definition 2. (μ,σ) *satisfies* χ*-consistency if there is a sequence of assessments* $\{(\mu_k,\sigma_k)\}\subseteq\Psi^{\chi}$ *such that* $\lim_{k\to\infty}(\mu_k,\sigma_k) = (\mu,\sigma)$*.*

[5]In the following, we will use $\mu^{\chi}(\cdot)$ to denote the belief system derived under χ-cursed Bayes' Rule. Also, note that both σ^{χ}_{-i} and μ^{χ} are induced by σ; that is, $\sigma^{\chi}_{-i}(\cdot) = \sigma^{\chi}_{-i}[\sigma](\cdot)$ and $\mu^{\chi}(\cdot) = \mu^{\chi}[\sigma](\cdot)$. For the ease of exposition, we drop $[\sigma]$ when it does not cause confusion.

For any $i \in N$, $\chi \in [0,1]$, σ, and $\theta \in \Theta$, let $\rho_i^{\chi}(h^T|h^t, \theta, \sigma_{-i}^{\chi}, \sigma_i)$ be player i's perceived conditional realization probability of terminal history $h^T \in \mathcal{H}^T$ at history $h^t \in \mathcal{H}\backslash\mathcal{H}^T$ if the type profile is θ and player i uses the behavioral strategy σ_i whereas perceives other players' using the cursed behavioral strategy σ_{-i}^{χ}. At every non-terminal history h^t, a χ-cursed player in χ-CSE will use χ-cursed Bayes' rule (Definition 1) to derive the posterior belief about the other players' types. Accordingly, a type θ_i player i's conditional expected payoff at history h^t is given by:

$$\mathbb{E}u_i(\sigma|h^t, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i(\theta_{-i}|h^t, \theta_i) \rho_i^{\chi}(h^T|h^t, \theta, \sigma_{-i}^{\chi}, \sigma_i) u_i(h^T, \theta_i, \theta_{-i}).$$

Definition 3. *An assessment (μ^*, σ^*) is a χ-cursed sequential equilibrium if it satisfies χ-consistency and $\sigma_i^*(h^t, \theta_i)$ maximizes $\mathbb{E}u_i(\sigma^*|h^t, \theta_i)$ for all i, θ_i, $h^t \in \mathcal{H}\backslash\mathcal{H}^T$.*

3 General Properties of χ-CSE

In this section, we characterize some general theoretical properties of χ-CSE. The first result is the existence of the χ-CSE. The definition of χ-CSE mirrors the definition of the sequential equilibrium by Kreps and Wilson (1982)—the only difference is that players in χ-CSE update their beliefs by χ-cursed Bayes' rule and best respond to χ-cursed (behavioral) strategies. Therefore, one can prove the existence of χ-CSE in a similar way as in the standard argument of the existence of sequential equilibrium.

Proposition 1. *For any $\chi \in [0,1]$ and any finite multi-stage game with observed actions, there is at least one χ-CSE.*

Proof. The proof follows a standard argument. See Appendix A for details. □

Let $\Phi(\chi)$ be the correspondence that maps $\chi \in [0,1]$ to the set of χ-CSE. Proposition 1 guarantees $\Phi(\chi)$ is non-empty for any $\chi \in [0,1]$. Because χ-cursed Bayes' rule changes continuously in χ, we can further prove in Proposition 2 that $\Phi(\chi)$ is an upper hemi-continuous correspondence.

Proposition 2. $\Phi(\chi)$ *is upper hemi-continuous with respect to χ.*

Proof. The proof follows a standard argument. See Appendix A for details. □

As shown in Corollary 1, a direct consequence of upper hemi-continuity is that every limit point of a sequence of χ-CSE when $\chi \to 0$ is a sequential equilibrium. This result bridges our behavioral equilibrium concept with standard equilibrium theory.

Corollary 1. *Every limit point of a sequence of χ-CSE with χ converging to 0 is a sequential equilibrium.*

Proof. By Proposition 2, we know $\Phi(\chi)$ is upper hemi-continuous at 0. Consider of a sequence of χ-CSE. As $\chi \to 0$, the limit point remains a CSE, which is a sequential equilibrium at $\chi = 0$. This completes the proof. $\square$

Finally, by a similar argument to Kreps and Wilson (1982), for any $\chi \in [0, 1]$, χ-CSE is also upper hemi-continuous with respect to payoffs. In other words, our χ-CSE preserves the continuity property of sequential equilibrium.

The next result is the characterization of a necessary condition for χ-CSE. As seen from Lemma 1, players update their beliefs more passively in χ-CSE than in the standard equilibrium—they put χ-weight on their beliefs formed in previous stage. To formalize this, we define the *χ-dampened updating property* in Definition 4. An assessment satisfies this property if at *any* non-terminal history, the belief puts at least χ weight on the belief in previous stage—both on and off the equilibrium path. In Proposition 3, we show that χ-consistency implies the χ-dampened updating property.

Definition 4. *An assessment (μ, σ) satisfies the χ-dampened updating property if for any $i \in N$, $\theta \in \Theta$ and $h^t = (h^{t-1}, a^t) \in \mathcal{H} \backslash \mathcal{H}^T$,*

$$\mu_i(\theta_{-i}|h^t, \theta_i) \geq \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i).$$

Proposition 3. *χ-consistency implies χ-dampened updating for any $\chi \in [0, 1]$.*

Proof. See Appendix A. $\square$

It follows that if assessment (μ, σ) satisfies the χ-dampened updating property, then for any player i, any history h^t and any type profile θ, player i's belief about θ_{-i} is bounded by

$$\chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) \leq \mu_i(\theta_{-i}|h^t, \theta_i) \leq 1 - \chi \sum_{\theta'_{-i} \neq \theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i).$$

One can see from this condition that when χ increases, the feasible range of $\mu_i(\theta_{-i}|h^t, \theta_i)$ shrinks, and the restriction on the belief system becomes more stringent. Moreover, if the history h^t is an off-path history of (μ, σ), then this condition characterizes the feasible set of off-path beliefs, which shrinks as χ increases.

An important implication of this observation is that $\Phi(\chi)$ is not lower hemi-continuous with respect to χ. The intuition is that for some χ-CSE that contains off-path histories, the off-path beliefs to support the equilibrium might not be χ-consistent for sufficiently large χ. In this case, the χ-CSE is not attainable by a sequence of χ_k-CSE where χ_k converges to χ from above, causing the lack of lower hemi-continuity.[6]

Lastly, another implication of χ-dampened updating property is that for each player i, history h^t and type profile θ, the belief $\mu_i(\theta_{-i}|h^t, \theta_i)$ has a lower bound that is *independent* of the strategy profile. The lower bound is characterized in Corollary 2. This result implies that when $\chi > 0$, $\mathcal{F}(\theta_{-i}|\theta_i) > 0$ implies $\mu_i(\theta_{-i}|h^t, \theta_i) > 0$ for all h^t, so that if prior beliefs are bounded away from zero, beliefs are always bounded away from 0 as well. In other words, when $\chi > 0$, because of the χ-dampened updating, beliefs will always have full support even if at off-path histories.

Corollary 2. *For any χ-consistent assessment (μ, σ), $i \in N$, $\theta \in \Theta$ and $h^t \in \mathcal{H} \backslash \mathcal{H}^T$,*

$$\mu_i(\theta_{-i}|h^t, \theta_i) \geq \chi^t \mathcal{F}(\theta_{-i}|\theta_i)$$

Proof. See Appendix A. □

If the game has only one stage, then the dampened updating property has no effect, in which case χ-CSE and χ-CE are equivalent solution concepts. This is formally stated and proved in Proposition 4.

Proposition 4. *For any one-stage game and for any χ, χ-CSE and χ-CE are equivalent.*

Proof. For any one-stage game, the only public history is the initial history $h_\emptyset$. Thus, in any χ-CSE, for each player $i \in N$ and type profile $\theta \in \Theta$, player i's belief about

[6] An example is provided in Section 4.1 (see Footnote 7).

other players' types at this history is

$$\mu_i(\theta_{-i}|h_\emptyset, \theta_i) = \mathcal{F}(\theta_{-i}|\theta_i).$$

Since the game has only one stage, the outcome is simply $a^1 = (a_1^1, \ldots, a_n^1)$, the action profile at stage 1. Moreover, given any behavioral strategy profile σ, player i believes a^1 will be the outcome with probability

$$\sigma_i(a_i^1|h_\emptyset, \theta_i) \times \left[\chi\bar{\sigma}_{-i}(a_{-i}^1|h_\emptyset, \theta_i) + (1-\chi)\sigma_{-i}(a_{-i}^1|h_\emptyset, \theta_{-i})\right].$$

Therefore, if σ is the behavioral strategy profile of a χ-CSE in an one-stage game, then for each player i, type $\theta_i \in \Theta_i$ and each $a_i^1 \in A_i(h_\emptyset)$ such that $\sigma_i(a_i^1|h_\emptyset, \theta_i) > 0$,

$$\begin{aligned} a_i^1 \in \underset{a_i^{1\prime} \in A_i(h_\emptyset)}{\operatorname{argmax}} \sum_{\theta_{-i} \in \Theta_{-i}} & \mathcal{F}(\theta_{-i}|\theta_i) \times \\ & \left\{ \sum_{a_{-i}^1 \in A_{-i}(h_\emptyset)} \left[\chi\bar{\sigma}_{-i}(a_{-i}^1|h_\emptyset, \theta_i) + (1-\chi)\sigma_{-i}(a_{-i}^1|h_\emptyset, \theta_{-i})\right] \right\} u_i(a_i^{1\prime}, a_{-i}^1, \theta_i, \theta_{-i}), \end{aligned}$$

which coincides with the maximization problem of χ-CE. This completes the proof. $\square$

From the proof of Proposition 4, one can see that in one-stage games players have *correct* perceptions about the average strategy of others. Therefore, the maximization problem of χ-CSE coincides with the problem of χ-CE. For general multi-stage games, because of the χ-dampened updating property, players will update beliefs incorrectly and thus their perceptions about other players' future moves can also be distorted.

4 Applications

In this section, we will explore χ-CSE in five applications of multi-stage games with observed actions, in order to illustrate the range of effects it can have and to show how it is different from the χ-CE and sequential equilibrium.

Our first application is the sender-receiver signaling game, which is practically the simplest possible multi-stage game. From our analysis, we will see both the theoretical and empirical implications of our χ-CSE.

4.1 Pooling Equilibria in Signaling Games

We first make a general observation about pooling equilibria in multi-stage games. Player j follows a *pooling strategy* if for every non-terminal history, h^t, all types of player j take the same action $a_j^{t+1} \in A_j(h^t)$. Conceptually, since every type of player j takes the same action, players other than j cannot make any inference about j's type from j's actions. A *pooling χ-CSE* is a χ-CSE where every player follows a pooling strategy. Hence, every player has correct beliefs about any other player's future move because every type of every player chooses the same action.

Since in any pooling χ-CSE, players can correctly anticipate other players' future moves no matter how cursed they are, one may naturally conjecture that a pooling χ-CSE is also a χ'-CSE for any $\chi' \in [0, 1]$. As shown by Eyster and Rabin (2005), this is true for one-stage Bayesian games: if a pooling strategy profile is a χ-cursed equilibrium, then it is also a χ'-cursed equilibrium for any $\chi' \in [0, 1]$. Surprisingly, this result does not extend to multi-stage games. Proposition 5 shows if a pooling behavioral strategy profile is a χ-CSE, then it remains a χ'-CSE only for $\chi' \leq \chi$, which is a weaker result than Eyster and Rabin (2005).

This result is driven by the χ-dampened updating property which restricts the set of off-path beliefs. As discussed above, when χ gets larger, the set of feasible off-path beliefs shrinks, eliminating some pooling χ-CSE.

Proposition 5. *A pooling χ-CSE is a χ'-CSE for $\chi' \leq \chi$.*

Proof. See Online Appendix. □

The proof strategy is similar to the one in Eyster and Rabin (2005) Proposition 3. Given a χ-CSE behavioral strategy profile, we can separate the histories into on-path and off-path histories. For on-path histories in a pooling equilibrium, since all types of players make the same decisions, players cannot make any inference about other players' types. Therefore, for on-path histories, their beliefs are the prior beliefs, which are independent of χ. On the other hand, for off-path histories, as shown in Proposition 3, a necessary condition for χ-CSE is that the belief system has to satisfy the χ-dampened updating property. When χ gets larger, this requirement becomes more stringent, and hence some pooling χ-CSE may break down.

Example 1 is a signaling game where the sender has only two types and two messages, and the receiver has only two actions. This example demonstrates the implication of Proposition 5 and shows the lack of lower hemi-continuity; i.e., it is possible for a pooling behavioral strategy profile to be a χ-CSE, but not a χ'-CSE for $\chi' > \chi$. We will also use this example to illustrate how the notion of cursedness in sequential cursed equilibrium proposed by Cohen and Li (2023) departs from CSE.

Example 1. The sender has two possible types drawn from the set $\Theta = \{\theta_1, \theta_2\}$ with $\Pr(\theta_1) = 1/4$. The receiver does not have any private information. After the sender's type is drawn, the sender observes his type and decides to send a message $m \in \{A, B\}$, or any mixture between the two. After that, the receiver decides between action $a \in \{L, R\}$ or any mixture between the two, and the game ends. The game tree is illustrated in Figure 1.

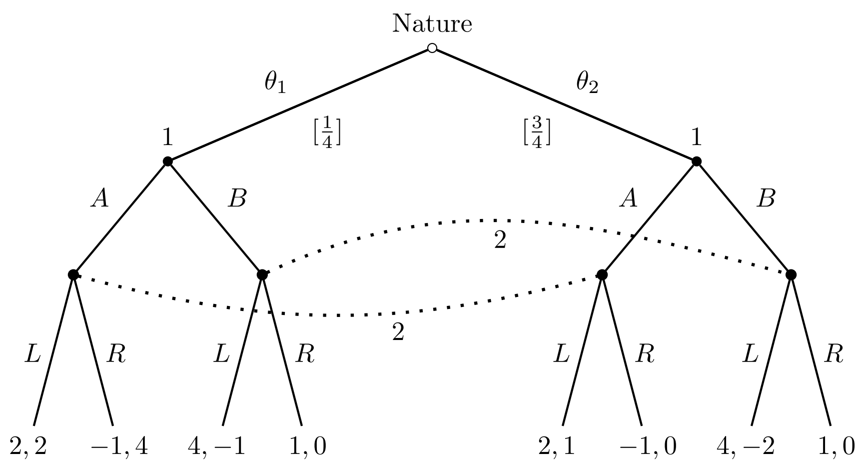

Figure 1: Game Tree for Example 1

If we solve for the χ-CE of the game (or the sequential equilibria), we find that there are two pooling equilibria for every value of χ. In the first pooling χ-CE, both sender types choose A; the receiver chooses L in response to A and R at the off-path history B. In the second pooling χ-CE, both sender types pool at B and the receiver chooses R at both histories. By Proposition 3 of Eyster and Rabin (2005), these two equilibria are in fact pooling χ-CE for all $\chi \in [0, 1]$. The intuition is that in a pooling χ-CE, players are not able to make any inference about other players' types from their actions because the average normal form strategy is the same as the type-conditional normal form strategy. Therefore, their beliefs are independent of χ, and hence a pooling χ-CE will still be an equilibrium for any $\chi \in [0, 1]$.

However, as summarized in Claim 1 below, the χ-CSE imposes *stronger* restrictions than χ-CE in this example, in the sense that when χ is sufficiently large, the second pooling equilibrium cannot be supported as a χ-CSE. The key reason is that when the game is analyzed in its normal form, the χ-dampened updating property shown in Proposition 3 does not have any bite, allowing both pooling equilibria to be supported as a χ-CE for any value of χ. Yet, in the χ-CSE analysis, the additional restriction of χ-dampened updating property eliminates some extreme off-path beliefs, and hence, eliminates the second pooling χ-CSE equilibrium for sufficiently large χ. For simplicity, we use a four-tuple $[(m(\theta_1), m(\theta_2)); (a(A), a(B))]$ to denote a behavioral strategy profile.

Claim 1. *In this example, there are two pure pooling χ-CSE, which are:*

1. $[(A, A); (L, R)]$ *is a pooling χ-CSE for any* $\chi \in [0, 1]$.
2. $[(B, B); (R, R)]$ *with* $\mu_2(\theta_1|A) \in \left[\frac{1}{3}, 1 - \frac{3}{4}\chi\right]$ *is a pooling χ-CSE if and only if* $\chi \le 8/9$.

Proof. See Online Appendix. ☐

From previous discussion, we know in general, the sets of χ-CSE and χ-CE are non-overlapping because of the nature of sequential distortion of beliefs in χ-CSE. Yet, a pooling χ-CSE is an exception. In a pooling χ-CSE, players can correctly anticipate others' future moves, so a pooling χ-CSE will mechanically be a pooling χ-CE. In cases such as this, we can find that χ-CSE is a *refinement* of χ-CE.[7]

Remark. This game is useful for illustrating some of the differences between the notions of "cursedness" in χ-CSE and the sequential cursed equilibrium ((χ_S, ψ_S)-SCE) proposed by Cohen and Li (2023). The first distinction is that the χ and χ_S

[7]Note that the χ-CSE correspondence $\Phi(\chi)$ is not lower hemi-continuous with respect to χ. To see this, we consider a sequence of $\{\chi_k\}$ where $\chi_k = \frac{8}{9} + \frac{1}{9k}$ for $k \ge 1$. From the analysis of Claim 1, we know $[(B, B); (R, R)] \notin \Phi(\chi_k)$ for any $k \ge 1$. However, in the limit where $\chi_k \to 8/9$, $[(B, B); (R, R)]$ with $\mu_2(\theta_1|A) = 1/3$ is indeed a CSE. That is, $[(B, B); (R, R)]$ is not approachable by this sequence of χ_k-CSE.

parameters capture substantively different sources of distortion in a player's beliefs about the other players' strategies. In χ-CSE, the degree of cursedness, χ, captures how much a player neglects the dependence of the other players' behavioral strategies on those players' (exogenous) *private information*, i.e, types, drawn by nature, and as a result, mistakenly treats different types as behaving the same with probability χ. In contrast, in (χ_S, ψ_S)-SCE, the cursedness parameter, χ_S, captures how much a player neglects the dependence of the other players' strategies on future moves of the others, or current moves that are unobserved because of simultaneous play. Thus, it is a neglect related to endogenous information. If player i *observes* a previous move by some other player j, then player i correctly accounts for the dependence of player j's chosen action on player j's private type, as would be the case in χ-CSE only at the boundary where $\chi = 0$.

