Revision Games

Part II: Applications and Robustness

By

Yuichiro Kamada and Michihiro Kandori

April 2018

CENTER FOR RESEARCH AND EDUCATION FOR POLICY EVALUATION
DISCUSSION PAPER NO. 13

CENTER FOR RESEARCH AND EDUCATION FOR POLICY EVALUATION (CREPE)
THE UNIVERSITY OF TOKYO
http://www.crepe.e.u-tokyo.ac.jp/

Revision Games*

Part II: Applications and Robustness

KAMADA, Yuichiro† and KANDORI, Michihiro‡

† Haas School of Business, University of California, Berkeley

‡ Faculty of Economics, University of Tokyo

First Version: November 25, 2008; This Version: April 19, 2017

Abstract

This paper illustrates a wide applicability of the theory of revision games. First, we present various applications of the revision game; exchange of goods, price competition, and election campaign. Those applications reveal how the possibility of cooperation, extent of cooperation, and the frequency and magnitude of revisions are related to key parameters, such as the marginal benefit and marginal cost of cooperation, the degree of product differentiation, and the office motivation of the electoral candidates. Second, we examine the robustness of our model and extension to the case of asynchronous revisions.

*This paper contains the results in a working paper previously distributed under the title "Revision Games." We thank the same set of people as in the companion paper, Kamada and Kandori (2017), for their helpful comments. Emily Her provided excellent research assistance. The second author (Kandori) acknowledges support by JSPS KAKENHI Grant Numbers JP25380224 and JP17K03617.

†Email: ykamada@berkeley.edu

‡Email: kandori@e.u-tokyo.ac.jp

1 Introduction

In revision games, players have stochastic (Poisson) opportunities to prepare and revise their actions before their final action is implemented at the deadline. In the companion paper (Kamada and Kandori, 2017), we show that cooperation is sustained under such a situation and prove the existence of the optimal trigger strategy equilibrium and provide its characterization. In particular, we give a necessary and sufficient condition for *cooperation* to emerge in revision games.

The present paper is devoted to studying applications of the revision-game model to various economic situations, and discussing robustness of cooperation. These analyses show wide applicability of the revision-games framework.

Under the optimal trigger strategy equilibrium, cooperation is sustained in the following manner: on the path of play, players prepare an action prescribed by a *plan* $x(t)$. Namely, when a revision opportunity arrives at time t, players are supposed to revise their actions to $x(t)$. If any player deviates from this instruction, both players revert to a Nash action in all future revision opportunities. The plan $x(t)$ has to be such that players prefer following it to deviating to a static best response. As a consequence of this incentive constraint, $x(t)$ starts with the fully collusive level (when the time to the deadline is sufficiently long) and gradually tends to the Nash action as the time t approaches the deadline.

We consider three applications, namely a good exchange game, price competition, and an election game, where the static Nash equilibrium is inefficient. In the good exchange game, players have a dominant action which is not to provide any good to the opponent; however, providing positive amounts to each other can Pareto-dominate such a situation. In the price competition, colluding at a high price Pareto-dominates the Nash price profile in which prices are so low that there is no further incentive for undercutting. In the election game, we assume policy-motivated candidates who would prefer a half-half lottery between the most ideal and worst policies to the sure chance of implementing Nash policies in the middle of the policy space.

In each of these applications, we solve for the optimal trigger strategy equilibrium plan by applying the general characterization in Kamada and Kandori (2017) that uses differential equations. We show that under the optimal plan, over time, the amount of exchange decreases, the price falls, and policies converge towards the middle. Furthermore, using the condition for cooperation that Kamada and Kandori

(2017) identify - called *Finite Time Condition*- we show that in the revision game of price competition, product differentiation is necessary and sufficient for nontrivial cooperation to arise. Also, in the revision game of the election game, office-motivation cannot be too strong relative to policy-motivation in order for a collusive course of actions to be possible in equilibrium. To obtain those results, we show and utilize a simple and useful lemma (Lemma 1) to judge the sustainability of cooperation.

To confirm the applicability of our model, we also examine the robustness of our model. Possibility of cooperation in our model hinges on the assumption that the action space and time are both continuous. In our analysis of robustness, we first formalize and prove this statement. Then, we focus on continuity of time and consider a discrete-time model. We demonstrate that, by introducing another realistic modification to the model, i.e., a small perturbation of the payoff function such that there is a small incentive to punish the opponent for deviation, a significant amount of cooperation can be achieved if and only if our continuous time model predicts cooperation.

Our final discussion is about asynchronous moves. We consider a model in which revision opportunities arrive independently across players, and show that our analysis under the synchronous revision carries over to the asynchronous case: If the arrival rates of the opportunities are common across the players and payoff function satisfies the separability condition, then the same action plan as the case with synchronous opportunities characterizes the optimal trigger strategy equilibrium.

Section 2 provides a recap of the general framework studied in Kamada and Kandori (2017). It also states and proves a lemma about sustainability of cooperation. In Section 3, we discuss the three applications of the model. Section 4 discusses robustness of cooperation. Section 5 concludes and discusses various follow-up papers of our project.

2 General Framework

In this section we recapitulate the general framework and the main result of Kamada and Kandori (2017).

Component game: Component game is a two-player normal-form game with players $i = 1, 2$. There is a common action set A that is a convex subset (an interval) of $\mathbb{R}$.

The payoff function is $\pi_i : A \times A \to \mathbb{R}$. We assume symmetry, i.e., $\pi_1(a, a') = \pi_2(a', a)$ for all $a, a' \in A$.

Revision game: Time continuously runs from $-T(< 0)$ to 0. At time $-T$, two players simultaneously choose their actions. During time in $(-T, 0)$, there is a Poisson process with arrival rate $\lambda > 0$, and at each arrival of the Poisson hit, two players simultaneously revise their actions, observing all the past events. At time 0, the action profile that is chosen at the last Poisson arrival is implemented, and the corresponding payoff profile is realized.

Assumptions and the finite time condition: We impose the following six assumptions throughout the paper.

- **A1:** A unique pure symmetric Nash equilibrium action a^N and the unique best symmetric action $a^* := \arg\max_{a \in A} \pi(a)$ exist, and $a^N < a^*$.[1]
- **A2:** The symmetric payoff $\pi(a)$ is strictly increasing for $a < a^*$.
- **A3:** $\pi_1(a_1, a_2)$ is continuous. Furthermore, $\max_{a_1} \pi_1(a_1, a_2)$ exists for all a_2, and therefore we can define the *gain from deviation* at a symmetric profile (a, a) by

$$d(a) := \max_{a_1} \pi_1(a_1, a) - \pi_1(a, a) \tag{1}$$

- **A4:** The gain from deviation $d(a)$ is strictly increasing on $[a^N, a^*]$ and non-decreasing for $a^* < a$.
- **A5:** The gain from deviation d (defined by (1)) is differentiable, and $d' > 0$ on $(a^N, a^*]$.
- **A6:** Function $f(x) := \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)}$ is Lipschitz continuous on $[a^N + \varepsilon, a^*]$ for any $\varepsilon \in (0, a^*]$,[2] where $\pi^N := \pi_i(a^N, a^N)$.

[1]This inequality is without loss of generality, and the case with $a^* < a^N$ can be analyzed in a symmetric manner.

[2]$f(x)$ is Lipschitz continuous on $[a^N + \varepsilon, a^*]$, if there exists a finite number $K \geq 0$ such that $\left|\frac{f(x)-f(y)}{x-y}\right| \leq K$ for all $x \neq y$ in $[a^N + \varepsilon, a^*]$

Also, the following condition is the key to distinguishing those component games with which a nontrivial equilibrium exists in the revision game and those with which there is no such equilibrium.

- **Finite Time Condition**

$$\lim_{a \downarrow a^N} \int_a^{a^*} \frac{1}{f(x)} dx < \infty. \tag{2}$$

Optimal trigger strategy equilibrium plan: A trigger strategy is characterized by its revision plan $x : [0, T] \to A$. Players start with initial action $x(T)$, and when a revision opportunity arrives at time $-t$, they choose action $x(t)$. If any player fails to follow that rule, then both players choose the Nash equilibrium action of the component game in all future revision opportunities. Formally, *the set of feasible plans* is:

$$X := \{x : [0, T] \to A \mid \pi \circ x \text{ is measurable}\}.$$

Given a feasible plan $x \in X$, the (trigger strategy) *incentive constraint* at time t is

$$\text{(IC(t)): } d(x(t))e^{-\lambda t} \leq \int_0^t \left(\pi(x(s)) - \pi^N\right) \lambda e^{-\lambda s} ds. \tag{3}$$

The set of *trigger strategy equilibrium plans* is:

$$X^* := \{x \in X \mid \text{IC(t) holds for all } t \in [0, T]\}.$$

A plan that achieves the highest ex ante expected payoff within X^* is referred to as *an optimal trigger strategy equilibrium plan.* The following result claims uniqueness of such a plan. In fact, if a plan x is an optimal trigger strategy equilibrium plan, then another plan y that does not coincide with x only for t's in a measure zero set also constitutes an optimal trigger strategy equilibrium plan. The uniqueness that we state below is modulo such multiplicity.[3]

Theorem 1 (Kamada and Kandori (2017)) *Suppose that A1-A6 hold.*

1. *The optimal trigger strategy equilibrium plan $\overline{x}(t)$ is the unique plan with the following properties: (i) it is continuous in t and departs a^N at $t = 0$ (i.e.,*

[3]See Proposition 1 in Kamada and Kandori (2017) for the detail.

$\overline{x}(t) = a^N$ *if and only if* $t = 0$*), (ii) for* $t > 0$*, it solves differential equation*

$$\frac{dx}{dt} = \frac{\lambda\left(d(x) + \pi(x) - \pi^N\right)}{d'(x)} =: f(x) \tag{4}$$

until $\overline{x}(t)$ *hits the optimal action* a^**, and (iii) if* $\overline{x}(t)$ *hits the optimal action* a^* *it stays there (i.e.,* $\bar{x}(t') = a^*$ *for some* $t' \le T$ *implies* $\bar{x}(t'') = a^*$ *for all* $t'' \in [t', T]$*).*

2. *The plan* $\bar{x}(t)$ *is nontrivial (i.e.,* $\bar{x}(t) \neq a^N$ *for some* t*) if and only if the Finite Time Condition (2) holds. Under the Finite Time Condition (2), if the time horizon* T *is large enough,* $\overline{x}(t)$ *always hits the optimal action* a^* *at a finite time*

$$t(a^*) := \lim_{a \downarrow a^N} \int_a^{a^*} \frac{1}{f(x)} dx. \tag{5}$$

3. *If* $\liminf_{a \downarrow a^N} \frac{d(a)}{\pi(a) - \pi^N} > 0$*, the Finite Time Condition fails and the unique trigger strategy equilibrium is to play the Nash action all the time:* $x(t) \equiv a^N$*.*

Part 2 of Theorem 1 implies the following lemma that we utilize in the three examples in the next section. Since $\frac{1}{f(x)}$ in the definition of the Finite Time Condition (2) is finite for $x \in (a^N, a^*]$, the following holds.

Lemma 1 *If* $\lim_{a \downarrow a^N} |\frac{1}{f(a)}|$ *exists and is finite, the Finite Time Condition (2) holds and the optimal trigger strategy equilibrium plan* $\overline{x}(t)$ *is nontrivial (i.e., not identically equal to the Nash action* a^N*).*

This lemma will be useful in judging sustainability of cooperation in applications. This is because, in many cases, we can use l'Hôpital's rule to show that the limit of $|\frac{1}{f(a)}|$ as $a \downarrow a^N$ is well-defined and is finite.

3 Applications

In this section, we use the general framework of revision games to analyze various economic applications. Specifically, we use the differential equation provided in Theorem 1 to analyze good exchange games, price competition with product differentiation, and an election model.

3.1 Good Exchange Game

We first present a simple model to illustrate how a revision game works. Suppose two players produce and exchange goods. Player i produces a_i units of goods, with production cost $c(a_i)$ and gives it to player $-i$, who enjoys benefit $b(a_{-i})$.[4] Formally, this component game has two players $i = 1, 2$ with a common action space $A = [0, \bar{a}]$ for some $\bar{a} > 0$, and their payoff function is

$$\pi_i(a_i, a_{-i}) = b(a_{-i}) - c(a_i),$$

where b and c are twice continuously differentiable and strictly increasing functions such that $b(0) = c(0) = 0$ and $b'(0) > 0$. Note that there is a dominant strategy Nash equilibrium action $a_i = 0$ and the Nash payoff is $\pi^N = 0$. We assume that the Nash equilibrium is inefficient, and there exists a unique optimal action $a^* > 0$ that maximizes the symmetric payoff $\pi(a) = \pi_i(a, a)$. We also assume that $\pi(a)$ is strictly increasing on $[0, a^*]$. With these assumptions, A1-A6 are satisfied. Let us call this class *good exchange games*.

The differential equation for the optimal trigger strategy equilibrium plan is[5]

$$\frac{dx}{dt} = f(x) = \lambda \frac{b(x)}{c'(x)}.$$

When is cooperation sustained? For example, the *public goods provision game*, where

$$\pi_i(a_i, a_{-i}) = (a_i + a_{-i}) - ra_i, \quad 1 < r < 2, \quad a_i, a_{-i} \in [0, \bar{a}]$$

can be regarded as a special case of the good exchange game with $b(a_{-i}) = a_{-i}$, $c(a_i) = (r - 1)a_i$ and $a^* = \bar{a}$. Part 3 of Theorem 1 implies that *no* cooperation is sustained in this special case. The impossibility of cooperation comes from the property that $c'(0) > 0$ (the Nash action $a^N = 0$ is a *corner solution*) and $b'(0)$ is finite. More generally, the good exchange game provides the following insights into when cooperation is sustainable.

Proposition 1 *In the revision game of the good exchange game, the possibility of*

[4]Alternatively, we can assume that players produce and exchange one unit of goods, and a_i represents the quality of goods produced by player i.

