

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Chen, Chia-Hui; Ishida, Junichiro

Working Paper A war of attrition with experimenting players

ISER Discussion Paper, No. 1014

Provided in Cooperation with: The Institute of Social and Economic Research (ISER), Osaka University

Suggested Citation: Chen, Chia-Hui; Ishida, Junichiro (2017) : A war of attrition with experimenting players, ISER Discussion Paper, No. 1014, Osaka University, Institute of Social and Economic Research (ISER), Osaka

This Version is available at: https://hdl.handle.net/10419/197688

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Discussion Paper No. 1014

A WAR OF ATTRITION WITH EXPERIMENTING PLAYERS

Chia-Hui Chen Junichiro Ishida

October 2017

The Institute of Social and Economic Research Osaka University 6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

A War of Attrition with Experimenting Players^{*}

Chia-Hui Chen[†] and Junichiro Ishida[‡]

October 20, 2017

Abstract

A standard incomplete-information war of attrition is extended to incorporate experimentation and private learning. We obtain a characterization of all equilibria in this extended setup and use this setup to illuminate a tradeoff between short-run and long-run gains of experimentation. The extension yields qualitative impacts on the strategic nature of the problem. The option value of experimentation serves as a credible commitment device to stay in the game, which is instrumental in inducing the other player to concede earlier. As a direct consequence, there may be an equilibrium in which the strictly less efficient player can get the better end of the deal, implying that slow learning can be a blessing in this type of competition. Our analysis gives insight into why an apparently inferior technology often survives in many standards competitions and more broadly offers implications for technology adoption and industry dynamics. We also show that there is a non-degenerate set of parameters that can support the Pareto-efficient allocation as an equilibrium outcome whereas it is never possible in the standard setup.

JEL Classification Number: D82, D83

Keywords: war of attrition, bandit problem, experimentation, standards competitions, private learning.

^{*}The first author acknowledges financial support from JSPS KAKENHI Grant-in-Aid (B) JP16H03615. The second author acknowledges financial support from JSPS KAKENHI Grant-in-Aid (S) JP15H05728, (A) JP20245031, (B) JP16H03615 and (C) JP24530196 as well as the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

[†]Institute of Economic Research, Kyoto University. Email: chchen@kier.kyoto-u.ac.jp

[‡]Institute of Social and Economic Research, Osaka University. Email: jishida@iser.osaka-u.ac.jp

1 Introduction

In many standards competitions, we observe instances where an apparently less efficient or less matured technology survives in the end as the defacto standard, as exemplified most notably by the video format war between Betamax and VHS.¹ Although standards competitions often take the form of a war of attrition, this observation seems to be at odds with theoretical predictions from some standard versions of it, suggesting that there may be a missing link that has been overlooked in the literature. In this paper, we attempt to fill this gap by extending a standard war of attrition to ask whether there are factors, other than sheer technological superiority, which affect the outcome of this type of economic competition in any significant way.

One observation stands out in search of this missing link: standards competitions typically take place in innovation-intensive industries, such as telecommunications, home VCR, audio, operating systems, high-speed rail and automated driving. In most of these cases, incoming firms, or standards, are initially still technologically premature but gradually improve their productive efficiency or profitability over time via experimentation. The fact that competing standards may innovate alters the strategic nature of competition at least in two ways. First, it is not only the current efficiency level but also the margin of experimentation, i.e., how much room for improvement a given technology is expected to possess, which can play a decisive role: a player with large room for improvement clearly has more incentive to experiment and hence less incentive to drop out. Second, the possibility of technological progress also implies that it is *ex ante* not clear how many standards can survive in a given market. Although the market for standards is typically not large enough to accommodate many inefficient ones, there are still cases where multiple standards attain enough efficiency to coexist for a substantial duration of time, as in the competition between Windows and Mac in operating systems or among Nintendo, SONY and Microsoft in video game consoles. When multiple standards emerge, the strategic nature of the problem tips over, for it is no longer a war of attrition with no point in waiting for the rival to drop out.

To describe this situation, we extend a two-player incomplete-information war of attrition to incorporate experimentation and private learning. The setup is quite standard, except that each player's type may change over time as a result of (learning-by-doing) experimentation. Each player's type, which is either weak or strong, is his own private information. If a player

¹There is naturally an extensive literature discussing whether and how a society gets "locked in" with an inefficient standard, mainly focusing on various issues arising from network effects (Katz and Shapiro [13], [14], Farrell and Saloner [9], [10], Arthur [2], Choi [7]). Also, see Katz and Shapiro [15] and Stango [24] for surveys.

is weak, however, he may achieve a technological breakthrough and become strong with some probability as long as he stays in the game. The probability of achieving a breakthrough depends on the unobservable state of nature which is not directly observable and must be uncovered via experimentation. Finally, we assume that a strong player can earn positive payoff regardless of the other player's presence, so that it is a dominant strategy for him to stay in the game indefinitely. It is this last assumption which gives rise to the *ex ante* uncertainty about how many players can survive in the end.

We obtain a characterization of all equilibria in this extended setup. The extension to incorporate experimentation is technically rather straightforward but yields some qualitative impacts on the strategic nature of the problem. In our extended setup, a weak player may stay in the game just because he may achieve a breakthrough tomorrow. The option value of experimentation thus serves as a credible commitment device to stay in the game, which is instrumental in inducing the other player to concede earlier. As a direct consequence of this force, we find an equilibrium in which the strictly less efficient player, but with more room for experimentation, can get the better end of the deal, i.e., slow learning can be a blessing. This stands in sharp contrast to the standard setup with no experimentation (hereafter, simply the standard war of attrition) where the equilibrium payoff is monotonic with respect to a player's "efficiency" level. Our analysis sheds new light on the tradeoff between short-run and long-run gains of experimentation from a previously unexplored channel and suggests some practical implications for technology choice and industry dynamics. On a more specific level, it gives insight into why an inferior technology often survives in standards competitions.

From a more technical point of view, it is also worth noting that there may exist an efficient equilibrium (in the sense of Pareto) in our extended setup, again in clear contrast to the standard war of attrition where there generically exists no efficient equilibrium. This is because the efficient allocation requires that any weak player exit immediately unless both of them are weak. A way to (approximately) achieve this allocation is that: one player, say player 1, exits with probability one at time 0 if he is weak; player 2 then exits with probability one an instant later if he is weak and player 1 did not exit. There is no such equilibrium, however, because if player 1 knows that player 2 will exit with strictly positive probability an instant later, he is always better off with waiting than exiting immediately. More generally, in any equilibrium of the standard setup, at most one player can exit with strictly positive probability, which immediately excludes the possibility of realizing the efficient allocation.

The situation changes rather drastically once the element of experimentation is introduced into the setup. In the model with experimentation, the efficient allocation typically involves delay, as it leads to better informed decisions. A player may exit with strictly positive probability when the option value of experimentation equals the flow cost, the timing of which may arrive differently between the players. This property is crucial for the construction of equilibrium in general and can be exploited to construct an equilibrium in which both of the players exit with strictly positive probability at some (different) points of the game. More importantly, due to this property, there exists a non-degenerate set of parameters that can support the efficient allocation as an equilibrium outcome.

Related literature: A war of attrition, which originates from theoretical biology, is well suited for various economic situations, such as standards competitions, where multiple parties compete for a fixed resource. In many economic applications, however, participating players may innovate and improve their "fitness" in a relatively short span of time. This could be an important departure from biological settings where genetic evolution occurs rather slowly and such a possibility is almost negligible.

An essential feature of our war of attrition is that there is a type of player who would never concede. Without the element of experimentation, our model becomes a variant of Fudenberg and Tirole [11] who analyze a war of attrition when there are types who would never concede. Among other things, they establish the uniqueness of equilibrium when some types may never exit – a property which is also retained in our setup – whereas there is typically a continuum of equilibria when there are no such types. Similarly, Ordover and Rubinstein [21] analyze a concession game in which one player is informed about the disagreement outcome; there, a player who knows that the disagreement outcome is favorable never concedes. In this vein, our model is also related to reputational models of bargaining, such as Abreu and Gul [1] and Kambe [12], which include a stubborn type who "irrationally" sticks to his own demand and never concedes.

Kim and Lee [16] analyze information acquisition in a war of attrition where the war terminates stochastically, in which case each player receives a payoff according to the underlying state of nature. In this setup, they consider a situation where each player can observe the state of nature by engaging in costly information acquisition. Our analysis shares an aspect with this work in that learning takes place in a war of attrition but differs in the way players acquire information: in theirs, information acquisition is done once and for all when a player incurs the cost whereas in ours, it is through time-consuming experimentation with information being revealed only gradually over time. This difference is crucial as the speed of learning, or the margin of experimentation, is one of the main focuses of our analysis.

Chen and Ishida [6] analyze a hierarchial model of experimentation in which a principal retains the authority to terminate the joint project while an agent focuses strictly on running the project assigned to him. A key departure of the analysis is the possibility that the principal may be informed about the project type, so that her termination decision becomes a signal of the project quality. In equilibrium, a principal who knows that the project is good never terminates the project, just like the strong type in the current setup who never exits. The model can thus be seen as a hybrid of experimentation and war of attrition models. The focus of that work is on characterizing the agent's effort dynamics and how they depend on factors such as the principal's ability to evaluate the eventual value of a project.

Several works examine the optimal timing of exit with information externalities among players. Chamley and Gale [5] consider a model of strategic investment in which there are N players, a random number n of whom have an investment option. Assuming that the value of investment depends positively on the random number n, there is an incentive to wait and see others' investment decisions. Decamps and Mariotti [8] also consider a setup in which the value of investment is common across the players while the cost is only privately known. Murto and Välimäki [20] analyze an exit game with private learning where each player receives a signal in each period which partially reveals his own type. In those models, the players' payoffs are positively correlated through the aggregate state of nature: a good (aggregate) state is a good news for all players in the game. In contrast, in our setup, the state of nature is individual-specific and a good state for one player is necessarily a bad news for the other.

Finally, one driving force of our model is the *ex ante* uncertainty about the number of players who can survive in the end. Some previous works explore this issue mainly by focusing on demand-side uncertainty. Most notably, Rob [22] analyzes entry dynamics of firms in a situation where uncertainty exists with respect to the "limit of the market" and it is costly to overshoot the market limit. See Vettas [25], Rob and Vettas [23], and Barbarino and Jovanovic [3] for various extensions of this approach.

2 Model

Environment: Consider a war of attrition between two players, indexed by i = 1, 2. The basic setup is a standard war of attrition under incomplete information, except that a player's type may change over time as a result of experimentation. Time is continuous, extending from zero to infinity, and each player decides whether to stay in the game or exit from it at each instant with a decision to exit irreversible and publicly observable. Below, we first outline the standard part and then describe how we augment it with the possibility of experimentation and private learning.

Standard war of attrition: A player is either weak or strong at each point in time. There

is asymmetry of information between the players where each player's type is his private information and cannot be observed by the other player. The flow payoff to a player is one if he is strong or the other player has exited from the game; otherwise, the flow payoff is $-\rho_i$ where $\rho_i \in (0, \infty)$ captures the flow cost of staying in the game for player *i*. This specification implies that there are two (equally effective) ways for a weak player to earn positive payoff, either to achieve a breakthrough or to wait for the other player to exit. The reservation payoff, which a player earns when he exits, is normalized at zero. Each player discounts payoffs by a common discount rate *r*. For the sake of exposition, we refer to this particular setup as the standard war of attrition in what follows.