In the context of pooling equilibria in sender-receiver signaling games, if $\chi_S = 1$, then in SCE the sender believes the receiver will respond the same way both on and off the equilibrium path. This distorts how the sender perceives the receiver's future action in response to an off-equilibrium path message. In χ-CSE, cursedness does not hinder the sender from correctly perceiving the receiver's strategy since the receiver only has one type. Take the strategy profile $[(A, A); (L, R)]$ for example, which is a pooling χ-CSE equilibrium for all $\chi \in [0, 1]$. However, with (χ_S, ψ_S)-SCE, a sender misperceives that the receiver, upon receiving the off-path message B, will, with probability χ_S, take the same action (L) as when receiving the on-path message A. If χ_S is sufficiently high, the sender will deviate to send B, which implies that $[(A, A); (L, R)]$ cannot be supported as an equilibrium when χ_S is sufficiently large ($\chi_S > 1/3$). The distortion induced by χ_S also creates an additional SCE if χ_S is sufficiently large: $[(B, B); (L, R)]$. To see this, if $\chi_S = 1$, then a sender incorrectly believes that the receiver will continue to choose R if the sender deviates to A, rather than switching to L, and hence B is optimal for both sender types. However, $[(B, B); (L, R)]$ is *not* a χ-CSE equilibrium for any $\chi \in [0, 1]$, or a χ-CE in the sense of Eyster and Rabin (2005), or a sequential equilibrium.

In the two possible pooling equilibria analyzed in the last paragraph, the second SCE parameter, ψ_S, does not have any effect, but the role of ψ_S can be illustrated in the context of the $[(B, B); (R, R)]$ sequential equilibrium. This second SCE parameter, ψ_S, is introduced to accommodate a player's possible failure to fully account for the informational content from *observed* events. The larger $(1-\psi_S)$ is, the greater extent a player neglects the informational content of observed actions. Although the parameter ψ_S has a similar flavor to $1-\chi$ in χ-CSE, it is different in a number of ways. In particular this parameter only has an effect via its interaction with χ_S and thus does not independently arise. In the two parameter model, the overall degree of cursedness is captured by the product, $\chi_S(1-\psi_S)$, and thus any cursedness effect of ψ_S is shut down when $\chi_S = 0$. For instance, under our χ-CSE, the strategy profile $[(B,B);(R,R)]$ can only be supported as an equilibrium when χ is sufficiently small. However, $[(B,B);(R,R)]$ can be supported as a (χ_S,ψ_S)-SCE even when $(1-\psi_S)=1$ as long as χ_S is sufficiently small. In fact, when $\chi_S = 0$, a (χ_S,ψ_S)-SCE is equivalent to sequential equilibrium regardless of the value of ψ_S.

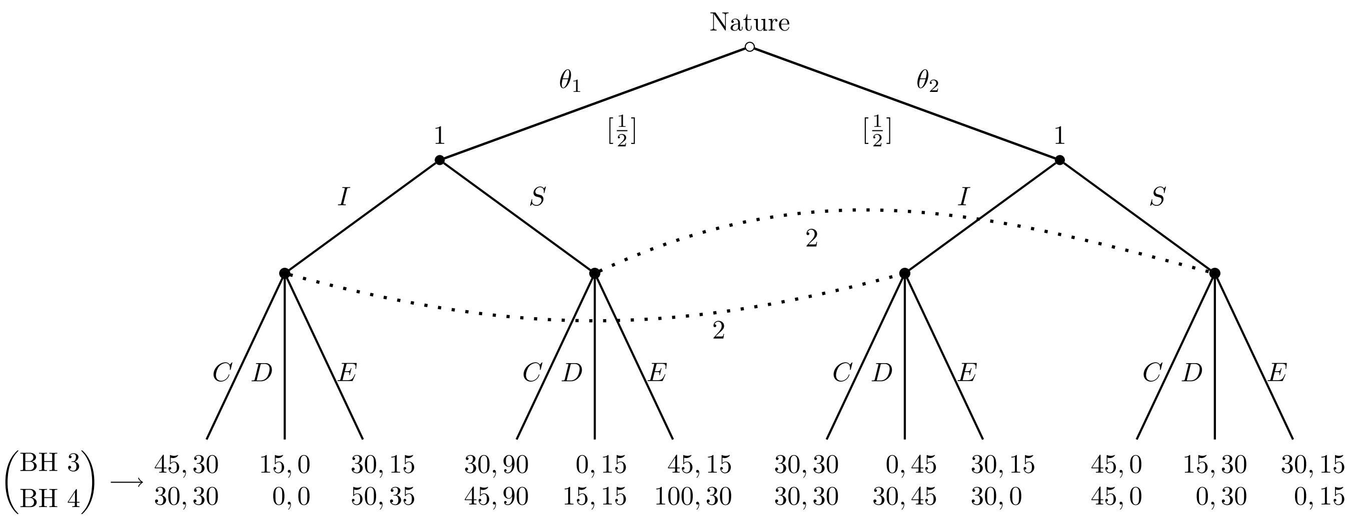

Figure 2: Game Tree for *BH 3* and *BH 4* in Brandts and Holt (1993)

Example 2. Here we analyze two signaling games that were studied experimentally by Brandts and Holt (1993) (*BH 3* and *BH 4*) and show that χ-CSE can help explain some of their findings. In both Game *BH 3* and *BH 4*, the sender has two possible types $\{\theta_1,\theta_2\}$ which are equally likely. There are two messages $m \in \{I, S\}$ available to the sender.[8] After seeing the message, the receiver chooses an action from $a \in$

[8] I stands for "**I**ntuitive" and S stands for "**S**equential but not intuitive", corresponding to the two pooling sequential equilibria of the two games.

$\{C, D, E\}$. The game tree and payoffs for both games are summarized in Figure 2.

In both games, there are two pooling sequential equilibria. In the first equilibrium, both sender types send message I, and the receiver will choose C in response to I and D in response to S. In the second equilibrium, both sender types send message S, and the receiver will choose D in response to I while choose C in response to S. Both are sequential equilibria, in both games, but only the first equilibrium where the sender sends I satisfies the intuitive criterion proposed by Cho and Kreps (1987).

Since the equilibrium structure is similar in both games, the sequential equilibrium and the intuitive criterion predict the behavior should be the same in both games. However, this prediction is strikingly rejected by the data. Brandts and Holt (1993) report that in the later rounds of the experiment, almost all type θ_1 senders send I in Game *BH 3* (97 %), and yet all type θ_1 senders send S in Game *BH 4* (100%). In contrast, type θ_2 senders behave similarly in both games—46.2% and 44.1% of type θ_2 senders send I in Games *BH 3* and *BH 4*, respectively. Qualitatively speaking, the empirical pattern reported by Brandts and Holt (1993) is that *sender type θ_1 is more likely to send I in Game BH 3 than Game BH 4 while sender type θ_2's behavior is insensitive to the change of games.*

To explain this finding, Brandts and Holt (1993) propose a descriptive story based on *naive receivers*. A naive receiver will think both sender types are equally likely, regardless of which message is observed. This naive reasoning will lead the receiver to choose C in both games. Given this naive response, a type θ_1 sender has an incentive to send I in Game *BH 3* and choose S in Game *BH 4*. (Brandts and Holt (1993), p. 284 – 285)

In fact, their story of naive reasoning echoes the logic of χ-CSE. When the receiver is fully cursed (or naive), he will ignore the correlation between the sender's action and type, causing him to not update the belief about the sender's type. Proposition 6 characterizes the set of χ-CSE of both games. Following the previous notation, we use a four-tuple $[(m(\theta_1), m(\theta_2)); (a(I), a(S))]$ to denote a behavioral strategy profile.

Proposition 6. *The set of χ-CSE of Game BH 3 and BH 4 are characterized as below.*

- *In Game BH 3, there are three pure χ-CSE:*

1. $[(I, I); (C, D)]$ *is a pooling* χ*-CSE if and only if* $\chi \leq 4/7$.
2. $[(S, S); (D, C)]$ *is a pooling* χ*-CSE if and only if* $\chi \leq 2/3$.
3. $[(I, S); (C, C)]$ *is a separating* χ*-CSE if and only if* $\chi \geq 4/7$.

- *In Game BH 4, there are three pure* χ*-CSE:*

1. $[(I, I); (C, D)]$ *is a pooling* χ*-CSE if and only if* $\chi \leq 4/7$.
2. $[(S, S); (D, C)]$ *is a pooling* χ*-CSE if and only if* $\chi \leq 2/3$.
3. $[(S, S); (C, C)]$ *is a pooling* χ*-CSE for any* $\chi \in [0, 1]$.

Proof. See Online Appendix. □

As noted earlier for Example 1, by Proposition 3 of Eyster and Rabin (2005), pooling equilibria (1) and (2) in games BH 3 and BH 4 survive as χ-CE for *all* $\chi \in [0, 1]$. Hence, Proposition 6 implies that χ-CSE refines the χ-CE pooling equilibria for larger values of χ. Moreover, χ-CSE actually eliminates *all* pooling equilibria in BH 3 if $\chi > 2/3$. Proposition 6 also suggests that for any $\chi \in [0, 1]$, sender type θ_2 will behave similarly in both games, which is qualitatively consistent with the empirical pattern. In addition, χ-CSE predicts that a highly cursed ($\chi > 2/3$) type θ_1 sender will send different messages in different games—highly cursed type θ_1 senders will send I and S in Games *BH 3* and *BH 4*, respectively. This is consistent with the empirical data.

4.2 A Public Goods Game with Communication

Our second application is a threshold public goods game with private information and pre-play communication, variations of which have been studied in laboratory experiments (Palfrey and Rosenthal, 1991; Palfrey et al., 2017). Here we consider the "unanimity" case where there are N players and the threshold is also N.

Each player i has a private cost parameter c_i, which is independently drawn from a uniform distribution on $[0, K]$ where $K > 1$. After each player's c_i is drawn, each player observes their own cost, but not the others' costs. Therefore, c_i is player i's private information and corresponds to θ_i in the general formulation.[9] The game consists of two stages. After the profile of cost parameters is drawn, the game will proceed to stage 1 where each player simultaneously broadcasts a public message $m_i \in \{0, 1\}$ without any cost or commitment. After all players observe the message profile from this first stage, the game proceeds to stage 2 which is a unanimity threshold public goods game. Player i has to pay the cost c_i if he contributes, but the public good will be provided only if all players contribute. The public good is worth a unit of payoff for every player. Thus, if the public good is provided, each player's payoff will be $1 - c_i$.

If there is no communication stage, the unique Bayesian Nash equilibrium is that no player contributes, which is also the unique χ-CE for any $\chi \in [0, 1]$. In contrast, with the communication stage, there exists an efficient sequential equilibrium where each player i sends $m_i = 1$ if and only if $c_i \leq 1$ and contributes if and only if all players send 1 in the first stage.[10] Since this is a private value game, the standard cursed equilibrium has no bite, and this efficient sequential equilibrium is also a χ-CE for all values of χ, by Proposition 2 of Eyster and Rabin (2005). In the following, we demonstrate that the prediction of χ-CSE is different from CE (and sequential equilibrium).

To analyze the χ-CSE, consider a collection of "cutoff" costs, $\{C_c^\chi, C_0^\chi, C_1^\chi, \ldots, C_N^\chi\}$. In the communication stage, each player communicates the message $m_i = 1$ if and only if $c_i \leq C_c^\chi$. In the second stage, if there are exactly $0 \leq k \leq N$ players sending $m_i = 1$ in the first stage, then such a player would contribute in the second stage if and only if $c_i \leq C_k^\chi$. A χ-CSE is a collection of these cost cutoffs such that the associated strategies are a χ-CSE for the public goods game with communication. The most efficient sequential equilibrium identified above for $\chi = 0$ corresponds to cutoffs with $C_0^0 = C_1^0 = \cdots = C_{N-1}^0 = 0$ and $C_c^0 = C_N^0 = 1$.

There are in fact multiple equilibria in this game with communication. In order to demonstrate how the cursed belief can distort players' behavior, here we will focus on the χ-CSE that is similar to the most efficient sequential equilibrium identified

[9] This application has a continuum of types. The framework of analysis developed for finite types is applied in the obvious way.

[10] One can think of the first stage as a poll, where players are asked the following question: "Are you willing to contribute if everyone else says they are willing to contribute?". The message $m_i = 1$ corresponds to a "yes" answer and the message $m_i = 0$ corresponds to a "no" answer.

above, where $C_0^\chi = C_1^\chi = \cdots = C_{N-1}^\chi = 0$ and $C_c^\chi = C_N^\chi$. The resulting χ-CSE is given in Proposition 7.

Proposition 7. *In the public goods game with communication, there is a χ-CSE where*

1. $C_0^\chi = C_1^\chi = \cdots = C_{N-1}^\chi = 0$, *and*
2. *there is a unique $C^*(N,K,\chi) \leq 1$ s.t. $C_c^\chi = C_N^\chi = C^*(N,K,\chi)$ that solves:*

$$C^*(N,K,\chi) - \chi\left[\frac{C^*(N,K,\chi)}{K}\right]^{N-1} = 1-\chi.$$

Proof. See Online Appendix. □

To provide some intuition, we sketch the proof by analyzing the two-person game, where the χ-CSE is characterized by four cutoffs $\{C_c^\chi, C_0^\chi, C_1^\chi, C_2^\chi\}$, with $C_0^\chi = C_1^\chi = 0$ and $C_c^\chi = C_2^\chi$. If players use the strategy that they would send message 1 if and only if the cost is less than C_c^χ, then by Lemma 1, at the history where both players send 1, player i's cursed posterior belief density would be

$$\mu_i^\chi(c_{-i}|\{1,1\}) = \begin{cases} \chi\cdot\left(\frac{1}{K}\right) + (1-\chi)\cdot\left(\frac{1}{C_c^\chi}\right) & \text{if } c_{-i} \leq C_c^\chi \\ \chi\cdot\left(\frac{1}{K}\right) & \text{if } c_{-i} > C_c^\chi. \end{cases}$$

Notice that cursedness leads a player to put some probability weight on a type that is not compatible with the history. Namely, for χ-cursed players, when seeing another player sending 1, they still believe the other player might have $c_{-i} > C_c^\chi$. When χ converges to 1, the belief simply collapses to the prior belief as fully cursed players never update their beliefs. On the other hand, when χ converges to 0, the belief converges to $1/C_c^\chi$, which is the correct Bayesian inference.

Given this cursed belief density, the optimal cost cutoff to contribute, C_2^χ, solves

$$C_2^\chi = \int_0^{C_2^\chi} \mu_i^\chi(c_{-i}|\{1,1\})dc_{-i}.$$

Finally, at the first stage cutoff equilibrium, the C_c^χ type of player would be indifferent between sending 1 and 0 at the first stage. Therefore, C_c^χ satisfies

$$0 = \left(\frac{C_c^\chi}{K}\right)\left\{-C_c^\chi + \int_0^{C_2^\chi} \mu_i^\chi(c_{-i}|\{1,1\})dc_{-i}\right\}.$$

After substituting $C_c^\chi = C_2^\chi$, we obtain the χ-CSE: $C_c^\chi = C_2^\chi = \frac{K-K\chi}{K-\chi}$.

From this expression, one can see that the cutoff C_c^χ (as well as C_2^χ) is decreasing in χ and K. When $\chi \to 0$, C_c^χ converges to 1, which is the cutoff of the sequential equilibrium. On the other hand, when $\chi \to 1$, C_c^χ converges to 0, so there is no possibility for communication when players are fully cursed. Similarly, when $K \to 1$, C_c^χ converges to 1, which is the cutoff of the sequential equilibrium, while $\lim_{K\to\infty} C_c^\chi = 1-\chi$.

These comparative statics results with respect to χ and K are not just a special property of the $N = 2$ case, but hold for all $N > 1$. Furthermore, there is a similar effect of increasing N that results in a lower cutoff (less effective communication). These properties of $C^*(N,K,\chi)$ are summarized in Corollary 3.

Corollary 3. *The efficient χ-CSE predicts the following comparative statics for all $N \geq 2$ and $K > 1$:*

1. $C^*(N,K,0) = 1$ *and* $C^*(N,K,1) = 0$.

2. $C^*(N,K,\chi)$ *is strictly decreasing in N, K, and χ for any $\chi \in (0,1)$.*

3. *For all $\chi \in [0,1]$,* $\lim_{N\to\infty} C^*(N,K,\chi) = \lim_{K\to\infty} C^*(N,K,\chi) = 1-\chi$.