[5]This is derived as $f(x) := \frac{\lambda\left(d(x)+\pi(x)-\pi^N\right)}{d'(x)} = \frac{\lambda(c(x)+(b(x)-c(x))-0)}{c'(x)}$.

cooperation depends on the marginal cost and benefit at the Nash action:

1. *When* $c'(0) = 0$ *and* $c''(0) > 0$*, there exists a nontrivial trigger strategy equilibrium plan.*

2. *When* $c'(0) > 0$*, there does not exist a nontrivial trigger strategy equilibrium plan if* $b'(0) < \infty$*.*

3. *When* $c'(0) > 0$*, there is* $b(\cdot)$ *with* $b'(0) = \infty$ *such that there exists a nontrivial trigger strategy equilibrium plan.*

Proof. The Finite Time Condition (2) is expressed as

$$\lim_{a \downarrow a^N} \int_a^{a^*} \frac{1}{f(x)} dx < \infty \iff \lim_{a \downarrow a^N} \int_a^{a^*} \frac{c'(x)}{\lambda b(x)} dx < \infty.$$

This is satisfied in Case (1). Since $\frac{c'(x)}{b(x)} < \infty$ for any $x > 0$, and $\lim_{x \downarrow a^N} \frac{c'(x)}{b(x)} = \lim_{x \downarrow a^N} \frac{c''(x)}{b'(x)} < \infty$ by l'Hôpital's rule. Hence, Lemma 1 implies that cooperation can be sustained by the optimal trigger strategy equilibrium plan. In Case (2), the sufficient condition for no cooperation (part 3 of Theorem 1)

$$\liminf_{x \downarrow a^N} \frac{d(x)}{\pi(x) - \pi^N} > 0$$

is satisfied. This is because $\frac{d(x)}{\pi(x) - \pi^N} = \frac{c(x)}{(b(x) - c(x)) - 0}$ and by l'Hôpital's rule

$$\lim_{x \downarrow 0} \frac{c(x)}{b(x) - c(x)} = \frac{c'(0)}{b'(0) - c'(0)} > 0.$$

In Case (3), cooperation is sustained if, for example, $b(x) = \sqrt{x}$ and $c(x) = x$. The integrand $\frac{1}{f(x)} = \frac{c'(x)}{\lambda b(x)} = \frac{1}{\lambda \sqrt{x}}$ in the Finite Time Condition (2) diverges to infinity as x tends to the Nash action 0, but it does so slowly enough in this example. As a result, the Finite Time Condition (2) holds because

$$\lim_{a \downarrow 0} \int_a^{a^*} \frac{1}{f(x)} dx = \left[\frac{2}{\lambda} x^{\frac{1}{2}} \right]_0^{a^*} = \frac{2}{\lambda} a^{*\frac{1}{2}} < \infty,$$

so the cooperation is sustained by the optimal trigger strategy equilibrium plan. ■

We present a simple example that admits a closed-form solution, and lets us evaluate how much cooperation can be sustained.

Linear benefit and quadratic cost

Proposition 2 *In the revision game of the good exchange game with* $b(a) = a$ *and* $c(a) = c \cdot a^2$ *where* $c > 0$ *is a constant, the optimal trigger strategy equilibrium plan,* $\bar{x}(t)$*, is characterized by*

$$\bar{x}(t) = \begin{cases} \frac{\lambda}{2c} t & \text{if} \quad t < t(a^*) \\ a^* = \frac{1}{2c} & \text{if} \quad t(a^*) \leq t \end{cases},$$

where $t(a^*) = \frac{1}{\lambda}$.

The plan characterized in Proposition 2 is depicted in Figure 1 for the case with $c = 1$. When $T \geq 1/\lambda$, the plan starts at the optimal action $a^* = 0.5$ and stays there until the time reaches $-1/\lambda$. After that, the prepared action decreases over time to reach the Nash action $a^N = 0$ at the deadline. The closed-form solution of the optimal trigger strategy equilibrium plan enables us to compute the expected payoff. Specifically, the next corollary shows that 74% of the fully collusive payoff can be sustained through the revision process, even though the players do not have a long-term relationship.

Corollary 1 *In the revision game of the good exchange game with* $b(a) = a$ *and* $c(a) = c \cdot a^2$ *where* $c > 0$ *is a constant, for any* $\lambda > 0$ *and* $T > t(a^*) = \frac{1}{\lambda}$*, the following are true, where* e *denotes the base of natural logarithms:*

1. *The expected payoff under the optimal trigger strategy equilibrium plan is* $\frac{1}{2ec}$.
2. *The ratio of the expected payoff under the optimal trigger strategy equilibrium plan to the fully collusive payoff is* $\frac{2}{e} \cong 0.74$*, independent of the value of* c.

The calculation is given in Appendix A.1.

3.2 Price Competition: Product Differentiation Affects Collusion

We consider the price-competition revision game, which captures the situation where firms revise their posted prices before the opening of the market/their stores. We

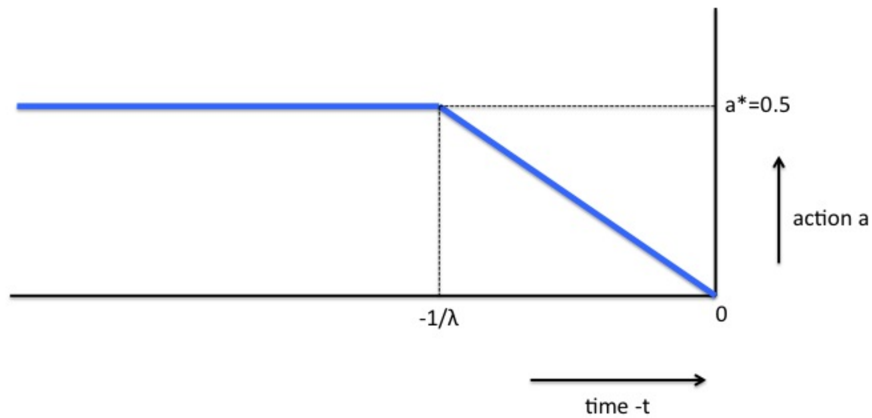

Figure 1: The optimal trigger strategy equilibrium plan $\bar{x}(t)$ for the good exchange game.

will show that firms' abilities to collude hinges on the degree of product differentiation. In particular, we show that *product differentiation is a necessary and sufficient condition for the sustainability of collusive prices.* This prediction is in stark contrast to the prediction of infinitely repeated games, in which for any level of product differentiation, sufficient patience guarantees sustainability of collusion.

To demonstrate this result, we need a model to accommodate various degrees of product differentiation, including no differentiation as a special case. A standard model with such a feature is the Hotelling's location model with price-setting firms.[6] It is illustrated in Figure 2. Figure 3 summarizes the main results. The figure shows the expected profits associated with the optimal trigger strategy equilibrium, joint profit maximization (full collusion), and the one-shot Nash equilibrium, all as a function of the level of product differentiation. In the model, consumers' transportation cost (c) relative to their value of the goods (v) measures the degree of product differentiation.

[6] An alternative would be to assume that firm i's demand is determined by $Q_i = a - bp_i + cp_{-i}$ under prices (p_i, p_{-i}) for some constants $a, b, c > 0$, but this specification does not nest the case of no differentiation as a special case.

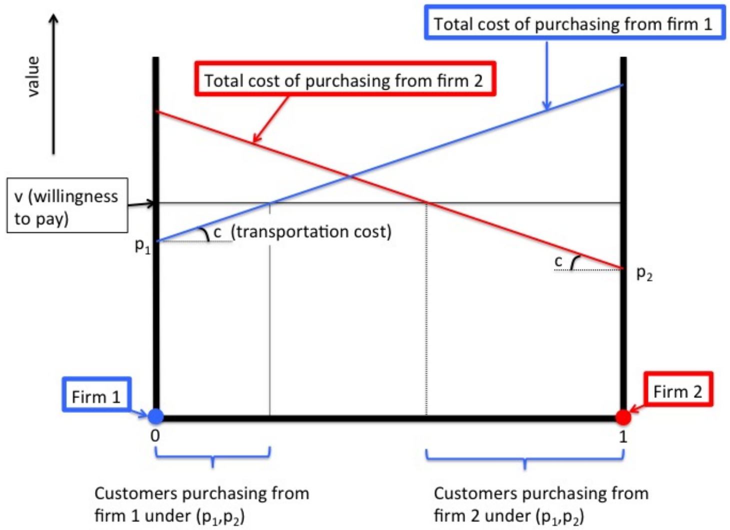

Figure 2: The component game for the Bertrand competition.

Degree of product differentiation ($h = \frac{c}{v}$)	0	.1	.2	.3	.5	$.6\dot{6}$
$\frac{\text{Expected payoff}}{\text{Fully collusive payoff}}$	0	.429	.641	.797	.963	1
$\frac{\text{Expected payoff - Nash payoff}}{\text{Fully collusive payoff - Nash payoff}}$	0	.362	.539	.686	.889	1

Table 1: Degrees of product differentiation and cooperation (the right-bottom entry (i.e., 1) is the limit value as $h \uparrow \frac{2}{3}$) when the horizon is long enough.

Note that, when $c/v = 2/3 = 0.67$, each firm becomes a local monopolist and the Nash equilibrium coincides with the fully collusive outcome. Table 1 shows that the sustainable level of collusion varies from zero (when there is no differentiation) to higher levels as the differentiation increases.

Model

There is a unit mass of buyers uniformly distributed over $[0, 1]$. Two firms $i = 1, 2$ are located at 0 and 1, respectively. A buyer at location $s \in [0, 1]$ receives payoff $v - c|s - s'| - p$ if she buys from a firm at s' with price p. If the buyer does not buy, her payoff is 0. When c is high enough (i.e., $c > \frac{2}{3}v$), each firm becomes a local monopolist and the joint profit maximization is achieved by the one-shot Nash equilibrium. Hence, we focus on the non-trivial case $c \in [0, \frac{2}{3}v)$.

Each buyer can purchase one unit at most, and decides her purchase behavior

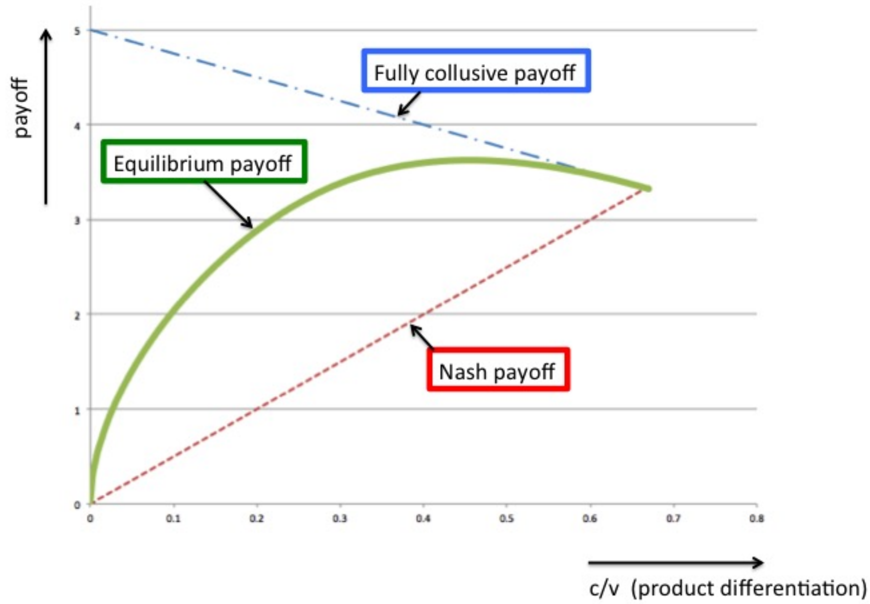

Figure 3: The expected payoffs under the optimal trigger strategy equilibrium plan for the Bertrand competition game: $v = 10$.

to maximize the payoff.[7] Each firm's marginal cost is normalized at 0. The payoff function for firm i is therefore given by:

$$\pi_i(p_i, p_{-i}) = p_i \times (\text{market share under } (p_i, p_{-i})).$$

Differential equations

It is straightforward to show that $p^N = c$ and $\pi^N = \frac{c}{2}$ hold in the Nash equilibrium, the fully collusive price is $p^* = v - \frac{c}{2}$, and the symmetric payoff function is $\pi(p) = \frac{p}{2}$.[8] The differential equation depends on the gain from deviation $d(p)$, and it takes on two functional forms for the following reasons. First, if c is high relative to the rival's price p, the static best response is to steal only *a part* of the buyers from the rival firm. Second, if c is relatively low, the static best response is to steal *all* the customers.

[7]If a buyer is indifferent between purchasing and not, she makes a purchase. If all buyers are indifferent between the two firms, then the firms equally split the market share. If a buyer is indifferent between purchasing from two firms, she mixes between them with equal probability.

[8]We provide a detailed explanation for these values in the Online Appendix.

Those two cases correspond to two functional forms of $d(p)$. For a high c, only the first case arises, while for a low c both cases can arise.

1. **High product differentiation:** $c \in (\frac{2}{7}v, \frac{2}{3}v)$.

   In this case, partial stealing of the customers from the rival is the myopic best reply and $d(p) = \frac{(p-c)^2}{8c}$ holds for all $p \in [p^N, p^*] = [c, v - \frac{c}{2}]$. A1-A6 are satisfied for this range of p as well, so Theorem 1 implies that the optimal trigger strategy equilibrium exists and its plan is a solution to the following differential equation:

$$\frac{dp}{dt} = f(p) = \lambda \frac{p + 3c}{2}. \tag{6}$$

   Since $\left|\frac{1}{f(p)}\right| < \infty$ holds for all $p \in (p^N, p^*]$ and $\lim_{p \downarrow c} \left|\frac{1}{f(p)}\right| < \infty$, Lemma 1 implies that the optimal trigger strategy equilibrium is nontrivial.

2. **Low product differentiation:** $c \in (0, \frac{2}{7}v]$.