Experimentation: We augment the above setup by introducing the hidden state of nature for each player which indirectly influences his type. More precisely, before the game begins, nature chooses the state for each player *i*, denoted by $\theta_i \in \{0, 1\}$. The realized state (θ_1, θ_2) is not observable to anyone and drawn independently with $P(\theta_i = 1) = \alpha_i$. Given the state, nature then chooses the initial type of each player where

$$P(\text{player } i \text{ weak } | \theta_i = 0) = 1, \ P(\text{player } i \text{ weak } | \theta_i = 1) = \eta_i.$$

The unconditional probability that player *i* is strong at time 0 is hence given by $(1 - \eta_i)\alpha_i$. We assume, however, that a player who is initially weak may achieve a breakthrough and become strong over time: if a weak player stays in the game for [t, t + dt), a breakthrough arrives with probability $\lambda_i \theta_i dt$. For clarity, we call $\lambda_i \in (0, \infty)$ the *learning rate* for player *i*. Note that in this specification, a player can be strong only if the underlying state is good $(\theta_i = 1)$.

Interpretation: Throughout the analysis, we in particular focus on two parameters, α_i and η_i , as they admit economically relevant interpretations. First, since a player can be strong only if the state is good, α_i indicates the upper bound of the probability that player *i* is strong, i.e., if player *i* stays in the game indefinitely, the probability that he is good converges to α_i . For this reason, we refer to α_i as the *asymptotic efficiency* of player *i*. Second, η_i measures how much room for experimentation player *i* has, which we refer to as the *margin of experimentation*. As a practical interpretation, we take this as capturing the "technological maturity" of player *i*, where a low η_i means that the player is in a technologically more advanced or matured stage with less room for experimentation. Note that the learning rate plays a similar, though not equivalent, role to the margin of experimentation since $1 - \eta_i$ can alternatively be regarded as the mass probability of achieving a breakthrough at time 0.

3 Analysis

3.1 Preliminaries

Given the payoff structure, it is a dominant strategy for a strong player to stay in the game indefinitely with no further decisions to make. We can thus focus on a weak player, provided that both players are still in the game. Throughout the analysis, we denote a generic player by i and "the other player" by j.

Let $p_i(t)$ denote the belief that the state for player *i* is good ($\theta_i = 1$) at time *t*, conditional on being weak. Since a breakthrough occurs at a constant rate, the belief depends only on time and is given by

$$p_i(t) = \frac{\alpha_i \eta_i e^{-\lambda_i t}}{Q_i(t)},$$

where $Q_i(t) := 1 - \alpha_i + \alpha_i \eta_i e^{-\lambda t}$. Note that the belief declines monotonically over time, i.e., no news is a bad news.

If a player is strong or the only player has exited, he earns a continuation payoff of $\frac{1}{r}$. As such, $r\rho_i$ measures the cost-benefit ratio of staying in the game for player *i*. If the learning rate is low relative to the cost-benefit ratio, a player has no incentive to experiment, and the game is effectively reduced to the standard war of attrition. To focus on relevant and interesting cases, therefore, we assume throughout the analysis that the option value of experimentation is sufficiently large.

Assumption 1 $\lambda_i > r\rho_i$.

3.2 The optimality conditions

The strategy for player i is G_i which specifies the unconditional distribution of exiting time. Define g_i as its corresponding density, which represents the rate at which player i exits. We assume that G_i is right continuous.

At each instant t, each player decides whether to stay or exit. If a player exits, he earns a continuation payoff of zero. In contrast, if a player stays and exits at the next instant, he can achieve a breakthrough with probability $p_{i,t}\lambda dt$ while the rival player exits with probability $\frac{g_j(t)dt}{1-G_j(t)}$; in either case, the player earns $\frac{1}{r}$. Since the flow cost is $\rho_i dt$, a necessary condition for player i to exit at t is

$$r\rho_i \ge p_i(t)\lambda_i + \frac{g_j(t)}{1 - G_j(t)},\tag{1}$$

which shows the key tradeoff between exiting now and waiting for an instant. If there is no possibility of a breakthrough ($\lambda_i = 0$), the model becomes the standard war of attrition where the hazard rate is constant over time.

If $\alpha_i \eta_i \geq \frac{r\rho_i}{\lambda_i}$, there exists τ_i such that

$$p_i(\tau_i) = \frac{\alpha_i \eta_i e^{-\lambda_i \tau_i}}{Q_i(\tau_i)} = \frac{r\rho_i}{\lambda_i},\tag{2}$$

which is the point at which the option value of experimentation equals the flow cost for player i. From (2), we obtain

$$\tau_i = \frac{1}{\lambda_i} \ln \left[\frac{(\lambda_i - r\rho_i)\alpha_i \eta_i}{r\rho_i (1 - \alpha_i)} \right].$$

Let $\tau_i = 0$ if $\frac{r\rho_i}{\lambda_i} > \alpha_j \eta_j$. For the subsequent analysis, we suppose $\tau_2 > \tau_1 \ge 0$ and define $\Delta_{\tau} := \tau_2 - \tau_1$. This is of course not without loss of generality, as it excludes the possibility that $\tau_1 = \tau_2$. We exclude the symmetric case because it is a knife-edge case exhibiting some peculiar properties and also because the analysis of this case is relatively straightforward (see section 4.1 for a brief discussion on the symmetric case).²

Assumption 2 $\tau_2 > \tau_1 \ge 0$.

We first establish some useful properties that must hold in any equilibrium.

Lemma 1 $G_i(t) = 0$ for $t < \tau_i$, i.e., player i never exits before τ_i .

Proof. For $t < \tau_i$, $\lambda p_i(t) > r\rho_i$ regardless of the other player's strategy G_j . As such, it is a dominant strategy to stay in the game.

Lemma 2 For $t > \tau_2$, G_i is given by

$$G_i(t) = 1 - (1 - G_i(\tau_2))\phi_j(t), \tag{3}$$

where $\phi_j(t) := e^{-r\rho_j(t-\tau_2)} \frac{Q_j(\tau_2)}{Q_j(t)}$, until no weak player remains in the game.

Proof. See Appendix.

These results mark two key departures from the standard war of attrition. First, there is a tipping point τ_i which can arrive strictly later than time 0. Second, the optimal strategy

²Although we only discuss the case where $\tau_1 = \tau_2 = 0$, we can essentially apply the same argument to any $\tau_1 = \tau_2 > 0$.

must take into account how the other player acquires information over time, which is captured by $\frac{Q_j(\tau_2)}{Q_j(t)}$. If there is no learning, i.e., $\lambda_j = 0$, then $\frac{Q_j(\tau_2)}{Q_j(t)} = 1$ for all t, and the model is reduced to the standard war of attrition as we have seen. If $\lambda_j > 0$, on the other hand, $\frac{Q_j(\tau_2)}{Q_j(t)}$ is strictly increasing in t, prompting the player to exit at an accelerating rate.

Given that a weak player must exit at some positive rate, there may be a point at which no weak player remains in the game. To compute this termination point, define $W_i(t)$ as the probability that player *i* is not strong (either is weak or has exited) at time *t* which is assumed to be left continuous.³ Note that if player *i* never exits, the probability that he is weak at *t* is $Q_i(t)$ which constitutes the lower bound for $W_i(t)$. Also, $W_i(t) \ge G_i(t)$ by definition. We will later discuss in detail how we construct this function W_i , but for now, we take that such a function exists and is well defined for any given *t*.

Let $T_i := \inf\{t : G_i(t) = W_i(t)\}$ denote the termination point for player *i*. If there exists $T_i < \infty$, $g_i(t) = 0$ and G_i stays constant for all $t > T_i$, i.e., any remaining player must be strong and never exits once the game reaches the termination point.

3.3 In the gap

For $t > \tau_2$, it is quite straightforward to obtain the continuation equilibrium which is characterized by the indifference condition (3). We also know that before τ_i , player *i* never exits, i.e., $G_i(t) = 0$ for $t < \tau_i$. What remains to be seen is the case in-between, i.e., for $t \in [\tau_1, \tau_2]$. The following statement clarifies what could possibly happen in this interval.

Lemma 3 In equilibrium, generically, one of the following must happen:⁴

- 1. Player 1 exits with strictly positive probability at τ_1 ;
- 2. Player 2 exits with strictly positive probability at τ_2 ;
- 3. Both players 1 and 2 exit with strictly positive probability at τ_1 and τ_2 , respectively.

Moreover, no player exits in (τ_1, τ_2) .

Proof. For $t \in [\tau_1, \tau_2)$, player 2 never exits by Lemma 1. Given this, it is never optimal for player 1 to exit at any $t \in (\tau_1, \tau_2)$, because he strictly prefers to exit at τ_1 rather than at any $t \in (\tau_1, \tau_2)$. To prove the lemma, therefore, it suffices to show that player 1 never exits with positive probability at τ_2 . If $G_2(\tau_2) = 0$, player 1 strictly prefers to exit at τ_1 than at τ_2 . If

³Note that G_i is assumed to be right continuous, so that we implicitly consider the timing structure in which $W_i(t)$ is realized first and then G_i is chosen.

⁴There is also a case where $\tau_1 = \tau_2$ and $T_1(0) = T_2(0)$, so that neither player exits with positive probability in equilibrium. Throughout the analysis, we ignore this non-generic case.

 $G_2(\tau_2) > 0$, player 1 strictly prefers to wait until $\tau_2 + dt$ than exiting at τ_2 . This means that G_1 can make a discrete jump only at τ_1 .

This result suggests that any equilibrium of this game is thoroughly characterized by a pair $(G_1(\tau_1), G_2(\tau_2))$ where $G_1(\tau_1) = G_1(\tau_2)$ by Lemma 3. Let $G_i^* := G_i(\tau_i), W_i^* := W_i(\tau_i)$ and $p_i^* := p_i(\tau_i)$ to save notation.

We now derive $W_i(t)$ to see how the termination point T_i is related to G_i^* . Note that $W_i(t) = Q_i(t)$ if player *i* has never exited. On the other hand, once player *i* starts exiting, we in general have $W_i(t) > Q_i(t)$ because $W_i(t)$ includes those who exited but would have become strong later had they stayed in the game. More precisely, for a player who exits at t', the probability that he would achieve a breakthrough by *t* if he stayed is $p_i(t')(1 - e^{-\lambda_i(t-t')})$. Since player *i* may exit with positive probability only at τ_i , we obtain

$$W_{i}(t) = \begin{cases} Q_{i}(t) & \text{for } t \in [0, \tau_{i}], \\ Q_{i}(t) + G_{i}^{*} p_{i}^{*} (1 - e^{-\lambda_{i}(t - \tau_{i})}) + \int_{\tau_{2}}^{t} g_{i}(s) p_{i}(s) (1 - e^{-\lambda_{i}(t - s)}) ds & \text{for } t \in (\tau_{i}, T_{i}], \end{cases}$$

where $g_i(t) = (r\rho_j - p_j(t)\lambda_j)(1 - G_i^*)\phi_j(t)$. For $t > T_i$, $W_i(t)$ stays constant at $G_i(T_i)$.