Proof. See Online Appendix. □

These properties are illustrated in Figure 3. The left panel illustrates the equilibrium condition for C^* in a graph where the horizontal axis is $C \in [0,K]$. We can rewrite the characterization of $C^*(N,K,\chi)$ in Proposition 7 as a solution for C to the following equation: $\frac{1-C}{\chi} = 1 - \left[\frac{C}{K}\right]^{N-1}$. The left panel displays the LHS of this equation, $\frac{1-C}{\chi}$, as the downward sloping line that connects the points $(0, \frac{1}{\chi})$ and $(1,0)$. The RHS is displayed for $N = 2$ and $N = 3$ by the two curves that connect the points $(0,1)$ and $(K,0)$. The equilibrium, $C^*(N,K,\chi)$, is given by the (unique) intersection of the LHS and RHS curves. It is easy to see from this graph

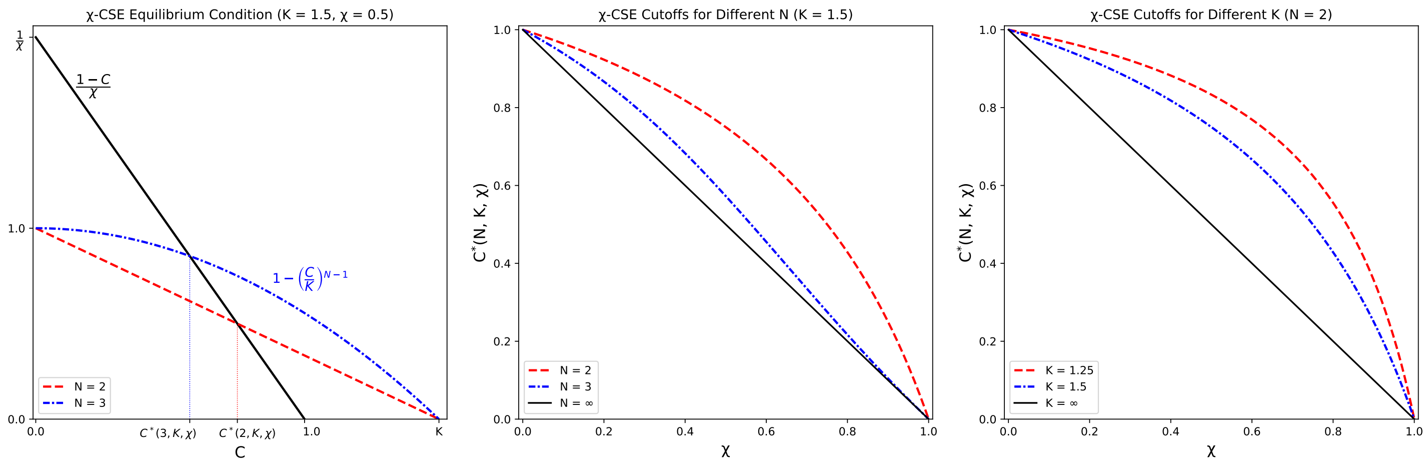

Figure 3: (Left) Illustration of the χ-CSE equilibrium condition when $K = 1.5$ and $\chi = 0.5$. (Middle) The χ-CSE cutoff $C^*(N, K, \chi)$ for $N = 2, 3$ and for $N \to \infty$ when $K = 1.5$. (Right) The χ-CSE cutoff $C^*(N, K, \chi)$ for $K = 1.25, 1.5$ and for $K \to \infty$ when $N = 2$.

that $C^*(N, K, \chi)$ is strictly decreasing in N, K, and χ. When N increases, the RHS increases for all $C \in (0, K)$, resulting in an intersection at a lower value of C. When K increases, again the RHS increases for all $C \in (0, K)$, and also the intercept of the RHS on the horizontal axis increases, leading to a similar effect; and when χ increases, the intercept of the LHS on the horizontal axis decreases, resulting in an intersection at a lower value of C. In addition, when N grows without bound, the RHS approaches a constant function equal to 1 for $C < K$, resulting in a limiting intersection at $C^*(\infty, K, \chi) = 1 - \chi$. This is illustrated in the middle panel of Figure 3, which graphs $C^*(2, 1.5, \cdot)$, $C^*(3, 1.5, \cdot)$, and $C^*(\infty, 1.5, \cdot)$. A similar effect occurs for $K \to \infty$, illustrated in the right panel of Figure 3, which displays $C^*(2, 1.25, \cdot)$, $C^*(2, 1.5, \cdot)$, and $C^*(2, \infty, \cdot)$.

An interesting takeaway of this analysis is that in the public goods game with communication, *cursedness limits information transmission*: χ-CSE predicts when players are more cursed (higher χ), it will be harder for them to effectively communicate in the first stage for efficient coordination in the second stage. Moreover, Corollary 3 shows this χ-CSE varies systematically with *all three parameters of the model*: N, K, *and* χ. In contrast, in the standard χ-CE, players best respond to the *average type-contingent strategy* rather than the average behavioral strategy. Since it is a private value game, players do not care about the distribution of types, only the distribution of actions. Thus, the prediction of standard CE coincides with the equilibrium prediction *for all values of* N, K, *and* χ. This seems behaviorally implausible and is also suggestive of an experimental design that varies the two parameters N and K, since the qualitative effects of changing these parameters are identified.

4.3 Reputation Building: The Centipede Game with Altruists

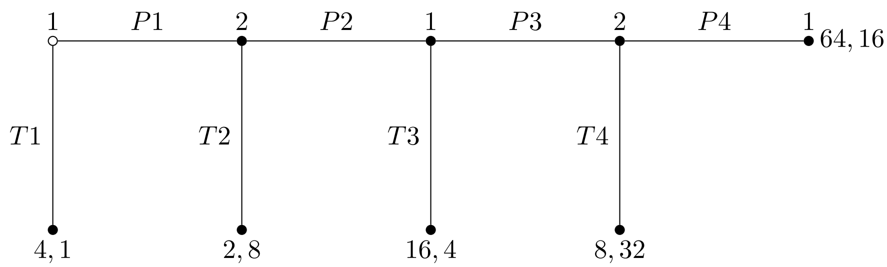

Figure 4: Four-stage Centipede Game

In order to further demonstrate the difference between χ-CE and χ-CSE, in this section we consider a variation of the centipede game with private information, as analyzed in McKelvey and Palfrey (1992) and Kreps (1990). This game is an illustration of reputation-building, where a selfish player imitates an altruistic type in order to develop a reputation for passing, which in turn entices the opponent to pass and leads to higher payoffs.

There are two players and four stages, and the game tree is shown in Figure 4. In stage one, player one can choose either Take ($T1$) or Pass ($P1$). If she chooses action $T1$, the game ends and the payoffs to players one and two are 4 and 1, respectively. If she chooses the action $P1$, the game continues and player two has a choice between take ($T2$) and pass ($P2$). If he chooses $T2$, the game ends and the payoffs to players one and two are 2 and 8, respectively. If he chooses $P2$, the game continues to the third stage where player one chooses between $T3$ and $P3$. Similar to the previous stages, if she chooses $T3$, the payoffs to players one and two are 16 and 4, respectively. If she chooses $P3$, the game proceeds to the last stage where player two chooses between

$T4$ and $P4$. If player two chooses $T4$ the payoffs are 8 and 32, respectively. If player two alternatively chooses $P4$, the payoffs are 64 and 16, respectively.

There are two types of player one, selfish and altruistic. Selfish players are assumed to have a utility function that is linear in their own payoff. Altruistic players are assumed to have a utility function that is linear in the sum of the two payoffs. For the sake of simplicity, we assume that player two has only one type, selfish. The common knowledge probability that player one is altruistic is α. Player one knows her own type, but player two does not. Therefore, player one's type is her private information. In the following, we will focus on the interesting case where $\alpha \leq 1/7$.[11]

Because this is a game of incomplete information with *private values,* the standard χ-CE is equivalent to the Bayesian Nash equilibrium of the game for all $\chi \in [0, 1]$, and yields the same take probabilities as the Bayesian equilibrium. Since altruistic player one wants to maximize the sum of the payoffs, it is optimal for her to always pass. The equilibrium behavior is summarized in Claim 2.

Claim 2. *In the Bayesian Nash equilibrium, selfish player one will choose $P1$ with probability $\frac{6\alpha}{1-\alpha}$ and choose $T3$ with probability 1; player two will choose $P2$ with probability $\frac{1}{7}$ and choose $T4$ with probability 1.*

Proof. See Online Appendix. □

It is useful to see exactly why, in this example (and more generally) the standard χ-CE is the same as the perfect Bayesian equilibrium. In particular, why it is not the case that cursed beliefs will change player two's updating process after observing $P1$ at stage one. Belief updating is *not* a property of the standard χ-CE as the analysis is in the strategic form, and thus is solved as a BNE of the game in the reduced normal form.[12] Table 1 summarizes the payoff matrices in the reduced normal form of centipede game for selfish and altruistic type.

It is easily verified that at the Bayesian Nash equilibrium, selfish player one would choose T_1 with probability $(1-7\alpha)/(1-\alpha)$ and choose P_1T_3 with probability $6\alpha/(1-\alpha)$, while player two would choose T_2 with probability $6/7$.

[11] If $\alpha > 1/7$, player two always chooses $P2$ in the second stage since the probability of encountering altruistic player one is sufficiently high. Selfish player one would thus chooses $P1$ in the first stage and choose $T3$ in the third stage.

[12] The analysis is similar for the unreduced normal form.

Table 1: Reduced Normal Form Centipede Game Payoff Matrix

selfish $(1-\alpha)$	T_2	P_2T_4	P_2P_4
T_1	$4,1$	$4,1$	$4,1$
P_1T_3	$2,8$	$16,4$	$16,4$
P_1P_3	$2,8$	$8,32$	$64,16$

altruistic (α)	T_2	P_2T_4	P_2P_4
T_1	$5,1$	$5,1$	$5,1$
P_1T_3	$10,8$	$20,4$	$20,4$
P_1P_3	$10,8$	$40,32$	$80,16$

To solve the standard χ-CE, let selfish player one choose T_1 with probability p and P_1T_3 with probability $1-p$. Let player two choose T_2 with probability q and P_2T_4 with probability $1-q$. Notice that for player two, P_2P_4 is a dominated strategy and given this, it is also sub-optimal for selfish player one to choose P_1P_3. In this case, selfish player one would choose T_1 if and only if

$$4 \geq 2q + 16(1-q) \iff q \geq 6/7,$$

implying that selfish player one's best response correspondence in the standard cursed analysis coincides with the Bayesian Nash equilibrium analysis. On the other hand, to solve for player two's best responses we need to first solve for the *perceived* strategy. When player two is χ-cursed, he would think that player one is using $\sigma_1^\chi(a|\theta)$ where $a \in \{T_1, P_1T_3, P_1P_3\}$ and $\theta \in \{\text{selfish}, \text{ altruistic}\}$. Player one's true strategy is given in Table 2.

Table 2: Player 1's True Strategy

	player one's type	
$\sigma_1(a|\theta)$	selfish	altruistic
T_1	p	0
P_1T_3	$1-p$	0
P_1P_3	0	1

In this case, player one's *average strategy* is simply:

$$\bar{\sigma}_1(T_1) = (1-\alpha)p, \;\; \bar{\sigma}_1(P_1T_3) = (1-\alpha)(1-p), \;\; \bar{\sigma}_1(P_1P_3) = \alpha.$$

By definition, $\sigma_1^\chi(a|\theta) = \chi\bar{\sigma}_1(a) + (1-\chi)\sigma_1(a|s)$ and hence we can find that $\sigma_1^\chi(a|\theta)$ is given in Table 3.

From player two's perspective, given any action profile, player two's expected

Table 3: Cursed Perception of Player 1's Strategy

	player one's type	
$\sigma_1^\chi(a\|\theta)$	selfish	altruistic
T_1	$p(1-\chi\alpha)$	$p\chi(1-\alpha)$
P_1T_3	$(1-p)(1-\chi\alpha)$	$(1-p)\chi(1-\alpha)$
P_1P_3	$\chi\alpha$	$1-\chi+\chi\alpha$

payoff is not affected by whether player one is selfish or altruistic. Hence, player two only cares about the *marginal* distribution of player one's actions. In this case, χ-cursed player two believes player one will choose $a \in \{T_1, P_1T_3, P_1P_3\}$ with probability $\bar{\sigma}_1(a)$. Therefore, it is optimal for player two to choose T_2 if and only if

$$\bar{\sigma}_1(T_1) + 8\left[1-\bar{\sigma}_1(T_1)\right] \geq \bar{\sigma}_1(T_1) + 4\bar{\sigma}_1(P_1T_3) + 32\bar{\sigma}_1(P_1P_3) \iff p \leq \frac{1-7\alpha}{1-\alpha},$$

implying player two's best responses in the standard cursed analysis also coincides with the Nash best responses. As a result, one concludes that standard χ-CE would make exactly the same prediction as the Bayesian Nash equilibrium regardless how cursed the players are.

In contrast, the χ-CSE will exhibit distortions to the *conditional beliefs* of player two, given that player one has passed, because player two incorrectly takes into account how player one's choice to pass depended on player one's private information. In particular, it is harder to build a reputation, since a selfish type will have to imitate altruists in such a way that the true posterior on altruistic type conditional on a pass is higher than in the perfect Bayesian equilibrium, because the updating by player two about player one's type is dampened relative to this true posterior due to cursedness. This distorted belief updating will result in *less* passing by player one compared to the Bayesian equilibrium. Formally, the χ-CSE is described in Proposition 8.

Proposition 8. *In the χ-CSE, selfish player one will choose $P1$ with probability q_1^χ and choose $T3$ with probability 1; player two will choose $P2$ with probability q_2^χ and choose $T4$ with probability 1 where*

$$q_1^\chi = \begin{cases} \left[\frac{7\alpha-7\alpha\chi}{1-7\alpha\chi} - \alpha\right] \Big/ (1-\alpha) & \text{if } \chi \leq \frac{6}{7(1-\alpha)} \\ 0 & \text{if } \chi > \frac{6}{7(1-\alpha)} \end{cases} \qquad \textit{and,}$$

$$q_2^\chi = \begin{cases} 1/7 & \text{if } \chi \leq \frac{6}{7(1-\alpha)} \\ 0 & \text{if } \chi > \frac{6}{7(1-\alpha)}. \end{cases}$$

Proof. See Online Appendix. □

In order to see how the cursedness affects the equilibrium behavior, here we focus on the case of $\chi \leq \frac{6}{7(1-\alpha)}$ where selfish player one and player two will both mix at stage one and two. Given selfish player one chooses $P1$ with probability q_1^χ, by Lemma 1, we know when the game reaches stage two, player two's belief about player one being altruistic becomes

$$\mu^\chi = \chi\alpha + (1-\chi)\left[\frac{\alpha}{\alpha + (1-\alpha)q_1^\chi}\right].$$

Here we see that when χ is larger, player two will update his belief more slowly. Therefore, in order to maintain indifference at the mixed equilibrium, selfish player one has to pass with *lower* probability so that $P1$ is a more informative signal to player two. As a result, to make player two indifferent between $T2$ and $P2$, the following condition must hold at the equilibrium:

$$\mu^\chi = \frac{1}{7} \iff q_1^\chi = \left[\frac{7\alpha - 7\alpha\chi}{1 - 7\alpha\chi} - \alpha\right] \Big/ (1-\alpha).$$

To conclude this section, in Figure 5, we plot the probabilities of choosing $P1$ and $P2$ at χ-CSE when there is a five percent chance that player one is an altruist (i.e., $\alpha = 0.05$). From our analysis above, we can find that both the standard equilibrium theory and χ-CE predict selfish player one chooses $P1$ with probability and player two chooses $P2$ with probability 0.14. Moreover, these probabilities are independent of χ. However, χ-CSE predicts when players are more cursed, selfish player one is less likely to choose $P1$. When players are sufficiently cursed ($\chi \geq 0.91$), selfish player one and player two will never pass—i.e., behave as if there were no altruistic players.

4.4 Sequential Voting over Binary Agendas

In this section, we apply the concept of χ-CSE to the model of strategic binary amendment voting with incomplete information studied by Ordeshook and Palfrey

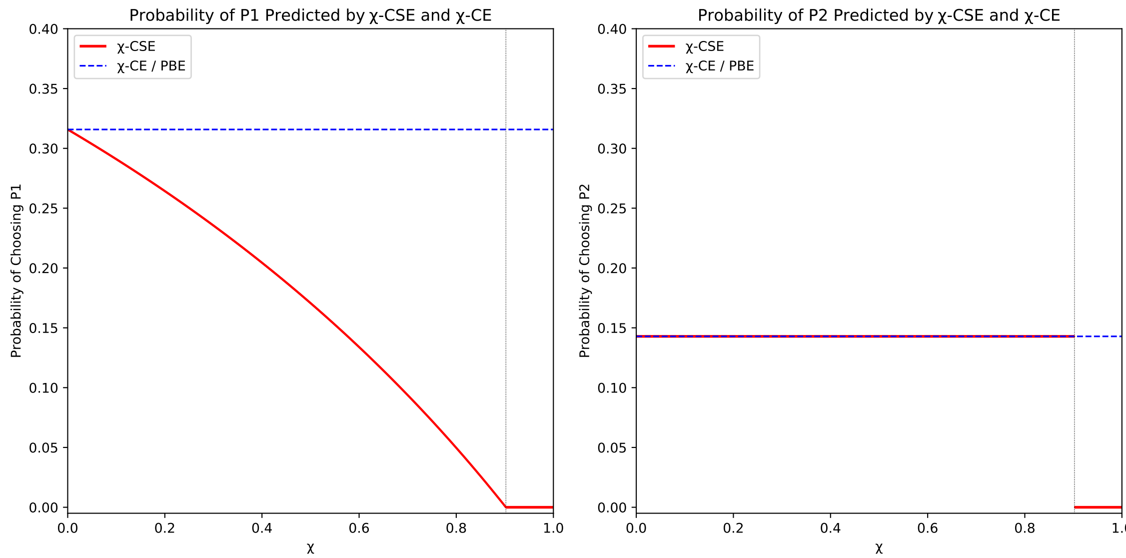

Figure 5: χ-CSE of the centipede game with altruistic players ($\alpha = 0.05$)

(1988). Let $N = \{1, 2, 3\}$ denote the set of voters. These three voters will vote over three possible alternatives in $X = \{a, b, c\}$. Voting takes place in a two-stage agenda. In the first stage, voters vote between a and b. In the second stage, voters vote between c and the majority rule winner of the first stage. The majority rule winner of the second stage is the outcome.

Each voter i has three possible private-value types where $\Theta \in \{\theta_1, \theta_2, \theta_3\}$ is the set of possible types. Each voter's type is independently drawn from a common prior distribution of types, p. In other words, the probability of a voter being type θ_k is p_k. Each voter's type is their own private information. Each voter has the same type-dependent payoff function, which is denoted by $u(x|\theta)$ for any $x \in X$ and $\theta \in \Theta$. We summarize the payoff function with the following table.

		x		
	$u(x\|\theta)$	a	b	c
	θ_1	1	v	0
θ	θ_2	0	1	v
	θ_3	v	0	1

Notice that $v \in (0, 1)$ is a parameter that measures the *intensity* of the second ranked outcome relative to the top ranked outcome. This intensity parameter, v, is assumed to be the same for all types of all voters. Because this is a game of private values, the standard χ-CE and the Bayesian Nash equilibrium coincide.

We use $a_i^1(\theta)$ to denote type θ voter i's action at stage 1. As is standard in majority voting games we will focus on the analysis of symmetric pure-strategy equilibria where voters do not use weakly dominated strategies. In other words, we will consider $a_i^t(\cdot) = a_j^t(\cdot)$ for all $i, j \in N$, and will drop the subscript.

In this PBE (and χ-CE) all voters will vote sincerely in equilibrium *except for* type θ_1 voters at stage 1. To see this, first note that voting insincerely in the last stage is dominated and thus eliminated, so all types of voters vote for their preferred alternative on the last ballot. Second, voting sincerely in both stages is a dominant strategy for a type θ_2 voter, who prefers any lottery between b and c to either a or c. Third, voting sincerely in both stages is also dominant for a type θ_3 voter in the sense that, in the event that neither of the other two voters are type θ_3, then any lottery between a and c is better than a vote between b and c since b (i.e., type θ_3's least preferred alternative) will win.[13]

The PBE (and χ-CE) prediction about a type θ_1 voter's strategy at stage 1 is summarized in the following claim.

