

On the Impossibility of Stability-Based Equilibria in Infinite Horizon: An Example*

Alexander Frug[†] Malachy James Gavan[‡]

May 16, 2026

Abstract

We show that stability-based equilibrium refinements may fail to generate renegotiation-proof equilibria in infinite-horizon settings. We study a refinement, called a dynamically consistent partition, that classifies equilibria according to whether they remain immune to future mutually beneficial revisions in every continuation game, allowing for incomplete information. Equilibria that survive this recursive test are interpreted as renegotiation-proof. By contrast, equilibria that lead to a continuation game in which players would jointly prefer to switch to a better equilibrium that survives the test from that point onward are not. We provide a concrete example in which a dynamically consistent partition exists in every finite-horizon truncation and selects a renegotiation-proof equilibrium, but fails to exist in the infinite-horizon limit.

Keywords: Dynamically Consistent Partition, Renegotiation-proof, Stability

*Frug acknowledges financial support from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S).

[†]Universitat Pompeu Fabra and Barcelona School of Economics; email: alexander.frug@upf.edu.

[‡]University of Liverpool; email: malachy.gavan@liverpool.ac.uk

1 Introduction

This paper shows that stability-based equilibrium refinements may fail to generate renegotiation-proof equilibria in an infinite-horizon setting. Using the notions of internal and external stability for equilibrium selection, we demonstrate an environment where, in every finite-horizon truncation, there exists a dynamically consistent partition of equilibria—of the original game and all its continuation games—into equilibria that are immune to renegotiation and those vulnerable to renegotiation. This leads to a natural (and non-empty) selection of renegotiation-proof equilibria of the original game. In the infinite-horizon limit, however, such a dynamically consistent partition of equilibria fails to exist, and the refinement therefore cannot select any renegotiation-proof equilibrium.

We consider a concept of renegotiation-proofness based on internal and external stability, in the spirit of [Asheim \(1997\)](#); [Bernheim and Ray \(1989\)](#) and [Xue \(2000\)](#). The underlying logic in these papers is closely related to the notion of stability, introduced by ([von Neumann and Morgenstern, 1953](#), Ch. XII) and further developed in [Greenberg \(1990\)](#)'s theory of social situations, and applied there to equilibria of dynamic games. We adapt it to allow for incomplete information.¹

The general logic behind such solution concepts is to rule out plays that the players would abandon in favor of “robust” mutually preferred alternatives later on. A formal way to capture this is to define dynamically consistent partitions of equilibria—of the whole game and of its continuation games—into “Good” sets (robust to renegotiation) and “Bad” sets (vulnerable to renegotiation). Intuitively, an equilibrium belongs to the Good set of a given game if there is no reachable continuation game, induced by the equilibrium, at which players would jointly prefer to switch to another equilibrium that is itself deemed Good (in that continuation). Conversely, an equilibrium is in the Bad set of any given game if, now or at some future point, the players can agree to switch to a better plan from the Good set of the relevant continuation game. The requirement of dynamic consistency ensures that this classification holds throughout the game.

Existence of these concepts has been shown in specific settings, for instance finite horizon environments, infinitely repeated games, and infinite games with dynasties of players (see for example, [Asheim \(1991b\)](#); [Asheim \(1997\)](#); [Bernheim and Ray \(1989\)](#); [Xue](#)

¹For some recent contributions to adapting cooperative-based solution concepts to incomplete information, see [Dutta and Vohra \(2005\)](#); [de Clippel \(2007\)](#); [Bloch and Dutta \(2009\)](#); [Salamanca \(2020\)](#).

(2000); Ales and Sleet (2014)). Our contribution is to show, through a concrete example, that a dynamically consistent partition may fail to exist in an infinite-horizon setting with incomplete information. While the method yields a natural classification of equilibria into renegotiation-proof and non-renegotiation-proof in every finite truncation of our game, this classification breaks down in a natural infinite extension. This result serves as a warning against applying such refinements in infinite-horizon environments.

We develop this result within a dynamic cheap-talk framework in which the sender undergoes exogenous gradual learning. As in standard cheap-talk models beginning with Crawford and Sobel (1982), the sender transmits non-verifiable messages about the state of the world, and the receiver takes an action based on those messages. In our setting, however, the sender acquires information gradually over time, according to a fixed exogenous learning process, and communicates repeatedly with the receiver who decides whether to take a single action (which would end the interaction) or wait (and continue to the next round of the sender's learning and communication), after each sender's report.²

To further illustrate the environment and the logic of the solution concept that will be introduced shortly, we now turn to an example.

Example 1. Consider the canonical uniform-quadratic case with constant sender bias $b = \frac{1}{8}$. That is, the sender and receiver have utility functions

$$u_S(a, \theta) = -(\theta + \frac{1}{8} - a)^2 \text{ and } u_R(a, \theta) = -(\theta - a)^2,$$

where $a \in [0, 1]$ is the action chosen by the receiver and $\theta \in [0, 1]$ is the uniformly distributed state of the world.

Suppose that there are two periods. Within each period, the sender first obtains information, then sends a message, and the receiver decides whether to take an action, thereby terminating the game, or delay. If the receiver takes no action $a \in [0, 1]$ by the end of period 2, then payoffs for all players are -1 .

Now consider the learning process where the sender is initially uninformed, then, at time t_1 , the sender learns whether the state is above or below 0.4, and finally, at time $t_2 > t_1$, the sender perfectly learns the state (see Figure 1). Formally, denoting by L_{t_i} the sender's information

²Cheap talk with exogenous gradual learning was previously studied by Frug (2022). Similar environments but with endogenous learning processes were studied by Ivanov (2015, 2016) and Frug (2016, 2018).

partition in period t_i , we have

$$L_{t_0} = \{[0, 1]\}, L_{t_1} = \{[0, 0.4), [0.4, 1]\}, \text{ and } L_{t_2} = \{\{\theta\} | \theta \in [0, 1]\}.$$




Figure 1: Learning process in Example 1.

We now consider the equilibria of this game. If there is no informative communication at t_1 and the players get to t_2 , from the usual [Crawford and Sobel \(1982\)](#) logic, the only receiver's information partitions (up to interval boundaries) that are consistent with equilibrium are babbling (which reveals no information) and the interval partition at the threshold 0.25:

$$\{[0, 1]\} \text{ and } \{[0, 0.25), [0.25, 1]\}.$$

However, the equilibrium preferred by both players ex ante is the one in which the sender fully reveals his information (only) in period t_1 , thereby inducing the receiver's information partition³

$$\{[0, 0.4), [0.4, 1]\}.$$

There are no equilibria with informative communication in both periods. Hence, following the literature, we refer to this as the "most informative" equilibrium outcome (note that since no information is transmitted in t_2 , whether the receiver chooses her best response action in t_1 or t_2 is immaterial).⁴

³To see that this is an equilibrium note that after communication, the receiver believes that the state is uniform on $[0, 0.4)$ or $[0.4, 1]$ and therefore chooses, following the corresponding reports, 0.2 and 0.7. As at t_1 the sender has no additional information beyond that, the expected payoff when sending the message to indicate $[0, 0.4)$ is $\int_0^{0.4} -(\theta + \frac{1}{8} - 0.2)^2 d\theta = -\frac{139}{12000}$, while misreporting would lead to a lower payoff of $\int_0^{0.4} -(\theta + \frac{1}{8} - 0.7)^2 d\theta = -\frac{739}{12000}$. Clearly, if the sender learns that $\theta \geq 0.4$, there is no incentive to misreport as the sender is upwards biased.