   In this case, the myopic best reply is full stealing of customers if the rival's price is above $\hat{p} := 3c$, and a partial stealing is optimal otherwise. Hence, the functional form of the gain from deviation changes at $\hat{p}$;

$$d(p) = \begin{cases} \frac{(p-c)^2}{8c} & \text{if} \quad p \leq \hat{p} \\ \frac{p}{2} - c. & \text{if} \quad \hat{p} \leq p \end{cases}.$$

   Assumptions A1-A6 are satisfied for $p \in [p^N, p^*] = [c, v - \frac{c}{2}]$,[9] so Theorem 1 implies that the optimal trigger strategy equilibrium exists and its plan is a solution to the following differential equation:

$$\frac{dp}{dt} = f(p) = \begin{cases} \lambda \frac{p+3c}{2} & \text{if} \quad p \leq \hat{p} \\ \lambda(2p - 3c). & \text{if} \quad \hat{p} \leq p \end{cases}.^{10} \tag{7}$$

[9] Since $\lim_{p \nearrow \hat{p}} \frac{d(p)-d(\hat{p})}{p-\hat{p}} = \frac{1}{2} = \lim_{p \searrow \hat{p}} \frac{d(p)-d(\hat{p})}{p-\hat{p}}$, d is differentiable at $p = \hat{p}$. Hence, A5 is satisfied. Note that π_i is not differentiable at $\hat{p} - c$; however, A5 only requires the differentiability of d, so we can still apply our theorem.

[10] From this differential equation, we can compute the optimal plan. Specifically, we first use the differential equation $\frac{dp}{dt} = \lambda \frac{p+3c}{2}$ for the region $[p^N, \hat{p})$ with the initial condition at the deadline given by the time-price pair $(0, c)$ (c is the Nash price p^N). Then, we consider the differential equation $\frac{dp}{dt} = \lambda(2p - 3c)$ for the region $[\hat{p}, p^*]$ with the initial condition given by the time-price pair $(t(\hat{p}), \hat{p})$, where we define $t(\hat{p}) := \lim_{p \nearrow \hat{p}} t(p)$ with $t(p)$ being the time at which the solution to the first differential equation is at price p. Note that, since the Finite Time Condition holds with $d(p) = \frac{(p-c)^2}{8c}$, $t(\hat{p})$ is finite.

Since $\frac{1}{f(p)} < \infty$ holds for all $p \in (p^N, p^*]$and $\lim_{p\downarrow c} \left|\frac{1}{f(p)}\right| < \infty$, Lemma 1 implies that the optimal trigger strategy equilibrium is nontrivial.

3. **No product differentiation: $c = 0$.**

   In this case, an infinitesimal price-cut can steal the entire unit mass of buyers. Thus, the supremum payoff from deviating from the price profile (p, p) is $p \times 1 = p$. This implies that $d(p) = p - \frac{p}{2} = \frac{p}{2}$.[11] A1-A2 and A4-A6 are satisfied for $p \in [p^N, p^*]$ (see footnote 11 regarding A3). Comparing the gain from deviation with the size of punishment, we have:

$$\liminf_{p\downarrow p^N} \frac{d(p)}{\pi(p) - \pi^N} = \liminf_{p\downarrow p^N} \frac{\frac{p}{2}}{\frac{p}{2} - 0} = 1 > 0.$$

   Hence, Part 3 of Theorem 1 implies that *no cooperation is sustained by the trigger strategy when there is no product differentiation.*

Summary and comparative statics

Overall, we obtain a conclusion that *a nontrivial collusive plan exists if and only if there is a product differentiation.* The intuition is as follows. If there is no product differentiation, each firm can steal the entire profit of the rival firm by an infinitesimal price-cut, whenever the current price is strictly higher than the marginal cost (which is 0 in this example). This is because all buyers switch to the deviating firm. Hence, if the current price p is not equal to the Nash price $p^N = c = 0$, the gain from deviation ($d(p) = \frac{p}{2}$) is *of the same order in magnitude* as the gain from cooperation ($\pi(p) - \pi^N = \frac{p}{2}$), however close p is to 0. This makes cooperation impossible. If there is a product differentiation, however, only a small fraction of buyers switch to the deviating firm as a result of a marginal price cut. As a result, the gain from deviation ($d(p) = \frac{(p-c)^2}{8c}$) near the Nash price ($p^N = c$) is of an order of magnitude smaller than the gain from cooperation ($\pi(p) - \pi^N = \frac{p-c}{2}$) and this makes cooperation possible.

[11] Technically speaking, d is not well-defined in the formula in A3 because there is no best response. In this paragraph, we use a modified definition of d in which we replace "max" in the definition with "sup." For a similar reason, $\pi_1(a_1, a_2)$ is not continuous, which is again a violation of A3. However, the proof of Theorem 1 only uses the fact that $\pi(x)$ is continuous and $d(x) = \sup_{a_1} \pi_1(a_1, a) - \pi_1(a, a)$ exists and is continuous. These conditions are satisfied in the case of $c = 0$ here, so the results in Theorem 1 still go through.

The differential equations in cases 1 and 2 above have closed-form solutions, and a formal description of the optimal trigger strategy equilibrium plan is described in the following:

Proposition 3 *In the price competition revision game, the optimal trigger strategy equilibrium plan, $\bar{p}(t)$, is characterized as follows:*

1. *If $c \in (\frac{2}{7}v, \frac{2}{3}v)$,*

$$\bar{p}(t) = \begin{cases} c\left(4e^{\lambda\frac{t}{2}} - 3\right) & \text{if} \quad t < t(p^*) \\ p^* = v - \frac{c}{2} & \text{if} \quad t(p^*) \leq t \end{cases},$$

where $t(p^*) = \frac{2}{\lambda}\ln\left(\frac{v}{4c} + \frac{5}{8}\right)$ is the time to achieve fully collusive price p^*.

2. *If $c \in (0, \frac{2}{7}v]$,*

$$\bar{p}(t) = \begin{cases} c\left(4e^{\lambda\frac{t}{2}} - 3\right) & \text{if} \quad t < \widehat{t} \\ c\left(\frac{8}{27}e^{2\lambda t} + \frac{3}{2}\right) & \text{if} \quad \widehat{t} \leq t < t(p^*) \\ p^* = v - \frac{c}{2} & \text{if} \quad t(p^*) \leq t \end{cases},$$

where $\widehat{t} = \frac{2}{\lambda}\ln\left(\frac{3}{2}\right)$ and $t(p^*) = \frac{3}{2\lambda}\ln\left(\frac{3}{2}\right) + \frac{1}{2\lambda}\ln\left(\frac{v}{c} - 2\right)$.

3. *If $c = 0$, $\bar{p}(t) = 0$ for all t.*

The parameter $\widehat{t}$ in Case 2 is the time to achieve the critical price $\hat{p}$ in (7), where the functional form of the gain from deviation (and therefore that of the differential equation) changes. The optimal plans are depicted in Figure 4 for $v = 10$ and $c = 0, 1, 2, 3,$ and 5. As c deceases, the Nash price p^N decreases and the fully collusive price p^* increases. The optimal plans are the curves connecting these two prices. Note that the optimal plan is the solution (the blue curve) to a single differential equation (6) when c is high ($c = 5$, 3). In contrast, when c is low ($c = 2$, 1), the optimal plan consists of the solutions to the two differential equations in (7) pasted together at the critical price level $\hat{p} := 3c$ (the blue and red curves pasted at the black dots). The figure shows that, as the degree of product differentiation goes down to zero, the expected number of price revisions increases. This can be seen from the fact that the optimal price path departs from the fully collusive level farther away from the deadline when the degree of product differentiation is smaller (i.e., when c is smaller).

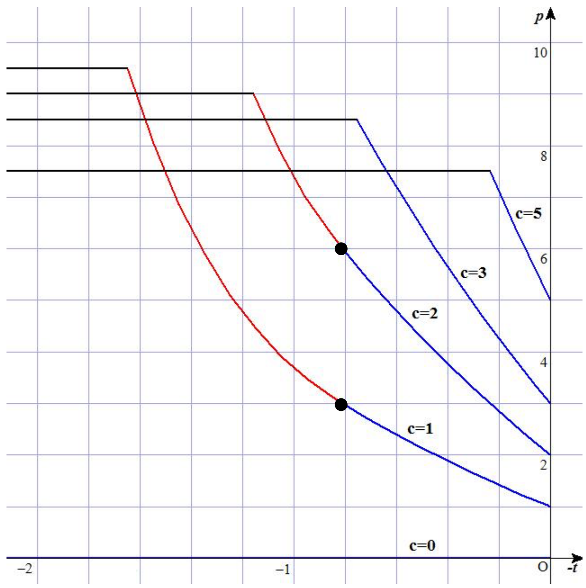

Figure 4: The optimal trigger strategy equilibrium plan $p(t)$ for the Bertrand competition game: $\lambda = 1$, $v = 10$. The black dots represent the time-price pairs $(\hat{t}, \hat{p})$ at which the two paths are pasted (cf. footnote 9).

However, at the limit (i.e., when $c = 0$), no price revision occurs even on the optimal trigger strategy equilibrium.

Let us calculate the expected payoff. So far, we have treated c (transportation cost) as the degree of product differentiation. What really matters, however, is the magnitude of c *relative to* v (willingness to pay). Hence, we use $h := \frac{c}{v}$ as the degree of product differentiation, and the expected payoff is characterized as follows.

Corollary 2 *In the price competition revision game, for any* $h = c/v \in [0, \frac{2}{3})$ *and* $T > t(p^*)$*, the expected payoff under the optimal trigger strategy equilibrium is*

$$\begin{cases} v\left(\frac{5}{2}h - \frac{1}{4}\left(\frac{(8h)^2}{2+5h}\right)\right) & \text{if} \quad h \in (\frac{2}{7}, \frac{2}{3}) \\ v\left(\frac{h}{2} + \frac{4h}{9}\left(\left(\frac{3}{2h} - 3\right)^{\frac{1}{2}}\right)\right) & \text{if} \quad h \in (0, \frac{2}{7}] \\ 0 & \text{if} \quad h = 0 \end{cases}.$$

The proof is given in Appendix A.2. Figure 3 at the beginning of this subsection shows the graph of the expected payoff and other benchmark payoffs. It shows that a significant degree of payoff improvement relative to the Nash equilibrium is achieved in the revision game.

The next corollary shows that the degree of collusion is increasing in the degree of product differentiation. Consider two measures of the degrees of collusion: First, $\bar{C}(h)$ is the expected payoff under the optimal trigger strategy equilibrium divided by the fully collusive payoff under h. Second, $\tilde{C}(h)$ is the expected payoff under the optimal trigger strategy equilibrium minus the Nash payoff, divided by the fully collusive payoff minus the Nash payoff under h. Note that $\tilde{C}(h)$ is a more conservative measure than $\bar{C}(h)$ because it measures the ratio of the payoff *increment* relative to the Nash payoff.

Corollary 3 *The two measures of the degree of collusion are strictly increasing in the degree of product differentiation:* $\bar{C}'(h) > 0$ *and* $\tilde{C}'(h) > 0$ *if* $T > t(p^*)$ *under* h.

Table 1 shows, for various degrees of product differentiation h, the values of $\bar{C}(h)$ and $\tilde{C}(h)$ when the horizon is long enough.[12] As Corollary 3 predicts, those ratios are increasing in h. The table shows that the opportunities of revising prices can provide high levels of collusion, under reasonable degrees of product differentiation. For example, if $h = .5$, on average, a buyer's willingness to pay for the worse good is 71.4% of that of the preferred good. For such a degree of product differentiation, 96% of fully collusive payoffs can be achieved in the revision game. Even under the more conservative measure of cooperation, 89% of the increment in the expected profit relative to the Nash profit is achieved through the revision game.

3.3 Election Campaign: Policy Platforms Gradually Converge

The celebrated median voter theorem shows that two parties offer identical policy platforms. In reality, however, two parties often start with quite different platforms, and during an election campaign they gradually converge. To explain such a phenomenon, we present a simple election model with policy-motivated candidates and show that *the policies converge over time towards the middle of the policy space*. A sample path of announced policies under the optimal trigger strategy equilibrium is depicted in Figure 5. It demonstrates that the policies become closer and closer to each other towards the end of the election campaign, but do not converge perfectly to the Nash policies even at the election date (i.e., at $t = 0$). In our model, we pa-

[12] We can calculate these values since the ratios are the values derived in the corollary to $\pi^* = \frac{1}{2}\left(v - \frac{c}{2}\right)$. We derive these ratios in Appendix A.2.

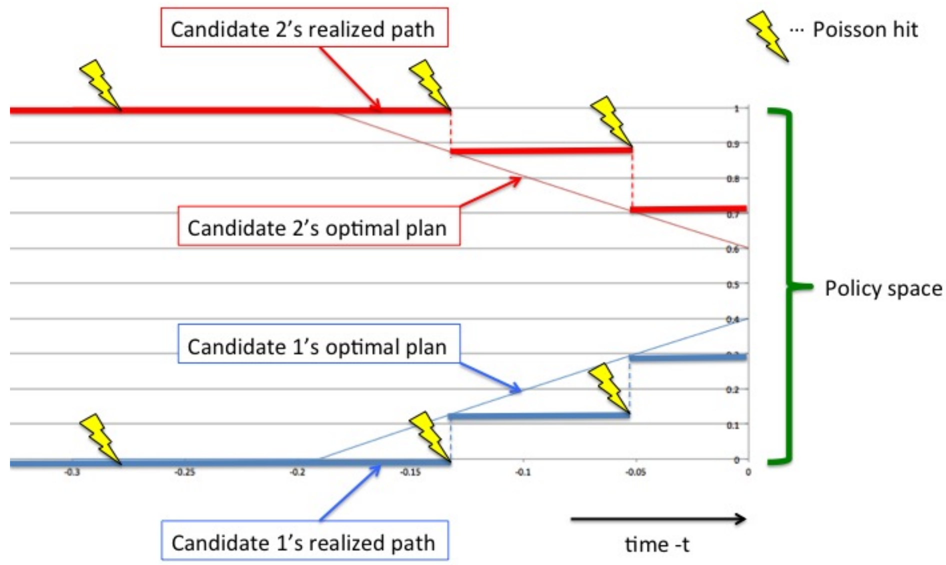

Figure 5: A realized policy path under the optimal trigger strategy equilibrium plan profile $(y_1(t), y_2(t))$ under $v = 0.9$ for the election game: $\lambda = 1$.

rameterize the degree of office motivation, and it turns out to have a non-monotonic effect on the expected number of the policy revisions and the size of a revision (conditional on there being a revision): it is positive and increasing when the degree of office motivation is not too large, while it suddenly drops to zero and stays constant above a certain threshold. The nonmonotonicity is due to the fact that the Finite Time Condition fails if the degree of office motivation is above the threshold. We will explain this point later.