By definition, $G_i(T_i) = W_i(T_i)$. Letting $\Phi_i(t) := (r\rho_j - p_j(t)\lambda_j)\phi_j(t)$, the termination point T_i must solve

$$1 - (1 - G_i^*)\phi_j(T_i) = Q_i(t) + G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) + (1 - G_i^*) \int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds,$$

Note that $\{G_i(s)\}_{s=\tau_2}^{T_i}$ is uniquely determined from (3) once we pin down G_i^* while $p_i(t)$ is independent of G_i . As such, the termination date is essentially determined by G_i^* , and we write $T_i(G_i^*)$ to indicate its dependence on G_i^* .

Lemma 4 There must exist a unique $T_i(G_i^*) < \infty$ for any $G_i^* \in [0, W_i^*]$. $T_i(G_i^*)$ is strictly decreasing in G_i^* with $T_i(W_i^*) = \tau_2$.

Proof. See Appendix.

For a given G_i^* , player *i* exits according to (3) from τ_2 to $T_i(G_i^*)$. At $T_i(G_i^*)$, the probability that player *i* is strong conditional on the continuation of the game reaches one, and hence $g_i(t) = 0$ for $t > T_i(G_i^*)$. We denote by T_i^* the equilibrium termination point. By standard argument, we can then show that $T_i^* = T_j^*$ must hold in equilibrium if the game continues beyond τ_2 , which implies that (G_1^*, G_2^*) must be chosen so that the two players stop exiting at the same time.

Lemma 5 $T_1^* = T_2^* \text{ if } \max\{T_1^*, T_2^*\} > \tau_2.$

Proof. Suppose on the contrary that $T_i^* > T_j^*$ while $T_i^* > \tau_2$. That is, only player *i* exits for $t \in (T_j^*, T_i^*)$. However, given that player *j* never exits, it is strictly better for player *i* to exit at max{ T_j^*, τ_i } rather than waiting until T_i^* , which is a contradiction.

We now consider player 1's problem at τ_1 , where his viable choice is either to exit with some positive probability at τ_1 or to wait at least until τ_2 . Suppose that player 1 waits until τ_2 . At time τ_2 , player 2 exits with probability G_2^* , in which case player 1 can earn a continuation payoff of $\frac{1}{r}$. If player 2 does not exit, which occurs with probability $1 - G_2^*$, player 1 stays in the game and earns positive payoff if and only if he achieves a breakthrough by τ_2 . Therefore, player 1 has an incentive to wait until τ_2 , rather than exiting at τ_1 , only if

$$p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_\tau})}{r + \lambda_1} + G_2^* e^{-r\Delta_\tau} (1 - p_1^* + p_1^* e^{-\lambda_1\Delta_\tau}) \ge (1 - p_1^*)(1 - e^{-r\Delta_\tau})\rho_1.$$
(4)

Clearly, player 1 has a stronger incentive to stay in the game when player 2 is more likely to exit at τ_2 . In fact, $G_2^* > 0$ is a necessary condition for $G_1^* < W_1^*$, as we will see below.

Lemma 6 For any $\Delta_{\tau} > 0$, (4) fails to hold if G_2^* is sufficiently small.

Proof. From (4), it suffices to show that

$$(1 - p_1^*)(1 - e^{-r\Delta_{\tau}})\rho_1 > p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_{\tau}})}{r + \lambda_1}.$$

Note that $p_1^* = \min\{\frac{r\rho_1}{\lambda_1}, p_0\}$ by definition. If $p_1^* = \frac{r\rho_1}{\lambda_1}$, this condition is reduced to

$$\frac{1 - e^{-r\Delta_{\tau}}}{r} > \frac{1 - e^{-(r+\lambda_1)\Delta_{\tau}}}{r+\lambda_1}$$

which holds for any $\Delta_{\tau} > 0$. Clearly, the same condition should also hold for any $p_0 < \frac{r\rho_1}{\lambda_1}$.

Which of the three possibilities emerges in equilibrium depends largely on the termination points (T_1, T_2) . What is in particular crucial is the termination point when the player does not exit with positive probability, i.e., $T_i(0)$.

Lemma 7 If $T_1(0) \ge T_2(0)$, player 1 must exit with positive probability at τ_1 , i.e., $G_1^* > 0$.

Proof. Suppose on the contrary that $G_1^* = 0$. This means that $G_2^* > 0$ by Lemma 3. In equilibrium, however, we must have $T_1(0) = T_2(G_2^*)$, which cannot be satisfied if $T_1(0) \ge T_2(0) > T_2(G_2^*)$.

4 Equilibrium characterization

4.1 Benchmark: no learning with fixed type

With the preceding lemmas, we are now ready to characterize the equilibrium of the entire game. Before we proceed further, however, we first consider a benchmark case where the players' types are fixed over time, i.e., $\lambda_1 = \lambda_2 = 0$, in order to single out the impact of experimentation in our setup. This particular example merits some independent attention as it is not a special case of our analysis which assumes $\tau_1 \neq \tau_2$. When the type is fixed over time, η_i is irrelevant and we simply assume $\eta_1 = \eta_2 = 0$ to save notation.

Given that $\lambda_1 = \lambda_2 = 0$, the model actually reduces to a standard war of attrition under incomplete information a la Fudenberg and Tirole (1986). Since the option value of experimentation is invariably zero, we have $\tau_1 = \tau_2 = 0$, implying that both players must exit at some positive rate from the outset. Moreover, generically, either one of the players must exit with strictly positive probability at time 0. To see this, T_i is now given by

$$1 - e^{-r\rho_j T_i} (1 - G_i(0)) = 1 - \alpha_i.$$

If $\alpha_i > 0$, then we can find a finite T_i that satisfies this condition for a given $G_i(0)$. Simple computation yields

$$T_i(G_i(0)) = \frac{\ln[1 - G_i(0)] - \ln \alpha_i}{r\rho_j}.$$

Although our argument thus far excludes the symmetric case where $\tau_1 = \tau_2$, we can apply essentially the same argument to this case to derive the equilibrium of this special case. First, we must still have the players stop exiting at the same time in equilibrium, i.e., $T_1(G_1(0)) = T_2(G_2(0))$. Second, we can also show by standard argument that the two players cannot exit with strictly positive probability at the same time, for there would always be an incentive to wait and exit an instant later. This suggests that player *i* must exit with strictly positive probability at time 0 if

$$\rho_j \ln \alpha_j > \rho_i \ln \alpha_i.$$

The equilibrium probability of immediate concession, $G_i(0)$, must satisfy

$$\frac{\ln[1 - G_i(0)] - \ln \alpha_i}{\rho_j} = -\frac{\ln \alpha_j}{\rho_i}.$$
(5)

It is well known that this is the unique equilibrium of this special case. Note that neither player exits with positive probability if and only if $T_1(0) = T_2(0)$ or $\rho_1 \ln \alpha_1 = \rho_2 \ln \alpha_2$.

Several remarks are in order. First, the equilibrium is unique for any given set of parameters, which stems from the fact that there is a termination point which must be reached in finite time. Although a war of attrition often admits a continuum of equilibria, Fudenberg and Tirole (1986) show that the equilibrium can be uniquely pinned down when there are types who never concede, such as the strong type in our setup. As we will see below, this uniqueness result is also retained in our model which incorporates experimentation.

Second, the equilibrium allocation is determined entirely by the asymptotic efficiency α_i and the flow cost ρ_i or more precisely by $\psi_i := \rho_i \ln \alpha_i$. Note that ψ_i is strictly increasing in α_i and decreasing in ρ_i with $\lim_{\alpha_i \to 1} \psi_i = \lim_{\rho_i \to 0} \psi_i = 0$. In particular, if the players share the same flow cost, the one who is expected to be stronger (a higher α_i) can induce the other player to give in and concede earlier. As a consequence, the player with a higher α_i (or more generally $\rho_i \ln \alpha_i$) always earns a higher payoff in equilibrium.

Finally, there exists no efficient equilibrium in the sense of Pareto. Since it is a pure waste of resource for two weak players to stay in the game simultaneously, the efficient allocation must have one player, say player 1, exit immediately if he is weak. Formally, this requires $G_1(0) = 1 - \alpha_1$. Clearly, this is not feasible for any $(\alpha_1, \alpha_2) \in (0, 1)^2$ and $(\rho_1, \rho_2) \in (0, \infty)^2$: from (5), $G_i(0) < 1 - \alpha_i$, so that the game continues beyond time 0 with a probability strictly larger than zero even when both players are weak. Intuitively, if $G_1(0) = 1 - \alpha_1$, player 1 can signal that he is strong by not exiting at time 0. Given this, player 2 must exit at the next instant if he is weak. If this is the case, however, player 1 no longer has an incentive to exit at time 0 even when he is weak because he strictly prefers to wait and see player 2's action.

4.2 Equilibrium with private learning

We now turn to the model which incorporates experimentation to see how this additional option alters the equilibrium allocation. Although there could be different types of equilibrium, of particular interest is whether there exists a pure-strategy equilibrium which implements the efficient allocation, i.e., $G_1^* = W_1^*$ and $G_2^* = W_2^*$.⁵ To see this possibility, note that if $G_1^* = W_1^*$, then $g_1(t) = 0$ for all $t > \tau_1$. Given this, player 2 has no incentive to continue beyond τ_2 if weak, and as such, $G_2^* = W_2^*$. This implies that the pure-strategy equilibrium exists if and only if

$$(1 - p_1^*)(1 - e^{-r\Delta_\tau})\rho_1 \ge p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_\tau})}{r+\lambda_1} + W_2^* e^{-r\Delta_\tau}(1 - p_1^* + p_1^* e^{-\lambda_1\Delta_\tau}).$$
(6)

⁵It is easy to verify that this allocation is Pareto efficient. To improve player 2's payoff, player 1 must exit earlier than τ_1 , but this necessarily lowers his payoff. The same argument applies to player 2.

In particular, if $\tau_1 > 0$, this condition becomes

$$r\rho_1\Big(\frac{1-e^{-r\Delta_\tau}}{r}-\frac{1-e^{-(r+\lambda_1)\Delta_\tau}}{r+\lambda_1}\Big) \ge W_2^*e^{-r\Delta_\tau}\frac{\lambda_1-r\rho_1(1-e^{-\lambda_1\Delta_\tau})}{\lambda_1-r\rho_1}.$$

Proposition 1 If (6) holds, there exists a unique equilibrium in which $G_1^* = W_1^*$ and $G_2^* = W_2^*$.

Proof. If (6) holds with strict inequality, it is a dominant strategy for player 1 to exit completely at τ_i . Given this, the pure-strategy equilibrium is clearly the only possible equilibrium. If (6) holds with equality, player 1 is indifferent. If $G_2^* < W_2^*$, however, (7) is violated, implying that $G_1^* = W_1^*$ must hold.