⁴Recall that, from the ex-ante perspective, the players rank equilibria identically and prefer those that induce lower expected variance. In particular, since both partitions, $\{[0, 0.4), [0.4, 1]\}$ and $\{[0, 0.25), [0.25, 1]\}$, consist of two intervals and the former is more "balanced", prior to any sender's learning, it is preferred by both players.

Let us consider the stability of coordinating on the most informative equilibrium outcome. Notice that this requires informative communication in t_1 and relies on no information being transmitted in t_2 . Now, suppose that the communication in period t_1 revealed that $\theta \geq 0.4$. Consider the point in time right after the sender reveals this information.

At that point, instead of choosing an action and terminating the interaction (or equivalently, postponing the action to the next period in which informative communication with the sender is not expected), the players may recognize that they are facing the following continuation game. They now share the same information – that the state is uniformly distributed on $[0.4, 1]$ – and understand that the sender is about to learn the state perfectly. They also understand that the two (Crawford-Sobel) equilibrium outcomes of this continuation game are babbling and the informative communication which leads to the receiver's information partition of $[0.4, 1]$ into

$$\{[0.4, 0.45), [0.45, 1]\}.$$

Both players (who are at that point equally (un)informed) strictly prefer the informative equilibrium outcome of the continuation game.

In analyzing static communication games à la [Crawford and Sobel \(1982\)](#), it is often assumed that the players will seek to avoid dominated equilibria. By the same token, if the players are sufficiently sophisticated to coordinate on a Pareto efficient equilibrium ex-ante, why would they ignore opportunities to reach Pareto improvements down the road? In other words, why not revise at the beginning of the continuation game? The problem is that if the players expect to revise their plan and coordinate on the most informative equilibrium of the continuation game, it changes the incentives for truth-telling in period t_1 . Specifically, the sender who learns that the state is in $[0, 0.4)$ will prefer to (mis)report that the state is in $[0.4, 1]$, planning to claim that it belongs to $[0.4, 0.45)$ in stage t_2 .

Therefore, informative communication at t_1 leads to a desire to jointly renegotiate future play. We need a solution concept that takes this into account since it seems unlikely that the players will follow a dominated equilibrium of the (continuation-) game just because they promised to do so in the past. □

To more concretely pin down the idea of joint dynamic revision, we adapt a solution concept similar to [Greenberg \(1990\)](#), [Asheim \(1997\)](#) and [Xue \(2000\)](#). We adapt the solution concept to our setting by allowing for incomplete information as follows.

We begin by considering all *continuation games*⁵ and identifying the set of equilibria of each such game and the (stochastic) transition from a given game into its continuation games.⁶ Then, for each such game, we consider a partition of the set of equilibria. Such a partition is interpreted as identifying equilibria that are robust to renegotiation, and must satisfy the following dynamic consistency requirements:

Internal Stability: an equilibrium included in the robust-to-renegotiation set of equilibria of a given game cannot lead to a continuation game where there is a Pareto improving equilibrium, which in itself is robust to renegotiation in that continuation game.

External Stability: any equilibrium not included in the robust-to-renegotiation set of equilibria of a given game must lead (with positive probability) to a continuation game where there is a Pareto improving equilibrium in the set of equilibria that are robust to renegotiation in that continuation.

We refer to such a partition for all possible games as a *Dynamically Consistent Partition, DCP*. An equilibrium is *renegotiation-proof* if it is selected by a DCP, i.e. belongs to the robust-to-renegotiation set of equilibria of the original game.

The behavioral principle we aim to capture is simple. At any point in time, the players seek to coordinate on a Pareto efficient equilibrium and, whenever possible, they are willing to replace the (continuation of the) equilibrium they currently plan to play in favor of a Pareto-improving equilibrium of the continuation game (they are currently at); however, they are willing to do so only if they expect that the new equilibrium will not be replaced later in favor of another Pareto-improving equilibrium (of some reachable continuation game).

First, we illustrate the solution concept in the finite horizon. For example, in the single-period information transmission game of [Crawford and Sobel \(1982\)](#), the solution concept selects all the Pareto efficient equilibria. Note that this set depends on the assumption of the sender's initial information. Consider the baseline [Crawford and Sobel \(1982\)](#)'s variant of Example 1 – i.e., same preferences but suppose the sender is fully informed from the outset. In this case, any equilibrium is renegotiation-proof since

⁵This includes the original game and all the games that begin after a finite non-terminal sequence of reports.

⁶We note that this transition is determined not only by the sender's learning process but also by the equilibrium played (e.g., the receiver's information in any continuation game under the babbling equilibrium of the original game is that the state is uniform between zero and one).

any equilibrium is top-ranked by some sender types. On the other hand, if the players begin the interaction before the sender knows the state, only the (most) informative equilibrium is selected. This is true in general for the uniform-quadratic specification since the players rank identically all receiver's information partitions from the ex-ante perspective. To illustrate the dynamic implications of renegotiation-proofness in our context, we now turn back to Example 1 (including the sender's learning process).

Example 1 (Revisited). *In the unique renegotiation-proof equilibrium, information is transmitted only in period t_2 , when the sender reveals whether the state is above or below 0.25.*

The previously established logic rules out informative communication in period t_1 (since it will lead to a continuation game where it will be renegotiated).

Now consider the communication in period t_2 if the equilibrium played reveals no information in period t_1 . Since no information has been revealed yet, and the communication in period t_2 takes place when the sender is perfectly informed, the equilibria of that continuation game are identical to Crawford-Sobel equilibria. That is, there is the babbling equilibrium and the informative equilibrium in which the receiver learns whether the state is below or above 0.25. The informative equilibrium is a priori preferred to the babbling by the receiver and all sender types prior to learning L_{t_2} (i.e., the sender "period t_1 type" whose belief is that the state is distributed uniformly between 0 and 0.4 and the one who believes the state is distributed uniformly between 0.4 and 1). Therefore, the informative equilibrium is renegotiation-proof in the continuation game that begins after no information is revealed in period t_1 .

Working backwards, consider the equilibrium of the entire game that begins with babbling in period t_1 and continues with the renegotiation-proof equilibrium described above for the continuation game of period t_2 . Refer to this equilibrium as "delayed-communication equilibrium."

Since the delayed-communication equilibrium continues with a renegotiation-proof equilibrium after the first round of (uninformative) communication, the players will not renegotiate at that point. Do they want to renegotiate it in period t_1 ? Even though, prior to any sender's learning, both players prefer the most informative equilibrium (where the receiver learns whether the state is above or below 0.4), the players will not replace the delayed-communication equilibrium in favor of the most-informative equilibrium of the entire game since they understand that they may have an incentive to revise it in the future. Hence, the delayed-communication equilibrium is a renegotiation-proof equilibrium of the entire game and it is unique since there are no other equilibria. \square

In the above example, the solution concept led to a delay in communication. This

arguably natural result is typical since, as time goes by, revising the course of play becomes more difficult since agreeing on Pareto improvements with a sender whose information is more refined requires more demanding constraints.⁷

The logic of constructing dynamically consistent partitions of equilibria for the continuation games satisfies some natural benchmarks for being well behaved. Specifically, whenever such a dynamically consistent partition exists, there exists at least one renegotiation-proof equilibrium of the original game.⁸ Hence, the result of the paper is not merely a result of the non-existence of renegotiation-proof equilibria but a deeper negative result on the existence of dynamically consistent partitions. In the example we provide, this result is driven *only* via taking the solution concept to the infinite horizon, as the solution concept is always non-empty in this setting with a finite horizon. Therefore, in our setting, the issue in the infinite horizon is not merely that no equilibrium can be supported as renegotiation-proof, but that no dynamically consistent partition exists from which such an equilibrium could be selected.