Claim 3. *The symmetric (undominated pure) PBE strategy for type θ_1 voters in the first stage can be characterized as follows.*

1. $a^1(\theta_1) = b$ *is a PBE strategy if and only if* $v \geq \frac{p_1}{p_1+p_2}$.
2. $a^1(\theta_1) = a$ *is a PBE strategy if and only if* $v \leq \frac{p_1}{p_1+p_3}$.

Proof. See Ordeshook and Palfrey (1988). □

Claim 3 shows that, if v is relatively large, only type θ_1 voting sophisticatedly for b instead of sincerely for a can be supported by a PBE. Conditional on being pivotal, voting for b in the first stage guarantees an outcome of b and thus guarantees getting v, while voting for a leads to a lottery between a and c. As a result, when v is sufficiently high, a type θ_1 voter will have an incentive to strategically vote for b to avoid the risk of having c elected in the last stage.

[13]When there is another type θ_3 voter, the first ballot does not matter since their most preferred alternative c will always win in the second stage.

The analysis of a cursed sequential equilibrium is different from the standard cursed equilibrium in strategic form because the cursedness affects belief updating over the stages of the game, and players anticipate future play of the game. Due to the dynamics and the anticipation of future cursed behavior, such cursed behavior at later stages of a game can feedback and affect strategic behavior earlier in the game.

In the context of the two-stage binary amendment strategic voting model, cursed behavior and belief updating mean that voters in the first stage use the expected cursed beliefs in the second stage to compute the continuation values in the two continuation games of the second stage, either a vote between a and c or a vote between b and c. Because they have a cursed understanding about the relationship between types and the voting behavior in the first stage, this affects their predictions about which alternative wins in the second stage, conditional on which alternative wins in the first stage.

It is noteworthy that, given any $\chi \in [0, 1]$, all voters will still vote sincerely in χ-CSE *except for* type θ_1 voters at stage 1. As implied by Proposition 4, a voter in the last stage would act as if solving a maximization problem of χ-CE but under an (incorrectly) updated belief. Therefore, we can follow the same arguments as solving for the undominated Bayesian equilibrium and conclude that type θ_2 and θ_3 voters as well as type θ_1 voters at stage 2 will vote sincerely under a χ-CSE.

Proposition 9 establishes that the set of parameters v and p that can support a χ-CSE in which type θ_1 voters vote sophisticatedly for b shrinks as χ increases.

Proposition 9. *If $a^1(\theta_1) = b$ can be supported by a symmetric χ-CSE, then it can also be supported by a symmetric χ'-CSE for all $\chi' \leq \chi$.*

Proof. See Online Appendix. □

The intuition behind strategic voting over agendas mainly comes from the information content of *hypothetical* pivotal events. However, a cursed voter does not (fully) take such information into consideration, and thus becomes overly optimistic about his favorite alternative a being elected in the second stage. Therefore, a type θ_1 voter has a stronger incentive to deviate from sophisticated voting to sincere voting in stage 1 as χ increases.

Interestingly, the set of v and p that can support a χ-CSE in which type θ_1 voters vote sincerely for a does not necessarily expand as the level of cursedness becomes higher, as characterized in Proposition 10.

Proposition 10. *Given p and $v \in (0,1)$, there exists $\tilde{\chi}(p,v)$ such that*

1. *If $v > \frac{p_1}{p_1+p_3}$, then $a^1(\theta_1) = a$ is a χ-CSE strategy if and only if $\chi \geq \tilde{\chi}(p,v)$;*
2. *If $v < \frac{p_1}{p_1+p_3}$, then $a^1(\theta_1) = a$ is a χ-CSE strategy if and only if $\chi \leq \tilde{\chi}(p,v)$.*

Proof. See Online Appendix. □

Thus, Proposition 10 shows that, when χ is sufficiently large, there are some values of (v,p) that cannot support sincere voting for type θ_1 voters under PBE (and χ-CE) but can support it under χ-CSE. Alternatively, there also exist some values of (v,p) that can support sincere voting under PBE but fail to support it under χ-CSE when χ is large.

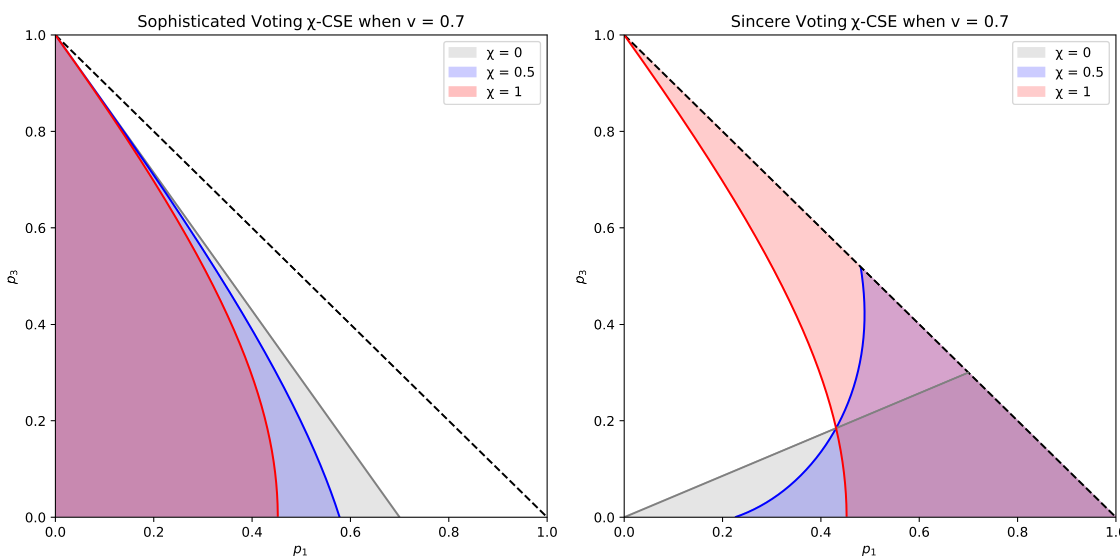

Figure 6: χ-CSE for Sophisticated (left) and Sincere (right) Voting When $v = 0.7$

To illustrate this, Figure 6 plots the set of p (fixing $v = 0.7$) that can support a χ-CSE for type θ_1 voters at stage 1 to vote sophisticatedly for b and sincerely for a. The left panel of Figure 6 shows that a sophisticated voting χ-CSE becomes harder to be supported as χ increases, as indicated by Proposition 9. For example, when $p \equiv (p_1, p_2, p_3) = (0.6, 0.3, 0.1)$, type θ_1 voters will not vote for second preferred alternative b if $\chi > 0.18$.

On the other hand, the right panel of Figure 6 shows that, while type θ_1 voters who sincerely vote for a at stage 1 cannot be supported under PBE when p_3 is large, they may emerge in a χ-CSE with sufficiently high χ. Also note that when p_2 is large, sincere voting by type θ_1 voters is no longer a χ-CSE with high χ. In such a sincere voting equilibrium, a fully rational type θ_1 voter knows there will be only one type θ_2 voter among the other two voters when being pivotal. As a result, whether to sincerely vote for a is determined by the ratio of p_1 to p_3. When p_3 is large, sincere voting at stage 1 will likely lead to zero payoff for type θ_1 voters and thus cannot be a PBE strategy. However, cursed type θ_1 voters will take the possibility of having two type θ_2 voters into account since they are not correctly conditioning on pivotality. As a result, when p_2 is large, sincere voting at stage 1 will likely lead to zero payoff for type θ_1 voters, and thus cannot be a χ-CSE strategy with high χ, while voting sophisticatedly for b can likely secure a payoff of v.

4.5 The Dirty Faces Game

The dirty faces game was first described by Littlewood (1953) to study the relationship between common knowledge and behavior.[14] There are several different variants of this game, but here we focus on a simplified version, the two-person dirty faces game, which was theoretically analyzed by Fudenberg and Tirole (1991a) and Lin (2022) and was experimentally studied by Weber (2001) and Bayer and Chan (2007).

Let $N = \{1, 2\}$ be the set of players. For each $i \in N$, let $x_i \in \{O, X\}$ represent whether player i has a clean face (O) or a dirty face (X). Each player's face type is independently and identically determined by a commonly known probability $p = \Pr(x_i = X) = 1 - \Pr(x_i = O)$. Once the face types are drawn, each player i can observe the other player's face x_{-i} but not their own face.[15] If there is at least one player with a dirty face, a public announcement of this fact is broadcast to both

[14] The dirty faces game has also been reframed as the "cheating wives puzzle" (Gamow and Stern, 1958), the "cheating husbands puzzle" (Moses et al., 1986), the "muddy children puzzle" (Barwise, 1981) and (Halpern and Moses, 1990), and the "red hat puzzle" (Hardin and Taylor, 2008).

[15] To fit into the framework, each player's "type" (their own private information) can be specified as "other players' faces." That is, $\theta_i = x_{-i}$.

players at the beginning of the game. Let $\omega \in \{0, 1\}$ denote whether there is an announcement or not. If there is an announcement ($\omega = 1$), all players are informed there is at least one dirty face but not the identities. When $\omega = 0$, it is common knowledge to both players that their faces are clean and the game becomes trivial. Hence, in the following, we will focus only on the interesting case where $\omega = 1$.

There are a finite number of $T \geq 2$ stages. In each stage, each player i simultaneously chooses $s_i \in \{U, D\}$. The game ends as soon as either player (or both) chooses D, or at the end of stage T in case neither player has chosen D. Actions are revealed at the end of each stage. Payoffs depend on own face types and action. If a player chooses D, he will get $\alpha > 0$ if he has a dirty face while receive -1 if he has a clean face. We assume that

$$p\alpha - (1-p) < 0 \iff 0 < \bar{\alpha} \equiv \frac{\alpha}{(1-p)(1+\alpha)} < 1, \tag{1}$$

where $p\alpha - (1-p)$ is the expected payoff of D when the belief of having a dirty face is p. Thus, Assumption (1) guarantees it is strictly dominated to choose D at stage 1 when observing a dirty face. In other words, players will be rewarded when correctly inferring the dirty face but penalized when wrongly claiming the dirty face.

The payoffs are discounted with a common discount factor $\delta \in (0, 1)$. To summarize, conditional on reaching stage t, each player's payoff function (which depends on their own face and action) can be written as:

$$u_i(s_i|t, x_i = X) = \begin{cases} \delta^{t-1}\alpha & \text{if } s_i = D \\ 0 & \text{if } s_i = U \end{cases} \quad \text{and} \quad u_i(s_i|t, x_i = O) = \begin{cases} -\delta^{t-1} & \text{if } s_i = D \\ 0 & \text{if } s_i = U. \end{cases}$$

Therefore, a two-person dirty faces game is defined by a tuple $\langle p, T, \alpha, \delta \rangle$.

Since the game ends as soon as some player chooses D, the information sets of the game can be specified by the face type the player observes and the stage number. Thus a behavioral strategy can be represented as:

$$\sigma : \{O, X\} \times \{1, \ldots, T\} \to [0, 1],$$

which is a mapping from information sets to the probability of choosing D, where

$\{O, X\}$ corresponds to a player's observation of the other player's face.

There is a unique Nash equilibrium. When observing a clean face, a player would immediately know his face is dirty. Hence, it is strictly dominant to choose D at stage 1 in this case. On the other hand, when observing a dirty face, because of Assumption (1), it is optimal for the player to choose U at stage 1. However, if the game proceeds to stage 2, the player would know his face is dirty because the other player would have chosen D at stage 1 if his face were clean and the game would not have reached stage 2. This result is independent of the payoffs, the timing, the discount factor, and the (prior) probability of having a dirty face. The only assumption for this argument is common knowledge of rationality.

Alternatively, when players are "cursed," they are not able to make perfect inferences from the other player's actions. Specifically, since a cursed player has incorrect perceptions about the relationship between the other player's actions and their private information after seeing the other player choose U in stage 1, a cursed player does not believe they have a dirty face for sure. At the extreme when $\chi = 1$, fully cursed players never update their beliefs. In the following, we will compare the predictions of the standard χ-CE and the χ-CSE. A surprising result is that there is always a *unique* χ-CE, but there can be multiple χ-CSE.

For the sake of simplicity, we will focus on the characterization of pure strategy equilibrium in the following analysis. Since the game ends when some player chooses D, we can equivalently characterize a stopping strategy as a mapping from the observed face type to a stage in $\{1, 2, \ldots, T, T+1\}$ where $T+1$ corresponds to the strategy of never stopping. Furthermore, both χ-CE and χ-CSE will be symmetric because if players were to stop at different stages, least one of the players would have a profitable deviation. Finally, we use $\hat{\sigma}^{\chi}(x_{-i})$ and $\tilde{\sigma}^{\chi}(x_{-i})$ to denote the equilibrium stopping strategies of χ-CE and χ-CSE, respectively.

We characterize the χ-CE in Proposition 11. Since χ-CE is defined for simultaneous move Bayesian games, to solve for the χ-CE, we need to look at the corresponding normal form where players simultaneously choose $\{1, 2, \ldots, T, T+1\}$ given the observed face type.

Proposition 11. *The χ-cursed equilibrium can be characterized as follows.*

1. *If $\chi > \bar{\alpha}$, the only χ-CE is that both players choose:*

$$\hat{\sigma}^{\chi}(O) = 1 \quad \text{and} \quad \hat{\sigma}^{\chi}(X) = T + 1.$$

2. *If $\chi < \bar{\alpha}$, the only χ-CE is that both players choosing*

$$\hat{\sigma}^{\chi}(O) = 1 \quad \text{and} \quad \hat{\sigma}^{\chi}(X) = 2.$$

Proof. See Online Appendix. □

Proposition 11 shows that χ-CE makes an extreme prediction—when observing a dirty face, players would either choose D at stage 2 (the equilibrium prediction) or never choose D. In addition, the prediction of χ-CE is unique for $\chi \neq \bar{\alpha}$. As characterized in Proposition 12, for extreme values of χ, the prediction of χ-CSE coincides with χ-CE. But for intermediate values of χ, there can be multiple χ-CSE.

Proposition 12. *The pure strategy χ-CSE can be characterized as follows.*

1. *$\tilde{\sigma}^{\chi}(O) = 1$ for all $\chi \in [0, 1]$.*

2. *Both players choosing $\tilde{\sigma}^{\chi}(X) = T + 1$ is a χ-CSE if and only if $\chi \geq \bar{\alpha}^{\frac{1}{T+1}}$.*

3. *Both players choosing $\tilde{\sigma}^{\chi}(X) = 2$ is a χ-CSE if and only if $\chi \leq \bar{\alpha}$.*

4. *For any $3 \leq t \leq T$, both players choosing $\tilde{\sigma}^{\chi}(X) = t$ is a χ-CSE if and only if*

$$\left(\frac{1-\kappa(\chi)}{1-p}\right)^{\frac{1}{t-2}} \leq \chi \leq \bar{\alpha}^{\frac{1}{t-1}} \quad \text{where}$$

$$\kappa(\chi) \equiv \frac{[(1+\alpha)(1+\delta\chi) - \alpha\delta] - \sqrt{[(1+\alpha)(1+\delta\chi) - \alpha\delta]^2 - 4\delta\chi(1+\alpha)}}{2\delta\chi(1+\alpha)}.$$

Proof. See Online Appendix. □

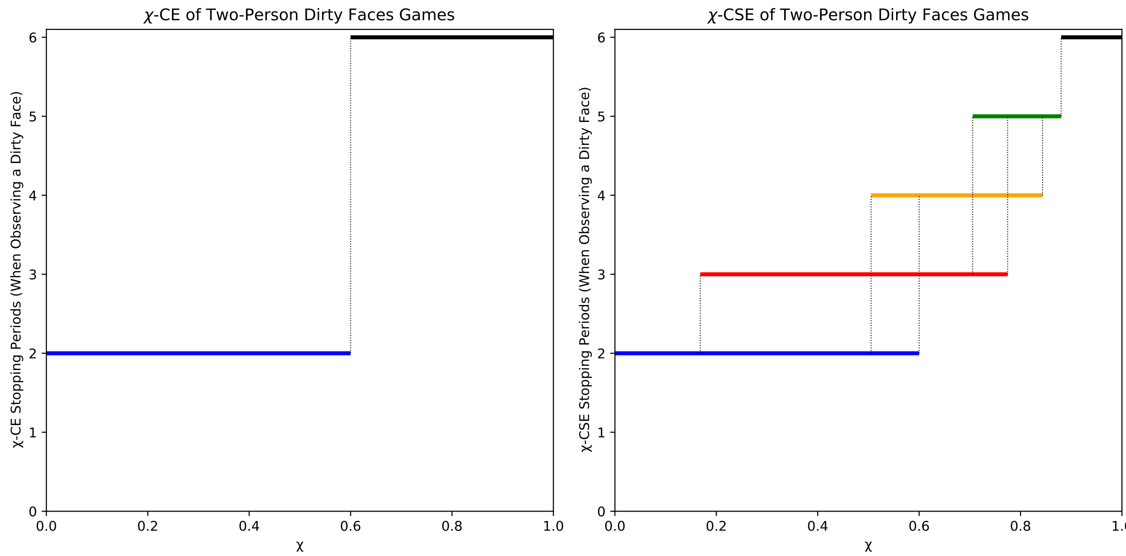

Figure 7: χ-CE vs. χ-CSE When $(\alpha, \delta, p, T) = \left(\frac{1}{4}, \frac{4}{5}, \frac{2}{3}, 5\right)$

Illustrative Example

In order to illustrate the sharp contrast between the predictions of χ-CE and χ-CSE, here we consider an illustrative example where $\alpha = 1/4$, $\delta = 4/5$, $p = 2/3$ and the horizon of the game is $T = 5$. As characterized by Proposition 11, χ-CE predicts players will choose $\hat{\sigma}^{\chi}(X) = 2$ if $\chi \leq \bar{\alpha} = 0.6$; otherwise, they will choose $\hat{\sigma}^{\chi}(X) = 6$, i.e., they never choose D when observing a dirty face. As demonstrated in the left panel of Figure 7, χ-CE is (generically) unique and it predicts players will either behave extremely sophisticated or unresponsive to the other's action at all.

In contrast, as characterized by Proposition 12, there can be multiple χ-CSE. As shown in the right panel of Figure 7, when $\chi \leq \bar{\alpha} = 0.6$, both players stopping at stage 2 is still an equilibrium, but it is not unique except for very low values of χ. For $0.168 \leq \chi \leq 0.505$, both players stopping at stage 3 is also a χ-CSE, and for $0.505 \leq \chi \leq 0.6$, there are three pure strategy χ-CSE where both players stop at stage 2, 3, or 4, respectively.