The crucial assumption to derive gradual convergence of policies is that *candidates have strong opinions about what the right policy should be.* In such a situation, a candidate's payoff sharply decreases as the implemented policy moves away from her bliss point, and once implemented policy is far away, where it is located does not matter so much. This can happen if the policy platforms concern such issues as same-sex marriage, abortion, or gun control. As we will explain below, the *convexity* of policy payoff captures this crucial assumption.

Model

The policy space is the interval $[0,1]$. There are two candidates, $i = 1, 2$, where candidate i chooses policy y_i. There is a continuum of voters, each of whom votes for one of the candidates (no abstention). The candidate who attracts a larger amount of votes wins the election, and each candidate is elected with probability $\frac{1}{2}$ in the case of ties. Each voter votes for the candidate whose policy is closer to her bliss point, and randomizes between the two candidates with equal probabilities when indifferent. The population distribution of the bliss points of voters is not fixed, but stochastic. The winner of the election, given platforms y_1 and y_2, is determined by the position of the median voter, which is a random variable and denoted by y^*.

As in the standard models of policy-motivated candidates, the candidates do not know the position of the median voter, while its distribution is common knowledge.[13] For simplicity, we assume that y^* is uniformly distributed over the policy space $[0,1]$. In that case, the winner of the election, denoted by random variable $w(y_1, y_2) \in \{1, 2\}$, is determined as

$$w(y_1, y_2) = 1 \text{ with probability } \begin{cases} \frac{y_1+y_2}{2} & \text{if} \quad y_1 < y_2 \\ \frac{1}{2} & \text{if} \quad y_1 = y_2 \\ 1 - \frac{y_1+y_2}{2} & \text{if} \quad y_2 < y_1 \end{cases} .$$

To see why, consider the case $y_1 < y_2$. Candidate 1 wins if a majority of voters have bliss points less than $\frac{y_1+y_2}{2}$. This happens when the position of median voter y^* is less than $\frac{y_1+y_2}{2}$. Under the uniform-distribution assumption, the probability of this event is just equal to $\frac{y_1+y_2}{2}$.

Candidate i's realized payoff is given by, for $w = w(y_1, y_2)$,

$$g_i(y_i, y_{-i}) = v \cdot \mathbb{I}_{\{i=w\}} + b(|y_w - \overline{y}_i|)$$

where $v \geq 0$ represents the value of winning *per se*, and $b(\cdot)$ is a "policy payoff" which depends on the distance between the winner's policy y_w and candidate i's "bliss point" that is denoted by $\overline{y}_i$. We assume that $\overline{y}_1 = 0$ and $\overline{y}_2 = 1$. That is, candidate 1 is "left wing" and candidate 2 is "right wing." For the moment, we also assume that

[13] Pioneering papers on policy-motivated candidates such as Alesina (1988) and Roemer (1994) employ such an assumption. Specifically, Alesina (1988) considers a model where every possible policy profile determines a non-degenerate probability of winning, and Roemer (1994) considers a situation where a voter distribution is chosen from a non-singleton set of distributions.

$v \in (\frac{1}{2}, 1]$ which will turn out to be the case in which cooperation can be sustained. We will discuss the case with $v \notin (\frac{1}{2}, 1]$ later.

Candidates first set their policies at time $-T$, and then make revisions over the time internal $(-T, 0)$ without knowing the realized position of the median voter. We interpret the revision phase as the time period for an election campaign. In the revision phase, candidates obtain opportunities to express their policy positions, for example, at an open policy debate on radio or television. At each opportunity, candidates can amend or revise their policy platforms (as is often the case in reality). At time 0 of the revision game, the election takes place, and the elected candidate is committed to implementing her finally-announced policy.[14]

Two key assumptions

On top of the above standard specification of the election model with policy-motivated candidates, we postulate two additional assumptions:

1. First, we assume that the policy payoff function $b(\cdot)$ is *convex*. Such policy preferences are especially relevant for issues that provoke strongly opinionated reactions (e.g. same-sex marriage, abortion, gun control, and so forth). This is because, for these policy issues, it is natural to assume that one's utility arising from policy preferences $b(|y_w - \overline{y}_i|)$ sharply decreases as the winner's policy (y_w) moves away from her bliss point ($\overline{y}_i$).[15] To capture this possibility in the simplest setting, we assume $b(z) = \max\{\frac{1}{2} - z, 0\}$. With this specification, a unique pure Nash equilibrium exists. It is not generally equal to $\frac{1}{2}$ and given by a symmetric policy profile around $\frac{1}{2}$. Since the winning probability for each candidate is always equal to $\frac{1}{2}$ under *any* symmetric profiles, a profile $(y_1, y_2) = (0, 1)$ Pareto-dominates this Nash policy profile when the latter is not $(0, 1)$, as is illustrated by Figure 6. This means that there is a potential room for cooperation in the revision game as long as the Nash profile is not $(0, 1)$.[16]

[14] This "policy announcement game" is proposed and analyzed in Kamada and Sugaya (2014), in which they analyze the case where candidates cannot announce inconsistent policies while they have an option not to specify their policies. Their analysis is based on an analogue of backward induction, and thus different from ours.

[15] See Osborne (1995) for a criticism on the use of concave utility functions for preferences over electoral policies. Kamada and Kojima (2014) discuss implications of convex voter utility functions.

[16] Note that, on the other hand, there would be no room for cooperation if b is concave, as traditionally assumed in the political science literature.

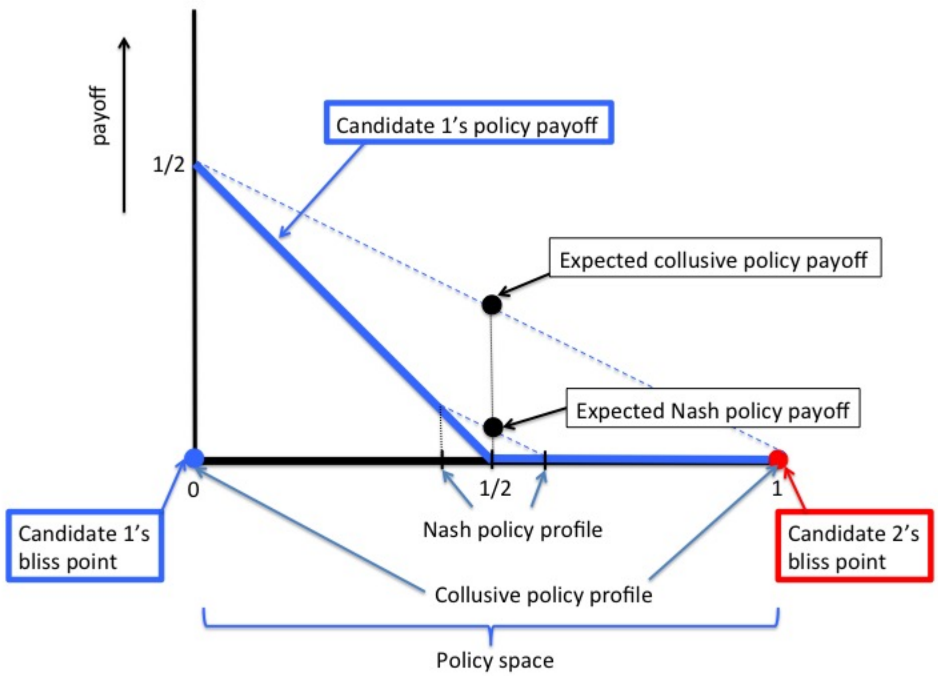

Figure 6: The policy preference term for the election game.

2. Second, we assume that candidate 1 chooses policy y_1 from $[0, \frac{1}{2}]$ and candidate 2 chooses policy y_2 from $[\frac{1}{2}, 1]$. The motivation behind this assumption is that candidate 1 is faithful to her party's identity ("left wing") so that she never wants to choose a "right wing" policy in $[\frac{1}{2}, 1]$, possibly because of reputational concerns. Symmetric explanation applies to candidate 2.[17]

Differential equation

The payoff functions are not symmetric as they are, but by redefining actions by

$$x_1 = y_1 \text{ and } x_2 = 1 - y_2,$$

[17]Technically, if each candidate can choose a policy from $[0, 1]$, then there does not exist a pure Nash equilibrium in the component game because a best response does not necessarily exist. We conjecture that, if we allowed for candidates to choose their policies from $[0, 1]$, there would exist a nontrivial equilibrium in which at each opportunity candidates mix across multiple policies. Kamada and Sugaya (2014) analyze a mixed strategy plan in the context of their model, and their simulation result shows quite complicated dynamics of mixing probabilities. For this reason, here we do not delve into such an analysis.

we can retain symmetry.[18] The new variable x_i measures the distance between her policy and her bliss point. Let $\pi_i(x_1, x_2)$ be the associated payoff.

The winning probability of candidate i is given by $\frac{1+x_i-x_{-i}}{2}$, and the unique Nash policy $x^N = \frac{2v-1}{2}$, where $v \in (\frac{1}{2}, 1]$ represents the value of winning.[19] Note that the Nash policy profile is $(y_1, y_2) = (\frac{2v-1}{2}, 1 - \frac{2v-1}{2})$, and (i) it is close to $(\frac{1}{2}, \frac{1}{2})$ when value of winning v is high (close to 1), and (ii) it approaches the bliss profile $(0, 1)$ as v decreases towards $\frac{1}{2}$. For any value of $v \in (\frac{1}{2}, 1]$, the fully collusive profile $(0, 1)$ is better than the Nash profile, because of the convexity of the policy payoff function. A straightforward calculation shows that $\pi(x) = \frac{1}{2}\left(v + \frac{1}{2} - x\right)$, $\pi^N = \frac{1}{2}$, and $d(x) = \frac{1}{8}(v - \frac{1}{2} - x)^2$.

One can check that A1-A6 are satisfied for $x \in [0, \frac{2v-1}{2}]$, so Theorem 1 implies that the optimal trigger strategy equilibrium exists and its plan is a solution to the following differential equation:

$$\frac{dx}{dt} = f(x) = \lambda \frac{2x - 2v - 7}{4}.$$

Now, $\left|\frac{1}{f(x)}\right| < \infty$ holds for all $x \in [0, \frac{2v-1}{2})$ and $\lim_{x \uparrow \frac{2v-1}{2}} \left|\frac{1}{f(x)}\right| < \infty$, Lemma 1 implies that the optimal trigger strategy equilibrium is nontrivial.

The cases with weak and strong office motivation

In the analysis so far, we assumed $v \in (\frac{1}{2}, 1]$. If $v \leq \frac{1}{2}$, the office-motivation is so weak that $(0, 1)$ is the Nash policy profile, so there is no room for cooperation. On the other hand, if $v > 1$, then the office motivation is so strong that the first-order condition does not hold at the Nash policy profile $(\frac{1}{2}, \frac{1}{2})$. Specifically, we have that

$$d(x) = \begin{cases} \frac{1}{2}(v-1)\left(\frac{1}{2} - x\right) & \text{if} \quad x > \frac{3}{2} - v \\ \frac{1}{8}(v - \frac{1}{2} - x)^2 & \text{if} \quad x \leq \frac{3}{2} - v \end{cases}.$$

[18] Assumption A1 stipulates that $a^N < a^*$ holds, and this is not satisfied in the current example. However, we can still use Theorem 1 by relabeling actions (for example, multiply -1 to each action). The only difference is that in the Finite Time Condition, we now have the action in the limit approaching the Nash action from below, instead of having it approaching from above as stated in condition (2).

[19] If the action space were $[0, 1]$ instead of $[0, \frac{1}{2}]$, then at the policy profile $(y_1, y_2) = (\frac{2v-1}{2}, 1 - \frac{2v-1}{2})$, candidate 1 would have an incentive to deviate to another policy $1 - \frac{2v-1}{2} - \epsilon$ for a small enough $\epsilon > 0$ (and the same is true for candidate 2).

Thus,

$$\liminf_{x\uparrow\frac{1}{2}} \frac{d(x)}{\pi(x)-\pi^N} = \liminf_{x\uparrow\frac{1}{2}} \frac{\frac{1}{2}(v-1)\left(\frac{1}{2}-x\right)}{\frac{1}{2}(v+\frac{1}{2}-x)-\frac{v}{2}} = \frac{1}{2}(v-1) > 0.$$

Hence, part 3 of Theorem 1 implies that the Finite Time Condition fails and the unique trigger strategy equilibrium is to set the policy at $x=\frac{1}{2}$ all the time.

Summary

The above differential equation has a closed-form solution as follows.

Proposition 4 *In the revision game of the election game, the optimal trigger strategy equilibrium plan, $(\bar{y}_1(t), \bar{y}_2(t))$, is characterized by the following:*

1. *If $v \in [0, \frac{1}{2}]$, then $\bar{y}_1(t) = \bar{y}_2(t) = 0$ for all t.*
2. *If $v \in (\frac{1}{2}, 1]$, then*

$$\bar{y}_1(t) = \begin{cases} \frac{7+2v-8\cdot e^{\frac{\lambda}{2}t}}{2} & \textit{if} \quad t < t(y_1^*) \\ y_1^* = 0 & \textit{if} \quad t(y_1^*) \le t \end{cases},$$

where $t(y_1^*) = \frac{2}{\lambda}(\ln(7+2v) - 3\ln 2)$ is the time to achieve full collusion and $\bar{y}_2(t) = 1 - y_1(t)$.

3. *If $v \in (1, \infty)$, then $\bar{y}_1(t) = \bar{y}_2(t) = \frac{1}{2}$ for all t.*

The above proposition shows that in the revision game of the election game, when the office motivation is not too large or too small, each candidate starts from announcing their most preferred policies. They just stick to their original announcements until a certain time ($t(y_1^*)$) before the election day, and then begin catering to the middle towards the end. Thus the model captures the well-observed phenomena of candidates changing their policy announcements, moving towards the middle when the election day is approaching. The plan characterized in Proposition 4 is depicted in Figure 7 for various values of parameter v. Notice that there is a discontinuity at $v = 1$, i.e., the limit of the optimal plan as $v \downarrow 1$ does not converge to the optimal plan at $v = 1$. The sample path Figure 5 presented at the beginning of this section corresponds to the case with $v = 0.9$.