While it was suspected in the literature that the logic of such solution concepts does not easily extend to the infinite horizon (see, for example, [Asheim \(1997\)](#) and [Asheim \(1991b\)](#)), to the best of our knowledge, our example where the solution concept leads to non-existence in the infinite horizon while always existing in a finite horizon is the first in the literature. Therefore we provide a warning that some caution should be taken when applying such concepts to infinite horizon settings.

Before turning to the model and analysis, we first briefly discuss the literature.

2 Literature Review

Several papers have introduced versions of renegotiation-proofness in the literature ([Bernheim and Ray, 1989](#); [Farrell and Maskin, 1989](#); [Asheim, 1997](#)). Our notion of renegotiation-proofness (or consistent partition) builds on the concept of [Asheim \(1997\)](#), and adapts it to incomplete information by allowing each information type to be treated

⁷Qualitatively, this means that, in “later” continuation games—when the sender’s information improves—fewer equilibria will fail to be renegotiation-proof (e.g., in a typical single-stage communication game where the sender is perfectly informed from the outset, all equilibria are renegotiation-proof since each such equilibrium is top-ranked from the perspective of some sender type).

⁸Note this is not the case in general when considering only the set of equilibria, see the literature review for further discussion.

as a player. [Asheim \(1997\)](#) relates his concept to coalition proofness of [Bernheim et al. \(1987\)](#) and the concept of consistent planning of [Strotz \(1955\)](#) and [Pollak \(1968\)](#), as well as other notions of renegotiation-proofness such as [Bernheim and Ray \(1989\)](#); [Farrell and Maskin \(1989\)](#). [Asheim \(1997\)](#) uses the notion of optimistic stability, building on the logic of [von Neumann and Morgenstern \(1953\)](#)'s stable set. In addition, our notion also extends [Xue \(2000\)](#)'s "weakly negotiation-proof Nash equilibrium" which is shown to exist in finite games of complete information.

It is well known that there may be no renegotiation-proof plan ([Hellwig and Leininger, 1987](#); [Asheim, 1997](#)). This can occur even in finite horizon. The logic is as follows. Suppose, for example, that there is a single equilibrium of the original (multistage) game, where there are multiple continuation equilibria (but, due to uniqueness, only one such equilibrium is consistent with the equilibrium of the whole game). However, it is possible that the continuation equilibrium consistent with the unique equilibrium of the whole game is not renegotiation-proof (in stage 2). Rather upon arriving at that continuation, another equilibrium is preferred by all active players. In that case, it can disqualify the original equilibrium and therefore there is no renegotiation-proof equilibrium.⁹ In contrast, our result will show that in some environments a *DCP* in itself does not exist, leading to nonexistence of such an equilibrium for a conceptually different reason.

When extending concepts of dynamic consistency to the infinite horizon, it is unclear what challenges can occur. Indeed, little is understood about this notion beyond the finite action finite horizon case. For example, in the discussion of optimistic stable standards of behavior (or a dynamically consistent partition) in a consistent planning environment with complete information, [Asheim \(1997\)](#) points out that "*with finite action but infinite horizon, no counterexample to existence is available.*" [Asheim \(1991b\)](#) shows, in the context of repeated games, that requiring stationarity of the set of Pareto-perfect norms—analogue to the dynamically consistent partition in our setting and likewise based on internal and external stability—may be inappropriate in the infinite horizon. His analysis highlights the additional conceptual and technical difficulties that arise when extending stability-based concepts beyond the finite horizon.

A related but distinct concept is that of the stable set, introduced as the "solution" by [von Neumann and Morgenstern \(1953, Ch. XII\)](#). [Greenberg \(1990\)](#)'s theory of social

⁹See example 6 of the working paper of [Asheim \(1997\)](#): [Asheim \(1991a\)](#) for further discussion of this.

situations built on the stable set logic to provide a general notion of stability that encompasses both the stable set and the dynamically consistent partitions considered in this paper. [von Neumann and Morgenstern \(1953, Ch. XII, 65:Y\)](#) show that when a dominance relation over a set of alternatives is strictly acyclic, in that there is no infinite chain of alternatives where each dominates the previous, then a unique abstract stable set exists. It is well known that the stable set may not exist ([Lucas, 1969](#)) – that is, it may not be possible to partition the set of outcomes into those that are stable and unstable, due to ‘odd cycles’ on the induced graph ([Richardson, 1953](#)). We show that the dynamically consistent partition may not exist, and the logic behind our construction is different from that used in [Richardson \(1953\)](#); [Lucas \(1969\)](#).

3 Definitions and preliminary analysis

Our starting point is the uniform quadratic specification of [Crawford and Sobel \(1982\)](#). That is, we assume that $\theta \sim U[0, 1]$ and that the players’ payoffs are given by:

$$u_S(a, \theta) = -(\theta + b - a)^2 \text{ and } u_R(a, \theta) = -(\theta - a)^2,$$

where a is the action chosen by the receiver and $b > 0$ is the sender’s bias.

We now introduce several preliminary definitions that will be needed later to define the notion of renegotiation-proofness in terms of dynamically consistent partitions. While the definitions make use of some specification-related assumptions (such as the uniformly distributed state of the world), we prefer to offer a sufficiently general form of the definitions (that may go slightly beyond what is strictly necessary for our example) to fully illustrate the logic of the concept.

The original reporting game is given by

$$G^0 = ([0, 1], \tau, L)$$

where $[0, 1]$ is the receiver’s initial information set, $\tau = (t_i)_{i=0}^T$, is an increasing sequence of times, where T is finite or infinite, and $L = (L_t)_{t \in \tau}$ is the sender’s learning process on $[0, 1]$ such that, for each $t \in \tau$, L_t is a partition of $[0, 1]$ into intervals and whenever $t < t'$, $L_{t'}$ refines L_t .

Within each $t \in \tau$, first the sender learns to which element of L_t the state belongs. He then chooses a message. Finally, after receiving the message for that period, the receiver can decide whether to take an action, a , or delay. If action a is taken at the end of time t , the game terminates and payoffs are the ones associated with expectation over u_S and u_R . So long as future periods exist, if the receiver decides to delay in a period, the game continues to the next period. In the case that the receiver never takes an action the players' payoffs are -1 .

Fix the original reporting game $G^0 = ([0, 1], \tau, L)$. For each date $t \in \tau$, let $\sigma(L_t)$ denote the set of non-empty unions of elements of L_t . For any $I^R \in \sigma(L_t)$, there exists a reporting behavior by the sender up to and including time t such that the induced continuation game consists of the receiver's information I^R (over which she has a uniform belief), the remaining time structure which is the tail of τ from date t onward, denoted by τ^t , and the sender's learning process that is the restriction of the original learning process to I^R from date t onward. Formally, the induced reporting game is

$$G^{t, I^R} = (I^R, \tau^t, L^{t, I^R}),$$

where

$$\tau^t = \{s \in \tau \mid s \geq t\} \quad \text{and} \quad L^{t, I^R} = (L_s |_{I^R})_{s \in \tau^t}.$$

We denote the family of reachable reporting games by

$$\Gamma(L) := \{G^{t, I^R} \mid t \in \tau, I^R \in \sigma(L_t)\}$$

Given a reporting game G (which can be the original game or any continuation game), denote by $EQ(G)$ the set of Bayesian Nash equilibria of G .¹⁰ Note that the set of equilibria is always non-empty as babbling is always an equilibrium. In addition, for every G , let L_G^{init} denote the sender's information partition at the beginning of G (which we denoted L_0 for the original game). Note that never taking an action is not consistent with equilibrium.