The existence of multiple χ-CSE in which both players stop at $t > 2$ highlights a player's learning process in a multi-stage game, which does not happen in strategic form cursed equilibrium. In the strategic form, a player has no opportunity to learn about the other player's type in middle stages. Thus, when level of cursedness is not low enough to support a χ-CE with stopping at stage 2, both players would never stop. However, in a χ-CSE of the multi-stage game, a cursed player would still learn about his own face being dirty as the game proceeds, even though he might not be confident enough to choose D at stage 2. If χ is not too large, the expected payoff of choosing D would eventually become positive at some stage before the last stage T.[16] For some intermediate values of χ, there might be multiple stopping stages which yield positive expected payoffs. In this case, the dirty faces game becomes a special type of coordination games where both players coordinate on stopping strategies, resulting in the existence of multiple χ-CSE.[17]

5 Concluding Remarks

In this paper, we formally developed Cursed Sequential Equilibrium, which extends the strategic form cursed equilibrium (Eyster and Rabin, 2005) to multi-stage games, and illustrated the new equilibrium concept with a series of applications. While the standard CE has no bite in private value games, we show that cursed beliefs can actually have significant consequences for dynamic private value games. In the private value games we consider, our cursed sequential equilibrium predicts (1) under-contribution caused by under-communication in the public goods game with communication, (2) low passing rate in the presence of altruistic players in the centipede game, and (3) less sophisticated voting in the sequential two-stage binary agenda game. We also illustrate the distinction between CE and CSE in some non-private value games. In simple signaling games, χ-CSE implies refinements of pooling equilibria that are not captured by traditional belief-based refinements (or χ-CE), and are qualitatively consistent with some experimental evidence. Lastly, we examine the dirty face game, showing that the CSE further expands the set of equilibrium and predicts stopping in middle stages of the game. We summarize our findings from these applications in Table 4.

The applications we consider are only a small sample of the possible dynamic

[16] The upper bound of the inequality in Proposition 12 characterizes the stages at which stopping yields positive expected payoffs.

[17] Note that players with low levels of cursedness would not coordinate on stopping at late stages since the discount factor shrinks the informative value of waiting (i.e., both choosing U). This result is characterized by the lower bound of the inequality in Proposition 12.

Table 4: Summary of Findings in Section 4

Application	Private Values	χ-CE vs. BNE	χ-CSE vs. χ-CE
Signaling Game	No	$\neq$	χ-CSE $\subset$ χ-CE
Public Goods Game	Yes	$=$	$\neq$
Centipede Game	Yes	$=$	$\neq$
Sequential Voting Game	Yes	$=$	$\neq$
Dirty Faces Game	Yes	$\neq$	$\neq$

games where CSE could be usefully applied. One prominent class of problems where it would be interesting to study the dynamic effects of cursedness is social learning. For example, in the standard information cascade model of Bikhchandani et al. (1992), we conjecture that the effect would be to delay the formation of an information cascade because players will partially neglect the information content of prior decision makers. Laboratory experiments report evidence that subjects underweight the information contained in prior actions relative to their own signal (Goeree et al., 2007). A related class of problems involves information aggregation through sequential voting and bandwagon effects (Callander, 2007; Ali et al., 2008; Ali and Kartik, 2012). A natural conjecture is that CSE will impede information transmission in committees and juries as later voters will under-appreciate the information content of the decisions by early voters. This would dampen bandwagon effects. The centipede example we studied suggests that CSE might have broader implications for behavior in reputation-building games, such as the finitely repeated prisoner's dilemma or entry deterrence games such as the chain store paradox.

The generalization of CE to dynamic games presented in this paper is limited in several ways. First, the CSE framework is formally developed for finite multi-stage games with observed actions. We do not extend CSE for games with continuous types but we do provide one application that shows how such an extension is possible. However, a complete generalization to continuous types (or continuous actions) would require more technical development and assumptions. We also assume that the number of stages is finite, and extending this to infinite horizon multi-stage games would be a useful exercise. Extending CSE to allow for imperfect monitoring in the form of private histories is another interesting direction to pursue. The SCE approach in Cohen and Li (2023) allows for cursedness with respect to both public and private endogenous information, which leads to some important differences from our CSE approach. In CSE, we find that subjects are limited in their ability to make correct inferences about hypothetical events, but the mechanism is different from SCE, which introduces a second free parameter that modulates cursedness with respect to hypothetical events. For a more detailed discussion of these and other differences and overlaps between CSE and SCE, see Fong et al. (2023).

As a final remark, our analysis of applications of χ-CSE suggests some interesting experiments. For instance, χ-CSE predicts in the public goods game with communication, when either the number of players (N) or the largest possible contribution cost (K) increases, pre-play communication will be less effective, while the prediction of sequential equilibrium and χ-CE is independent of N and K. In other words, in an experiment where N and K are manipulated, significant treatment effects in this direction would provide evidence supporting χ-CSE over χ-CE. Also, χ-CSE makes qualitatively testable predictions in the sequential voting games and the dirty faces games, which have not been extensively studied in laboratory experiments. In the sequential voting game, it would be interesting to test how sensitive strategic (vs. sincere) voting behavior is to preference intensity (v) and the type distribution. In the dirty faces game, it would be interesting to design an experiment to identify the extent to which deviations from sequential equilibrium are related to the coordination problem that arises in χ-CSE.

References

Ali, N., J. K. Goeree, N. Kartik, and T. R. Palfrey (2008): "Information Aggregation in Standing and Ad Hoc Committees," *American Economic Review*, 98, 181–186.

Ali, N. and N. Kartik (2012): "Herding with Collective Preferences," *Economic Theory*, 51, 601–626.

Barwise, J. (1981): "Scenes and other situations," *Journal of Philosophy*, 78, 369–397.

Bayer, R. C. and M. Chan (2007): "The dirty faces game revisited," Tech. rep., University of Adelaide, School of Economics.

Bazerman, M. H. and W. F. Samuelson (1983): "I won the auction but don't want the prize," *Journal of Conflict Resolution*, 27, 618–634.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 100, 992–1026.

Brandts, J. and C. A. Holt (1993): "Adjustment patterns and equilibrium selection in experimental signaling games," *International Journal of Game Theory*, 22, 279–302.

Callander, S. (2007): "Bandwagons and Momentum in Sequential Voting," *Review of Economic Studies*, 74, 653–684.

Camerer, C. F., S. Nunnari, and T. R. Palfrey (2016): "Quantal Response and Non-equilibrium Beliefs Explain Overbidding in Maximum Value Auctions," *Games and Economic Behavior*, 98, 243–263.

Capen, E. C., R. V. Clapp, and W. M. Campbell (1971): "Competitive bidding in high-risk situations," *Journal of Petroleum Technology*, 23, 641–653.

Carrillo, J. and T. R. Palfrey (2009): "The Compromise Game: Two-Sided Adverse Selection in the Laboratory," *American Economic Journal: Microeconomics*, 1, 151–181.

Carrillo, J. D. and T. R. Palfrey (2011): "No trade," *Games and Economic Behavior*, 71, 66–87.

Cho, I.-K. and D. M. Kreps (1987): "Signaling games and stable equilibria," *Quarterly Journal of Economics*, 102, 179–221.

Cohen, S. and S. Li (2023): "Sequential Cursed Equilibrium," *arXiv preprint arXiv:2212.06025v3 First version posted December 13, 2022.*

Dyer, D., J. H. Kagel, and D. Levin (1989): "A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis," *Economic Journal*, 99, 108–115.

Eckel, C. C. and C. A. Holt (1989): "Strategic Voting in Agenda Controlled Committee Experiments," *American Economic Review*, 79, 763–773.

Eyster, E. and M. Rabin (2005): "Cursed equilibrium," *Econometrica*, 73, 1623–1672.

Fong, M.-J., P.-H. Lin, and T. R. Palfrey (2023): "A Note on Cursed Sequential Equilibrium and Sequentia Cursed Equilibrium," *arXiv preprint arXiv:2304.05515.*

Forsythe, R. E., R. M. Isaac, and T. R. Palfrey (1989): "Theories and Tests of Blind Bidding in Sealed-bid Auctions," *Rand Journal of Economics*, 20, 214–238.

Fudenberg, D. and J. Tirole (1991a): *Game theory*, MIT press.

——— (1991b): "Perfect Bayesian equilibrium and sequential equilibrium," *Journal of Economic Theory*, 53, 236–260.

Gamow, G. and M. Stern (1958): "Forty unfaithful wives," *Puzzle Math*, 20–23.

Goeree, J. K., T. R. Palfrey, B. W. Rogers, and R. D. McKelvey (2007): "Self-correcting information cascades," *Review of Economic Studies*, 74, 733–762.

Guarnaschelli, S., R. D. McKelvey, and T. R. Palfrey (2000): "An experimental study of jury decision rules," *American Political Science Review*, 94, 407–423.

Halpern, J. Y. and Y. Moses (1990): "Knowledge and common knowledge in a distributed environment," *Journal of the ACM (JACM)*, 37, 549–587.

Hardin, C. S. and A. D. Taylor (2008): "An introduction to infinite hat problems," *The Mathematical Intelligencer*, 30, 20–25.

Holt, C. A. and R. Sherman (1994): "The loser's curse," *American Economic Review*, 84, 642–652.

Ivanov, A., D. Levin, and M. Niederle (2010): "Can Relaxation of Beliefs Rationalize the Winner's Curse," *Econometrica*, 78, 1435–1452.

Jehiel, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81–104.

Jehiel, P. and F. Koessler (2008): "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 62, 533–557.

Kagel, J. H. and D. Levin (1986): "The winner's curse and public information in common value auctions," *American Economic Review*, 894–920.

——— (2009): "Common value auctions and the winner's curse," in *Common Value Auctions and the Winner's Curse*, Princeton University Press.

Kagel, J. H., D. Levin, R. C. Battalio, and D. J. Meyer (1989): "First-price common value auctions: bidder behavior and the "Winner's Curse"," *Economic Inquiry*, 27, 241–258.

Kreps, D. M. (1990): *A course in microeconomic theory*, Princeton university press.

Kreps, D. M. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

Levine, M. E. and C. R. Plott (1977): "Agenda influence and its implications," *Virginia Law Review*, 561–604.

Lin, P.-H. (2022): "Cognitive Hierarchies in Multi-Stage Games of Incomplete Information," *arXiv preprint arXiv:2208.11190*.

Lin, P.-H. and T. R. Palfrey (2022): "Cognitive Hierarchies in Extensive Form Games," *Caltech Social Science Working Paper*.

Lind, B. and C. R. Plott (1991): "The winner's curse: experiments with buyers and with sellers," *American Economic Review*, 81, 335–346.

Littlewood, J. E. (1953): *A Mathematician's Miscellany*, London, England: Meuthen & Co. Ltd.

McKelvey, R. D. and T. R. Palfrey (1992): "An experimental study of the centipede game," *Econometrica*, 803–836.

——— (1998): "Quantal response equilibria for extensive form games," *Experimental Economics*, 1, 9–41.

Moses, Y., D. Dolev, and J. Y. Halpern (1986): "Cheating husbands and other stories: a case study of knowledge, action, and communication," *Distributed Computing*, 1, 167–176.

Ordeshook, P. C. and T. R. Palfrey (1988): "Agendas, strategic voting, and signaling with incomplete information," *American Journal of Political Science*, 441–466.

Palfrey, T. R. and H. Rosenthal (1991): "Testing for effects of cheap talk in a public goods game with private information," *Games and Economic Behavior*, 3, 183–220.

Palfrey, T. R., H. Rosenthal, and N. Roy (2017): "How Cheap Talk Enhances Efficiency in Threshold Public Goods Games," *Games and Economic Behavior*, 101, 234–259.

Plott, C. R. and M. E. Levine (1978): "A Model of Agenda Influence on Committee Decisions," *American Economic Review*, 68, 146–160.

Rogers, B. W., T. R. Palfrey, and C. F. Camerer (2009): "Heterogeneous quantal response equilibrium and cognitive hierarchies," *Journal of Economic Theory*, 144, 1440–1467.

Samuelson, W. F. and M. H. Bazerman (1985): "Negotiation under the winner's curse," *Research in experimental economics*, 3, 105–138.

Søvik, Y. (2009): "Strength of dominance and depths of reasoning—An experimental study," *Journal of Economic Behavior & Organization*, 70, 196–205.

Weber, R. A. (2001): "Behavior and learning in the "dirty faces" game," *Experimental Economics*, 4, 229–242.

Appendix A Omitted Proofs of Section 2 and 3

Proof of Lemma 1: By definition 1, for any $(\mu, \sigma) \in \Psi^{\chi}$, any history h^{t-1}, any player i and any type profile $\theta = (\theta_i, \theta_{-i})$,

$$
\begin{aligned}
&\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)[\chi \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1-\chi)\sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i})] \\
&= \chi \underbrace{\left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i) \right]}_{=1} \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1-\chi) \underbrace{\left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i}) \right]}_{=\bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i)} \\
&= \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i).
\end{aligned}
$$

Therefore, since $(\mu, \sigma) \in \Psi^{\chi}$, with some rearrangement, it follows that

$$
\begin{aligned}
\mu_i(\theta_{-i}|h^t, \theta_i) &= \frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i}, \theta_i)}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)\sigma^{\chi}_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i}, \theta_i)} \\
&= \frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i)[\chi \bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i) + (1-\chi)\sigma_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i})]}{\bar{\sigma}_{-i}(a^t_{-i}|h^{t-1}, \theta_i)} \\
&= \chi \mu_i(\theta_{-i}|h^{t-1}, \theta_i) + (1-\chi) \left[\frac{\mu_i(\theta_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a^t_{-i}|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a^t_{-i}|h^{t-1}, \theta'_{-i})} \right].
\end{aligned}
$$

Proof of Proposition 1: The proof is similar to the proof for sequential equilibrium and proceeds in three steps. First, for any finite multi-stage games with observed actions, Γ, we construct an ϵ-perturbed game Γ^{ϵ} that is identical to Γ but every player in every information set has to play any available action with probability at least ϵ. Second, we defined a cursed best-response correspondence for Γ^{ϵ} and prove that the correspondence has a fixed point by Kakutani's fixed point theorem. Finally, in step 3, we use a sequence of fixed points in perturbed games, with ϵ converging to 0, where the limit of this sequence is a χ-CSE.

Step 1: Let Γ^ϵ be a game identical to Γ but for each player $i \in N$, player i must play any available action in every information set $\mathcal{I}_i = (\theta_i, h^t)$ with probability at least ϵ where $\epsilon < \frac{1}{\sum_{j=1}^n |A_j|}$. Let $\Sigma^\epsilon = \times_{j=1}^n \Sigma_j^\epsilon$ be set of feasible behavioral strategy profiles for players in the perturbed game Γ^ϵ. For any behavioral strategy profile $\sigma \in \Sigma^\epsilon$, let $\mu^\chi(\cdot) \equiv (\mu_i^\chi(\cdot))_{i=1}^n$ be the belief system induced by σ via χ-cursed Bayes' rule. That is, for each player $i \in N$, information set $\mathcal{I}_i = (\theta_i, h^t)$ where $h^t = (h^{t-1}, a^t)$ and type profile $\theta_{-i} \in \Theta_{-i}$,

$$\mu_i^\chi(\theta_{-i}|h^t, \theta_i) = \chi \mu_i^\chi(\theta_{-i}|h^{t-1}, \theta_i) + (1-\chi)\left[\frac{\mu_i^\chi(\theta_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i})}{\sum_{\theta'_{-i} \in \Theta_{-i}} \mu_i^\chi(\theta'_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a_{-i}^t|h^{t-1}, \theta'_{-i})}\right].$$

Notice that the χ-cursed Bayes' rule is only defined on the family of multi-stage games with observed actions. As σ is fully mixed, the belief system is uniquely pinned down.

Finally, let $B^\epsilon : \Sigma^\epsilon \rightrightarrows \Sigma^\epsilon$ be the cursed best response correspondence which maps any behavioral strategy profile $\sigma \in \Sigma^\epsilon$ to the set of ϵ-constrained behavioral strategy profiles $\tilde{\sigma} \in \Sigma^\epsilon$ that are best replies given the belief system $\mu^\chi(\cdot)$.

Step 2: Next, fix any $0 < \epsilon < \frac{1}{\sum_{j=1}^n |A_j|}$ and show that B^ϵ has a fixed point by Kakutani's fixed point theorem. We check the conditions of the theorem:

1. It is straightforward that Σ^ϵ is compact and convex.

2. For any $\sigma \in \Sigma^\epsilon$, as $\mu^\chi(\cdot)$ is uniquely pinned down by χ-cursed Bayes' rule, it is straightforward that $B^\epsilon(\sigma)$ is non-empty and convex.

3. To verify that B^ϵ has a closed graph, take any sequence of ϵ-constrained behavioral strategy profiles $\{\sigma^k\}_{k=1}^\infty \subseteq \Sigma^\epsilon$ such that $\sigma^k \to \sigma \in \Sigma^\epsilon$ as $k \to \infty$, and any sequence $\{\tilde{\sigma}^k\}_{k=1}^\infty$ such that $\tilde{\sigma}^k \in B^\epsilon(\sigma^k)$ for any k and $\tilde{\sigma}^k \to \tilde{\sigma}$. We want to prove that $\tilde{\sigma} \in B^\epsilon(\sigma)$.

Fix any player $i \in N$ and information set $\mathcal{I}_i = (\theta_i, h^t)$. For any $\sigma \in \Sigma^\epsilon$, recall that $\sigma_{-i}^\chi(\cdot)$ is player i's χ-cursed perceived behavioral strategies of other players induced by σ. Specifically, for any type profile $\theta \in \Theta$, non-terminal history h^{t-1} and action profile $a_{-i}^t \in A_{-i}(h^{t-1})$,

$$\sigma_{-i}^\chi(a_{-i}^t|h^{t-1}, \theta_{-i}, \theta_i) = \chi \bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) + (1-\chi)\sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i}).$$

Additionally, recall that $\rho_i^\chi(\cdot)$ is player i's belief about the terminal nodes (conditional on the history and type profile), which is also induced by σ. Since $\mu^\chi(\cdot)$ is continuous in σ we have thaat $\sigma_{-i}^\chi(\cdot)$ and $\rho_i^\chi(\cdot)$ are also continuous in σ.