The closed-form of the optimal plan enables us to obtain implications about the observed behavior. To state those implications formally, let us introduce a few pieces

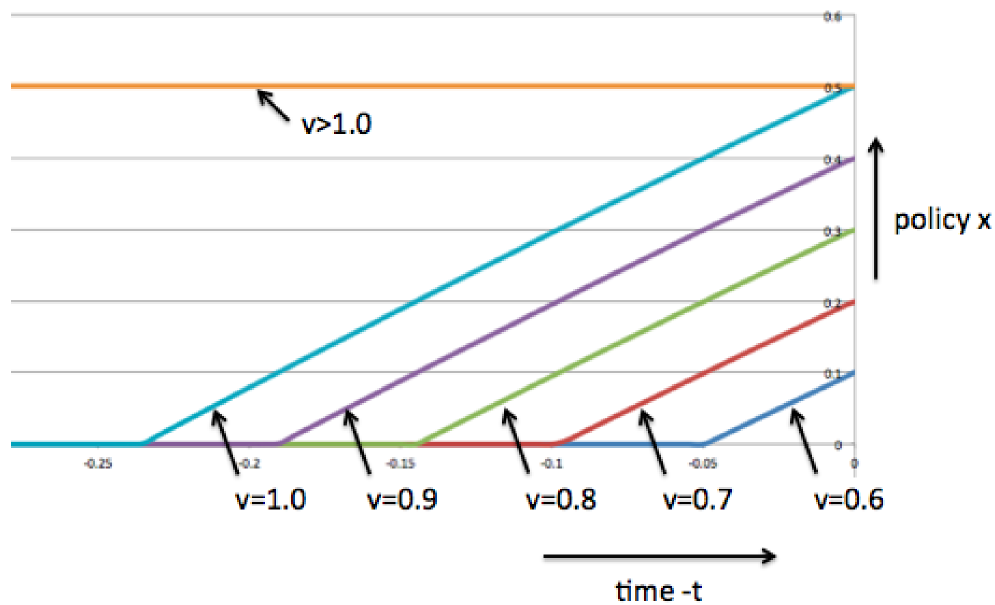

Figure 7: The optimal trigger strategy equilibrium plans $y_1(t)$ for various values of v for the election game: $\lambda = 1$.

of notation to measure the magnitude of policy revisions. Fix a component game with parameter v and consider the optimal trigger strategy equilibrium. For $t < T$, let $F_t^v(a)$ denote the cumulative distribution function of action a at the beginning of time $-t$ before revisions can possibly take place at that moment. That is, it represents the probability distribution of the last action prepared strictly before time $-t$. Also, in order to make clear the dependence of the optimal plan on v, denote by x^v the optimal plan. Let $T > t(y_1^*)$ (defined in Proposition 4: $-t(y_1^*)$ is the time when the optimal plan departs from the optimal action). For each $t \in [0, T)$, define

$$\Delta_t(v) := \int_0^{\frac{1}{2}} (x^v(t) - a) dF_t^v(a).$$

That is, $\Delta_t(v)$ measures the *expected size of the policy change* at time $-t$ when the component game has parameter v, conditional on there being a revision opportunity at $-t$. Finally, for the component game with parameter v, let $N(v)$ be the expected number of policy changes under the optimal trigger strategy equilibrium.

Corollary 4 *In the revision game of the election game, the following hold.*

1. *Suppose that* $v \in (\frac{1}{2}, 1]$ *and* $T > t(y_1^*)$ *under* v *(defined in Proposition 4).*

   (a) *Then, the expected number of policy changes under the optimal trigger strategy equilibrium is strictly increasing in* v*:* $N'(v) > 0$.

   (b) *For each* $t \in [0, t(y_1^*))$*, the expected policy change is strictly increasing:* $\Delta_t'(v) > 0$.

2. *Suppose that* $v \notin (\frac{1}{2}, 1]$*. Then, there is no policy change under the optimal trigger strategy equilibrium plan.*

The proof is in Appendix A.4. The corollary implies that the magnitude of policy changes, measured by the expected number of changes and the conditional size of each change, is greater when office motivation is larger. This is because a larger office motivation makes it difficult to keep candidates away from catering to the middle.

The difference in the behavior has an implication on the candidates' payoffs for different levels of office motivation. To see this, first let us calculate the equilibrium payoffs.

Corollary 5 *In the revision game of the election game, if* $v \in (\frac{1}{2}, 1]$ *and* $T > t(y_1^*)$ *(defined in Proposition 4), the expected payoff under the optimal trigger strategy equilibrium plan is* $\frac{5}{2} - \frac{16}{7+2v}$.[20]

The calculation is given in Appendix A.3. Figure 8 depicts the expected payoffs stated in part 1 of Corollary 5. As mentioned in the discussion of Figure 7, there is a discontinuity of the optimal plan at $v = 1$. This results in the discontinuity at $v = 1$ in Figure 8.

Define the measures of collusion $\bar{C}(v)$ and $\tilde{C}(v)$ as in the last subsection. That is, we use $\bar{C}(v)$ to denote the expected payoff under the optimal trigger strategy equilibrium divided by the fully collusive payoff under v. Also, $\tilde{C}(v)$ is the expected payoff under the optimal trigger strategy equilibrium minus the Nash payoff, divided by the fully collusive payoff minus the Nash payoff under v.

[20] As in the case of Bertrand competition, Appendix A.3 derives the ratio of this payoff to the fully collusive payoff.

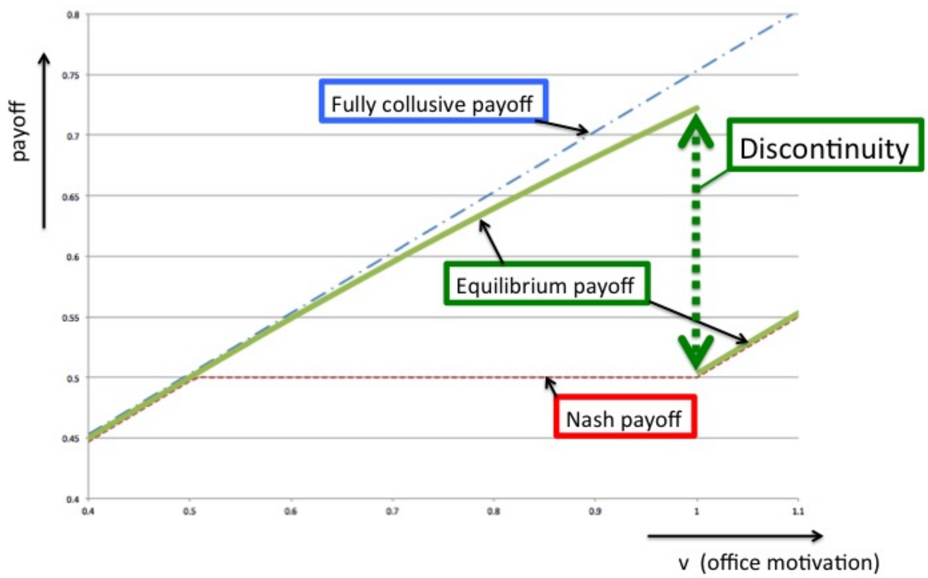

Figure 8: The expected payoffs under the optimal trigger strategy equilibrium plan for the election campaign game.

Corollary 6 *In the revision game of the election game, the two measures of the degree of collusion are strictly decreasing in the levels of the strength of office motivation:* $\bar{C}'(h) < 0$ *and* $\tilde{C}'(h) < 0$ *if* $T > t(y_1^*)$ *under* v *(defined in Proposition 4).*

Hence, whenever the optimal plan is nontrivial, the degree of collusion is strictly decreasing in the degree of office motivation, under the two measures. The reason is, again, that a larger office motivation makes it difficult to keep candidates away from catering to the middle.

4 Robustness of Cooperation

4.1 Discrete Time Model and Small Willingness to Punish

There are two key features to sustain cooperation in revision games. First, because revision opportunities arrive by a Poisson process in continuous time, *there is no "last revision opportunity"*: even when a revision opportunity arrives very close to the deadline, there is a positive probability that another revision (where punishment is imposed) is possible in the future. Second, even when punishment is fairly mild, some degree of (or "a little bit" of) cooperation is possible. This is guaranteed by the *continuous action spaces (and smooth payoffs)* in our model. If either of those two features were absent, no cooperation would be possible in our model. This begs the question of the robustness of our results.

Consider the application of *one* of the following modifications to our model. We continue assuming that the component game has a unique and pure Nash equilibrium.

1. Time is discrete and finite, $-K\Delta, ..., -2\Delta, -\Delta, 0$. In each period, players can revise their actions with some constant probability $\gamma > 0$.[21]

2. It takes time $\epsilon > 0$ to react to revised actions.[22]

3. Action space of each player is finite.

If one of those modification is applied, no cooperation is possible; the only equilibrium is that each player always chooses the Nash action of the component game. For Modification 1, the usual backward induction argument works: In period 0, players play Nash actions because there is no punishment in the future, and the same is true for any previous periods. Modification 2 is similar. In time interval $(-\epsilon, 0]$, there is no future punishment and therefore the Nash actions are chosen. Given this, the same is true for $(-2\epsilon, \epsilon]$, and so on.

Lastly, consider Modification 3. For simplicity, let us first consider trigger strategy equilibria in *pure* strategies. Take any non-Nash action profile (a_1, a_2).[23] There should be at least one player who can gain by deviating from this profile, and denote the gain from deviation by $\underline{d}(a_1, a_2) > 0$. Since actions are finite, the minimum of $\underline{d}(a_1, a_2)$ over all non-Nash pure profiles, denoted by $\underline{d} > 0$, exists. Also, let

$$P = \max_{i=1,2} \left[\left(\max_{(a_1,a_2)} \pi_i(a_1, a_2) \right) - \left(\min_{(a'_1,a'_2)} \pi_i(a'_1, a'_2) \right) \right]$$

[21]This is a discrete-time approximation of our model: when $K = \frac{T}{\Delta}$, $\gamma = \lambda\Delta$, and $\Delta \to 0$, we obtain our Poisson model.

[22]That is, if players revise at time $-t$ and another revision opportunity arrives at $-s \in (-t, -t+\epsilon)$, players cannot react to the action profile $a(t)$.

[23]There is no such profile if and only if each player's action space is a singleton. However, in such a case, it is obvious that there is a unique subgame-perfect equilibrium, and it is a repetition of the unique Nash equilibrium of the component game.

be the maximal punishment that can be physically imposed. Note that P is finite because the action space is finite for each player.[24] Recall that $e^{-\lambda t}$ is the probability of no revision when the remaining time is t, and define $t^* > 0$ by the unique solution to the equality $e^{-\lambda t^*}\underline{d} = \left(1 - e^{-\lambda t^*}\right) P$. At any time $-t \in (-t^*, 0]$, if a non-Nash action profile is chosen, the gain from deviation of some player is at least $e^{-\lambda t}\underline{d}$, while the physically possible future punishment is at most $\left(1 - e^{-\lambda t}\right) P < e^{-\lambda t}\underline{d}$. Hence, no cooperation is possible in $(-t^*, 0]$. Repeating the same argument, we can show that no cooperation is possible in $(-nt^*, -(n-1)t^*]$ for all $n = 2, 3, ...$ One can show that an analogous conclusion continues to hold even if we allow for mixed strategies if the unique Nash equilibrium is strict. Formalizing this idea, however, is a complicated task. We relegate it to Appendix A.5.

These observations suggest that the possibility of cooperation in revision games might not be robust to certain changes of our assumptions. In what follows, we closely examine this issue. First, note that the preceding observations are partly similar to those for the robustness of cooperation in infinitely repeated games. If the stage game has a unique Nash equilibrium, no cooperation is sustained in a finitely repeated game. In reality human players eventually die, and this observation seems to suggest that no cooperation would be possible among human players in a repeated interaction. A common defense for the theory of infinitely repeated games states, however, that even if human players eventually die, at any point in time they may face some possibility of future interaction (and therefore the theory applies to such a situation). Similarly in revision games, if the deadline is "soft," even if the deadline is eventually implemented, at any point in time players may face some possibility of future revisions. In reality deadlines are often soft, and cooperation would be sustained when we add this feature to the discrete-time model (Modification 1) or the reaction-time model (Modification 2).

What if the deadline is not soft? When the deadline is "firm," if a revision opportunity comes very close to the deadline, no more revision would be possible in reality. Modifications 1 or 2 formulate this fact, and one may argue that adding those modifications would make the model more realistic and our central conclusion (sustainability of cooperation) would fail. This is a valid concern and we take it seriously.

It is true that our results are not robust to the realistic modifications (1 or 2),

[24] Again, $P = 0$ if and only if the action space is a singleton for each player.

provided that we keep our assumption of complete selfishness. However, if we also incorporate another realistic aspect of human beings that we have a tendency to punish a deviator (if it is not too costly), then our conclusion would be restored. Willingness to punish a deviator is a well-documented fact in behavioral economics (for example, see Fehr and Gachter (2002)). Our model shows that the tendency to punish a deviator near the deadline *need not be so strong* to sustain substantial cooperation. The interaction between gain from deviation and benefit of cooperation builds up: *in the revision process eventually players can sustain substantial levels of cooperation, when they have a fairly mild incentive to punish a deviator near the deadline.* We view that this is the take-home message of our model, which can be applicable to real-life problems.

To make this point more formally, we offer a numerical analysis of the discrete-time model (Modification 1).[25] As we have argued, no cooperation is sustained in this model if players are completely selfish. We examine what happens if players have an ε-incentive to punish a deviator. More specifically, we consider the following two component games, both of which have action space $[0, 1]$.

- Model 1 : $\pi_i(a_i, a_{-i}) = 2a_{-i} - \max\{a_i^2 - \varepsilon, 0\}$. (Cooperation is sustained in the revision game with $\varepsilon = 0$.)
- Model 2 : $\pi_i(a_i, a_{-i}) = 2a_{-i} - \max\{a_i - \varepsilon, 0\}$. (Cooperation cannot be sustained in the revision game with $\varepsilon = 0$.)

Under complete selfishness ($\varepsilon = 0$), the Nash action is 0, while the optimal action is 1 in both models. When $\varepsilon > 0$, the set of Nash actions is $[0, \varepsilon]$. Hence, in the last revision opportunity, players can play $a_i = \varepsilon$ if players have cooperated in the past and choose $a_i = 0$ if a deviation occurred. This is how "a mild incentive to punish" is formulated, and the incentive to punish is represented by ε.