For the purposes of renegotiation, it is sufficient to restrict attention to the points of the game that occur after communication in a given period and before the receiver

¹⁰Formally, our definition only allows for pure reporting strategies. The results would not change with a more general definition. However, this would come at the cost of heavier notation and therefore is omitted.

takes an action or delays.¹¹

Comparing Equilibria. For $e', e'' \in EQ(G)$, we say that e' *Pareto dominates* e'' (denote $e' \succ e''$) if e' is preferred over e'' by the receiver and *all* sender's L_G^{init} elements. Some of the preferences may be weak but at least one strict preference is required.

Equilibrium Transition Between Games. For $G_1, G_2 \in \Gamma(L)$ and $e \in EQ(G_1)$, we write

$$G_1 \xrightarrow{e} G_2$$

if the game G_2 is induced with positive probability by the reporting strategy played under equilibrium e of G_1 .

Renegotiation-Proof Equilibria. Next, for all $G \in \Gamma(L)$, we define a partition of $EQ(G)$ into equilibria that are immune to renegotiation and those that are not, that is consistent across games.

A dynamically consistent partition (DCP) with respect to \succ , is a collection of subsets $RP(G) \subseteq EQ(G)$, for all $G \in \Gamma(L)$ that satisfies:

- *Internal Stability.* For each $e \in RP(G)$, for each \tilde{G} with $G \xrightarrow{e} \tilde{G}$ and for each $\tilde{e} \in RP(\tilde{G})$, $\tilde{e} \not\succeq e|_{\tilde{G}}$.
- *External Stability.* For each $e \notin RP(G)$, there is \tilde{G} with $G \xrightarrow{e} \tilde{G}$ and there is $\tilde{e} \in RP(\tilde{G})$ with $\tilde{e} \succ e|_{\tilde{G}}$.

Given a DCP, we refer to an equilibrium $e \in RP(G^0)$ as renegotiation-proof (of the entire game).

While it was shown for other settings the existence of partitions into “Good” and “Bad” equilibria analogous to our DCP does not guarantee the non-emptiness of the set of “Good” equilibria in the original game (see the literature review) in our setting, however, if a DCP exists, a renegotiation-proof equilibrium of the entire game exists (i.e. $RP(G^0) \neq \emptyset$).

¹¹Assume that at that point in time the receiver believes the state is uniform on I . For any later point in time where the receiver has the same information, the sender's information is more refined. As will become clear from the definition of renegotiation proofness, if the players have a joint incentive to renegotiate given an information structure, they would also have an incentive to renegotiate when the information structure is coarser.

Proposition 1. *If DCP exists, a renegotiation-proof equilibrium of the entire game G^0 exists.*

We note that a similar result is given in [Asheim \(1997\)](#) (Proposition 2.1). However, as our solution concept is for incomplete information and only compares equilibria (rather than all paths, as Asheim’s concept does), we cannot use that result directly.

We relegate the proof, and all remaining proofs, to the appendix.

Our first observation is that the solution concept is well-behaved for finite horizon.

Proposition 2. *If T is finite, a DCP exists.*

4 Infinite Horizon

The definitions in Section 3 (and Proposition 1) allow for finite as well as infinite horizon. In this section, we provide a concrete counterexample to the existence of DCP for infinite horizon.

Example 2. *Consider the uniform-quadratic case with the sender’s constant bias $b = \frac{1}{12}$. That is, $\theta \sim U[0, 1]$ and the sender and receiver’s utility functions are, respectively,*

$$u_S(a, \theta) = -\left(\theta + \frac{1}{12} - a\right)^2 \text{ and } u_R(a, \theta) = -(\theta - a)^2.$$

Consider the increasing sequence $\{x_j\}_{j=1}^\infty$, where $x_j = \frac{1}{3} \cdot \frac{j}{j+1}$, and suppose the sender’s learning process is given by $L_{t_0} = \{[0, 1]\}$, and for $j \in \mathbb{N}$,

$$L_{t_j} = \{\{\theta\}_{\theta \in [0, x_j]}, (x_j, 1]\}.$$

This means that the sender is initially uninformed but if the state is below x_j (and only then), the sender will learn the state perfectly within the first j periods, and since $x_j \rightarrow \frac{1}{3}$, states (weakly) above $\frac{1}{3}$ will always belong to the unique non-degenerate interval for each j .¹²

¹²We note that a crucial aspect of the learning process in our example is that the sequence is that the learning is “upwards” (the sender learns the state perfectly below a threshold, and the sequence of thresholds is increasing). As a result, informative communication in a given period will rely on no




Figure 2: Learning process in Example 2.

We begin by characterizing the set of equilibria in this game. First, like in most cheap talk environments, (a set of) uninformative equilibria (babbling) exist, which we denote e_0 . In addition, for each $j \in \mathbb{N}$, there exists a class of payoff-equivalent equilibria, denoted e_j , in which the sender reports at time t_j either “low” or “high”, where low means $\theta \leq x_j$ and high means $\theta > x_j$, and in all other periods the reports are uninformative. Formally, $e_j = \{e_{j,j'} | j' \geq j\}$ is the set of equilibria, differing only in the period in which the receiver takes the action, $t_{j'}$, after informative communication has occurred in time j (if there is any).

To see that strategies of this form are equilibria note that, upon learning that the state is weakly below x_j , the sender strictly prefers $a_{L_j} = \frac{x_j}{2}$ over $a_{H_j} = \frac{x_j+1}{2}$ (which are the receiver’s best responses to learning the state is below or above x_j , respectively). To see this, note that for $\theta \leq x_j$,

$$u_S(a_{L_j}, \theta) = -\left(\theta + \frac{1}{12} - \frac{x_j}{2}\right)^2 > -\left(\theta + \frac{1}{12} - \frac{x_j+1}{2}\right)^2 = u_S(a_{H_j}, \theta),$$

which can be seen by first noticing that the inequality holds for the threshold type $\theta = x_j$ for all $x_j < \frac{1}{3}$, from which the same inequality follows for types $\theta < x_j$. In addition, a sender who learns that the state is in $(x_j, 1]$ strictly prefers to report truthfully (“high”) inducing a_{H_j} over any lower action, since the sender’s ideal action is $a_{H_j} + \frac{1}{12}$ and the sender’s payoff is increasing in the receiver’s action for any $a < a_{H_j}$.

Finally, there cannot be more than a single round of informative communication. Such communication would induce two actions below $\frac{1}{3}$ and would incentivize types who, according to such an equilibrium candidate, must induce the lowest action to deviate and induce a higher action.

communication in future periods. If instead the learning was “downwards” (the sender learns the state perfectly *above* a threshold, and the sequence of thresholds is *decreasing*) the players would have no incentives for future renegotiation.

Thus, the set of equilibria in the game G^0 is

$$EQ(G^0) = \{e_{j,j'} \mid j, j' \in \{0\} \cup \mathbb{N}, j' \geq j\}.$$

If no informative communication has been made by time t_i , the continuation game is $G^{t_i, [0,1]}$. The equilibria of this game are those where no informative communication takes place, but the action is taken later than t_i , and those where informative communication will take place in some period $t_j \geq t_i$, which will be followed by the receiver's action in some future period,

$$EQ(G^{t_i, [0,1]}) = \{e_{0,j} \mid j \geq i\} \cup \{e_{j,j'} \mid j, j' \in \mathbb{N}, j' \geq j \geq i\}.$$

A convenient feature of this example is the comparison of equilibria. Specifically, regardless of when the receiver's action is taken

$$e_{j,j'} \succ e_{i,i'} \iff j > i.$$

That is, any equilibrium in a class that provides information later is preferred by both the sender and the receiver to equilibria that provide information earlier. This is due to the lower expected variance induced by later partitions as, under the uniform quadratic specification assumed throughout, both players rank equilibria identically according to a decreasing order of the expected variance of the induced receiver's information partition.