We further define

$$\mathcal{S}_{\mathcal{I}_i}^k \equiv \left\{\sigma_i' \in \Sigma_i^\epsilon : \sigma_i'(\,\cdot\,|\mathcal{I}_i) = \tilde{\sigma}_i^k(\,\cdot\,|\mathcal{I}_i)\right\}, \mathcal{S}_{\mathcal{I}_i} \quad \equiv \left\{\sigma_i' \in \Sigma_i^\epsilon : \sigma_i'(\,\cdot\,|\mathcal{I}_i) = \tilde{\sigma}_i(\,\cdot\,|\mathcal{I}_i)\right\}.$$

Since $\tilde{\sigma}^k \in B^\epsilon(\sigma^k)$, for any $\sigma_i' \in \Sigma_i^\epsilon$, we can obtain that

$$\begin{aligned}\max_{\sigma_i'' \in \mathcal{S}_{\mathcal{I}_i}^k} &\left\{\sum_{\theta_{-i}\in\Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma^k](\theta_{-i}|h^t,\theta_i)\rho_i^\chi(h^T|h^t,\theta,\sigma_{-i}^\chi[\sigma^k],\sigma_i'')u_i(h^T,\theta_i,\theta_{-i})\right\} \\ &\geq \sum_{\theta_{-i}\in\Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma^k](\theta_{-i}|h^t,\theta_i)\rho_i^\chi(h^T|h^t,\theta,\sigma_{-i}^\chi[\sigma^k],\sigma_i')u_i(h^T,\theta_i,\theta_{-i}).\end{aligned}$$

By continuity, as we take limits on both sides, we can obtain that

$$\begin{aligned}\max_{\sigma_i'' \in \mathcal{S}_{\mathcal{I}_i}} &\left\{\sum_{\theta_{-i}\in\Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma](\theta_{-i}|h^t,\theta_i)\rho_i^\chi(h^T|h^t,\theta,\sigma_{-i}^\chi[\sigma],\sigma_i'')u_i(h^T,\theta_i,\theta_{-i})\right\} \\ &\geq \sum_{\theta_{-i}\in\Theta_{-i}} \sum_{h^T \in \mathcal{H}^T} \mu_i^\chi[\sigma](\theta_{-i}|h^t,\theta_i)\rho_i^\chi(h^T|h^t,\theta,\sigma_{-i}^\chi[\sigma],\sigma_i')u_i(h^T,\theta_i,\theta_{-i}).\end{aligned}$$

Therefore, $\tilde{\sigma} \in B^\epsilon(\sigma)$, so by Kakutani's fixed point theorem, B^ϵ has a fixed point.
Step 3: For any ϵ, let σ^ϵ be a fixed point of B^ϵ and μ^ϵ be the belief system induced by σ^ϵ via χ-cursed Bayes' rule. We combine these two components and let $(\mu^\epsilon, \sigma^\epsilon)$ be the induced assessment. We now consider a sequence of $\epsilon \to 0$, where $\{(\mu^\epsilon, \sigma^\epsilon)\}$ is the corresponding sequence of assessments.

By compactness and the finiteness of Γ, the Bolzano-Weierstrass theorem guarantees the existence of a convergent subsequence of the assessments. As $\epsilon \to 0$, let $(\mu^\epsilon, \sigma^\epsilon) \to (\mu^*, \sigma^*)$. By construction, the limit assessment (μ^*, σ^*) satisfies χ-consistency and sequential rationality. Hence, (μ^*, σ^*) is a χ-CSE. ■

Proof of Proposition 2: To prove $\Phi(\chi)$ is upper hemi-continuous in χ, consider any sequence of $\{\chi_k\}_{k=1}^\infty$ such that $\chi_k \to \chi^* \in [0,1]$, and any sequence of CSE, $\{(\mu^k, \sigma^k)\}$, such that $(\mu^k, \sigma^k) \in \Phi(\chi_k)$ for all k. Let (μ^*, σ^*) be the limit assessment, i.e., $(\mu^k, \sigma^k) \to (\mu^*, \sigma^*)$. We need to show that $(\mu^*, \sigma^*) \in \Phi(\chi^*)$.

To simplify notation, for any player $i \in N$, any information set $\mathcal{I}_i = (h^t, \theta_i)$, any

$\sigma_i' \in \Sigma_i$, and any $\sigma \in \Sigma$, the expected payoff under the belief system $\mu^\chi(\cdot)$ induced by σ is denoted as:

$$\mathbb{E}_{\mu^\chi[\sigma]}\left[u_i(\sigma_i', \sigma_{-i}|h^t, \theta_i)\right] \equiv \sum_{\theta_{-i}\in\Theta_{-i}} \sum_{h^T\in\mathcal{H}^T} \mu_i^\chi(\theta_{-i}|h^t, \theta_i)\rho_i^\chi(h^T|h^t, \theta, \sigma_{-i}^\chi, \sigma_i')u_i(h^T, \theta_i, \theta_{-i}).$$

Suppose $(\mu^*, \sigma^*) \notin \Phi(\chi^*)$. Then there exists some player $i \in N$, some information set $\mathcal{I}_i = (h^t, \theta_i)$, some $\sigma_i' \in \Sigma_i$, and some $\epsilon > 0$ such that

$$\mathbb{E}_{\mu^{\chi^*}[\sigma^*]}\left[u_i(\sigma_i', \sigma_{-i}^*|h^t, \theta_i)\right] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]}\left[u_i(\sigma_i^*, \sigma_{-i}^*|h^t, \theta_i)\right] > \epsilon. \tag{A}$$

Since $\mu^\chi(\cdot)$ is continuous in χ, it follows that for any strategy profile σ, $\sigma_{-i}^\chi(\cdot)$ and $\rho_i^\chi(\cdot)$ are both continuous in χ. As a result, there exists a sufficiently large M_1 such that for every $k \geq M_1$,

$$\left|\mathbb{E}_{\mu^{\chi^k}[\sigma^k]}\left[u_i(\sigma_i^k, \sigma_{-i}^k|h^t, \theta_i)\right] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]}\left[u_i(\sigma_i^*, \sigma_{-i}^*|h^t, \theta_i)\right]\right| < \frac{\epsilon}{3}. \tag{B}$$

Similarly, there exists a sufficiently large M_2 such that for every $k \geq M_2$,

$$\left|\mathbb{E}_{\mu^{\chi^k}[\sigma^k]}\left[u_i(\sigma_i', \sigma_{-i}^k|h^t, \theta_i)\right] - \mathbb{E}_{\mu^{\chi^*}[\sigma^*]}\left[u_i(\sigma_i', \sigma_{-i}^*|h^t, \theta_i)\right]\right| < \frac{\epsilon}{3}. \tag{C}$$

Therefore, for any $k \geq \max\{M_1, M_2\}$, inequalities (A), (B) and (C) imply:

$$\mathbb{E}_{\mu^{\chi^k}[\sigma^k]}\left[u_i(\sigma_i', \sigma_{-i}^k|h^t, \theta_i)\right] - \mathbb{E}_{\mu^{\chi^k}[\sigma^k]}\left[u_i(\sigma_i^k, \sigma_{-i}^k|h^t, \theta_i)\right] > \frac{\epsilon}{3},$$

implying that σ_i' is a profitable deviation for player i at information set $\mathcal{I}_i = (h^t, \theta_i)$, which contradicts $(\mu^k, \sigma^k) \in \Phi(\chi_k)$. Therefore, $(\mu^*, \sigma^*) \in \Phi(\chi^*)$, as desired. ■

Proof of Proposition 3: Fix any $\chi \in [0, 1]$ and let (μ, σ) be a χ-consistent assessment. We prove the result by contradiction. Suppose (μ, σ) does not satisfy χ-dampened updating property. Then there exists $i \in N$, $\tilde{\theta} \in \Theta$ and a non-terminal history h^t such that $\mu_i(\theta_{-i}|h^t, \tilde{\theta}_i) < \chi\mu_i(\theta_{-i}|h^{t-1}, \tilde{\theta}_i)$. Since (μ, σ) is χ-consistent, there exists a sequence $\{(\mu^k, \sigma^k)\} \subseteq \Psi^\chi$ such that $(\mu^k, \sigma^k) \to (\mu, \sigma)$ as $k \to \infty$. By

Lemma 1, we know for this $i, \tilde{\theta}$ and h^t,

$$\mu_i^k(\tilde{\theta}_{-i}|h^t, \tilde{\theta}_i) = \chi \mu_i^k(\tilde{\theta}_{-i}|h^{t-1}, \tilde{\theta}_i) + (1-\chi) \left[\frac{\mu_i^k(\tilde{\theta}_{-i}|h^{t-1}, \tilde{\theta}_i)\sigma_{-i}^k(a_{-i}^t|h^{t-1}, \tilde{\theta}_{-i})}{\sum_{\theta'_{-i}} \mu_i^k(\theta'_{-i}|h^{t-1}, \tilde{\theta}_i)\sigma_{-i}^k(a_{-i}^t|h^{t-1}, \theta'_{-i})} \right]$$
$$\geq \chi \mu_i^k(\tilde{\theta}_{-i}|h^{t-1}, \tilde{\theta}_i).$$

As we take the limit $k \to \infty$ on both sides, we can obtain that

$$\mu_i(\tilde{\theta}_{-i}|h^t, \tilde{\theta}_i) = \lim_{k\to\infty} \mu_i^k(\tilde{\theta}_{-i}|h^t, \tilde{\theta}_i) \geq \lim_{k\to\infty} \chi \mu_i^k(\tilde{\theta}_{-i}|h^{t-1}, \tilde{\theta}_i) = \chi \mu_i(\tilde{\theta}_{-i}|h^{t-1}, \tilde{\theta}_i),$$

which yields a contradiction. ■

Proof of Corollary 2: We prove the statement by induction on t. For $t = 1$, by Proposition 3,

$$\mu_i(\theta_{-i}|h^1, \theta_i) \geq \chi \mu_i(\theta_{-i}|h_\emptyset, \theta_i) = \chi \mathcal{F}(\theta_{-i}|\theta_i).$$

Next, suppose there is t' such that the statement holds for all $1 \leq t \leq t' - 1$. At stage t', by Proposition 3 and the induction hypothesis, we can find that

$$\mu_i(\theta_{-i}|h^{t'}, \theta_i) \geq \chi \mu_i(\theta_{-i}|h^{t'-1}, \theta_i) \geq \chi \left[\chi^{t'-1} \mathcal{F}(\theta_{-i}|\theta_i) \right] = \chi^{t'} \mathcal{F}(\theta_{-i}|\theta_i). ■$$

Online Appendix: Omitted Proofs of Section 4

4.1 Pooling Equilibria in Signaling Games

Proof of Proposition 5

Let the assessment (μ, σ) be a pooling χ-CSE. We want to show that for any $\chi' \leq \chi$, the assessment (μ, σ) is also a χ'-CSE. Consider any non-terminal history h^{t-1}, any player i, any $a_i^t \in A_i(h^{t-1})$ and any $\theta \in \Theta$. We can first observe that

$$\begin{aligned}
\bar{\sigma}_{-i}(a_{-i}^t|h^{t-1}, \theta_i) &= \sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)\sigma_{-i}(a_{-i}^t|h^{t-1}, \theta'_{-i}) \\
&= \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i}) \left[\sum_{\theta'_{-i}} \mu_i(\theta'_{-i}|h^{t-1}, \theta_i)\right] \\
&= \sigma_{-i}(a_{-i}^t|h^{t-1}, \theta_{-i})
\end{aligned}$$

where the second equality holds because σ is a pooling behavioral strategy profile, so σ_{-i} is independent of other players' types. For this pooling χ-CSE, let G^σ be the set of on-path histories and $\tilde{G}^\sigma$ be the set of off-path histories. We can first show that for every $h \in G^\sigma$, $i \in N$ and $\theta \in \Theta$,

$$\mu_i(\theta_{-i}|h, \theta_i) = \mathcal{F}(\theta_{-i}|\theta_i).$$

This can be shown by induction on t. For $t = 1$, any $h^1 = (h_\emptyset, a^1)$ and any $\theta \in \Theta$, by Lemma 1, we can obtain that

$$\begin{aligned}
\mu_i(\theta_{-i}|h^1, \theta_i) =& \chi\mu_i(\theta_{-i}|h_\emptyset, \theta_i) + (1-\chi)\left[\frac{\mu_i(\theta_{-i}|h_\emptyset, \theta_i)\sigma_{-i}(a_{-i}^1|h_\emptyset, \theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^1|h_\emptyset, \theta_i)}\right] \\
=& \chi\mathcal{F}(\theta_{-i}|\theta_i) + (1-\chi)\mathcal{F}(\theta_{-i}|\theta_i)\underbrace{\left[\frac{\sigma_{-i}(a_{-i}^1|h_\emptyset, \theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^1|h_\emptyset, \theta_i)}\right]}_{=1} \\
=& \mathcal{F}(\theta_{-i}|\theta_i).
\end{aligned}$$

Now, suppose there is t' such that the statement holds for $1 \leq t \leq t'-1$. At stage t' and $h^{t'} = (h^{t'-1}, a^{t'}) \in G^\sigma$, by Lemma 1 and the induction hypothesis, we can again obtain that the posterior belief is the prior belief

$$
\begin{aligned}
\mu_i(\theta_{-i}|h^{t'},\theta_i) =& \chi\mu_i(\theta_{-i}|h^{t'-1},\theta_i) + (1-\chi)\left[\frac{\mu_i(\theta_{-i}|h^{t'-1},\theta_i)\sigma_{-i}(a_{-i}^{t'}|h^{t'-1},\theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^{t'}|h^{t'-1},\theta_i)}\right] \\
=& \chi\mathcal{F}(\theta_{-i}|\theta_i) + (1-\chi)\mathcal{F}(\theta_{-i}|\theta_i)\underbrace{\left[\frac{\sigma_{-i}(a_{-i}^{t'}|h^{t'-1},\theta_{-i})}{\bar{\sigma}_{-i}(a_{-i}^{t'}|h^{t'-1},\theta_i)}\right]}_{=1} \\
=& \mathcal{F}(\theta_{-i}|\theta_i).
\end{aligned}
$$

Therefore, we have shown that players will not update their beliefs at every on-path information set, so the belief system is independent of χ. Finally, for any off-path history $h^t \in \tilde{G}^\sigma$, by Proposition 3, we can find that the belief system satisfies for any $\theta \in \Theta$,

$$\mu_i(\theta_{-i}|h^t,\theta_i) \geq \chi\mu_i(\theta_{-i}|h^{t-1},\theta_i) \geq \chi'\mu_i(\theta_{-i}|h^{t-1},\theta_i),$$

implying that when $\chi' \leq \chi$, μ will still satisfy the dampened updating property. Therefore, (μ,σ) remains a χ'-CSE. This completes the proof. ■

Proof of Claim 1

First observe that after player 1 chooses B, it is strictly optimal for player 2 to choose R for all beliefs $\mu_2(\theta_1|B)$, and after player 1 chooses A, it is optimal for player 2 to choose L if and only if

$$2\mu_2(\theta_1|A) + [1-\mu_2(\theta_1|A)] \geq 4\mu_2(\theta_1|A) \iff \mu_2(\theta_1|A) \leq 1/3.$$

Equilibrium 1.

If both types of player 1 choose A, then $\mu_2(\theta_1|A) = 1/4$, so it is optimal for player 2 to choose L. Given $a(A) = L$ and $a(B) = R$, it is optimal for both types of player 1 to choose A as $2 > 1$. Hence $m(\theta_1) = m(\theta_2) = A$, $a(A) = L$ and $a(B) = R$ is a pooling χ-CSE for any $\chi \in [0,1]$.

Equilibrium 2.

In order to support $m(\theta_1) = m(\theta_2) = B$ to be an equilibrium, player 2 has to choose R at the off-path information set A, which is optimal if and only if $\mu_2(\theta_1|A) \geq$

1/3. In addition, by Proposition 3, we know in a χ-CSE, the belief system satisfies

$$\mu_2(\theta_2|A) \geq \frac{3}{4}\chi \iff \mu_2(\theta_1|A) \leq 1 - \frac{3}{4}\chi.$$

Therefore, the belief system has to satisfy that $\mu_2(\theta_1|A) \in \left[\frac{1}{3}, 1 - \frac{3}{4}\chi\right]$, which requires $\chi \leq 8/9$.

Finally, it is straightforward to verify that for any $\mu \in \left[\frac{1}{3}, 1 - \frac{3}{4}\chi\right]$, $\mu_2(\theta_1|A) = \mu$ satisfies χ-consistency. Suppose type θ_1 player 1 chooses A with probability p and type θ_2 player 1 chooses A with probability q where $p, q \in (0, 1)$. Given this behavioral strategy profile for player 1, by Lemma 1, we have:

$$\mu_2(\theta_1|A) = \frac{1}{4}\chi + (1 - \chi)\left[\frac{p}{p + 3q}\right].$$

In other words, as long as (p, q) satisfies

$$q = \left[\frac{4 - 4\mu - 3\chi}{12 - 3\chi}\right] p,$$

we can find that $\mu_2(\theta_1|A) = \mu$. Therefore, if $\{(p^k, q^k)\} \to (0, 0)$ such that

$$q^k = \left[\frac{4 - 4\mu - 3\chi}{12 - 3\chi}\right] p^k,$$

then $\mu_2^k(\theta_1|A) = \mu$ for all k. Hence, $\lim_{k\to\infty} \mu_2^k(\theta_1|A) = \mu$, suggesting that $\mu_2(\theta_1|A) = \mu$ is indeed χ-consistent. This completes the proof. ■

Proof of Proposition 6

Here we provide a characterization of χ-CSE of Game 1 and Game 2. For the analysis of both games, we denote $\mu_I \equiv \mu_2(\theta_1|m = I)$ and $\mu_S \equiv \mu_2(\theta_1|m = S)$.

Analysis of Game BH 3.

At information set S, given μ_S, the expected payoffs of C, D, E are $90\mu_S$, $30-15\mu_S$ and 15, respectively. Therefore, for any μ_S, E is never a best response. Moreover, C is the best response if and only if $90\mu_S \geq 30 - 15\mu_S$ or $\mu_S \geq 2/7$. Similarly, at information set I, given μ_I, the expected payoffs of C, D, E are 30, $45 - 45\mu_I$ and

15, respectively. Therefore, E is strictly dominated, and C is the best response if and only if $30 \geq 45 - 45\mu_I$ or $\mu_I \geq 1/3$. Now we consider four cases.