The proposition implies that there exists a non-degenerate set of parameters which support the efficient allocation as an equilibrium outcome: among other things, (6) is satisfied if Δ_{τ} is sufficiently large. This stands in clear contrast to the standard setup where the efficient allocation can never be attained. If τ_1 and τ_2 are sufficiently close to each other, on the other hand, (6) fails to hold, in which case the termination points (T_1, T_2) play a crucial role in determining the equilibrium allocation. Given that (6) does not hold, Lemma 6 implies that there must exist $\overline{G}_2 \in (0, W_2^*)$ such that

$$(1-p_1^*)(1-e^{-r\Delta_{\tau}})\rho_1 = p_1^* \frac{(\lambda_1 - r\rho_1)(1-e^{-(r+\lambda_1)\Delta_{\tau}})}{r+\lambda_1} + \overline{G}_2 e^{-r\Delta_{\tau}}(1-p_1^* + p_1^* e^{-\lambda_1\Delta_{\tau}}).$$
 (7)

When (7) holds, player 1 is indifferent between exiting at τ_1 and at τ_2 (or after).

Proposition 2 Suppose that (6) does not hold. Then, there exists a unique equilibrium in which

1. $G_1^* \in (0, W_1^*)$ and $G_2^* = \overline{G}_2 \in (0, W_2^*)$ such that $T_1(G_1^*) = T_2(\overline{G}_2)$ if $T_1(0) > T_2(\overline{G}_2)$; 2. $G_1^* = 0$ and $G_2^* \in (0, W_2^*)$ such that $T_1(0) = T_2(G_2^*)$ if $T_2(\overline{G}_2) \ge T_1(0)$.

Proof. Suppose that player 2 exits at τ_2 with probability \overline{G}_2 , in which case player 1 is indifferent. If $T_1(0) > T_2(\overline{G}_2)$, then player 1 can pick G_1^* which satisfies $T_1(G_1^*) = T_2(\overline{G}_2)$. If $T_2(\overline{G}_2) \ge T_1(0)$, then G_2^* must be high enough to satisfy $T_1(0) = T_2(G_2^*)$. Given this, player 1 has no incentive to exit at τ_1 , i.e., $G_1^* = 0$.

We can also show that no other equilibrium exists. Observe that if $G_1^* = W_1^*$, then $G_2^* = W_2^*$, which gives player 1 an incentive to deviate and stay in the game. This implies that $G_1^* < W_1^*$. If $G_2^* = W_2^*$, on the other hand, player 1 has a strict incentive to stay in the game, i.e., $G_1^* = 0$. This cannot constitute an equilibrium, however, because

 $T_1(0) > T_2(W_2^*) = \tau_2$. Given that $G_1^* \in [0, W_1^*)$ and $G_2^* \in (0, W_2^*)$, the indifference condition and the termination condition must be satisfied, giving a unique pair (G_1^*, G_2^*) .

Propositions 1 and 2 characterize all equilibria in this setup. Remarkably, player 2 must exit with some positive probability in any equilibrium of this game. The reason for this is that if $G_2^* = 0$, then player 1 has no incentive to stay in the game if he is weak. Given that $G_1^* = W_1^*$, however, player 2 also has no incentive to stay in the game beyond τ_2 if he is weak. Therefore, $G_2^* = 0$ cannot be sustained as an equilibrium, and player 2 must exit with positive probability at τ_2 .

4.3 Discussion: experimentation versus private learning

Our model incorporates experimentation and private learning into an otherwise standard war of attrition, meaning that our learning process consists of two distinct elements. The first necessary element is *private learning* where each player's type changes over time when a onetime breakthrough arrives. The second element, which turns out to be equally important, is *experimentation* where each player cannot directly observe his own state of nature. If we take out the first element, it is a model with fixed types which has already been discussed in section 4.1. Here, we examine the role of the second element by assuming that each player *i* can (privately) observe his own state θ_i .

Consider first the problem faced by a player who knows that the state is good. For this player, the belief $p_i(t)$ stays invariably at one for all t, regardless of the history of the game. Under the assumption that $\lambda_i > r\rho_i$, it is a dominant strategy to stay in the game, i.e., a player who knows that the state is good never exits.

Given this strategy, what should a player do if he knows that the state is bad? Since $p_{i,t} = 0$, player *i* exits at *t* if

$$r\rho_i \ge \frac{g_j(t)}{1 - G_j(t)},$$

which is the same condition as in section 4.1. Integrating both sides, the equilibrium strategy is given by

$$G_i(t) = 1 - (1 - G_i(0))e^{-r\rho_j t},$$

until $G_i(t) = 1 - \alpha_i$. We can then follow the same procedure to obtain $T_i(0)$ and $G_i(0)$ and derive the equilibrium.

This argument shows that when the state is directly observable, the model is essentially equivalent to the benchmark case with fixed types, provided that λ_i is large enough to satisfy

Assumption 1. In other words, a player's behavior does not depend on whether a player is currently strong or not, as long as he knows that the state is good. This shows that although the element of private learning is clearly indispensable, it is not enough by itself to alter the equilibrium allocation in a qualitative sense; all of our main results can be obtained only when it is combined with the element of experimentation.

5 Example: one-sided learning

5.1 The asymptotic efficiency and the equilibrium payoff

An interesting special case of our model is where only one player can experiment and learn while the other player's type is fixed over time. The case with one-sided learning illuminates the role of experimentation and how it can affect the equilibrium payoffs in a clear way. We in particular show that slow learning can be a blessing, as it provides a credible commitment device to stay in the game and is hence instrumental in inducing the other player to concede earlier. This stands in sharp contrast to the standard setup in which the stronger player always earns a higher payoff.

Now suppose that $\eta_1 = 0$ and $\eta_2 = 1$, so that player 1 has no margin of experimentation. To assure that this is the only source of heterogeneity, we fix other parameters at the same level: $\alpha_i = \alpha$, $\lambda_i = \lambda$ and $\rho_i = \rho$ for i = 1, 2. Under this specification, the probability that player 1 is strong stays constant at α , leaving no room for experimentation regardless of λ . In contrast, the probability that player 2 is strong at t is $\alpha(1 - e^{-\lambda t})$ which is strictly lower than α at any finite t, i.e., player 2 is strictly less efficient than player 1.

Player 1 knows that if he is weak at time 0, he will always be weak. Since $\tau_1 = 0$, player 1 must decide whether to exit immediately or wait until $\tau_2 := \frac{1}{\lambda} \ln[\frac{(\lambda - r\rho)\alpha}{r\rho(1-\alpha)}]$. Given $p_{1,0} = \alpha \eta_1 = 0$, from Proposition 1, player 1 exits immediately if and only if

$$(1 - e^{-r\tau_2})\rho \ge W_2^* e^{-r\tau_2} = (1 - \alpha + \alpha e^{-\lambda\tau_2})e^{-r\tau_2}.$$

which can be written as

$$\left(1 - \left(\frac{r\rho(1-\alpha)}{(\lambda-r\rho)\alpha}\right)^{\frac{r}{\lambda}}\right)\rho \ge (1-\alpha)\left(\frac{r\rho(1-\alpha)}{(\lambda-r\rho)\alpha}\right)^{\frac{r}{\lambda}}\frac{\lambda}{\lambda-r\rho}.$$
(8)

If (8) is satisfied, player 1, if he is weak, exits immediately in the unique equilibrium. As a consequence, player 1 earns positive payoff if and only if he is strong at time 0. Let $\Pi_i(G_i^*, G_j^*)$ denote player *i*'s expected equilibrium payoff. The equilibrium payoff for player 1 is then given by

$$r\Pi_1(W_1^*, W_2^*) = \alpha.$$

In contrast, player 2 can earn positive payoff from the outset with probability $1 - \alpha$. With the remaining probability, player 1 is strong and stays in the game indefinitely. Given this, player 2 waits until τ_2 and exits once and for all if he cannot achieve a breakthrough by then. The equilibrium payoff for player 2 is hence given by

$$r\Pi_2(W_2^*, W_1^*) = 1 - \alpha + \alpha \left(\alpha \frac{(\lambda - r\rho)(1 - e^{-(r+\lambda)\tau_2})}{r+\lambda} - (1 - \alpha)(1 - e^{-r\tau_2})\rho \right) > 1 - \alpha.$$

It is immediate to see that player 2, who is strictly inefficient, can earn a higher payoff if $0.5 \ge \alpha$. The question is then if (8) can be satisfied for some $\alpha \in (0, 0.5]$. The following statement shows that this is indeed possible when experimentation is less costly, provided that the players are relatively impatient.

Proposition 3 For any $\alpha \in (0, 0.5]$, if $r \ge \lambda > 0$, there exists $\overline{\rho}(\alpha, \frac{r}{\lambda}) \in (0, \frac{\alpha\lambda}{r})$ such that there exists an equilibrium in which the strictly inefficient player earns a higher payoff for $\rho < \overline{\rho}(\alpha, \frac{r}{\lambda})$.

Proof. A necessary condition for this statement is $\tau_2 > 0$ which can be written as $\alpha > \frac{r\rho}{\lambda}$. To prove this, therefore, it suffices to show that there exists some $\alpha \in (0, 0.5]$ which satisfies both $\alpha > \frac{r\rho}{\lambda}$ and (8). To this end, fix $\rho = \frac{\alpha\lambda}{(1+x)r}$ for some given $\alpha \in (0, 0.5]$, so that $\alpha > \frac{r\rho}{\lambda} = \frac{\alpha}{1+x}$ holds for any x > 0. Then, (8) can be written as

$$\alpha \left(1 + (1 + \frac{x}{1-\alpha})^{-\frac{r}{\lambda}}\right) \frac{(1 + \frac{x}{1-\alpha})^{1+\frac{r}{\lambda}}}{(1+x)^2} \ge \frac{\lambda}{r},$$

which holds for a sufficiently large x for any $\alpha \in (0, 0.5]$ if $\frac{r}{\lambda} \ge 1$. It is clear from this that the threshold value can be obtained as a function of α and $\frac{r}{\lambda}$.

The impact of a decrease in ρ is not immediately clear, as it yields two opposing effects: on one hand, it lowers the cost of staying in the game, rendering player 1 less willing to exit immediately; on the other hand, it can also lower the cost of experimentation for player 2, allowing him to extend τ_2 . Intuitively, the latter effect dominates the former when player 1 is less patient and the benefit of waiting (until τ_2) is smaller. The proposition suggests that for any $r \ge \lambda > 0$, we can find ρ that is low enough to let player 1 concede immediately.

There is another way to interpret this result. Consider an alternative specification in which $\eta_1 = \eta_2 = 1$ but $\lambda_1 > \lambda_2 = \lambda > 0$. The case with one-sided learning can then be seen as the limit case where $\lambda_1 \to \infty$. In this sense, Proposition 3 implies that slow learning can be a blessing when outlasting the opponent is the major concern. With a high λ_i , the belief $p_i(t)$ goes down and the option value of experimentation dissipates rather quickly; in

fact, as $\lambda_i \to \infty$, $p_i(t) \to 0$ for any arbitrarily small t > 0, thereby approximating a situation where $\eta_i = 0$. In contrast, if λ_i is small, the belief $p_i(t)$ goes down only gradually, and the option value of experimentation is hence kept relatively high over time. Due to this effect, slow learning can actually be beneficial, as it induces the other player to concede early even if he is strictly more efficient in the asymptotic sense.

5.2 Implications

Standards competitions: One application which fits our framework particularly well is standards competitions, mainly on two accounts. First, standards competitions are a typical example of war of attrition in that the market can accommodate only a few standards in the end largely due to network externalities. Second, it is also important to note that many of those competitions have historically taken place in innovation-intensive industries. To analyse a standards competition as a war of attrition, it seem imperative to take into account the possibility of innovation and technological progress.