Now suppose that some equilibrium in the class e_k is played. After the report at t_k , two continuation games can be induced: $G^{t_k, [0, x_k]}$ and $G^{t_k, (x_k, 1]}$ after the "low" and "high" reports respectively. Note that the only equilibrium outcome of the low continuation game is the completely uninformative (babbling).¹³ Whereas the equilibria of $G^{t_k, (x_k, 1]}$ are given by

$$EQ(G^{t_k, (x_k, 1]}) = \{e_{j,j'}^k \mid j, j' \in \mathbb{N}, j' \geq j \geq k\},$$

Here, we write $e_{j,j'}^k$ for the equilibrium where communication about the interval $(x_k, 1]$ (which was revealed under e_k) occurs at t_j and the receiver's action is taken at $t_{j'} \geq t_j$. Moreover, as before, note that

$$e_{j,j'}^k \succ e_{i,i'}^k \iff j > i.$$

The equilibria in $EQ(G^{t_i, (x_k, 1]})$ take an analogous form to $EQ(G^{t_i, [0,1]})$.

¹³This is as 4 times the sender's bias is longer than the length of the interval $[0, x_k]$.

The next figure illustrates the induced games when an equilibrium in the class e_1 is played in the original game G^0 . If the realized state exceeds $x_1 = \frac{1}{6}$, the induced continuation game is $G^{t_1, (x_1, 1]}$, where the interval remains long enough to sustain future informative communication in equilibrium (of the continuation game). By contrast, if the realized state is at most $\frac{1}{6}$, the induced continuation game is $G^{t_1, [0, x_1]}$, which consists of a short interval on which the sender is fully informed (in which case, babbling is the only equilibrium of the continuation game).




Figure 3: Continuation games induced when equilibrium e_1 is played.

Note that once the high continuation game is reached, the game has a structure analogous to the original game.

Claim 1. A DCP does not exist.

The proof of Claim 1 proceeds by contradiction and uses the recursive structure of the continuation games. Assume that a DCP exists. Since equilibrium classes are strictly ordered by Pareto dominance, internal stability implies that the renegotiation-proof equilibria (of the original game) must all belong to a single class, e_j for some $j \in \mathbb{N}$. Nothing from the next class, e_{j+1} , therefore can also be renegotiation-proof.

External stability then implies that, following any equilibrium in e_{j+1} , there exists a later period t_k , $k > j + 1$, such that babbling is renegotiation-proof at the continuation game of an equilibrium (of the original game) where informative communication occurs only in t_k and the information revealed is that $\theta \in (x_k, 1]$. This is shown in the first step of the proof.

The second step then includes a backwards propagation argument: if an equilibrium with later communication is renegotiation-proof in some continuation game but not in the original game, then external stability implies that there is an equilibrium with even later communication that is renegotiation-proof at an earlier stage. Iterating this eventually implies that some element of the class e_l , with $l > j$, must be renegotiation-

proof (in the original game), contradicting the initial claim that the only renegotiation-proof equilibria belong to e_j . Hence, no DCP can exist.

Remark: The learning process and the bias selected in Example 2 have a number of desired technical features to lay the ground for our non-existence result. In addition to the “upward” learning mentioned earlier, recall that for the selected bias, $b = \frac{1}{12}$, the unique informative Crawford-Sobel equilibrium induces the receiver’s information partition around $\frac{1}{3}$. Since from the ex-ante perspective, the players wish to obtain the most “balanced” equilibrium information partition, they are motivated to coordinate on late communication (without renegotiating it in the future). A possibly undesirable feature of our example is that there is not only no renegotiation-proof equilibrium, but there is also no ex ante-efficient equilibrium that can be played in finite time. However, we emphasize that the non-existence of the DCP does not rely on the efficient equilibrium not being achievable.

To see this, extend the state space to $[-\frac{2}{3}, 1]$ and add, before the learning in Example 2, another period of learning. Specifically, suppose the sender is initially uninformed, but in period “ t_1 ” he learns the state perfectly if it is negative and learns that $\theta \in [0, 1]$ if it is non-negative. From t_1 onward, conditional on $\theta \in [0, 1]$, the continuation of the learning process is identical to that in Example 2. There is then an equilibrium in which the sender only reveals whether the state is negative or not. Moreover, this is the ex-ante efficient equilibrium. Nonetheless, the same issue as in Example 2 is inherited after revealing that $\theta \in [0, 1]$ and therefore the DCP does not exist.

4.1 Restoring Existence: ϵ -DCP

In the above example a DCP does not exist as any arbitrarily small Pareto improvement was sufficient for the players to renegotiate. A classical weakening of similar notions is to require that all players must weakly benefit, and at least one must benefit by at least $\epsilon > 0$ to have an incentive to renegotiate (Asheim, 1992; Greenberg, 1990). We now define ϵ -renegotiation-proof equilibria and show that it restores existence.

An ϵ -dynamically consistent partition (ϵ -DCP) with respect to \succ is defined exactly as before, except that internal and external stability only treat as relevant those Pareto improvements for which, in addition, the continuation payoff of either the receiver

or at least one sender L_G^{init} element increases by more than ϵ . Thus, small Pareto improvements are ignored for the purposes of renegotiation. Let $RP^\epsilon(G)$ denote the set selected by the ϵ -DCP at G , and say that $e \in EQ(G^0)$ is ϵ -renegotiation-proof if $e \in RP^\epsilon(G^0)$.

For Example 2, it is useful to focus on the subclass of equilibria in which the receiver takes an action immediately, in the same period that informative communication occurs. In the notation introduced above, these are $e_{j,j}$. These are the natural candidates for ϵ -renegotiation-proofness, as once informative communication has taken place the game ends before further learning can generate sender types with stronger incentives to initiate a switch to a new equilibrium.

For j sufficiently large, under $e_{j,j}$, the gains from any possible renegotiation are less than ϵ . This generates a tail of continuation equilibria that are themselves immune to renegotiation and, in turn, rules out equilibria with informative communication occurring too early. Thus, the ϵ -adaptation selects only equilibria that are approximately ex-ante efficient. This gives a cutoff date starting from which informative communication is ϵ -renegotiation-proof, which pushes further into the future as ϵ tends to zero.

Claim 2. *For any $\epsilon > 0$ there is some $n^\epsilon \in \mathbb{N}$, where $\lim_{\epsilon \rightarrow 0} n^\epsilon = \infty$, such that the set of ϵ -renegotiation-proof equilibria:*

1. *includes all equilibria $e_{j,j}$ with $j \geq n^\epsilon$ and*
2. *excludes all equilibria with informative communication at time t_j for $j < n^\epsilon$.*

We conclude by noting that the focus on equilibria where the receiver takes an action immediately after informative communication is essential. At that point the receiver's and sender's information above the communicated cut-off, x_j , is the same. Further, the future cut-offs x_k are "close" to x_j so the effects on the expected payoffs of any future refinements of $(x_j, 1]$ are negligible at this point. If instead the action is delayed, further sender learning generates additional sender types before the game ends. Suppose the sender learns the state in $(x_j, x_{j+1}]$. Conditional on that information, the benefit for this sender type of refining the receiver's information is above ϵ (if ϵ is sufficiently small) for such a type. This would bring back the non-existence as in Claim 1. Thus, restoring existence requires selecting equilibria in which informative communication is followed immediately by action.