Time is discrete $t = 0, 1, 2, ..., T$, and t represents the remaining time until the deadline (thus $t = 0$ is the last revision opportunity). In each period, revision opportunity arrives with probability γ. In the continuous time model, a revision opportunity arrives approximately with probability $\lambda\Delta$ in a small time interval Δ (λ is the Poisson arrival rate). To compare the discrete and continuous time models, we set $\gamma = \lambda\Delta$. We assume that $\lambda = 1$, $\Delta = 0.1$, and therefore $\gamma = 0.1$.

[25] Similar arguments can be made for Modifications 2 and 3 as well.

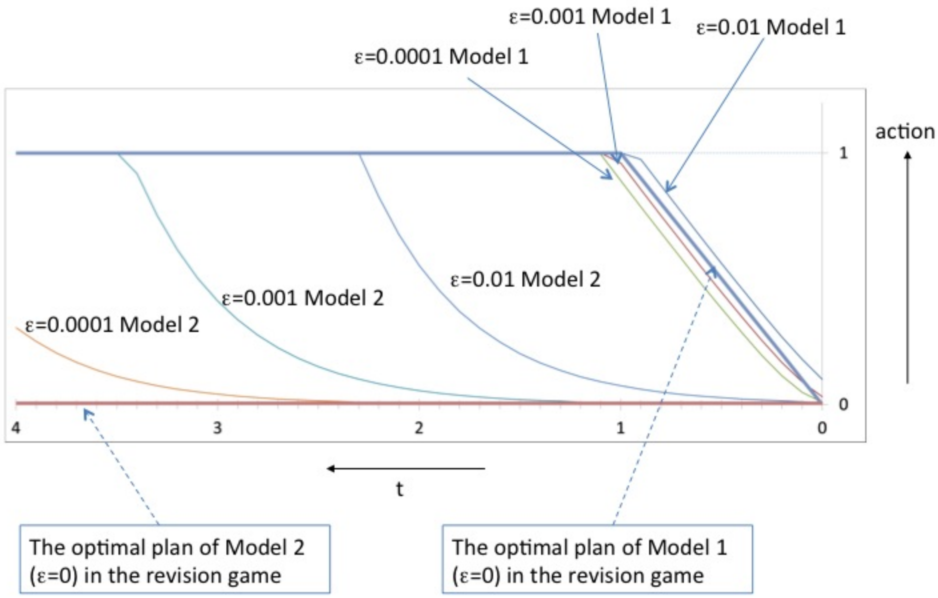

Figure 9: Finite Time Condition

The continuous time models with selfish players ($\varepsilon = 0$) have the following properties. In Model 1, the optimal plan is $x(t) = t$ until it hits the optimal action 1.[26] In Model 2, part 3 of Theorem 1 implies that the Finite Time Condition (2) fails and therefore the optimal plan does not exhibit any cooperation: $x(t) = 0$ for all t. Those predictions of revision games are depicted by the thick lines in Figure 9.

The discrete time models with ε-incentive to punish can be numerically solved backwards, and the solutions are depicted by thin lines in the figure (the discrete points are interpolated). As we can see, substantial cooperation is sustained even with tiny incentive to punish in Model 1. In contrast, in Model 2, cooperation becomes harder and harder to sustain as the incentive to punish decreases. To see the significance of the difference, we computed the expected payoffs under both models and the probability that the best action is implemented at the deadline. The result is summarized in Table 2.

To summarize, the prediction of the revision game is relevant in more realistic situations where (i) no more revision is possible near the deadline, but (ii) players

[26]This is the solution to the differential equation $\frac{dx}{dt} = \lambda \frac{d+\pi-\pi^N}{d'} = \frac{x^2+(2x-x^2)}{2x} = 1$.

ϵ	.0001	.001	.01
Expected payoff under Model 1	.432	.462	.521
Expected payoff under Model 2	.0159	.00531	.178
Probability of best action under Model 1	.282	.314	.349
Probability of best action under Model 2	.00786	.0250	.0886

Table 2: The expected payoffs for sufficiently large T and the probabilities of the best action in Models 1 and 2

have a mild incentive to punish a deviator. When cooperation is possible in our stylized model of a revision game, a fairly mild incentive to punish a deviator can sustain substantial cooperation in a more realistic situation with the properties (i) and (ii).

4.2 Asynchronous Revisions

The present paper and Kamada and Kandori (2017) have focused on synchronous revision games, which is an important first step towards the understanding of revision games. However, in some real-life situations, revisions may not be synchronized. In this subsection we consider a simple case of asynchronous revision games in which arrival rates of the two players are the same, and show that the main results of Kamada and Kandori (2017) carry over to that setting. A comprehensive analysis on the case when the arrival rates are heterogeneous can be found in Kamada and Kandori (2012).

Consider a component game with two players $i = 1, 2$. Let $\lambda_1 > 0$ and $\lambda_2 > 0$ be player 1 and 2's arrival rates, respectively. We assume that players observe all the past events in the revision game, including when revision opportunities arrived to the opponent (so i can see if j has actually followed the equilibrium action plan), and analyze the optimal symmetric trigger strategy equilibrium. Assume that *the payoff function is additively separable with respect to each player's action*. Specifically, we consider payoff functions of the following form: For each $i = 1, 2$,

$$\pi_i(a_i, a_{-i}) = b(a_{-i}) - c(a_i), \tag{8}$$

where $a_i, a_{-i} \in A$, where $A = [0, \bar{a}]$ for a finite $\bar{a}$ or $A = [0, \infty)$. We also assume that $b(0) = c(0) = 0$ and that π satisfies A1-A6. Notice that there is a unique Nash equilibrium, $(a_1, a_2) = (0, 0)$. Let a^* be the (unique) maximizer of $b(a) - c(a)$. The good exchange game in Section 3.1 fits this framework.

In general, player i's revision plan depends not only on the timing of revision but also on the opponent's action that is fixed at the time of revision (hence, a revision plan is represented by a function $x_i(t, a_{-i})$, where a_{-i} is the fixed action of the opponent at revision time $-t$). If the payoff is separable across players' actions, as we will formally show below, we can effectively ignore the dependence of the action plan with respect to the opponent's action. However, if the payoff function is not additively separable with respect to each player's action (as in the Cournot duopoly game), the dependence of revision plans on the opponent's action cannot be ignored. As a consequence, the analysis would be much more complicated than given in what follows. For example, it is not necessarily an optimal deviation to play the best response against the opponent's current action.[27]

Specifically, for each $i = 1, 2$, let $x_i(t)$ be player i's action plan at time $-t$. Fixing the opponent's action a_j, player i's payoff from cooperation plan at time $-t$ is

$$e^{-\lambda_j t} b(a_j) + \int_0^t b(x_j(\tau))\lambda_j e^{-\lambda_j \tau} d\tau - \left(e^{-\lambda_i t} c(x_i(t)) + \int_0^t c(x_i(\tau))\lambda_i e^{-\lambda_i \tau} d\tau \right). \quad (9)$$

On the other hand, using a Nash reversion, i's payoff from defection at time $-t$ is

$$e^{-\lambda_j t} b(a_j).$$

Hence, the incentive compatibility condition for player i at time $-t$ is:

$$e^{-\lambda_i t} c(x_i(t)) \leq \int_0^t \left(b(x_j(\tau))\lambda_j e^{-\lambda_j \tau} - c(x_i(\tau))\lambda_i e^{-\lambda_i \tau} \right) d\tau. \quad (10)$$

Notice that this condition does not depend on a_j, the fixed action of the opponent. This is the sense in which we said "we can effectively ignore the dependence of action plan with respect to the opponent's action." The intuition for this is simple: Whether or not player i cooperates at time $-t$, the only case where the opponent's fixed action matters in either case is when the opponent j will not have any further opportunity in the future. This happens with the same probability in the two cases, and by

[27] Kamada and Kandori (2012) demonstrate that even in the case of separable payoff functions, many complications and subtleties arise.

separability what player i is preparing does not affect the payoff from j's fixed action, $b(a_j)$.

In the case of homogeneous arrival rates, there is a simple characterization of the optimal symmetric trigger strategy equilibrium. To see this, substitute $\lambda_1 = \lambda_2 = \lambda$ in the incentive compatibility condition (10), and observe that the right hand side of the resulting condition can be simplified and is identical to the incentive compatibility condition (3) that we obtained in the main analysis. This gives us the following proposition, which implies that the results in Kamada and Kandori (2017) apply to the case of asynchronous revisions if the arrival rates are homogeneous.

Proposition 5 *The optimal trigger strategy equilibrium plans are identical under synchronous revisions with arrival rate λ and under asynchronous revisions with arrival rates λ_1 and λ_2, when the component-game payoff is separable as in (8) and the arrival rates are equal $\lambda = \lambda_1 = \lambda_2$.*

We end this subsection with two remarks. First, although the optimal plan is the same as for the case with synchronous revisions, the probability distribution of action profiles at the deadline is different. This is because two players' actions are perfectly correlated under synchronous revisions, while they are independent under asynchronous revisions. However, by additive separability, the expected payoffs stay the same even when the revisions are asynchronous.

Second, when arrival rates are heterogeneous, however, the simple characterization in the above proposition no longer applies because the right hand side of equation (10) is no longer identical to that of (3). Consequently, we need to work with two distinct incentive constraints (for the two players) simultaneously, which complicates the analysis. Kamada and Kandori (2012) deal with such a case.

5 Concluding Remarks

This paper presented wide applicability of the revision-games framework of Kamada and Kandori (2017) by considering economic applications and providing robustness analysis. As applications, we considered a good exchange game, price competition, and election game. For our robustness discussion, we focussed on timing of revisions. Specifically, we considered a variant of our model to a discrete-time setting, and also to the case with asynchronous revisions.

While we have been circulating the earlier versions of the present paper, various recent papers have analyzed revision games. Given the growing volume of such research, here we provide a brief summary of recent works, and provide our view on the future direction of research in revision games.

The first strand of such follow-up papers is concerned with uniqueness of equilibrium (and comparative statics on such a unique equilibrium) in asynchronous revision games with discrete actions. Calcagno et al. (2014) consider revision games with a finite action space and assume that revision opportunities arrive independently across players (asynchronous revision). As we show in this paper, the basic logic to sustain cooperation when the action space is continuous does not work when actions are finite. Calcagno et al. (2014) show that a quite different mechanism can operate in the finite-action case: a player's ability to commit sometimes provides an equilibrium selection when the component game has multiple equilibria. Ishii and Kamada (2011) generalize the selection results of Calcagno et al. (2014) allowing for a hybrid of synchronous and asynchronous moves. Romm (2014) examines the effect of reputation in the setting of Calcagno et al. (2014). Kamada and Sugaya (2014) introduce the first model of dynamic election campaigns into the literature on election by using a variant of the revision-games framework. Gensbittel et al. (2016) study revision games with zero-sum component games.

There are some attempts to obtain general properties of revision games. Moroni (2015) and Lovo and Tomala (2015) show existence of equilibria in their respective generalizations of revision games.

Roy (2014) conducts laboratory experiments that is related to our quantity revision game. The paper shows that the experimental results exhibit some important features of the trigger-strategy equilibrium that we identify in the present paper. This suggests that our model not only provides a theoretical possibility but also captures some mechanisms of cooperation via the revisions of actions in reality.

We suggest several possible directions for future research. First, in the continuation project, we investigate the case of asynchronous revision (Kamada and Kandori, 2012) and show that cooperation is still possible in such a setting. Second, we used trigger strategy equilibria to sustain cooperation, in which players revert to Nash actions upon deviation. Although this class of strategies is a natural starting point of analysis, a severer punishment might be possible. In our continuation work, we consider harsher punishment schemes than Nash reversion. Another direction would be to find a way to model the revision phase more generally. Revision games model the friction inherent in the revision phase as Poisson processes. Iijima and Kasahara (2016) also consider a model in which players revise their actions before the deadline, at which the component-game payoffs materialize once and for all. The friction they consider is the cost of revisions, i.e., players can make revisions at any time but there are costs associated to such revisions- the idea is reminiscent of Caruana and Einav (2008a,b).[28] It would be interesting to understand how and why the difference in the types of frictions in the revision phase implies a difference in the predictions. Finally, besides Romm (2014) that we mentioned above, a recent paper by Hopenhyan and Saeedi (2016) also incorporates imperfect information in a model like revision games (they study a dynamic auction with evolving valuations). More research in this direction may prove fruitful.

References

[1] Alesina, A (1998): "Credibility and Policy Convergence in a Two-Party System with Rational Voters," *American Economic Review*, 78, 796-805.

[2] Calcagno, R., Y. Kamada, S. Lovo and T. Sugaya (2014): "Asynchronicity and Coordination in Common and Opposing Interest Games," *Theoretical Economics*, 9(2): 409-434.

[3] Caruana, G. and L. Einav (2008a): "A Theory of Endogenous Commitment," *Review of Economic Studies*, 75: 99-116.

[4] Caruana, G. and L. Einav (2008b): "Production Targets," *RAND Journal of Economics*, 39(4): 990-1217.

[5] Fehr, E. and S. Gachter (2002): "Altruistic Punishment in Humans," *Nature*, 415: 137-140.

[6] Gensbittel, F., S. Lovo, J. Renault and T. Tomala (2016): "Zero-sum Revision Games," mimeo.

[28] The commitment due to random revision opportunities in our model implies equilibrium strategies quite different from those of Caruana and Einav (2008a). For example, close to the deadline, a player would like to switch to a static best response in our model, while in Caruana and Einav's (2008a) model she would not do so due to a high switching cost.

[7] Hopenhyan, H. and M. Saeedi (2016): "Bidding Dynamics in Auctions," mimeo.

[8] Iijima, R. and A. Kasahara (2016): "Equilibrium Uniqueness with Gradual and Noisy Adjustment," mimeo.

[9] Ishii, Y. and Y. Kamada (2011): "The Effect of Correlated Inertia on Coordination," mimeo.

[10] Kamada Y. and M. Kandori (2012): "Asynchronous Revision Games," mimeo.

[11] Kamada Y. and M. Kandori (2017): "Revision Games I: Theory," mimeo.