Concerning standards competitions, a puzzle that has attracted much attention in the literature is that we often observe instances where what appears to be an inferior technology survives and emerges as the defacto standard. One of the most famous anecdotes of this sort is perhaps the video format war between Betamax and VHS. In this new market for home VCR, Sony – the creator of Betamax – was the predecessor and quickly captured the entire market share by the end of 1975 before VHS was introduced in the following year. Sony at the time also possessed a more matured technology due to its prior involvement in the U-matic system – one of the first cassette-based tape formats. In fact, many recognized Betamax as the superior technology to VHS, with its better resolution, superior sound, more stable image, faster response to operating keys, and so on. Despite all these advantages, however, Betamax somehow faded out of popularity, forcing SONY to retreat from the format effectively by the end of the 1980s.

Many explanations have been suggested to account for why Betamax lost in the war,⁶ but our analysis sheds new light on this issue from a previously unexplored channel, by providing a mechanism through which a technologically less matured standard can outlast other possibly more matured predecessors in this type of competition. Although a lot remains to be seen at this point,⁷ we argue that the margin of experimentation, which has largely been ignored,

⁶The literature generally focuses on SONY's pricing and marketing strategies. Many also argue that Betamax's shorter recording time was a crucial deciding factor. See, e.g., Casumano et al. [4] and Liebowitz and Margolis [18].

⁷One important caveat is that a war of attrition is typically characterized by a mixed-strategy equilibrium, so almost any outcome, when it is seen in isolation, is consistent with theory. It is also emphasized in the literature that network externalities may give rise to the multiplicity of equilibrium, with some more efficient

can be a powerful force that can dictate the outcome of a standards competition, thereby offering some empirical implications.

Technology adoption: From a broader perspective, our analysis yields an important implication for technology adoption where there are choices over how to develop a particular standard. To put this idea in context, suppose that there are two available technologies, a (more matured) fast technology λ_F and a (less matured) slow technology λ_S where $\lambda_F > \lambda_S$. Suppose further that a player can choose which technology to adopt at time 0 (to simplify the argument, suppose that only one of the players has this choice). Our analysis then implies that there are cases where it is optimal to deliberately choose the slow technology λ^S , which takes more time to develop, even when the two technologies are perfectly equivalent in every other dimension.

Of course, in reality, technology adoption often involves a tradeoff between short-run and long-run gains. Although we consider an extreme case where the two technologies differ only in the learning rate for the sake of illustration, it is more realistic to assume that some technologies take more time to develop but are more efficient in the long run. One way to capture this possibility is to consider technologies that differ also in the asymptotic efficiency, i.e., a fast technology (α_F, λ_F) and a slow technology (α_S, λ_S) where $\alpha_S > \alpha_F$. This type of tradeoff is often present in the choice between exploitation (of old ideas) and exploration (of new ideas): it is extensively discussed in the innovation literature that exploration inherently entails more uncertainty and is more time-consuming but potentially brings about a bigger benefit (March [19]). Our analysis then implies that there arises an additional value of exploration in particular, and long-run innovation strategies in general, when a firm faces a war-of-attrition type of competition.

Industry dynamics: Consider yet another variant of our model where one player, say player 1, enters the game earlier at time $-\tilde{t}$, and the other player, player 2, joins later at time 0. Suppose that the two players are homogeneous in every aspect, except for the timing of entry. Then, having been in existence for a duration \tilde{t} of time, the probability that player 1 is weak at time 0 is $1 - \alpha + \alpha \eta e^{-\lambda \tilde{t}}$. Letting $\eta_1 := \eta e^{-\lambda \tilde{t}}$ and $\eta_2 = \eta$, therefore, the game is equivalent to the one in which the players differ in η_i . Taken at face value, our model thus predicts that younger firms, with less matured technologies, tend to outperform older counterparts, thereby generating and accelerating product life cycle. Although this hypothesis is hard to test empirically due to the lack of information about innovation and also to the difficulty in

than others, and an inefficient equilibrium can at least in theory emerge due to the lack of coordination. Even in this case, one still needs an explanation for why concerned parties coordinate on the wrong side. Our framework building on a war of attrition can provide a way.

controlling for the asymptotic efficiency,⁸ our analysis suggests a testable implication linking firm age and survival.

6 Conclusion

In this paper, we provide a simple extension of a standard war of attrition to incorporate the possibility that players may innovate through experimentation. We find that this extension yields some qualitative impacts on the strategic nature of the problem. Most notably, we show that the option value of experimentation provides a credible commitment device which affects the equilibrium allocation in a non-trivial way. As a direct consequence of this effect, there arises an equilibrium in which the strictly less efficient player, but with larger room for experimentation, can earn a higher payoff in equilibrium. Our analysis provide a plausible mechanism through which an inferior technology survives and becomes the de facto standard in a standards competition, and more broadly offers implications for technology adoption and industry dynamics.

References

- [1] Dilip Abreu and Faruk Gul. Bargaining and reputation. *Econometrica*, 68:85–117, 2000.
- [2] Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99:116–131, 1989.
- [3] Alessandro Barbarino and Boyan Jovanovic. Shakeouts and market crashes. International Economic Review, 48:385–420, 2007.
- [4] Michael A. Casumano, Yiorgos Mylonadis, and Richard S. Rosenbloom. Strategic maneuvering and mass-market dynamics: The triumph of vhs over beta. *Business History Review*, 66:51–94, 1992.
- [5] Christophe Chamley and Douglas Gale. Information revelation and strategic delay in a model of investment. *Econometrica*, 62:1065–1085, 1994.
- [6] Chia-Hui Chen and Junichiro Ishida. Hierarchical experimentation. ISER Discussion Paper No. 949, 2015.

⁸A more specific empirical challenge is that the timing of entry is most likely to be endogenous: some firms enter early because of their technological superiority (Klepper and Simons [17]), which might be translated into a higher α_i in our setup; as such, we might expect some correlation between the timing of entry and the asymptotic efficiency.

- [7] Jay P. Choi. Irreversible choice of uncertain technologies with network externalities. *RAND Journal of Economics*, 25:382–401, 1994.
- [8] Jean-Paul Decamps and Thomas Mariotti. Investment timing and learning externalities. Journal of Economic Theory, 118:80–102, 2004.
- [9] Joseph Ferrell and Garth Saloner. Standadization, compatibility, and innovation. RAND Journal of Economics, 16:70–83, 1985.
- [10] Joseph Ferrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncementa, and predation. *American Economic Review*, 76:940–955, 1986.
- [11] Drew Fudenberg and Jean Tirole. A theory of exit in duopoly. *Econometrica*, 54:943–960, 1986.
- [12] Shinsuke Kambe. Bargaining with imperfect commitment. Games and Economic Behavior, 28:217–237, 1999.
- [13] Michael Katz and Carl Shapiro. Network externalities, competition and compatibility. American Economic Review, 75:424–440, 1985.
- [14] Michael Katz and Carl Shapiro. Technology adoption in the presence of network externalities. Journal of Political Economy, 94:822–841, 1986.
- [15] Michael Katz and Carl Shapiro. Systems competition and network effects. Journal of Economic Perspectives, 8:93–115, 1994.
- [16] Kyungmin Kim and Frances Z. Xu Lee. Information acquisition in a war of attrition. American Economic Journal: Microeconomics, 6:37–78, 2014.
- [17] Steven Klepper and Kenneth L. Simons. Industry shakeouts and technological change. International Journal of Industrial Organization, 23:23–43, 2005.
- [18] Stan J. Liebowitz and Stephen E. Margolis. Path dependence, lock-in, and history. Journal of Law, Economics, and Organization, 11:205–226, 1995.
- [19] James March. Exploration and exploitation in organizational learning. Organizational Science, 2:71–87, 1991.
- [20] Pauli Murto and Juuso Välimäki. Learning and information aggregation in an exit game. *Review of Economic Studies*, 78:1426–1461, 2011.

- [21] Janusz A. Ordover and Ariel Rubinstein. A sequential concession game with asymmetric information. Quarterly Journal of Economics, 101:879–888, 1986.
- [22] Rafael Rob. Learning and capacity expansion under demand uncertainty. Review of Economic Studies, 58:655–675, 1991.
- [23] Rafael Rob and Nikolaos Vettas. Foreign direct investment and exports with growing demand. *Review of Economic Studies*, 70:629–648, 2003.
- [24] Victor Stango. The economics of standards wars. Review of Network Economics, 3:1–19, 2004.
- [25] Nikolaos Vettas. Demand and supply in new markets: Diffusion with bilateral learning. RAND Journal of Economics, 29:215–233, 1998.

Appendix

Proof of Lemma 2. We first establish that G_i must be strictly increasing and continuous, which implies that the players must be held indifferent. We can show that G_i must be increasing by standard argument. Suppose otherwise, i.e., there exists an interval [t', t''] such that $G_i(t)$ is constant for $t \in [t', t'']$ and increases for t > t''. Then, it is strictly better for player j to exit at t rather at any point in an interval shortly after t''. However, if $g_j(t) = 0$ for $t \in [t'', t'' + \varepsilon]$, it is strictly better for player i to exit at t'' rather than in this interval, which is a contradiction. This means that G_i must be strictly increasing for $t > \tau$ until until no weak player remains in the game.

To show that G_i must be continuous, suppose that G_i makes a discrete jump at some $t > \tau_2$. Then, the other player strictly prefers to wait and exit at t + dt rather than at any time shortly before t, which contradicts that G_i is strictly increasing.

These results suggest that in any continuation equilibrium after τ_2 , a weak player must be held indifferent, i.e.,

$$\frac{g_i(t)}{1 - G_i(t)} = r\rho_j - p_j(t)\lambda_j$$

Integrating both sides yields

$$\ln[1 - G_i(\tau_2)] - \ln[1 - G_i(t)] = r\rho_j(t - \tau_2) + \ln Q_j(t) - \ln Q_j(\tau_2).$$

Alternatively, this can be written as

$$G_i(t) = 1 - e^{-r\rho_j(t-\tau_2)} (1 - G_i(\tau_2)) \frac{Q_j(\tau_2)}{Q_j(t)},$$

which shows the necessary optimality condition that must be satisfied for $t > \tau_2$.

Proof of Lemma 4. Since G_i is a distribution function, it must be weakly increasing. Note also that $W_i(t)$ is strictly decreasing by definition, given that a player can only switch from weak to strong, but not the other way around. Moreover, for any $G_i^* \in [0, W_i^*]$,

$$\lim_{t \to \infty} G_i(t) = 1 > \lim_{t \to \infty} W_i(t).$$

there must exist a unique $T_i(G_i^*) < \infty$.