5 Conclusion

To conclude, we have shown that a dynamically consistent partition may fail to exist in the infinite horizon, despite existing in every finite truncation of the same environment. While related work had already suggested that extending stability-based concepts beyond finite-horizon settings is delicate, to the best of our knowledge we provide the first example in which non-existence arises only in the infinite-horizon. More broadly, while our analysis is developed in a cheap-talk environment, related stability-based ideas have also appeared in other applied settings, including dynamic collusion problems (see, for instance, [Farrell and Maskin \(1989\)](#)). Our analysis therefore suggests that additional care may be required when extending such concepts to infinite-horizon settings, especially in settings in which learning occurs dynamically, e.g., through changing market conditions. We therefore see the contribution of the paper mainly at the conceptual level, through improving our understanding of the logic of related concepts when taken to the infinite horizon.

At the same time, the ϵ -adaptation shows that once small gains from renegotiation are ignored, existence may be restored. The paper therefore provides both a cautionary example for applying stability-based refinements in infinite-horizon environments and an illustration of how an appropriately weakened notion of renegotiation may restore existence.

Proofs

Proof of Proposition 1:

Proof. Recall that $EQ(G^0) \neq \emptyset$ as $B \in EQ(G^0)$, where B denotes a completely uninformative (babbling) equilibrium. Assume by contradiction that $RP(G^0) = \emptyset$. In particular $B \notin RP(G^0)$, therefore, by external stability, there is a minimal date $t \in \tau$ and some $\tilde{G} \in \Gamma(L)$ such that $G^0 \xrightarrow{B} \tilde{G}$, where $\tilde{G} = G^{t,[0,1]}$, and there is $\tilde{e} \in RP(\tilde{G})$ with $\tilde{e} \succ B|_{\tilde{G}}$.

Notice that if babbling is played before date t , the continuation game \tilde{G} is reached with certainty. Consider the reporting strategy that is composed of babbling until time t , and then follows the reporting strategy in $\tilde{e} \in EQ(\tilde{G})$. It is immediate that this reporting strategy is consistent with equilibrium. Denote it by $e \in EQ(G^0)$. By assumption $RP(G^0) = \emptyset$, and therefore $e \notin RP(G^0)$.

Since $\tilde{e} \in RP(\tilde{G})$, by internal stability there cannot exist any continuation game $G' \in \Gamma(L)$ reached under e weakly after date t , together with some $e' \in RP(G')$, such that $e' \succ e|_{G'}$. Indeed, from date t onward, the equilibrium e coincides with \tilde{e} , and any such Pareto-improving renegotiation-proof equilibrium would contradict internal stability of \tilde{e} . Since $e \notin RP(G^0)$, external stability therefore implies that there must exist some continuation game $G' \in \Gamma(L)$ reached under e at some earlier date $s < t$, and some $e' \in RP(G')$, such that $e' \succ e|_{G'}$.

Recall that all elements of L_t^{init} agree that \tilde{e} is at least as good as transmitting no information to the receiver. Since L_s^{init} is a coarsening of L_t^{init} , it is also true that all elements of L_s^{init} agree that \tilde{e} is at least as good as transmitting no information to the receiver. Since $e' \succ e|_{G'}$ and $e|_{G'}$ is itself completely uninformative up to date s , it follows that all elements of L_s^{init} agree that e' is at least as good as transmitting no information to the receiver. Since e' is informative, the receiver prefers e' over the completely uninformative equilibrium as well. To conclude, notice that $G^0 \xrightarrow{B} G'$, and therefore we get a contradiction to the minimality of t because $e' \succ B|_{G'}$. \square

Proof of Proposition 2:

Proof. Since the reporting horizon is finite, the dynamically consistent partition $RP(G) \subseteq EQ(G)$ can be constructed backwards.

For notational convenience, for every continuation game $G \in \Gamma(L)$ beginning at the final date, define $R(G) = EQ(G)$ and let $P(G) \subseteq R(G)$ denote the set of all Pareto efficient equilibria within $R(G)$. Define $RP(G) = P(G)$.

Assume that $RP(\tilde{G}) \subseteq EQ(\tilde{G})$ has been defined for all continuation games $\tilde{G} \in \Gamma(L)$ beginning strictly after date t , and let $G \in \Gamma(L)$ begin at date t . Define $R(G) \subseteq EQ(G)$ to be the set of all equilibria with the property that if $e \in R(G)$, then, for every $\tilde{G} \in \Gamma(L)$ with $G \xrightarrow{e} \tilde{G}$ and \tilde{G} beginning strictly after t , it holds that $e|_{\tilde{G}} \in RP(\tilde{G})$. Let $P(G) \subseteq R(G)$ be the set of all Pareto efficient equilibria within $R(G)$. Define $RP(G) = P(G)$. It is straightforward that by this construction, we obtain a unique dynamic partition $RP(G) \subseteq EQ(G)$. In what follows, we verify that it is a dynamically consistent partition.

Let $e \in RP(G)$ for some $G \in \Gamma(L)$. Let $\tilde{G} \in \Gamma(L)$ such that $G \xrightarrow{e} \tilde{G}$ and let $\tilde{e} \in RP(\tilde{G})$. If $\tilde{e} \succ e|_{\tilde{G}}$, it follows that $e|_{\tilde{G}} \notin P(\tilde{G})$, thus $e|_{\tilde{G}} \notin RP(\tilde{G})$, contradicting our presumption that $e \in RP(G)$. Therefore internal stability is satisfied.

On the other hand, let $G \in \Gamma(L)$ and $e \notin RP(G)$. If $e \in R(G)$, then since $e \notin P(G)$ there is $e' \in RP(G)$ with $e' \succ e$, which establishes external stability. Otherwise, if $e \notin R(G)$, there is some continuation game $\tilde{G}_1 \in \Gamma(L)$ with $G \xrightarrow{e} \tilde{G}_1$ and $e|_{\tilde{G}_1} \notin RP(\tilde{G}_1)$. If $e|_{\tilde{G}_1} \in R(\tilde{G}_1)$, then since $e|_{\tilde{G}_1} \notin P(\tilde{G}_1)$, there is $e'_1 \in RP(\tilde{G}_1)$ with $e'_1 \succ e|_{\tilde{G}_1}$. Otherwise, there is some continuation game $\tilde{G}_2 \in \Gamma(L)$ with $\tilde{G}_1 \xrightarrow{e|_{\tilde{G}_1}} \tilde{G}_2$ and $e|_{\tilde{G}_2} \notin RP(\tilde{G}_2)$.

$RP(\tilde{G}_2)$. Continuing inductively, since the horizon is finite, after finitely many steps we reach some continuation game \tilde{G}_k such that $e|_{\tilde{G}_k} \in R(\tilde{G}_k)$. Thus, $e|_{\tilde{G}_k} \notin P(\tilde{G}_k)$ and therefore there is $e'_k \in RP(\tilde{G}_k)$ with $e'_k \succ e|_{\tilde{G}_k}$. Therefore external stability is satisfied. \square

Proof of Claim 1:

Proof. Assume that there exists a DCP: for all $G \in \Gamma(L)$ there exists $RP(G) \subseteq EQ(G)$ satisfying the conditions of internal and external stability. By Proposition 1, $RP(G^0) \neq \emptyset$. Further, by internal stability and the strict Pareto ordering of the equilibrium classes $\{e_j\}_{j \geq 0}$, there is some j such that $RP(G^0) \subseteq e_j$.

We separate the remainder of the proof into two Steps. The first step shows that if under a renegotiation-proof equilibrium informative communication must occur at t_j , then there exists a later period t_k such that babbling is renegotiation-proof at the continuation game of an equilibrium where informative communication occurs in t_k and the information revealed is that $\theta \in (x_k, 1]$, i.e., $G^{t_k, (x_k, 1]}$.