Case 1 $[m(\theta_1) = I, m(\theta_2) = S]$**:**

By Lemma 1, $\mu_I = 1-\chi/2$ and $\mu_S = \chi/2$. Moreover, since $\mu_I = 1-\chi/2 \geq 1/2$ for any χ, player 2 will choose C at information set I. To support this equilibrium, player 2 has to choose C at information set S. In other words, $[(I, S); (C, C)]$ is separating χ-CSE if and only if $\mu_S \geq 2/7$ or $\chi \geq 4/7$.

Case 2 $[m(\theta_1) = S, m(\theta_2) = I]$**:**

By Lemma 1, $\mu_I = \chi/2$ and $\mu_S = 1 - \chi/2$. Because $\mu_S \geq 1 - \chi/2 \geq 1/2$, it is optimal for player 2 to choose C at information set S. To support this as an equilibrium, player 2 has to choose D at information set I. Yet, in this case, type θ_2 player 1 will deviate to S. Therefore, this profile cannot be supported as an equilibrium.

Case 3 $[m(\theta_1) = I, m(\theta_2) = I]$**:**

Since player 1 follows a pooling strategy, player 2 will not update his belief at information set I, i.e., $\mu_I = 1/2$. χ-dampened updating property implies $\chi/2 \leq \mu_S \leq 1-\chi/2$. Since $\mu_I > 1/3$, player 2 will choose C at information set I. To support this profile to be an equilibrium, player 2 has to choose D at information set S, and hence, it must be the case that $\mu_S \leq 2/7$. Coupled with the requirement from χ-dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_S \leq 2/7$. That is, $[(I, I); (C, D)]$ is pooling χ-CSE if and only if $\chi/2 \leq 2/7$ or $\chi \leq 4/7$.

Case 4 $[m(\theta_1) = S, m(\theta_2) = S]$**:**

Similar to the previous case, since player 1 follows a pooling strategy, player 2 will not update his belief at information set S, i.e., $\mu_S = 1/2$. Also, the χ-dampened updating property suggests $\chi/2 \leq \mu_I \leq 1 - \chi/2$. Because $\mu_S > 2/7$, it is optimal for player 2 to choose C at information set S. To support this as an equilibrium, player 2 has to choose D at information set I. Therefore, it must be that $\mu_I \leq 1/3$. Combined with the requirement of χ-dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_I \leq 1/3$. As a result, $[(S, S); (D, C)]$ is a pooling χ-CSE if and only if $\chi \leq 2/3$.

Analysis of Game BH 4.

At information set I, given μ_I, the expected payoffs of C, D, E are 30, $45 - 45\mu_I$ and $35\mu_I$. Hence, D is the best response if and only if $\mu_I \leq 1/3$ while E is the best response if $\mu_I \geq 6/7$. For $1/3 \leq \mu_I \leq 6/7$, C is the best response. On the other hand, since player 2's payoffs at information set S are the same as in Game 1, player 2 will adopt the same decision rule—player 2 will choose C if and only if $\mu_S \geq 2/7$, and choose D if and only if $\mu_S \leq 2/7$. Now, we consider the following four cases.

Case 1 $[m(\theta_1) = I, m(\theta_2) = S]$**:**

In this case, by Lemma 1, $\mu_I = 1 - \chi/2$ and $\mu_S = \chi/2$. To support this profile to be an equilibrium, player 2 has to choose E and C at information set I and S, respectively. To make it profitable for player 2 to choose E at information set I, it must be that:

$$\mu_I = 1 - \chi/2 \geq 6/7 \iff \chi \leq 2/7.$$

On the other hand, player 2 will choose C at information set S if and only if $\chi/2 \geq 2/7$ or $\chi \geq 4/7$, which is not compatible with the previous inequality. Therefore, this profile cannot be supported as an equilibrium.

Case 2 $[m(\theta_1) = S, m(\theta_2) = I]$**:**

In this case, by Lemma 1, $\mu_I = \chi/2$ and $\mu_S = 1 - \chi/2$. To support this as an equilibrium, player 2 has to choose D at both information sets. Yet, $\mu_S = 1 - \chi/2 > 2/7$, implying that it is not a best reply for player 2 to choose D at information set S. Hence this profile also cannot be supported as an equilibrium.

Case 3 $[m(\theta_1) = I, m(\theta_2) = I]$**:**

Since player 1 follows a pooling strategy, player 2 will not update his belief at information set I, i.e., $\mu_I = 1/2$. The χ-dampened updating property implies $\chi/2 \leq \mu_S \leq 1 - \chi/2$. Because $1/3 < \mu_I = 1/2 < 6/7$, player 2 will choose C at information set I. To support this profile as an equilibrium, player 2 has to choose D at information set S, and hence, it must be the case that $\mu_S \leq 2/7$. Coupled with the requirement of χ-dampened updating, the off-path belief has to satisfy $\chi/2 \leq \mu_S \leq 2/7$. That is, $[(I, I); (C, D)]$ is pooling χ-CSE if and only if $\chi/2 \leq 2/7$ or $\chi \leq 4/7$.

Case 4 $[m(\theta_1) = S, m(\theta_2) = S]$**:**

Similar to the previous case, since player 1 follows a pooling strategy, player 2 will not update his belief at information set S, i.e., $\mu_S = 1/2$. Also, the χ-dampened updating property implies $\chi/2 \le \mu_I \le 1 - \chi/2$. Because $\mu_S > 2/7$, it is optimal for player 2 to choose C at information set S. To support this as an equilibrium, player 2 can choose either C or D at information set I.

Case 4.1: To make it a best reply for player 2 to choose D at information set I, it must be that $\mu_I \le 1/3$. Combined with the requirement from χ-dampened updating, the off-path belief has to satisfy $\chi/2 \le \mu_I \le 1/3$. As a result, $[(S, S); (D, C)]$ is a pooling χ-CSE if and only if $\chi \le 2/3$.

Case 4.2: To make it a best reply for player 2 to choose C at information set I, it must be that $1/3 \le \mu_I \le 6/7$. Combined with the requirement from χ-dampened updating, the off-path belief has to satisfy

$$\max\left\{\frac{1}{2}\chi,\ \frac{1}{3}\right\} \le \mu_I \le \min\left\{\frac{6}{7}, 1 - \frac{1}{2}\chi\right\}.$$

For any $\chi \in [0, 1]$, one can find μ_I that satisfies both inequalities. Hence $[(S, S); (C, C)]$ is a pooling χ-CSE for any χ.

This completes the analysis of Game BH 3 and Game BH 4. ■

4.2 A Public Goods Game with Communication

Proof of Proposition 7

To prove this set of cost cutoffs form a χ-CSE, we need to show that there is no profitable deviation for any type at any subgame. First, at the second stage where there are exactly $0 \le k \le N-1$ players sending 1 in the first stage, since no players will contribute, setting $C_k^\chi = 0$ is indeed a best response. At the subgame where all N players send 1 in the first stage, we use $\mu_i^\chi(c_{-i}|N)$ to denote player i's cursed belief density. By Lemma 1, the cursed belief about all other players having a cost lower than c is simply:

$$
\begin{aligned}
F^{\chi}(c) &\equiv \int_{\{c_j \leq c,\ \forall j \neq i\}} \mu_i^{\chi}(c'_{-i}|N) dc'_{-i} \\
&= \begin{cases} \chi \left(c/K\right)^{N-1} + (1-\chi)\left(c/C_c^{\chi}\right)^{N-1} & \text{if } c \leq C_c^{\chi} \\ 1-\chi+\chi\left(c/C_c^{\chi}\right)^{N-1} & \text{if } c > C_c^{\chi}, \end{cases}
\end{aligned}
$$

and C_N^{χ} is the solution of the fixed point problem of $C_N^{\chi} = F^{\chi}(C_N^{\chi})$.

Moreover, in equilibrium, C_c^{χ} type of players would be indifferent between sending 1 and 0 in the communication stage. Thus, given C_N^{χ}, C_c^{χ} is the solution of the following equation

$$
0 = \left(\frac{C_c^{\chi}}{K}\right)^{N-1} \left[-C_c^{\chi} + F^{\chi}(C_N^{\chi})\right].
$$

As a result, we obtain that in equilibrium, $C_c^{\chi} = C_N^{\chi} = F^{\chi}(C_N^{\chi}) \leq 1$ and denote this cost cutoff by $C^*(N, K, \chi)$. Substituting it into $F^{\chi}(c)$, gives:

$$
C^*(N, K, \chi) - \chi \left[\frac{C^*(N, K, \chi)}{K}\right]^{N-1} = 1 - \chi.
$$

In the following, we show that for any $N \geq 2$ and χ, the cutoff $C^*(N, K, \chi)$ is unique.

Case 1: When $N = 2$, the cutoff $C^*(2, K, \chi)$ is the unique solution of the linear equation

$$
C^*(2, K, \chi) - \chi \left[\frac{C^*(2, K, \chi)}{K}\right] = 1 - \chi \iff C^*(2, K, \chi) = \frac{K - K\chi}{K - \chi}.
$$

Case 2: For $N \geq 3$, we define the function $h(y) : [0, 1] \to \mathbb{R}$ where

$$
h(y) = y - \chi \left(\frac{y}{K}\right)^{N-1} - (1 - \chi).
$$

It suffices to show that $h(y)$ has a unique root in $[0, 1]$. When $\chi = 0$, $h(y) = y - 1$ which has a unique root at $y = 1$. In the following, we will focus on the case where $\chi > 0$. Since $h(y)$ is continuous, $h(0) = -(1 - \chi) < 0$ and $h(1) = \chi \left[1 - (1/K)^{N-1}\right] > 0$, there exists a root $y^* \in (0, 1)$ by the intermediate value theorem. Moreover, as we take the second derivative, we can find that for any $y \in (0,1)$,

$$h''(y) = -\left(\frac{\chi}{K^{N-1}}\right)(N-1)(N-2)y^{N-3} < 0,$$

implying that $h(y)$ is strictly concave in $[0,1]$. Furthermore, $h(0) < 0$ and $h(1) > 0$, so the root is unique, as illustrated in the left panel of Figure 3. This completes the proof. ■

Proof of Corollary 3

By Proposition 7, we know the cutoff $C^*(N,K,\chi) \leq 1$ and it satisfies

$$C^*(N,K,\chi) - \chi\left[\frac{C^*(N,K,\chi)}{K}\right]^{N-1} = 1 - \chi.$$

Therefore, when $\chi = 0$, the condition becomes $C^*(N,K,0) = 1$. In addition, when $\chi = 1$, the condition becomes

$$C^*(N,K,1) - \left[\frac{C^*(N,K,1)}{K}\right]^{N-1} = 0,$$

implying $C^*(N,K,1) = 0$.

For $\chi \in (0,1)$, to prove $C^*(N,K,\chi)$ is strictly decreasing in N, K and χ, we consider a function $g(y;N,K,\chi) : (0,1) \to \mathbb{R}$ where $g(y;N,K,\chi) = y - \chi[y/K]^{N-1}$. For any $y \in (0,1)$ and fix any K and χ, we can observe that when $N \geq 2$,

$$g(y;N+1,K) - g(y;N,K) = -\chi\left[\frac{y}{K}\right]^{N} + \chi\left[\frac{y}{K}\right]^{N-1} > 0,$$

so $g(\cdot;N,K,\chi)$ is strictly increasing in N. Therefore, the cutoff $C^*(N,K,\chi)$ is strictly decreasing in N. Similarly, for any $y \in (0,1)$ and fix any N and χ, observe that when $K > 1$,

$$\frac{\partial g}{\partial K} = \chi(N-1)\left(\frac{y^{N-1}}{K^N}\right) > 0,$$

which implies that cutoff $C^*(N,K,\chi)$ is also strictly decreasing in K. For the comparative statics of χ, we can rearrange the equilibrium condition where

$$\frac{1 - C^*(N, K, \chi)}{\chi} = 1 - \left[\frac{C^*(N, K, \chi)}{K}\right]^{N-1}.$$

Since LHS is strictly decreasing in χ, the equilibrium cutoff is also strictly decreasing in χ. Finally, taking the limit on both sides of the equilibrium condition, we obtain:

$$\lim_{N\to\infty} C^*(N, K, \chi) = \lim_{K\to\infty} C^*(N, K, \chi) = 1 - \chi.$$

This completes the proof. ■

4.3 The Centipede Game with Altruistic Types

Proof of Claim 2

By backward induction, we know selfish player two will choose $T4$ for sure. Given that player two will choose $T4$ at stage four, it is optimal for selfish player one to choose $T3$. Now, suppose selfish player one will choose $P1$ with probability q_1 and player two will choose $P2$ with probability q_2. Given this behavioral strategy profile, player two's belief about the other player being altruistic at stage two is:

$$\mu = \frac{\alpha}{\alpha + (1 - \alpha)q_1}.$$

In this case, it is optimal for selfish player two to pass if and only if

$$32\mu + 4(1 - \mu) \geq 8 \iff \mu \geq \frac{1}{7}.$$

At the equilibrium, selfish player two is indifferent between $T2$ and $P2$. If not, say $32\mu + 4(1 - \mu) > 8$, player two will choose $P2$. Given that player two will choose $P2$, it is optimal for selfish player one to choose $P1$, which makes $\mu = \alpha$ and $\alpha > 1/7$. However, we know $\alpha \leq 1/7$ which yields a contradiction. On the other hand, if $32\mu + 4(1 - \mu) < 8$, then it is optimal for player two to choose $T2$ at stage two. As a result, selfish player one would choose $T1$ at stage one, causing $\mu = 1$. In this case, player two would deviate to choose $P2$, which again yields a contradiction. To summarize, in equilibrium, player two has to be indifferent between $T2$ and $P2$, i.e.,

$\mu = 1/7$. As we rearrange the equality, we can obtain that

$$\frac{\alpha}{\alpha + (1-\alpha)q_1^*} = \frac{1}{7} \iff q_1^* = \frac{6\alpha}{1-\alpha}.$$

Finally, since the equilibrium requires selfish player one to mix at stage one, selfish player one has to be indifferent between $P1$ and $T1$. Therefore,

$$4 = 16q_2^* + 2(1-q_2^*) \iff q_2^* = \frac{1}{7}.$$

This completes the proof. ■

Proof of Proposition 8

By backward induction, we know selfish player two will choose $T4$ for sure. Given this, it is optimal for selfish player one to choose $T3$. Now, suppose selfish player one will choose $P1$ with probability q_1 and player two will choose $P2$ with probability q_2. Given this behavioral strategy profile, by Lemma 1, player two's *cursed* belief about the other player being altruistic at stage 2 is:

$$\mu^\chi = \chi\alpha + (1-\chi)\left[\frac{\alpha}{\alpha + (1-\alpha)q_1}\right].$$

In this case, it is optimal for player two to pass if and only if

$$32\mu^\chi + 4(1-\mu^\chi) \geq 8 \iff \mu^\chi \geq \frac{1}{7}.$$

We can first show that in equilibrium, it must be that $\mu^\chi \leq 1/7$. If not, then it is strictly optimal for player two to choose $P2$. Therefore, it is optimal for selfish player one to choose $P1$ and hence $\mu^\chi = \alpha \leq 1/7$, which yields a contradiction. In the following, we separate the discussion into two cases.

Case 1: $\chi \leq \frac{6}{7(1-\alpha)}$

In this case, we argue that player two is indifferent between $P2$ and $T2$. If not, then $32\mu^\chi + 4(1-\mu^\chi) < 8$ and it is strictly optimal for player two to choose $T2$. This would cause selfish player one to choose $T1$ and hence $\mu^\chi = 1-(1-\alpha)\chi$. This yields a contradiction because

$$\mu^{\chi} = 1 - (1-\alpha)\chi < \frac{1}{7} \iff \chi > \frac{6}{7(1-\alpha)}.$$

Therefore, in this case, player two is indifferent between $T2$ and $P2$ and thus,

$$\begin{aligned}
\mu^{\chi} = \frac{1}{7} &\iff \chi\alpha + (1-\chi)\left[\frac{\alpha}{\alpha + (1-\alpha)q_1^{\chi}}\right] = \frac{1}{7} \\
&\iff \chi + \frac{1-\chi}{\alpha + (1-\alpha)q_1^{\chi}} = \frac{1}{7\alpha} \\
&\iff \alpha + (1-\alpha)q_1^{\chi} = (1-\chi)\Big/\left[\frac{1}{7\alpha} - \chi\right] \\
&\iff q_1^{\chi} = \left[\frac{7\alpha - 7\alpha\chi}{1 - 7\alpha\chi} - \alpha\right]\Big/(1-\alpha).
\end{aligned}$$

Since the equilibrium requires selfish player one to mix at stage 1, selfish player one has to be indifferent between $P1$ and $T1$. Therefore,

$$4 = 16q_2^{\chi} + 2(1 - q_2^{\chi}) \iff q_2^{\chi} = \frac{1}{7}.$$

Case 2: $\chi > \frac{6}{7(1-\alpha)}$

In this case, we know for any $q_1^{\chi} \in [0,1]$,

$$\mu^{\chi} = \chi\alpha + (1-\chi)\left[\frac{\alpha}{\alpha + (1-\alpha)q_1^{\chi}}\right] \leq 1 - (1-\alpha)\chi < \frac{1}{7},$$

implying that it is strictly optimal for player two to choose $T2$, and hence it is strictly optimal for selfish player one to choose $T1$ at stage 1. This completes the proof. ■

4.4 Sequential Voting over Binary Agendas

Proof of Proposition 9

If $a^1(\theta_1) = b$ and all other types of voters as well as type θ_1 at stage 2 vote sincerely, voter i's χ-cursed belief in the second stage upon observing $a^1_{-i} = (a, b)$ is

$$\mu_i^{\chi}(\theta_{-i}|a_{-i}^1=(a,b))=\begin{cases} p_1p_3\chi+\frac{p_1}{p_1+p_2}(1-\chi) & \text{if } \theta_{-i}=(\theta_3,\theta_1)\\ p_2p_3\chi+\frac{p_2}{p_1+p_2}(1-\chi) & \text{if } \theta_{-i}=(\theta_3,\theta_2)\\ p_kp_l\chi & \text{otherwise.}\end{cases}$$

As mentioned in Section 4.4, a voter would act as if he perceives the other voters' (behavioral) strategies correctly in the last stage. However, misunderstanding the link between the other voters' types and actions would distort a voter's belief updating process. In other words, a voter would perceive the strategies correctly but form beliefs incorrectly. As a result, the continuation value of the a vs c subgame to a type θ_1 voter is simply the voter's χ-cursed belief, conditional on being pivotal, about there being at least one type θ_1 voter among his opponents. Similarly, the continuation value of the b vs c subgame is equal to the voter's conditional χ-cursed belief about there being at least one type θ_1 or θ_2 voter among his opponents multiplied by v. Therefore, the continuation values to a type θ_1 voter in the two possible subgames of the second stage are (let $\tilde{p}_2\equiv\frac{p_1}{p_1+p_2}$):

$$\begin{aligned} a \text{ vs } c: &\quad \chi\left(1-(1-p_1)^2\right)+(1-\chi)\tilde{p}_2\\ b \text{ vs } c: &\quad \left(1-p_3^2\chi\right)v \end{aligned}$$

It is thus optimal for a type θ_1 voter to vote for b in the first stage if

$$\begin{aligned} &\chi\left(1-(1-p_1)^2\right)+(1-\chi)\tilde{p}_2\leq\left(1-p_3^2\chi\right)v\\ \iff &[2p_1-p_1^2-\tilde{p}_2+p_3^2v]\chi\leq v-\tilde{p}_2 \end{aligned} \tag{2}$$

Notice that the statement would automatically hold when $\chi=0$. In the following, we want to show that given v and p, if condition (2) holds for some $\chi\in(0,1]$, then it will hold for all $\chi'\leq\chi$. As $\chi>0$, we can rewrite condition (2) as

$$2p_1-p_1^2-\tilde{p}_2+p_3^2v\leq\frac{v-\tilde{p}_2}{\chi}. \tag{2'}$$

Case 1: $v-\tilde{p}_2<0$.