To show that T_i is decreasing in G_i^* , observe first that

$$\frac{\partial G_i(T_i)}{\partial T_i}\frac{\partial T_i}{\partial G_i^*} + \frac{\partial G_i(T_i)}{\partial G_i^*} = \frac{\partial W_i(T_i)}{\partial T_i}\frac{\partial T_i}{\partial G_i^*} + \frac{\partial W_i(T_i)}{\partial G_i^*}$$

Since $\frac{\partial G_i(T_i)}{\partial T_i} - \frac{\partial W_i(T_i)}{\partial T_i} > 0$, we need to show

$$\frac{\partial T_i}{\partial G_i^*} < 0 \ \Leftrightarrow \ \frac{\partial G_i(T_i)}{\partial G_i^*} > \frac{\partial W_i(T_i)}{\partial G_i^*},$$

which can be written as

$$\phi_j(T_i) > p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) - \int_{\tau_2}^{T_i} \Phi_j(s) (1 - e^{-\lambda(T_i - s)}) ds.$$
(9)

Note that $W_i(T_i) = G_i(T_i)$ by definition, i.e.,

$$1 - (1 - G_i^*)\phi_j(T_i) = Q_i(t) + G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) + (1 - G_i^*) \int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds,$$

which can be written as

$$\int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds = \frac{1 - Q_i(t) - G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)})}{1 - G_i^*} - \phi_j(T_j).$$

Plugging this into (9) yields

$$\frac{1 - Q_i(t) - G_i^* p_i^* (1 - e^{-\lambda_i (T_i - \tau_i)})}{1 - G_i^*} > p_i^* (1 - e^{-\lambda_i (T_i - \tau_2)}),$$

which is simplified to

$$1 - Q_i(T_i) > p_i^* (1 - e^{-\lambda_i(T_i - \tau_2)}).$$

This condition can be written as

$$\alpha_i(1-\eta_i e^{-\lambda_i T_i}) > \frac{\alpha_i \eta_i e^{-\lambda_i T_i}}{1-\alpha_i + \alpha_i \eta_i e^{-\lambda_i T_i}} (1-e^{-\lambda_i (T_i-\tau_2)}),$$

which holds because $\alpha_i \geq \frac{\alpha_i \eta_i e^{-\lambda_i T_i}}{1-\alpha_i + \alpha_i \eta_i e^{-\lambda_i T_i}}$ and $1 - \eta_i e^{-\lambda_i T_i} > 1 - e^{-\lambda_i (T_i - \tau_2)}$.

A War of Attrition with Experimenting Players^{*}

Chia-Hui Chen[†] and Junichiro Ishida[‡]

October 20, 2017

Abstract

A standard incomplete-information war of attrition is extended to incorporate experimentation and private learning. We obtain a characterization of all equilibria in this extended setup and use this setup to illuminate a tradeoff between short-run and long-run gains of experimentation. The extension yields qualitative impacts on the strategic nature of the problem. The option value of experimentation serves as a credible commitment device to stay in the game, which is instrumental in inducing the other player to concede earlier. As a direct consequence, there may be an equilibrium in which the strictly less efficient player can get the better end of the deal, implying that slow learning can be a blessing in this type of competition. Our analysis gives insight into why an apparently inferior technology often survives in many standards competitions and more broadly offers implications for technology adoption and industry dynamics. We also show that there is a non-degenerate set of parameters that can support the Pareto-efficient allocation as an equilibrium outcome whereas it is never possible in the standard setup.

JEL Classification Number: D82, D83

Keywords: war of attrition, bandit problem, experimentation, standards competitions, private learning.

^{*}The first author acknowledges financial support from JSPS KAKENHI Grant-in-Aid (B) JP16H03615. The second author acknowledges financial support from JSPS KAKENHI Grant-in-Aid (S) JP15H05728, (A) JP20245031, (B) JP16H03615 and (C) JP24530196 as well as the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

[†]Institute of Economic Research, Kyoto University. Email: chchen@kier.kyoto-u.ac.jp

[‡]Institute of Social and Economic Research, Osaka University. Email: jishida@iser.osaka-u.ac.jp

1 Introduction

In many standards competitions, we observe instances where an apparently less efficient or less matured technology survives in the end as the defacto standard, as exemplified most notably by the video format war between Betamax and VHS.¹ Although standards competitions often take the form of a war of attrition, this observation seems to be at odds with theoretical predictions from some standard versions of it, suggesting that there may be a missing link that has been overlooked in the literature. In this paper, we attempt to fill this gap by extending a standard war of attrition to ask whether there are factors, other than sheer technological superiority, which affect the outcome of this type of economic competition in any significant way.

One observation stands out in search of this missing link: standards competitions typically take place in innovation-intensive industries, such as telecommunications, home VCR, audio, operating systems, high-speed rail and automated driving. In most of these cases, incoming firms, or standards, are initially still technologically premature but gradually improve their productive efficiency or profitability over time via experimentation. The fact that competing standards may innovate alters the strategic nature of competition at least in two ways. First, it is not only the current efficiency level but also the margin of experimentation, i.e., how much room for improvement a given technology is expected to possess, which can play a decisive role: a player with large room for improvement clearly has more incentive to experiment and hence less incentive to drop out. Second, the possibility of technological progress also implies that it is *ex ante* not clear how many standards can survive in a given market. Although the market for standards is typically not large enough to accommodate many inefficient ones, there are still cases where multiple standards attain enough efficiency to coexist for a substantial duration of time, as in the competition between Windows and Mac in operating systems or among Nintendo, SONY and Microsoft in video game consoles. When multiple standards emerge, the strategic nature of the problem tips over, for it is no longer a war of attrition with no point in waiting for the rival to drop out.

To describe this situation, we extend a two-player incomplete-information war of attrition to incorporate experimentation and private learning. The setup is quite standard, except that each player's type may change over time as a result of (learning-by-doing) experimentation. Each player's type, which is either weak or strong, is his own private information. If a player

¹There is naturally an extensive literature discussing whether and how a society gets "locked in" with an inefficient standard, mainly focusing on various issues arising from network effects (Katz and Shapiro [13], [14], Farrell and Saloner [9], [10], Arthur [2], Choi [7]). Also, see Katz and Shapiro [15] and Stango [24] for surveys.

is weak, however, he may achieve a technological breakthrough and become strong with some probability as long as he stays in the game. The probability of achieving a breakthrough depends on the unobservable state of nature which is not directly observable and must be uncovered via experimentation. Finally, we assume that a strong player can earn positive payoff regardless of the other player's presence, so that it is a dominant strategy for him to stay in the game indefinitely. It is this last assumption which gives rise to the *ex ante* uncertainty about how many players can survive in the end.

We obtain a characterization of all equilibria in this extended setup. The extension to incorporate experimentation is technically rather straightforward but yields some qualitative impacts on the strategic nature of the problem. In our extended setup, a weak player may stay in the game just because he may achieve a breakthrough tomorrow. The option value of experimentation thus serves as a credible commitment device to stay in the game, which is instrumental in inducing the other player to concede earlier. As a direct consequence of this force, we find an equilibrium in which the strictly less efficient player, but with more room for experimentation, can get the better end of the deal, i.e., slow learning can be a blessing. This stands in sharp contrast to the standard setup with no experimentation (hereafter, simply the standard war of attrition) where the equilibrium payoff is monotonic with respect to a player's "efficiency" level. Our analysis sheds new light on the tradeoff between short-run and long-run gains of experimentation from a previously unexplored channel and suggests some practical implications for technology choice and industry dynamics. On a more specific level, it gives insight into why an inferior technology often survives in standards competitions.

From a more technical point of view, it is also worth noting that there may exist an efficient equilibrium (in the sense of Pareto) in our extended setup, again in clear contrast to the standard war of attrition where there generically exists no efficient equilibrium. This is because the efficient allocation requires that any weak player exit immediately unless both of them are weak. A way to (approximately) achieve this allocation is that: one player, say player 1, exits with probability one at time 0 if he is weak; player 2 then exits with probability one an instant later if he is weak and player 1 did not exit. There is no such equilibrium, however, because if player 1 knows that player 2 will exit with strictly positive probability an instant later, he is always better off with waiting than exiting immediately. More generally, in any equilibrium of the standard setup, at most one player can exit with strictly positive probability, which immediately excludes the possibility of realizing the efficient allocation.

The situation changes rather drastically once the element of experimentation is introduced into the setup. In the model with experimentation, the efficient allocation typically involves delay, as it leads to better informed decisions. A player may exit with strictly positive probability when the option value of experimentation equals the flow cost, the timing of which may arrive differently between the players. This property is crucial for the construction of equilibrium in general and can be exploited to construct an equilibrium in which both of the players exit with strictly positive probability at some (different) points of the game. More importantly, due to this property, there exists a non-degenerate set of parameters that can support the efficient allocation as an equilibrium outcome.

Related literature: A war of attrition, which originates from theoretical biology, is well suited for various economic situations, such as standards competitions, where multiple parties compete for a fixed resource. In many economic applications, however, participating players may innovate and improve their "fitness" in a relatively short span of time. This could be an important departure from biological settings where genetic evolution occurs rather slowly and such a possibility is almost negligible.

An essential feature of our war of attrition is that there is a type of player who would never concede. Without the element of experimentation, our model becomes a variant of Fudenberg and Tirole [11] who analyze a war of attrition when there are types who would never concede. Among other things, they establish the uniqueness of equilibrium when some types may never exit – a property which is also retained in our setup – whereas there is typically a continuum of equilibria when there are no such types. Similarly, Ordover and Rubinstein [21] analyze a concession game in which one player is informed about the disagreement outcome; there, a player who knows that the disagreement outcome is favorable never concedes. In this vein, our model is also related to reputational models of bargaining, such as Abreu and Gul [1] and Kambe [12], which include a stubborn type who "irrationally" sticks to his own demand and never concedes.

Kim and Lee [16] analyze information acquisition in a war of attrition where the war terminates stochastically, in which case each player receives a payoff according to the underlying state of nature. In this setup, they consider a situation where each player can observe the state of nature by engaging in costly information acquisition. Our analysis shares an aspect with this work in that learning takes place in a war of attrition but differs in the way players acquire information: in theirs, information acquisition is done once and for all when a player incurs the cost whereas in ours, it is through time-consuming experimentation with information being revealed only gradually over time. This difference is crucial as the speed of learning, or the margin of experimentation, is one of the main focuses of our analysis.

Chen and Ishida [6] analyze a hierarchial model of experimentation in which a principal retains the authority to terminate the joint project while an agent focuses strictly on running the project assigned to him. A key departure of the analysis is the possibility that the principal may be informed about the project type, so that her termination decision becomes a signal of the project quality. In equilibrium, a principal who knows that the project is good never terminates the project, just like the strong type in the current setup who never exits. The model can thus be seen as a hybrid of experimentation and war of attrition models. The focus of that work is on characterizing the agent's effort dynamics and how they depend on factors such as the principal's ability to evaluate the eventual value of a project.

Several works examine the optimal timing of exit with information externalities among players. Chamley and Gale [5] consider a model of strategic investment in which there are N players, a random number n of whom have an investment option. Assuming that the value of investment depends positively on the random number n, there is an incentive to wait and see others' investment decisions. Decamps and Mariotti [8] also consider a setup in which the value of investment is common across the players while the cost is only privately known. Murto and Välimäki [20] analyze an exit game with private learning where each player receives a signal in each period which partially reveals his own type. In those models, the players' payoffs are positively correlated through the aggregate state of nature: a good (aggregate) state is a good news for all players in the game. In contrast, in our setup, the state of nature is individual-specific and a good state for one player is necessarily a bad news for the other.

Finally, one driving force of our model is the *ex ante* uncertainty about the number of players who can survive in the end. Some previous works explore this issue mainly by focusing on demand-side uncertainty. Most notably, Rob [22] analyzes entry dynamics of firms in a situation where uncertainty exists with respect to the "limit of the market" and it is costly to overshoot the market limit. See Vettas [25], Rob and Vettas [23], and Barbarino and Jovanovic [3] for various extensions of this approach.