[12] Kamada, Y. and F. Kojima (2014): "Voter Preferences, Polarization, and Electoral Policies," *American Economic Journal: Microeconomics*, 6(4): 203-236.

[13] Kamada Y. and N. Muto (2015): "Multi-Agent Search with Deadline," mimeo.

[14] Kamada, Y. and T. Sugaya (2014): "Valence Candidates and Ambiguous Platforms in Policy Announcement Games," mimeo.

[15] Lovo, S. and T. Tomala (2015): "Markov Perfect Equilibria in Stochastic Revision Games," mimeo.

[16] Moroni, S. (2015): "Existence of Trembling Hand Equilibrium in Revision Games with Imperfect Information," mimeo.

[17] Osborne, Martin J. (1995): "Spatial Models of Political Competition Under Plurality Rule: A Survey of Some Explanations of The Number of Candidates and The Positions They Take," *Canadian Journal of Economics*, 2, 261-301.

[18] Roemer, J. E. (1994): "A theory of policy differentiation in single issue electoral politics," *Social Choice and Welfare*, 11, 355-380.

[19] Romm, A. (2014): "Building Reputation at the Edge of the Cliff," mimeo.

[20] Roy, N. (2014): "Stochastic Revision promotes Cooperation: Experimental Evidence," mimeo.

A Appendix

A.1 Calculation of Expected Payoffs for the Good Exchange Game

For $b(a) = a$ and $c(a) = c \cdot a^2$ with $c > 0$, the expected payoff under the optimal trigger strategy equilibrium can be calculated as follows:

$$\int_0^{t(a^*)} \left(\bar{x}(t) - c(\bar{x}(t))^2\right) \lambda e^{-\lambda t} dt + e^{-\lambda t(a^*)} \left(a^* - c(a^*)^2\right)$$
$$= \int_0^{\frac{1}{\lambda}} \left(\frac{\lambda}{2c} t - c\left(\frac{\lambda}{2c} t\right)^2\right) \lambda e^{-\lambda t} dt + e^{-\lambda \frac{1}{\lambda}} \left(\frac{1}{2c} - c\left(\frac{1}{2c}\right)^2\right) = \frac{1}{2ce}.$$

On the other hand, the fully collusive payoff is $\frac{1}{2c} - c\left(\frac{1}{2c}\right)^2 = \frac{1}{4c}$. Thus, the ratio between these two values is $\frac{\frac{1}{2ce}}{\frac{1}{4c}} = \frac{2}{e} \cong 0.74$.

A.2 Calculation of Expected Payoffs for the Bertrand Competition

<u>The case of high product differentiation:</u>

Suppose $c \in (\frac{2}{7}v, \frac{2}{3}v)$. The expected payoff under the optimal trigger strategy equilibrium can be calculated as follows:

$$\int_0^{t(p^*)} \lambda e^{-\lambda t} \frac{\bar{p}(t)}{2} dt + e^{-\lambda t(p^*)} \frac{p^*}{2} = \int_0^{\frac{2}{\lambda} \ln\left(\frac{v}{4c} + \frac{5}{8}\right)} \lambda e^{-\lambda t} \frac{c\left(4e^{\lambda \frac{t}{2}} - 3\right)}{2} dt + e^{-\lambda \frac{2}{\lambda} \ln\left(\frac{v}{4c} + \frac{5}{8}\right)} \frac{\left(v - \frac{c}{2}\right)}{2}$$
$$= v\left(\frac{5}{2}h - \frac{1}{4}\left(\frac{(8h)^2}{2+5h}\right)\right).$$

On the other hand, the fully collusive payoff is $\frac{v - \frac{c}{2}}{2} = v\frac{1 - \frac{h}{2}}{2}$. The ratio between these two values is:

$$\bar{C}(h) = \frac{v\left(\frac{5}{2}h - \frac{1}{4}\left(\frac{(8h)^2}{2+5h}\right)\right)}{v\frac{1-\frac{h}{2}}{2}} = \frac{10h(2+5h) - 64h^2}{(2+5h)(2-h)} = \frac{2h(10-7h)}{(2+5h)(2-h)}.$$

Differentiating this with respect to h, we obtain:

$$\bar{C}'(h) = \frac{(h+10)(2-3h)}{(2+5h)^2(2-h)^2}.$$

Note that this is strictly positive whenever $h < \frac{2}{3}$. Thus, $\bar{C}(h)$ is strictly increasing in h.

Next, the ratio of the payoff increments is:

$$\tilde{C}(h) = \frac{v\left(\frac{5}{2}h - \frac{1}{4}\left(\frac{(8h)^2}{2+5h}\right)\right) - \frac{c}{2}}{v\frac{1-\frac{h}{2}}{2} - \frac{c}{2}} = \frac{8h}{5h+2}.$$

This is strictly increasing in h.

The case of low product differentiation:

Suppose $c \in (0, \frac{2}{7}v]$. The expected payoff under the optimal trigger strategy equilibrium can be calculated as follows:

$$\begin{aligned}
&\int_0^{t^1} \lambda e^{-\lambda t}\frac{\bar{p}(t)}{2}dt + \int_{t^1}^{t^2} \lambda e^{-\lambda t}\frac{\bar{p}(t)}{2}dt + e^{-\lambda t^2}\frac{p^*}{2} \\
&= \int_0^{\frac{2}{\lambda}\ln\left(\frac{3}{2}\right)} \lambda e^{-\lambda t}\frac{c\left(4e^{\lambda\frac{t}{2}} - 3\right)}{2}dt + \int_{\frac{2}{\lambda}\ln\left(\frac{3}{2}\right)}^{\frac{3}{2\lambda}\ln\left(\frac{3}{2}\right)+\frac{1}{2\lambda}\ln\left(\frac{v}{c}-2\right)} \lambda e^{-\lambda t}\frac{c\left(\frac{8}{27}e^{2\lambda t} + \frac{3}{2}\right)}{2}dt \\
&\quad + e^{-\lambda\left(\frac{3}{2\lambda}\ln\left(\frac{3}{2}\right)+\frac{1}{2\lambda}\ln\left(\frac{v}{c}-2\right)\right)}\frac{v-\frac{c}{2}}{2} = v\left(\frac{h}{2} + \frac{4h}{9}\left(\frac{3}{2h} - 3\right)^{\frac{1}{2}}\right).
\end{aligned}$$

The fully collusive payoff can be calculated as before, and thus the ratio of the expected payoffs is:

$$\bar{C}(h) = \frac{v\left(\frac{h}{2} + \frac{4h}{9}\left(\left(\frac{3}{2h} - 3\right)^{\frac{1}{2}}\right)\right)}{\frac{v\left(1-\frac{h}{2}\right)}{2}} = \frac{2h\left(1 + \frac{8h}{9}\sqrt{\frac{3}{2h}(1-2h)}\right)}{2-h}.$$

Differentiating, we get

$$\bar{C}'(h) = \frac{4\sqrt{3}\left(9\sqrt{\frac{2(1-2h)}{3h}} + 8h^2 - 34h + 12\right)}{9(2-h)^2\sqrt{\frac{2(1-2h)}{h}}}.$$

Since $-34h + 12 > 0$ whenever $h \in (0, \frac{2}{7})$, this is strictly positive for all $h \in (0, \frac{2}{7})$. Hence, $\bar{C}(h)$ is strictly increasing in h.

Next, the ratio of the payoff increments is:

$$\tilde{C}(h) = \frac{v\left(\frac{h}{2} + \frac{4h}{9}\left(\left(\frac{3}{2h} - 3\right)^{\frac{1}{2}}\right)\right) - \frac{c}{2}}{\frac{v\left(1-\frac{h}{2}\right)}{2} - \frac{c}{2}} = \frac{16\sqrt{3h(1-2h)}}{9\sqrt{2}(2-3h)}.$$

Differentiating, we get:

$$\tilde{C}'(h) = \frac{8\sqrt{3}(2-5h)}{9(2-3h)^2\sqrt{2h(1-2h)}}.$$

This is strictly positive whenever $h \in (0, \frac{2}{7})$. Hence, $\tilde{C}(h)$ is strictly increasing in h.

A.3 Calculation of Expected Payoffs for the Election Campaign Game

For $v \in (\frac{1}{2}, 1]$, The expected payoff under the optimal trigger strategy equilibrium can be calculated as follows:

$$\int_0^{t(y_1^*)} \frac{v + \frac{1}{2} - \bar{y}_1(t)}{2} \lambda e^{-\lambda t} + e^{-\lambda t(y_1^*)} \frac{v + \frac{1}{2}}{2}$$

$$= \int_0^{\frac{2}{\lambda}(\ln(7+2v) - 3\ln 2)} \frac{v + \frac{1}{2} - \frac{7+2v-8\cdot e^{\frac{\lambda}{2}t}}{2}}{2} \lambda e^{-\lambda t} dt + e^{-\lambda \frac{2}{\lambda}(\ln(7+2v)-3\ln 2)} \frac{v + \frac{1}{2}}{2} = \frac{5}{2} - \frac{16}{7+2v}.$$

On the other hand, the fully collusive payoff is $\frac{v+\frac{1}{2}}{2}$. The ratio between these two values is:

$$\bar{C}(v) = \frac{\frac{5}{2} - \frac{16}{7+2v}}{\frac{v+\frac{1}{2}}{2}} = \frac{5(7+2v) - 32}{(7+2v)(v+\frac{1}{2})} = \frac{2(3+10v)}{(7+2v)(2v+1)}.$$

Differentiating this with respect to v, we obtain:

$$\bar{C}'(v) = 4\frac{(1-2v)(10v+11)}{(7+2v)^2(2v+1)^2}.$$

This is strictly negative whenever $v > \frac{1}{2}$. Hence, $\bar{C}(v)$ is strictly decreasing in v.

Next, the ratio of the payoff increments is:

$$\tilde{C}(v) = \frac{\frac{5}{2} - \frac{16}{7+2v} - \frac{1}{2}}{\frac{v+\frac{1}{2}}{2} - \frac{1}{2}} = \frac{8}{7+2v}.$$

This is strictly decreasing in v.

A.4 Proof for Corollary 4

Proof. Part 2 follows directly from the formula of the optimal plan provided in Proposition 4, so we provide the proof for Part 1. Part 1a follows because $t^v = \frac{2}{\lambda}(\ln(7+2v) - 3\ln 2)$ by Proposition 4 and this is strictly increasing in v for $v \in (\frac{1}{2}, 1]$.

To prove Part 1b, fix v and $\tilde{v}$ with $\frac{1}{2} < v < \tilde{v} \leq 1$. By a change of variables (or simply by the definition of $\Delta_t(v)$), we have

$$\Delta_t(v) = \int_0^T (x^v(t) - x^v(s))\lambda e^{-\lambda s} ds.$$

Suppose first that $t < t^v$. Then, by the formula of the optimal plan provided in Proposition 4, we obtain

$$\Delta_t(v) = \int_{s=t}^{s=t^v} 4(e^{\frac{\lambda}{2}s} - e^{\frac{\lambda}{2}t})\lambda e^{-\lambda s} ds + \int_{s=t^v}^{s=t^{\tilde{v}}} (x^v(t) - 0)\lambda e^{-\lambda s} ds + e^{-\lambda(t^{\tilde{v}} - t)} x^v(t),$$

and

$$\Delta_t(\tilde{v}) = \int_{s=t}^{s=t^v} 4(e^{\frac{\lambda}{2}s} - e^{\frac{\lambda}{2}t})\lambda e^{-\lambda s} ds + \int_{s=t^v}^{s=t^{\tilde{v}}} (x^{\tilde{v}}(t) - x^{\tilde{v}}(s))\lambda e^{-\lambda s} ds + e^{-\lambda(t^{\tilde{v}} - t)} x^{\tilde{v}}(t).$$

By the definition of t^v, we have that $x^v(t) - 0 \leq x^v(t) - \frac{7+2v-8\cdot e^{\frac{\lambda}{2}s}}{2}$ for $s \geq t^v$. Since $x^v(t) - \frac{7+2v-8\cdot e^{\frac{\lambda}{2}s}}{2} = 4(e^{\frac{\lambda}{2}s} - e^{\frac{\lambda}{2}t}) = x^{\tilde{v}}(t) - x^{\tilde{v}}(s)$ for $s \in [t^v, t^{\tilde{v}}]$, it must be the case that $x^v(t) - 0 \leq x^{\tilde{v}}(t) - x^{\tilde{v}}(s)$ for $s \in [t^v, t^{\tilde{v}}]$. Also, $x^v(t) < x^{\tilde{v}}(t)$ since $t < t^{\tilde{v}}$. Since $t < t^{\tilde{v}}$ (so $e^{-\lambda(t^{\tilde{v}} - t)} > 0$), these two facts imply that $\Delta_t(v) < \Delta_t(\tilde{v})$ holds when $t < t^v$.

Next, suppose that $t \geq t^v$. Then,

$$\Delta_t(v) = \int_{s=t}^{s=t^{\tilde{v}}} (x^v(t) - 0)\lambda e^{-\lambda s} ds + e^{-\lambda(t^{\tilde{v}} - t)} x^v(t),$$

and

$$\Delta_t(\tilde{v}) = \int_{s=t}^{s=t^{\tilde{v}}} (x^{\tilde{v}}(t) - x^{\tilde{v}}(s))\lambda e^{-\lambda s} ds + e^{-\lambda(t^{\tilde{v}}-t)} x^{\tilde{v}}(t).$$

By the definition of t^v, we have that $x^v(t) - 0 = 0$ for $s \geq t^v$. Since $t < t^{\tilde{v}}$, $x^{\tilde{v}}(t) - x^{\tilde{v}}(s) \geq 0$ for $s \in [t, t^{\tilde{v}}]$. Hence, it must be the case that $x^v(t) - 0 \leq x^{\tilde{v}}(t) - x^{\tilde{v}}(s)$ for $s \in [t, t^{\tilde{v}}]$. Also, $x^v(t) < x^{\tilde{v}}(t)$ since $t < t^{\tilde{v}}$. Since $t < t^{\tilde{v}}$ (so $e^{-\lambda(t^{\tilde{v}}-t)} > 0$), these two facts imply that $\Delta_t(v) < \Delta_t(\tilde{v})$ holds when $t \geq t^v$, completing the proof. ■

A.5 No Cooperation under Finite Component Games

In this section we consider finite component games with a unique Nash equilibrium that is strict, and show that no cooperation is possible. In order to formalize what we mean by this under a general strategy space, we first introduce notations and terminology.