In this case, we want to show that voting b in the first stage is never optimal for type θ_1 voter. That is, we want to show condition (2') never holds for $v < \tilde{p}_2$. To see this, we can first observe that the RHS is strictly increasing in χ. Therefore, it suffices to show

$$2p_1 - p_1^2 - \tilde{p}_2 + p_3^2 v > v - \tilde{p}_2.$$

This is true because

$$\begin{aligned} 2p_1 - p_1^2 - \tilde{p}_2 + p_3^2 v - (v - \tilde{p}_2) &= 2p_1 - p_1^2 - (1 - p_3^2)v \\ &> 2p_1 - p_1^2 - (1 + p_3)p_1 = p_1 p_2 \geq 0 \end{aligned}$$

where the second inequality holds as $v < \frac{p_1}{p_1+p_2}$.

Case 2: $v - \tilde{p}_2 \geq 0$.

Since the RHS of condition (2') is greater or equal to 0, it will weakly increase as χ decreases. Thus, if condition (2') holds for some $\chi \in (0,1]$, it will also hold for all $\chi' \leq \chi$. This completes the proof. ■

Proof of Proposition 10

Assuming that all voters vote sincerely in both stages, voter i's χ-cursed belief in the second stage upon observing $a^1_{-i} = (a, b)$ is

$$\mu_i^\chi(\theta_{-i} | a^1_{-i} = (a,b)) = \begin{cases} p_1 p_2 \chi + \frac{p_1}{p_1+p_3}(1-\chi) & \text{if } \theta_{-i} = (\theta_1, \theta_2) \\ p_2 p_3 \chi + \frac{p_3}{p_1+p_3}(1-\chi) & \text{if } \theta_{-i} = (\theta_3, \theta_2) \\ p_k p_l \chi & \text{otherwise.} \end{cases}$$

Similar to the proof of Proposition 9, the continuation values to a type θ_1 voter in the two possible subgames of the second stage are (let $\tilde{p}_3 \equiv \frac{p_1}{p_1+p_3}$):

$$\begin{aligned} a \text{ vs } c: &\quad \chi\left(1 - (1-p_1)^2\right) + (1-\chi)\tilde{p}_3 \\ b \text{ vs } c: &\quad \left(1 - p_3^2 \chi\right) v \end{aligned}$$

Thus, it is optimal for a type θ_1 voter to vote for a in the first stage if

$$\begin{aligned}
&\chi\left(1-(1-p_1)^2\right)+(1-\chi)\tilde{p}_3 \geq \left(1-p_3^2\chi\right)v \\
\iff &\chi\left(2p_1-p_1^2-\tilde{p}_3+p_3^2 v\right) \geq v-\tilde{p}_3.
\end{aligned} \tag{3}$$

Case 1: $v-\tilde{p}_3>0$.

In this case, we want to show that given p and v, there exists $\tilde{\chi}$ such that condition (3) holds if and only if $\chi \geq \tilde{\chi}$. Let $\tau \equiv 2p_1-p_1^2-\tilde{p}_3+p_3^2 v$. If $\tau>0$, then condition (3) holds if and only if $\chi \geq \tilde{\chi} \equiv \frac{v-\tilde{p}_3}{\tau}$. On the other hand, if $\tau \leq 0$, condition (3) will not hold for all $\chi \in [0,1]$ and hence we can set $\tilde{\chi}=2$.

Case 2: $v-\tilde{p}_3 \leq 0$.

In this case, we want to show that given p and v, there exists $\tilde{\chi}$ such that condition (3) holds if and only if $\chi \leq \tilde{\chi}$. If $\tau<0$, then condition (3) holds if and only if $\chi \leq \frac{v-\tilde{p}_3}{\tau}$ where the RHS is greater or equal to 0. On the other hand, if $\tau \geq 0$, then condition (3) will hold for any $\chi \in [0,1]$ and hence we can again set $\tilde{\chi}=2$. This completes the proof. ■

4.5 The Dirty Faces Game

Proof of Proposition 11

When observing a clean face, a player will know that he has a dirty face immediately. Therefore, choosing 1 (i.e., choosing D at stage 1) when observing a clean face is a strictly dominant strategy. In other words, for any $\chi \in [0,1]$, $\hat{\sigma}^{\chi}(O)=1$.

The analysis of the case where the player observes a dirty face is separated into two cases.

Case 1: $\chi > \bar{\alpha}$

In this case, we show that $\hat{\sigma}^{\chi}(X)=T+1$ is the only χ-CE. If not, suppose $\hat{\sigma}^{\chi}(X)=t$ where $t \leq T$ can be supported as a χ-CE. We can first notice that $\hat{\sigma}^{\chi}(X)=1$ cannot be supported as a χ-CE because it is strictly dominated to choose 1 when observing a dirty face. For $2 \leq t \leq T$, given the other player $-i$ chooses

$\hat{\sigma}^\chi(X) = t$, we can find player $-i$'s *average strategy* is

$$\bar{\sigma}_{-i}(j) = \begin{cases} 1-p & \text{if} \quad j=1 \\ p & \text{if} \quad j=t \\ 0 & \text{if} \quad j \neq 1, t. \end{cases}$$

Therefore, the other player $-i$'s χ-cursed strategy is:

$$\sigma^\chi_{-i}(j|x_i = O) = \begin{cases} \chi(1-p) + (1-\chi) & \text{if} \quad j=1 \\ \chi p & \text{if} \quad j=t \\ 0 & \text{if} \quad j \neq 1, t, \end{cases} \qquad \text{and}$$

$$\sigma^\chi_{-i}(j|x_i = X) = \begin{cases} \chi(1-p) & \text{if} \quad j=1 \\ \chi p + (1-\chi) & \text{if} \quad j=t \\ 0 & \text{if} \quad j \neq 1, t. \end{cases}$$

In this case, given (player i perceives that) player $-i$ chooses the χ-cursed strategy, player i's expected payoff to choose $2 \leq j \leq t$ when observing a dirty face is:

$$(1-p)\left[-\delta^{j-1}\chi p\right] + p\left\{\delta^{j-1}\alpha\left[\chi p + (1-\chi)\right]\right\} = p\delta^{j-1}\underbrace{\left[\alpha - \chi(1+\alpha)(1-p)\right]}_{<0 \iff \chi > \bar{\alpha}} < 0.$$

Hence, given the other player chooses t when observing a dirty face, it is strictly dominated to choose any $j \leq t$. Therefore, the only χ-CE is $\hat{\sigma}^\chi(X) = T+1$.

Case 2: $\chi < \bar{\alpha}$

In this case, we want to show that $\hat{\sigma}^\chi(X) = 2$ is the only χ-CE. If not, suppose $\hat{\sigma}(X) = t$ for some $t \geq 3$ can be supported as a χ-CE. We can again notice that since when observing a dirty face, it is strictly dominated to choose 1, 1 is never a best response. Given player $-i$ chooses $\hat{\sigma}^\chi(X) = t$, by the same calculation as in **Case 1**, the expected payoff to choose $2 \leq j \leq t$ is:

$$p\delta^{j-1}\underbrace{\left[\alpha - \chi(1+\alpha)(1-p)\right]}_{>0 \iff \chi < \bar{\alpha}} > 0,$$

which is decreasing in j. Therefore, the best response to $\hat{\sigma}^{\chi}(X) = t$ is to choose 2 when observing a dirty face. As a result, the only χ-CE in this case is $\hat{\sigma}^{\chi}(X) = 2$. This completes the proof. ■

Proof of Proposition 12

When observing a clean face, the player would know that his face is dirty. Thus, choosing D at stage 1 is a strictly dominant strategy, and $\tilde{\sigma}^{\chi}(O) = 1$ for all $\chi \in [0, 1]$. On the other hand, the analysis for the case where the player observes a dirty face consists of several steps.

Step 1: Assume that both players choosing D at some stage $\bar{t}$. We claim that at stage $t \leq \bar{t}$, the cursed belief $\mu^{\chi}(X|t, X) = 1 - (1-p)\chi^{t-1}$. We can prove this by induction on t. At stage $t = 1$, the belief about having a dirty face is simply the prior belief p. Hence this establishes the base case. Now suppose the statement holds for any stage $1 \leq t \leq t'$ (and $t' < \bar{t}$). At stage $t' + 1$, by Lemma 1,

$$\begin{aligned}
\mu^{\chi}(X|t'+1, X) &= \chi \mu^{\chi}(X|t', X) + (1-\chi) \\
&= \chi \left[1 - (1-p)\chi^{t'-1}\right] + (1-\chi) \\
&= 1 - (1-p)\chi^{t'}
\end{aligned}$$

where the second equality holds by the induction hypothesis. This proves the claim.

Step 2: Given the cursed belief computed in the previous step, the expected payoff to choose D at stage t is:

$$\begin{aligned}
\mu^{\chi}(X|t, X)\alpha - [1 - \mu^{\chi}(X|t, X)] &= \left[1 - (1-p)\chi^{t-1}\right]\alpha - \left[(1-p)\chi^{t-1}\right] \\
&= \alpha - (1-p)(1+\alpha)\chi^{t-1},
\end{aligned}$$

which is increasing in t. Notice that at the first stage, the expected payoff is $\alpha - (1-p)(1+\alpha) < 0$ by Assumption (1), so choosing U at stage 1 is strictly dominated. Furthermore, the player would choose U at every stage when observing a dirty face if and only if

$$\mu^{\chi}(X|T,X)\alpha - [1-\mu^{\chi}(X|T,X)] \le 0 \iff \alpha - (1-p)(1+\alpha)\chi^{T-1} \le 0$$
$$\iff \chi \ge \bar{\alpha}^{\frac{1}{T+1}}.$$

As a result, both players choosing $\tilde{\sigma}^{\chi}(X) = T+1$ is a χ-CSE if and only if $\chi \ge \bar{\alpha}^{\frac{1}{T+1}}$.

Step 3: In this step, we show both players choosing $\tilde{\sigma}^{\chi}(X) = 2$ is a χ-CSE if and only if $\chi \le \bar{\alpha}$. We can notice that given the other player chooses D at stage 2, the player would know stage 2 would be the last stage regardless of his face type. Therefore, it is optimal to choose D at stage 2 as long as the expected payoff of D at stage 2 is positive. Consequently, both players choosing $\tilde{\sigma}^{\chi}(X) = 2$ is a χ-CSE if and only if

$$\mu^{\chi}(X|2,X)\alpha - [1-\mu^{\chi}(X|2,X)] \ge 0 \iff \alpha - (1-p)(1+\alpha)\chi$$
$$\iff \chi \le \bar{\alpha}.$$

Step 4: Given the other player chooses $\tilde{\sigma}^{\chi}(X) > t$, as the game reaches stage t, the belief about the other player choosing U at stage t is:

$$\underbrace{\mu^{\chi}(X|t,X)}_{\text{prob. of dirty}} \; [\chi\mu^{\chi}(X|t,X) + (1-\chi)]$$
$$+ \underbrace{[1-\mu^{\chi}(X|t,X)]}_{\text{prob. of clean}} [\chi\mu^{\chi}(X|t,X)] = \mu^{\chi}(X|t,X).$$

Furthermore, we denote the expected payoff of choosing D at stage t as

$$\mathbb{E}\left[u^{\chi}(D|t,X)\right] \equiv \mu^{\chi}(X|t,X)\alpha - (1-\mu^{\chi}(X|t,X)).$$

In the following, we claim that for any stage $2 \le t \le T-2$, given the other player will stop at some stage later than stage $t+2$ or never stop, if it is optimal to choose

U at stage $t+1$, then it is also optimal for you to choose U at stage t. That is,

$$\begin{aligned}
&\mathbb{E}\left[u^{\chi}(D|t+1, X)\right] < \delta\mu^{\chi}(X|t+1, X)\mathbb{E}\left[u^{\chi}(D|t+2, X)\right] \\
\Longrightarrow\ &\mathbb{E}\left[u^{\chi}(D|t, X)\right] < \delta\mu^{\chi}(X|t, X)\mathbb{E}\left[u^{\chi}(D|t+1, X)\right].
\end{aligned}$$

To prove this claim, first observe that

$$\begin{aligned}
&\mathbb{E}\left[u^{\chi}(D|t+1, X)\right] < \delta\mu^{\chi}(X|t+1, X)\mathbb{E}\left[u^{\chi}(D|t+2, X)\right] \\
\iff\ &(1+\alpha)\mu^{\chi}(X|t+1, X) - 1 < \delta\mu^{\chi}(X|t+1, X)\left[(1+\alpha)\mu^{\chi}(X|t+2, X) - 1\right].
\end{aligned}$$

After rearrangement, the inequality is equivalent to

$$\delta\chi\left[\mu^{\chi}(X|t+1, X)\right]^2 + \left[\delta(1-\chi) - \frac{\delta}{1+\alpha} - 1\right]\mu^{\chi}(X|t+1, X) + \frac{1}{1+\alpha} > 0.$$

Consider a function $F : [0, 1] \to \mathbb{R}$ where

$$F(y) = \delta\chi y^2 + \left[\delta(1-\chi) - \frac{\delta}{1+\alpha} - 1\right]y + \frac{1}{1+\alpha}.$$

Since $\mu^{\chi}(X|j, X) = 1 - (1-p)\chi^{j-1}$ is increasing in j, it suffices to complete the proof of the claim by showing there exists a unique $y^* \in (0, 1)$ such that F is single-crossing on $[0, 1]$ where $F(y^*) = 0$, $F(y) < 0$ for all $y > y^*$, and $F(y) > 0$ for all $y < y^*$. Because F is continuous and

- $F(0) = \frac{1}{1+\alpha} > 0$,
- $F(1) = \delta\chi + \left[\delta(1-\chi) - \frac{\delta}{1+\alpha} - 1\right] + \frac{1}{1+\alpha} = -\frac{\alpha(1-\delta)}{1+\alpha} < 0$.

By intermediate value theorem, there exists a $y^* \in (0, 1)$ such that $F(y^*) = 0$. Moreover, y^* is the unique root of F on $[0, 1]$ because F is a strictly convex parabola and $F(1) < 0$. This establishes the claim.

Step 5: For any $3 \leq t \leq T$, in this step, we find the conditions to support both players choosing $\tilde{\sigma}^{\chi}(X) = t$ as a χ-CSE. We can first notice that both players choosing $\tilde{\sigma}^{\chi}(X) = t$ is a χ-CSE if and only if

1. $\mathbb{E}\left[u^{\chi}(D|t, X)\right] \geq 0$

2. $\mathbb{E}\left[u^\chi(D|t-1,X)\right] \leq \delta\mu^\chi(X|t-1,X)\mathbb{E}\left[u^\chi(D|t,X)\right]$.

Condition 1 is necessary because if it fails, then it is better for the player to choose U at stage t and get at least 0. Condition 2 is also necessary because if the condition doesn't hold, it would be profitable for the player to choose D before stage t. Furthermore, these two conditions are jointly sufficient to support $\tilde{\sigma}^\chi(X) = t$ as a χ-CSE by the same argument as **step 3**. From condition 1, we can obtain that

$$\begin{aligned}\mathbb{E}\left[u^\chi(D|t,X)\right] \geq 0 &\iff (1+\alpha)\mu^\chi(X|t,X) - 1 \geq 0 \\ &\iff 1-(1-p)\chi^{t-1} \geq \frac{1}{1+\alpha} \iff \chi \leq \bar{\alpha}^{\frac{1}{t-1}}.\end{aligned}$$

In addition, by the calculation of **step 4**, we know

$$\mathbb{E}\left[u^\chi(D|t-1,X)\right] \leq \delta\mu^\chi(X|t-1,X)\mathbb{E}\left[u^\chi(D|t,X)\right] \iff F\left(\mu^\chi(X|t-1,X)\right) \geq 0,$$

which is equivalent to

$$\begin{aligned}\mu^\chi(X|t-1,X) &\leq \frac{\left[1+\frac{\delta}{1+\alpha}-\delta(1-\chi)\right] - \sqrt{\left[1+\frac{\delta}{1+\alpha}-\delta(1-\chi)\right]^2 - 4\delta\chi\left(\frac{1}{1+\alpha}\right)}}{2\delta\chi} \\ &= \frac{\left[(1+\alpha)(1+\delta\chi)-\alpha\delta\right] - \sqrt{\left[(1+\alpha)(1+\delta\chi)-\alpha\delta\right]^2 - 4\delta\chi(1+\alpha)}}{2\delta\chi(1+\alpha)} \equiv \kappa(\chi).\end{aligned}$$

Therefore, condition 2 holds if and only if

$$1-(1-p)\chi^{t-2} \leq \kappa(\chi) \iff \chi \geq \left(\frac{1-\kappa(\chi)}{1-p}\right)^{\frac{1}{t-2}}.$$

In summary, both players choosing $\tilde{\sigma}^\chi(X) = t$ is a χ-CSE if and only if

$$\left(\frac{1-\kappa(\chi)}{1-p}\right)^{\frac{1}{t-2}} \leq \chi \leq \bar{\alpha}^{\frac{1}{t-1}}.$$

This completes the proof. ■