2 Model

Environment: Consider a war of attrition between two players, indexed by i = 1, 2. The basic setup is a standard war of attrition under incomplete information, except that a player's type may change over time as a result of experimentation. Time is continuous, extending from zero to infinity, and each player decides whether to stay in the game or exit from it at each instant with a decision to exit irreversible and publicly observable. Below, we first outline the standard part and then describe how we augment it with the possibility of experimentation and private learning.

Standard war of attrition: A player is either weak or strong at each point in time. There

is asymmetry of information between the players where each player's type is his private information and cannot be observed by the other player. The flow payoff to a player is one if he is strong or the other player has exited from the game; otherwise, the flow payoff is $-\rho_i$ where $\rho_i \in (0, \infty)$ captures the flow cost of staying in the game for player *i*. This specification implies that there are two (equally effective) ways for a weak player to earn positive payoff, either to achieve a breakthrough or to wait for the other player to exit. The reservation payoff, which a player earns when he exits, is normalized at zero. Each player discounts payoffs by a common discount rate *r*. For the sake of exposition, we refer to this particular setup as the standard war of attrition in what follows.

Experimentation: We augment the above setup by introducing the hidden state of nature for each player which indirectly influences his type. More precisely, before the game begins, nature chooses the state for each player *i*, denoted by $\theta_i \in \{0, 1\}$. The realized state (θ_1, θ_2) is not observable to anyone and drawn independently with $P(\theta_i = 1) = \alpha_i$. Given the state, nature then chooses the initial type of each player where

$$P(\text{player } i \text{ weak } | \theta_i = 0) = 1, \ P(\text{player } i \text{ weak } | \theta_i = 1) = \eta_i.$$

The unconditional probability that player *i* is strong at time 0 is hence given by $(1 - \eta_i)\alpha_i$. We assume, however, that a player who is initially weak may achieve a breakthrough and become strong over time: if a weak player stays in the game for [t, t + dt), a breakthrough arrives with probability $\lambda_i \theta_i dt$. For clarity, we call $\lambda_i \in (0, \infty)$ the *learning rate* for player *i*. Note that in this specification, a player can be strong only if the underlying state is good $(\theta_i = 1)$.

Interpretation: Throughout the analysis, we in particular focus on two parameters, α_i and η_i , as they admit economically relevant interpretations. First, since a player can be strong only if the state is good, α_i indicates the upper bound of the probability that player *i* is strong, i.e., if player *i* stays in the game indefinitely, the probability that he is good converges to α_i . For this reason, we refer to α_i as the *asymptotic efficiency* of player *i*. Second, η_i measures how much room for experimentation player *i* has, which we refer to as the *margin of experimentation*. As a practical interpretation, we take this as capturing the "technological maturity" of player *i*, where a low η_i means that the player is in a technologically more advanced or matured stage with less room for experimentation. Note that the learning rate plays a similar, though not equivalent, role to the margin of experimentation since $1 - \eta_i$ can alternatively be regarded as the mass probability of achieving a breakthrough at time 0.

3 Analysis

3.1 Preliminaries

Given the payoff structure, it is a dominant strategy for a strong player to stay in the game indefinitely with no further decisions to make. We can thus focus on a weak player, provided that both players are still in the game. Throughout the analysis, we denote a generic player by i and "the other player" by j.

Let $p_i(t)$ denote the belief that the state for player *i* is good ($\theta_i = 1$) at time *t*, conditional on being weak. Since a breakthrough occurs at a constant rate, the belief depends only on time and is given by

$$p_i(t) = \frac{\alpha_i \eta_i e^{-\lambda_i t}}{Q_i(t)},$$

where $Q_i(t) := 1 - \alpha_i + \alpha_i \eta_i e^{-\lambda t}$. Note that the belief declines monotonically over time, i.e., no news is a bad news.

If a player is strong or the only player has exited, he earns a continuation payoff of $\frac{1}{r}$. As such, $r\rho_i$ measures the cost-benefit ratio of staying in the game for player *i*. If the learning rate is low relative to the cost-benefit ratio, a player has no incentive to experiment, and the game is effectively reduced to the standard war of attrition. To focus on relevant and interesting cases, therefore, we assume throughout the analysis that the option value of experimentation is sufficiently large.

Assumption 1 $\lambda_i > r\rho_i$.

3.2 The optimality conditions

The strategy for player i is G_i which specifies the unconditional distribution of exiting time. Define g_i as its corresponding density, which represents the rate at which player i exits. We assume that G_i is right continuous.

At each instant t, each player decides whether to stay or exit. If a player exits, he earns a continuation payoff of zero. In contrast, if a player stays and exits at the next instant, he can achieve a breakthrough with probability $p_{i,t}\lambda dt$ while the rival player exits with probability $\frac{g_j(t)dt}{1-G_j(t)}$; in either case, the player earns $\frac{1}{r}$. Since the flow cost is $\rho_i dt$, a necessary condition for player i to exit at t is

$$r\rho_i \ge p_i(t)\lambda_i + \frac{g_j(t)}{1 - G_j(t)},\tag{1}$$

which shows the key tradeoff between exiting now and waiting for an instant. If there is no possibility of a breakthrough ($\lambda_i = 0$), the model becomes the standard war of attrition where the hazard rate is constant over time.

If $\alpha_i \eta_i \geq \frac{r\rho_i}{\lambda_i}$, there exists τ_i such that

$$p_i(\tau_i) = \frac{\alpha_i \eta_i e^{-\lambda_i \tau_i}}{Q_i(\tau_i)} = \frac{r\rho_i}{\lambda_i},\tag{2}$$

which is the point at which the option value of experimentation equals the flow cost for player i. From (2), we obtain

$$\tau_i = \frac{1}{\lambda_i} \ln \left[\frac{(\lambda_i - r\rho_i)\alpha_i \eta_i}{r\rho_i (1 - \alpha_i)} \right].$$

Let $\tau_i = 0$ if $\frac{r\rho_i}{\lambda_i} > \alpha_j \eta_j$. For the subsequent analysis, we suppose $\tau_2 > \tau_1 \ge 0$ and define $\Delta_{\tau} := \tau_2 - \tau_1$. This is of course not without loss of generality, as it excludes the possibility that $\tau_1 = \tau_2$. We exclude the symmetric case because it is a knife-edge case exhibiting some peculiar properties and also because the analysis of this case is relatively straightforward (see section 4.1 for a brief discussion on the symmetric case).²

Assumption 2 $\tau_2 > \tau_1 \ge 0$.

We first establish some useful properties that must hold in any equilibrium.

Lemma 1 $G_i(t) = 0$ for $t < \tau_i$, i.e., player i never exits before τ_i .

Proof. For $t < \tau_i$, $\lambda p_i(t) > r\rho_i$ regardless of the other player's strategy G_j . As such, it is a dominant strategy to stay in the game.

Lemma 2 For $t > \tau_2$, G_i is given by

$$G_i(t) = 1 - (1 - G_i(\tau_2))\phi_j(t), \tag{3}$$

where $\phi_j(t) := e^{-r\rho_j(t-\tau_2)} \frac{Q_j(\tau_2)}{Q_j(t)}$, until no weak player remains in the game.

Proof. See Appendix.

These results mark two key departures from the standard war of attrition. First, there is a tipping point τ_i which can arrive strictly later than time 0. Second, the optimal strategy

²Although we only discuss the case where $\tau_1 = \tau_2 = 0$, we can essentially apply the same argument to any $\tau_1 = \tau_2 > 0$.

must take into account how the other player acquires information over time, which is captured by $\frac{Q_j(\tau_2)}{Q_j(t)}$. If there is no learning, i.e., $\lambda_j = 0$, then $\frac{Q_j(\tau_2)}{Q_j(t)} = 1$ for all t, and the model is reduced to the standard war of attrition as we have seen. If $\lambda_j > 0$, on the other hand, $\frac{Q_j(\tau_2)}{Q_j(t)}$ is strictly increasing in t, prompting the player to exit at an accelerating rate.

Given that a weak player must exit at some positive rate, there may be a point at which no weak player remains in the game. To compute this termination point, define $W_i(t)$ as the probability that player *i* is not strong (either is weak or has exited) at time *t* which is assumed to be left continuous.³ Note that if player *i* never exits, the probability that he is weak at *t* is $Q_i(t)$ which constitutes the lower bound for $W_i(t)$. Also, $W_i(t) \ge G_i(t)$ by definition. We will later discuss in detail how we construct this function W_i , but for now, we take that such a function exists and is well defined for any given *t*.

Let $T_i := \inf\{t : G_i(t) = W_i(t)\}$ denote the termination point for player *i*. If there exists $T_i < \infty$, $g_i(t) = 0$ and G_i stays constant for all $t > T_i$, i.e., any remaining player must be strong and never exits once the game reaches the termination point.

3.3 In the gap

For $t > \tau_2$, it is quite straightforward to obtain the continuation equilibrium which is characterized by the indifference condition (3). We also know that before τ_i , player *i* never exits, i.e., $G_i(t) = 0$ for $t < \tau_i$. What remains to be seen is the case in-between, i.e., for $t \in [\tau_1, \tau_2]$. The following statement clarifies what could possibly happen in this interval.

Lemma 3 In equilibrium, generically, one of the following must happen:⁴

- 1. Player 1 exits with strictly positive probability at τ_1 ;
- 2. Player 2 exits with strictly positive probability at τ_2 ;
- 3. Both players 1 and 2 exit with strictly positive probability at τ_1 and τ_2 , respectively.

Moreover, no player exits in (τ_1, τ_2) .

Proof. For $t \in [\tau_1, \tau_2)$, player 2 never exits by Lemma 1. Given this, it is never optimal for player 1 to exit at any $t \in (\tau_1, \tau_2)$, because he strictly prefers to exit at τ_1 rather than at any $t \in (\tau_1, \tau_2)$. To prove the lemma, therefore, it suffices to show that player 1 never exits with positive probability at τ_2 . If $G_2(\tau_2) = 0$, player 1 strictly prefers to exit at τ_1 than at τ_2 . If

³Note that G_i is assumed to be right continuous, so that we implicitly consider the timing structure in which $W_i(t)$ is realized first and then G_i is chosen.

⁴There is also a case where $\tau_1 = \tau_2$ and $T_1(0) = T_2(0)$, so that neither player exits with positive probability in equilibrium. Throughout the analysis, we ignore this non-generic case.

 $G_2(\tau_2) > 0$, player 1 strictly prefers to wait until $\tau_2 + dt$ than exiting at τ_2 . This means that G_1 can make a discrete jump only at τ_1 .

This result suggests that any equilibrium of this game is thoroughly characterized by a pair $(G_1(\tau_1), G_2(\tau_2))$ where $G_1(\tau_1) = G_1(\tau_2)$ by Lemma 3. Let $G_i^* := G_i(\tau_i), W_i^* := W_i(\tau_i)$ and $p_i^* := p_i(\tau_i)$ to save notation.