Finite component game: Let A_i be the finite set of actions for player $i = 1, 2$. $A = A_1 \times A_2$. Player i's payoff function is $\pi_i : A \to \mathbb{R}$. We extend the domain of π_i in the usual manner, by writing $\pi_i(a_i, \alpha_{-i})$, $\pi_i(\alpha)$, and so forth. A mixed action profile $\alpha = (\alpha_1, \alpha_2) \in \Delta(A_1) \times \Delta(A_2)$ is a Nash equilibrium if $\pi_i(\alpha) \geq \pi_i(\alpha'_i, \alpha_{-i})$ for all $\alpha'_i \in \Delta(A_1)$. We say that action a_i is an ϵ-best response to $\alpha_{-i} \in \Delta(A_{-i})$ if $\left(\max_{a'_i} \pi_i(a'_i, \alpha_{-i})\right) - \pi_i(a_i, \alpha_{-i}) \leq \epsilon$.

Histories, strategies, expected payoffs, and equilibrium: At time $-t$, a generic history h_k of the revision game at the k'th opportunity is written as:

$$h_k = \left(t, (t^l, a^l)_{l \in \mathbb{N}, l<k}\right) \in \mathbb{R}_+ \times (\mathbb{R}_+ \times A)^{k-1} := H_k.$$

The interpretation is that $-t$ (the first element of h_k) is the time at which the current opportunity (i.e., the k'th opportunity) arrives, $-t^l$ is the time at which l'th opportunity has arrived, and a^l is the action profile that is taken at that opportunity. Note that we ignore the events under which infinitely many opportunities have arrived in the past, because such events have probability zero under Poisson processes.

The set H_k is endowed with a Borel sigma-algebra (H_k can be seen as a subset of $\mathbb{R}_+ \times (\mathbb{R}_+ \times \mathbb{Z})^{k-1}$ because A is finite). The set of all strategies is $H = \bigcup_{k \in \mathbb{N}} H_k$, and we assume that this is endowed with a sigma-algebra induced by the sigma-algebras on H_k (note that H is a countable union of H_k's). Player i's behavioral strategies is function $\sigma_i : H \to \Delta(A_i)$ that is Borel measurable with respect to this sigma-algebra.[29] The space of i's behavioral strategies is Σ_i. The measurability ensures that there is a well-defined payoff function $u_i : \Sigma_i \times \Sigma_{-i} \times H \to \mathbb{R}$, where $u_i(\sigma|h)$ is interpreted to be player i's payoff induced by the strategy profile σ conditional on the history h.

We define a subgame-perfect equilibrium to be a strategy profile σ^* such that for all $h \in H$, $u_i(\sigma^*|h) \geq u_i(\sigma'_i, \sigma^*_{-i}|h)$ for all $\sigma'_i \in \Sigma_i$.

Now we are ready to state the main result of this section. Recall that A is finite.

Proposition 6 *Fix a component game $G = (A, \pi)$. Suppose that G has a unique Nash equilibrium $a^* \in A$ and it is strict. Then, the revision game of G has a unique subgame-perfect equilibrium σ^*, and it satisfies $\sigma^*_i(h)(a^*_i) = 1$ for every $h \in H$ and $i = 1, 2$.*

The idea of the proof is that, if the deadline is close enough, any mixed action profile played under any on- and off-path histories of any subgame perfect equilibria must be close enough to the unique Nash equilibrium. But then the only best response to such a distribution is the unique pure Nash action by the strictness of the equilibrium. We use this to show that for a small time interval close to the deadline, only the Nash action can be played. We can use this logic to implement backward induction. The formalization of backward induction follows the "continuous-time backward induction" of Calcagno et al. (2015).

Proof. Fix a component game $G = (A, \pi)$ such that (i) A is finite, (ii) G has a unique Nash equilibrium $a^* \in A$, and (iii) a^* is a strict Nash equilibrium of G. Fix $T < \infty$. Fix a subgame-perfect strategy profile $\sigma \in \Sigma$. We will show that $\sigma_i(h)(a^*_i) = 1$ for any $h \in H$ and $i = 1, 2$.

To see this, suppose to the contrary that there is a nonempty set $H^N \subseteq H$ such that $h \in H^N$ implies there exists player i such that $\sigma_i(h)(a^*_i) < 1$. Let $H_t = \bigcup_{k \in \mathbb{N}} \left(\{t\} \times (\mathbb{R}_+ \times A)^k\right) \subseteq H$ be the set of histories at time $-t$, and let

$$S = \sup_{H^N \cap \left(\bigcup_{s \in [0,t)} H_s\right) = \emptyset} t$$

[29]This formulation follows that of Kamada and Muto (2015).

be the supremum of t such that at every time after $-t$, no player i plays any action $a_i \neq a_i^*$ at any history.

The rest of the proof is divided into four steps. Step 1 defines the maximal amount of feasible reward $\bar{\pi}_i(t)$. Since A is finite, this is finite, and converges to zero as $t \downarrow S$. Step 2 shows that any action that can be taken at t must be a static $\bar{\pi}_i(t)$-best response. Step 3 studies a property of the component game, and shows that if all pure actions played under α are ϵ-best responses for small $\epsilon > 0$ then α must be the unique strict Nash equilibrium. These steps imply that the only action profile that can be taken close to time S is the unique strict Nash equilibrium of the component game. Formally, we will show that there exists $\delta > 0$ such that $\sigma_i(h)(a_i^*) = 1$ for any $h_t \in H_t$ with $t \in [S, \min\{S+\delta, T\}]$ and $i = 1, 2$. Step 4 says that such a conclusion is a contradiction to either (i) the assumption that H^N is nonempty or (ii) the definition of S.[30]

Step 1: Defining the maximal reward $\bar{\pi}_i(t)$.

Consider an opportunity at time $-t$. Let

$$\bar{\pi}_i(t) := \frac{1 - e^{-\lambda(t-S)}}{e^{-\lambda(t-S)}} \left[\left(\max_{a \in A} \pi_i(a) \right) - \left(\min_{a' \in A} \pi_i(a') \right) \right]$$

for $t \in [S, T]$ and $i = 1, 2$. Since A is finite, this is well-defined. Notice that $\bar{\pi}_i(t)$ measures the maximal amount of reward that i can receive after time $-t$.

Step 2: All actions in the support must be $\bar{\pi}_i(t)$-best response at time $-t$.

For any fixed $\epsilon > 0$, let $BR_i^{\epsilon}(\alpha_{-i}) \in A_i$ be the set of i's *pure* ϵ-best responses to $-i$'s mixed action $\alpha_{-i} \in \Delta(A_{-i})$. We show that, for every $a_i \in \mathrm{supp}(\sigma_i(h_t))$,

$$a_i \in BR_i^{\bar{\pi}_i(t)}(\sigma_{-i}(h_t)) \quad \text{for all} \quad h_t \in H_t.$$

To see this, suppose the contrary. Then, there exists $a_i \in \mathrm{supp}(\sigma_i(h_t))$ such that

$$\left(\max_{a_i' \in A_i} \pi_i(a_i', \sigma_{-i}(h_t)) \right) - \pi_i(a_i, \sigma_{-i}(h_t)) > \bar{\pi}_i(t).$$

[30] This is due to the following observations: First, if $S = T$, then the assumption that H^N is nonempty implies that $\sigma_i(h_T)(a_i^*) < 1$. Hence, the conclusion that $\sigma_i(h_t)(a_i^*) = 1$ for $t \in [S, \min\{S+\delta, T\}] = \{T\}$ is a contradiction. Second, if $S < T$, then it is obvious that contradiction is obtained.

By the definition of $\bar{\pi}_i(t)$, this implies that

$$e^{-\lambda(t-S)}\left[\left(\max_{a_i'\in A_i}\pi_i(a_i',\sigma_{-i}(h_t))\right)-\pi_i(a_i,\sigma_{-i}(h_t))\right]>(1-e^{-\lambda(t-S)})\left[\left(\max_{a\in A}\pi_i(a)\right)-\left(\min_{a'\in A}\pi_i(a')\right)\right],$$

or

$$(1-e^{-\lambda S})\pi(a^*)+e^{-\lambda S}\left[e^{-\lambda(t-S)}\left(\max_{a_i'\in A_i}\pi_i(a_i',\sigma_{-i}(h_t))\right)+(1-e^{-\lambda(t-S)})\left(\min_{a'\in A}\pi_i(a')\right)\right]$$

$$>(1-e^{-\lambda S})\pi(a^*)+e^{-\lambda S}\left[e^{-\lambda(t-S)}\pi_i(a_i,\sigma_{-i}(h_t))+(1-e^{-\lambda(t-S)})\left(\max_{a\in A}\pi_i(a)\right)\right].$$

The left hand side of this inequality is the maximum expected payoff from deviation from σ_i at history h_t, assuming the severest feasible punishment during time interval $(-t,,S]$. The right hand side is the maximum expected payoff from taking action a_i, assuming the best feasible reward during time interval $(-t,,S]$. Notice that, by assumption, the action profile played during time interval $(-S,0]$ is a^* with probability 1, hence the first term of each side. These observations imply that the inequality implies that assigning positive probability to action a_i cannot be a best response at h_t. Hence, i has a profitable deviation from σ_i to σ_i', where $\sigma_i'(h_{t'}') = \sigma_i(h_{t'}')$ for all $h_{t'}' \in H \setminus \{h_t\}$ and $\sigma_i'(h_t)(a_i'') = 1$ where $a_i'' \in \arg\max_{a_i'\in A_i}\pi_i(a_i',\sigma_{-i}(h_t))$. This contradicts the optimality of σ_i at h_t.

Step 3: If all actions in the support are ϵ-best response then it is α^*.

Define $\alpha_i^* \in \Delta(A_i)$ for each $i=1,2$ by $\alpha_i^*(a_i^*)=1$ for each $i=1,2$. Now we show that there exists $\bar{\epsilon}>0$ such that for all $\epsilon<\bar{\epsilon}$, if $a_i \in BR_i^{\epsilon}(\alpha_{-i})$ for every $a_i \in \mathrm{supp}(\alpha_i)$ for each $i=1,2$, then $\alpha=\alpha^*$.

To see this, suppose the contrary. Then, since A_i is finite for each $i=1,2$, there exists player i, action $\tilde{a}_i \neq a_i^*$ and a sequence $(\epsilon^k,\alpha^k)_{k=1}^{\infty}$ such that (i) $\epsilon_k \to 0$ as $k\to\infty$, (ii) $\tilde{a}_i \in BR_i^{\epsilon^k}(\alpha_{-i}^k)$ for all k, and (iii) $a_j' \in BR_j^{\epsilon^k}(\alpha_{-j}^k)$ for all $a_j' \in \mathrm{supp}(\alpha_j^k)$ for all j and for all k. There are two cases to consider.

1. Suppose first that α^k converges to α^*. Then, the following two claims must be true:

   (a) By the definition of convergence, for any $\delta>0$, there exists $\bar{k}$ such that for all $k>\bar{k}$, $|\alpha^k-\alpha^*|<\delta$.

(b) Since a^* is a strict Nash equilibrium, there exist $\epsilon > 0$ and $\delta > 0$ such that $|\alpha - \alpha^*| < \delta$ implies $a_i \notin BR_i^\epsilon(\alpha_{-i})$ if $a_i \neq a_i^*$.

Combining 1 and 2 above, we have that there exist $\epsilon > 0$ and $\bar{k}$ such that $k > \bar{k}$ implies $\tilde{a}_i \notin BR_i^\epsilon(\alpha_{-i}^k)$. This contradicts (ii).

2. Suppose next that α^k does not converge to α^*. Then, since $\Delta(A)$ is compact, there exists α'^* and a subsequence $(\hat{\alpha}^l)_{l=1}^\infty$ of the sequence $(\alpha^k)_{k=1}^\infty$ (i.e., there exists a strictly increasing function $g : \mathbb{N} \to \mathbb{N}$ such that $\hat{\alpha}^l = \alpha^{g(l)}$) that converges to α'. Then, the following two claims must be true:

(a) By the definition of convergence, for any $\delta > 0$, there exists $\bar{l}$ such that for all $l > \bar{l}$, $|\hat{\alpha}^l - \alpha'| < \delta$.

(b) Since a^* is a unique Nash equilibrium, there exist $\epsilon > 0$ and $\delta > 0$ such that $|\hat{\alpha}^l - \alpha'| < \delta$ implies there exists j and $a_j' \in \text{supp}(\hat{\alpha}_j^l)$ such that $a_j' \notin BR_j^\epsilon(\hat{\alpha}_{-j}^l)$.

Combining 1 and 2 above, we have that there exist $\epsilon > 0$ and $\bar{l}$ such that $l > \bar{l}$ implies there exists j and $a_j' \in \text{supp}(\hat{\alpha}_j^l)$ such that $a_j' \notin BR_j^\epsilon(\hat{\alpha}_{-j}^l)$. This contradicts (iii).

Since these two cases are exhaustive, we have now shown that there exists $\bar{\epsilon} > 0$ such that for all $\epsilon < \bar{\epsilon}$, if $a_i \in BR_i^\epsilon(\alpha_{-i})$ for every $a_i \in \text{supp}(\alpha_i)$ for each $i = 1, 2$, then $\alpha = \alpha^*$.

Step 4: Backward induction implies uniqueness.

First, notice that $\bar{\pi}_i(t)$ converges continuously to 0 as $t \downarrow S$ for each $i = 1, 2$. By Step 2, this implies that for any $\epsilon > 0$, there exists $\delta > 0$ such that for all $t \in [S, S+\delta]$, for every $a_i \in \text{supp}(\sigma_i(h_t))$,

$$a_i \in BR_i^\epsilon(\sigma_{-i}(h_t))$$

holds for all $h_t \in H_t$. Then, by choosing $\epsilon > 0$ strictly less than the $\bar{\epsilon}$ identified in Step 3, Step 3 shows that $\sigma_i(h)(a_i^*) = 1$ for any $h_t \in H_t$ with $t \in [S, S + \delta]$ and $i = 1, 2$. This is the desired claim, completing the proof. ■