Step 1 *If $RP(G^0) \subseteq e_j$, then there exists $j' \geq k > j + 1$ such that $e_{k, j'}^k \in RP(G^{t_k, (x_k, 1]})$.*

Since $RP(G^0) \cap e_{j+1} = \emptyset$, the equilibrium $e_{j+1, j+1}$, in which informative communication occurs at t_{j+1} and the receiver takes the action immediately, is not in $RP(G^0)$. By external stability, there is some continuation game \bar{G} with $G^0 \rightarrow_{e_{j+1, j+1}} \bar{G}$ and some $\bar{e} \in RP(\bar{G})$ such that $\bar{e} \succ e_{j+1, j+1}|_{\bar{G}}$.

The continuation game \bar{G} cannot be one that is also reached under any equilibrium selected in $RP(G^0)$. To see this, let \hat{e}_j be an equilibrium in $RP(G^0) \subseteq e_j$. If \bar{G} were reachable under \hat{e}_j , in this case \bar{e} would be a renegotiation-proof equilibrium in a continuation game reached under \hat{e}_j and would strictly dominate the continuation prescribed there, contradicting internal stability and therefore contradicting $\hat{e}_j \in RP(G^0)$. Therefore \bar{G} must occur strictly after the histories at which \hat{e}_j differs from $e_{j+1, j+1}$.

Moreover, \bar{G} cannot lie on the continuation after the report that $\theta \leq x_{j+1}$ at t_{j+1} , since all equilibria of this continuation game, $G^{t_{j+1}, [0, x_{j+1}]}$, are babbling and therefore payoff equivalent. Therefore \bar{G} must lie on the high branch after the report at t_{j+1} . Hence the babbling continuation prescribed there by $e_{j+1, j+1}$, denoted in the continuation-game notation by $e_{j+1, j+1}^{j+1}$ (recall that the superscript denotes the initial continuation game induced), cannot belong to $RP(G^{t_{j+1}, (x_{j+1}, 1]})$ (as otherwise $e_{j+1, j+1} \in RP(G^0)$, a contradiction).

By Proposition 1, applied to the continuation game $G^{t_{j+1}, (x_{j+1}, 1]}$, the set $RP(G^{t_{j+1}, (x_{j+1}, 1]})$ must therefore include an informative equilibrium. I.e. there exist integers $j' \geq k > j + 1$ such that $e_{k, j'}^{j+1} \in RP(G^{t_{j+1}, (x_{j+1}, 1]})$.

Consider now the continuation induced under $e_{k,j'}^{j+1}$ after the high report at t_k . This continuation game is $G^{t_k, (x_k, 1]}$, and the continuation prescribed there by $e_{k,j'}^{j+1}$ is the babbling equilibrium $e_{k,j'}^k$. By definition of DCP it must therefore be that $e_{k,j'}^k \in RP(G^{t_k, (x_k, 1]})$, as required. This proves step 1.

Having established that there is a continuation game $G^{t_k, (x_k, 1]}$ on which babbling is renegotiation-proof, we now show that this implies that we can find a renegotiation-proof equilibrium of the original game with an even later date of communication. This, in turn, will establish a contradiction and prove the claim.

Step 2 For $j' \geq k > j + 1$ such that $RP(G^0) \cap e_k = \emptyset$ and the babbling equilibrium $e_{k,j'}^k$ is in $RP(G^{t_k, (x_k, 1]})$, there exists some $l > k$ such that $RP(G^0) \cap e_l \neq \emptyset$.

Consider the equilibrium (of the original game, G^0) $e_{k,j'}$. Since this equilibrium is not in $RP(G^0)$, by external stability, there is some continuation game \tilde{G}_{i_1} reached under $e_{k,j'}$ and some $\tilde{e}_{k_1} \in RP(\tilde{G}_{i_1})$ such that $\tilde{e}_{k_1} \succ e_{k,j'}|_{\tilde{G}_{i_1}}$. This continuation game \tilde{G}_{i_1} cannot occur after the high report at t_k , i.e., after revealing that $\theta \in (x_k, 1]$, since the continuation prescribed there by $e_{k,j'}$ is the babbling equilibrium, $e_{k,j'}^k$, and $e_{k,j'}^k \in RP(G^{t_k, (x_k, 1]})$.

Therefore, \tilde{G}_{i_1} must be reached at some earlier date t_{i_1} with $i_1 < k$. Since Pareto dominance in this environment is equivalent to informative communication occurring strictly later, \tilde{e}_{k_1} must belong to some class e_{k_1} with $k_1 > k$.

If $i_1 = 0$, and therefore $t_{i_1} = t_0$, we have that $\tilde{G}_{i_1} = G^0$, and therefore $RP(G^0) \cap e_{k_1} \neq \emptyset$, proving the claim as $l = k_1 > k$. If $i_1 > 0$ and $\tilde{e}_{k_1} \notin RP(G^0)$, while $\tilde{e}_{k_1} \in RP(\tilde{G}_{i_1})$, we apply the same argument to the equilibrium \tilde{e}_{k_1} (where the babbling continuation of \tilde{e}_{k_1} is in $RP(G^{t_{i_1}, (x_{i_1}, 1]})$ by definition of DCP). This yields some continuation game \tilde{G}_{i_2} reached at an earlier date t_{i_2} with $i_2 < i_1$, together with some selected equilibrium $\tilde{e}_{k_2} \in RP(\tilde{G}_{i_2})$ such that \tilde{e}_{k_2} belongs to a class e_{k_2} with $k_2 > k_1$.

Continuing inductively, we obtain a finite and strictly decreasing sequence of continuation dates $k > i_1 > i_2 > \dots > i_{l^*}$ and a strictly increasing sequence of communication classes $k < k_1 < k_2 < \dots < k_{l^*} = l$. As it is finite we conclude that $i_{l^*} = 0$. At the final stage, the corresponding equilibrium \tilde{e}_l belongs to the equilibria with informative communication in period t_l with $l > k$, and it is in $RP(G^0)$. This is a contradiction to $RP(G^0) \subseteq e_j$.

□

Proof of Claim 2:

Proof. We proceed in two steps.

Step 1 *Immediate receiver's action after communication that occurs sufficiently late is ϵ -renegotiation-proof in the continuation game.*

Fix $\epsilon > 0$. For each $k \in \mathbb{N}$, consider the continuation game $G^{t_k, (x_k, 1]}$ reached immediately after the high communication at time t_k . Since at that point there is a single sender type, the expected gain from further communication in the continuation game is bounded (for both that sender type and the receiver) by the difference between the payoff of no further communication on $(x_k, 1]$ and splitting the receiver's information set $(x_k, 1]$ into $(x_k, 1/3]$ and $(1/3, 1]$. This gain is the reduction in the conditional expected variance:

$$\Delta_k := \text{Var}(\theta \mid (x_k, 1]) - \mathbb{E}(\text{Var}(\theta) \mid \{(x_k, 1/3], (1/3, 1]\}) = \frac{1}{18(k+1)}.$$

Denote by n^ϵ the lowest integer k such that $\frac{1}{18(k+1)} \leq \epsilon$. It follows that for every $j \geq n^\epsilon$, $e_{j,j}^j \in RP^\epsilon(G^{t_j, (x_j, 1]})$. Recall that $e_{j,j}^j$ represents an equilibrium of the continuation game $G^{t_j, (x_j, 1]}$, right after the sender reveals that $\theta > x_j$, where the receiver chooses an action immediately, prior to further sender's learning.

Step 2 *Using equilibria with sufficiently late informative communication and immediate action to characterize the ϵ -DCP.*

The next step is to show that these equilibria that communicate at time t_j , $j \geq n^\epsilon$ and end in immediate receiver action remain ϵ -renegotiation-proof when viewed from earlier high continuation games. That is, if informative communication is postponed until some t_j with $j \geq n^\epsilon$, then the corresponding equilibrium $e_{j,j}^k$ is also renegotiation-proof in every earlier high continuation game $G^{t_k, (x_k, 1]}$ with $k < j$.

Establishing this serves two purposes. First, taking $k = 0$ shows that equilibria with sufficiently delayed informative communication, where the receiver action is taken immediately after informative communication, are themselves ϵ -renegotiation-proof in the original game. Second, these equilibria can then be used to rule out equilibria in which informative communication occurs before t_{n^ϵ} .

The key observation is that every equilibrium in the class e_k induces the same continuation game $G^{t_k, (x_k, 1]}$. Moreover, the equilibrium that ends immediately after communication has no further continuation games (as no more sender learning takes place). Therefore, among the continuation games reached under equilibria in e_k , this is the continuation in which the gain from renegotiation is strongest. Consequently, if the babbling continuation in $G^{t_k, (x_k, 1]}$ for any $k < n^\epsilon$ can be ruled out by ϵ -internal stability, then every equilibrium in the class e_k is ruled out as renegotiation-proof in the original game, G^0 , as well.

Step 2.1 *Equilibria with sufficiently late informative communication and immediate action remain ϵ -renegotiation-proof in earlier high continuation games.*

To show that sufficiently delayed equilibria remain ϵ -renegotiation-proof when viewed from earlier high continuation games, consider an earlier high continuation game $G^{t_k, (x_k, 1]}$ together with the equilibrium $e_{j,j}^k$, where $j \geq n^\epsilon$ and $j > k$. That is, in the continuation, informative communication is postponed until t_j , after which the receiver takes the action immediately.

By the step 1, both continuations induced after communication at t_j are already ϵ -renegotiation-proof. We now show that renegotiating earlier, at the start of $G^{t_k, (x_k, 1]}$, also cannot generate a gain greater than ϵ .

Relative to the continuation game beginning at t_j , the equilibrium $e_{j,j}^k$ already induces the conditional mean action on the additional interval $(x_k, x_j]$. Therefore no other equilibrium can improve on $e_{j,j}^k$ on this interval for the receiver.

Any potential gain from renegotiation can therefore arise only on the upper interval $(x_j, 1]$. Conditional on reaching $(x_j, 1]$, the maximal possible gain is exactly Δ_j . However, from the perspective of $G^{t_k, (x_k, 1]}$, this continuation occurs only with conditional probability $\frac{1-x_j}{1-x_k} < 1$. Therefore the total gain from renegotiation is strictly smaller than Δ_j for the receiver, and hence strictly smaller than ϵ .

Since at common information the sender's payoff differs from the receiver's only by a constant, the gain is identical for both players. Therefore, excluding $e_{j,j}^k$ would violate ϵ -external stability, and so $e_{j,j}^k \in RP^\epsilon(G^{t_k, (x_k, 1]})$.

Step 2.2 *Equilibria with sufficiently late informative communication and immediate action are ϵ -renegotiation-proof in the original game.*

Taking $k = 0$, this immediately implies that $e_{j,j} \in RP^\epsilon(G^0)$.

Step 2.3 *Equilibria with informative communication before t_{n^ϵ} are ruled out.*

Now let us consider an equilibrium in e_k , which induces the babbling continuation on $G^{t_k, (x_k, 1]}$, where $k < n^\epsilon$. By the definition of n^ϵ , $\Delta_k > \epsilon$. As the bound Δ_k is tight, and by establishing that for all $j > n^\epsilon$, $e_{j,j}^k \in RP^\epsilon(G^{t_k, (x_k, 1]})$, there exists some sufficiently high j' such that $e_{j',j'}^k \in RP^\epsilon(G^{t_k, (x_k, 1]})$, while the gain from this renegotiation from the babbling continuation equilibrium is strictly higher than ϵ .

Therefore by ϵ -internal stability equilibria that continue with babbling in this continuation must be ruled out. This allows us to conclude that no equilibrium of the original game that provides informative communication earlier than t_{n^ϵ} is renegotiation-proof.

Combining the previous steps, every equilibrium $e_{j,j}$ with $j \geq n^\epsilon$ belongs to $RP^\epsilon(G^0)$, while no equilibrium with informative communication before t_{n^ϵ} does. Finally, by definition of n^ϵ , we have $n^\epsilon \rightarrow \infty$ as $\epsilon \rightarrow 0$.

□

References

- Ales, L. and Sleet, C. (2014). Revision proofness. *Journal of Economic Theory*, 152:324–355.
- Asheim, G. (1991a). Individual and collective time consistency. *No 1991-69, Discussion Paper from Tilburg University, Center for Economic Research*.
- Asheim, G. B. (1991b). Extending renegotiation-proofness to infinite horizon games. *Games and Economic Behavior*, 3(3):278–294.
- Asheim, G. B. (1992). A unique solution to n-person sequential bargaining. *Games and Economic Behavior*, 4(2):169–181.
- Asheim, G. B. (1997). Individual and collective time-consistency. *The Review of Economic Studies*, 64(3):427–443.
- Bernheim, B. D., Peleg, B., and Whinston, M. D. (1987). Coalition-proof nash equilibria i. concepts. *Journal of Economic Theory*, 42(1):1–12.
- Bernheim, B. D. and Ray, D. (1989). Collective Dynamic Consistency in Repeated Games. *Games and Economic Behavior*, 1(4):295–326.
- Bloch, F. and Dutta, B. (2009). Correlated equilibria, incomplete information and coalitional deviations. *Games and Economic Behavior*, 66(2):721–728.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451.
- de Clippel, G. (2007). The type-agent core for exchange economies with asymmetric information. *Journal of Economic Theory*, 135(1):144–158.
- Dutta, B. and Vohra, R. (2005). Incomplete information, credibility and the core. *Mathematical Social Sciences*, 50(2):148–165.
- Farrell, J. and Maskin, E. (1989). Renegotiation in Repeated Games. *Games and Economic Behavior*, 1(4):327–360.
- Frug, A. (2016). Dynamic cheap talk with static informational control. *Economics Letters*, 143:118–120.
- Frug, A. (2018). Strategic gradual learning and information transmission. *Journal of Economic Theory*, 177:594–615.
- Frug, A. (2022). On strategic transmission of gradually arriving information. *Available at SSRN 4279000*.

- Greenberg, J. (1990). *The theory of social situations: an alternative game-theoretic approach*. Cambridge University Press.
- Hellwig, M. and Leininger, W. (1987). On the existence of subgame-perfect equilibrium in infinite-action games of perfect information. *Journal of Economic Theory*, 43(1):55–75.
- Ivanov, M. (2015). Dynamic information revelation in cheap talk. *The BE Journal of Theoretical Economics*, 15(2):251–275.
- Ivanov, M. (2016). Dynamic learning and strategic communication. *International Journal of Game Theory*, 45:627–653.
- Lucas, W. F. (1969). The proof that a game may not have a solution. *Transactions of the American Mathematical Society*, 137:219–229.
- Pollak, R. A. (1968). Consistent planning. *The Review of Economic Studies*, 35(2):201–208.
- Richardson, M. (1953). Solutions of irreflexive relations. *Annals of Mathematics*, 58(3):573–590.
- Salamanca, A. (2020). A generalization of the harsanyi ntu value to games with incomplete information. *International Journal of Game Theory*, 49(1):195–225.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180.
- von Neumann, J. and Morgenstern, O. (1953). *Theory of Games and Economic Behaviour*. Princeton University Press.
- Xue, L. (2000). Negotiation-proof nash equilibrium. *International Journal of Game Theory*, 29:339–357.