We now derive $W_i(t)$ to see how the termination point T_i is related to G_i^* . Note that $W_i(t) = Q_i(t)$ if player *i* has never exited. On the other hand, once player *i* starts exiting, we in general have $W_i(t) > Q_i(t)$ because $W_i(t)$ includes those who exited but would have become strong later had they stayed in the game. More precisely, for a player who exits at t', the probability that he would achieve a breakthrough by *t* if he stayed is $p_i(t')(1 - e^{-\lambda_i(t-t')})$. Since player *i* may exit with positive probability only at τ_i , we obtain

$$W_{i}(t) = \begin{cases} Q_{i}(t) & \text{for } t \in [0, \tau_{i}], \\ Q_{i}(t) + G_{i}^{*} p_{i}^{*} (1 - e^{-\lambda_{i}(t - \tau_{i})}) + \int_{\tau_{2}}^{t} g_{i}(s) p_{i}(s) (1 - e^{-\lambda_{i}(t - s)}) ds & \text{for } t \in (\tau_{i}, T_{i}], \end{cases}$$

where $g_i(t) = (r\rho_j - p_j(t)\lambda_j)(1 - G_i^*)\phi_j(t)$. For $t > T_i$, $W_i(t)$ stays constant at $G_i(T_i)$.

By definition, $G_i(T_i) = W_i(T_i)$. Letting $\Phi_i(t) := (r\rho_j - p_j(t)\lambda_j)\phi_j(t)$, the termination point T_i must solve

$$1 - (1 - G_i^*)\phi_j(T_i) = Q_i(t) + G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) + (1 - G_i^*) \int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds,$$

Note that $\{G_i(s)\}_{s=\tau_2}^{T_i}$ is uniquely determined from (3) once we pin down G_i^* while $p_i(t)$ is independent of G_i . As such, the termination date is essentially determined by G_i^* , and we write $T_i(G_i^*)$ to indicate its dependence on G_i^* .

Lemma 4 There must exist a unique $T_i(G_i^*) < \infty$ for any $G_i^* \in [0, W_i^*]$. $T_i(G_i^*)$ is strictly decreasing in G_i^* with $T_i(W_i^*) = \tau_2$.

Proof. See Appendix.

For a given G_i^* , player *i* exits according to (3) from τ_2 to $T_i(G_i^*)$. At $T_i(G_i^*)$, the probability that player *i* is strong conditional on the continuation of the game reaches one, and hence $g_i(t) = 0$ for $t > T_i(G_i^*)$. We denote by T_i^* the equilibrium termination point. By standard argument, we can then show that $T_i^* = T_j^*$ must hold in equilibrium if the game continues beyond τ_2 , which implies that (G_1^*, G_2^*) must be chosen so that the two players stop exiting at the same time.

Lemma 5 $T_1^* = T_2^* \text{ if } \max\{T_1^*, T_2^*\} > \tau_2.$

Proof. Suppose on the contrary that $T_i^* > T_j^*$ while $T_i^* > \tau_2$. That is, only player *i* exits for $t \in (T_j^*, T_i^*)$. However, given that player *j* never exits, it is strictly better for player *i* to exit at max{ T_j^*, τ_i } rather than waiting until T_i^* , which is a contradiction.

We now consider player 1's problem at τ_1 , where his viable choice is either to exit with some positive probability at τ_1 or to wait at least until τ_2 . Suppose that player 1 waits until τ_2 . At time τ_2 , player 2 exits with probability G_2^* , in which case player 1 can earn a continuation payoff of $\frac{1}{r}$. If player 2 does not exit, which occurs with probability $1 - G_2^*$, player 1 stays in the game and earns positive payoff if and only if he achieves a breakthrough by τ_2 . Therefore, player 1 has an incentive to wait until τ_2 , rather than exiting at τ_1 , only if

$$p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_\tau})}{r + \lambda_1} + G_2^* e^{-r\Delta_\tau} (1 - p_1^* + p_1^* e^{-\lambda_1\Delta_\tau}) \ge (1 - p_1^*)(1 - e^{-r\Delta_\tau})\rho_1.$$
(4)

Clearly, player 1 has a stronger incentive to stay in the game when player 2 is more likely to exit at τ_2 . In fact, $G_2^* > 0$ is a necessary condition for $G_1^* < W_1^*$, as we will see below.

Lemma 6 For any $\Delta_{\tau} > 0$, (4) fails to hold if G_2^* is sufficiently small.

Proof. From (4), it suffices to show that

$$(1 - p_1^*)(1 - e^{-r\Delta_{\tau}})\rho_1 > p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_{\tau}})}{r + \lambda_1}.$$

Note that $p_1^* = \min\{\frac{r\rho_1}{\lambda_1}, p_0\}$ by definition. If $p_1^* = \frac{r\rho_1}{\lambda_1}$, this condition is reduced to

$$\frac{1 - e^{-r\Delta_{\tau}}}{r} > \frac{1 - e^{-(r+\lambda_1)\Delta_{\tau}}}{r+\lambda_1}$$

which holds for any $\Delta_{\tau} > 0$. Clearly, the same condition should also hold for any $p_0 < \frac{r\rho_1}{\lambda_1}$.

Which of the three possibilities emerges in equilibrium depends largely on the termination points (T_1, T_2) . What is in particular crucial is the termination point when the player does not exit with positive probability, i.e., $T_i(0)$.

Lemma 7 If $T_1(0) \ge T_2(0)$, player 1 must exit with positive probability at τ_1 , i.e., $G_1^* > 0$.

Proof. Suppose on the contrary that $G_1^* = 0$. This means that $G_2^* > 0$ by Lemma 3. In equilibrium, however, we must have $T_1(0) = T_2(G_2^*)$, which cannot be satisfied if $T_1(0) \ge T_2(0) > T_2(G_2^*)$.

4 Equilibrium characterization

4.1 Benchmark: no learning with fixed type

With the preceding lemmas, we are now ready to characterize the equilibrium of the entire game. Before we proceed further, however, we first consider a benchmark case where the players' types are fixed over time, i.e., $\lambda_1 = \lambda_2 = 0$, in order to single out the impact of experimentation in our setup. This particular example merits some independent attention as it is not a special case of our analysis which assumes $\tau_1 \neq \tau_2$. When the type is fixed over time, η_i is irrelevant and we simply assume $\eta_1 = \eta_2 = 0$ to save notation.

Given that $\lambda_1 = \lambda_2 = 0$, the model actually reduces to a standard war of attrition under incomplete information a la Fudenberg and Tirole (1986). Since the option value of experimentation is invariably zero, we have $\tau_1 = \tau_2 = 0$, implying that both players must exit at some positive rate from the outset. Moreover, generically, either one of the players must exit with strictly positive probability at time 0. To see this, T_i is now given by

$$1 - e^{-r\rho_j T_i} (1 - G_i(0)) = 1 - \alpha_i.$$

If $\alpha_i > 0$, then we can find a finite T_i that satisfies this condition for a given $G_i(0)$. Simple computation yields

$$T_i(G_i(0)) = \frac{\ln[1 - G_i(0)] - \ln \alpha_i}{r\rho_j}.$$

Although our argument thus far excludes the symmetric case where $\tau_1 = \tau_2$, we can apply essentially the same argument to this case to derive the equilibrium of this special case. First, we must still have the players stop exiting at the same time in equilibrium, i.e., $T_1(G_1(0)) = T_2(G_2(0))$. Second, we can also show by standard argument that the two players cannot exit with strictly positive probability at the same time, for there would always be an incentive to wait and exit an instant later. This suggests that player *i* must exit with strictly positive probability at time 0 if

$$\rho_j \ln \alpha_j > \rho_i \ln \alpha_i.$$

The equilibrium probability of immediate concession, $G_i(0)$, must satisfy

$$\frac{\ln[1 - G_i(0)] - \ln \alpha_i}{\rho_j} = -\frac{\ln \alpha_j}{\rho_i}.$$
(5)

It is well known that this is the unique equilibrium of this special case. Note that neither player exits with positive probability if and only if $T_1(0) = T_2(0)$ or $\rho_1 \ln \alpha_1 = \rho_2 \ln \alpha_2$.

Several remarks are in order. First, the equilibrium is unique for any given set of parameters, which stems from the fact that there is a termination point which must be reached in finite time. Although a war of attrition often admits a continuum of equilibria, Fudenberg and Tirole (1986) show that the equilibrium can be uniquely pinned down when there are types who never concede, such as the strong type in our setup. As we will see below, this uniqueness result is also retained in our model which incorporates experimentation.

Second, the equilibrium allocation is determined entirely by the asymptotic efficiency α_i and the flow cost ρ_i or more precisely by $\psi_i := \rho_i \ln \alpha_i$. Note that ψ_i is strictly increasing in α_i and decreasing in ρ_i with $\lim_{\alpha_i \to 1} \psi_i = \lim_{\rho_i \to 0} \psi_i = 0$. In particular, if the players share the same flow cost, the one who is expected to be stronger (a higher α_i) can induce the other player to give in and concede earlier. As a consequence, the player with a higher α_i (or more generally $\rho_i \ln \alpha_i$) always earns a higher payoff in equilibrium.

Finally, there exists no efficient equilibrium in the sense of Pareto. Since it is a pure waste of resource for two weak players to stay in the game simultaneously, the efficient allocation must have one player, say player 1, exit immediately if he is weak. Formally, this requires $G_1(0) = 1 - \alpha_1$. Clearly, this is not feasible for any $(\alpha_1, \alpha_2) \in (0, 1)^2$ and $(\rho_1, \rho_2) \in (0, \infty)^2$: from (5), $G_i(0) < 1 - \alpha_i$, so that the game continues beyond time 0 with a probability strictly larger than zero even when both players are weak. Intuitively, if $G_1(0) = 1 - \alpha_1$, player 1 can signal that he is strong by not exiting at time 0. Given this, player 2 must exit at the next instant if he is weak. If this is the case, however, player 1 no longer has an incentive to exit at time 0 even when he is weak because he strictly prefers to wait and see player 2's action.

4.2 Equilibrium with private learning

We now turn to the model which incorporates experimentation to see how this additional option alters the equilibrium allocation. Although there could be different types of equilibrium, of particular interest is whether there exists a pure-strategy equilibrium which implements the efficient allocation, i.e., $G_1^* = W_1^*$ and $G_2^* = W_2^*$.⁵ To see this possibility, note that if $G_1^* = W_1^*$, then $g_1(t) = 0$ for all $t > \tau_1$. Given this, player 2 has no incentive to continue beyond τ_2 if weak, and as such, $G_2^* = W_2^*$. This implies that the pure-strategy equilibrium exists if and only if

$$(1 - p_1^*)(1 - e^{-r\Delta_\tau})\rho_1 \ge p_1^* \frac{(\lambda_1 - r\rho_1)(1 - e^{-(r+\lambda_1)\Delta_\tau})}{r+\lambda_1} + W_2^* e^{-r\Delta_\tau}(1 - p_1^* + p_1^* e^{-\lambda_1\Delta_\tau}).$$
(6)

⁵It is easy to verify that this allocation is Pareto efficient. To improve player 2's payoff, player 1 must exit earlier than τ_1 , but this necessarily lowers his payoff. The same argument applies to player 2.

In particular, if $\tau_1 > 0$, this condition becomes

$$r\rho_1\Big(\frac{1-e^{-r\Delta_\tau}}{r}-\frac{1-e^{-(r+\lambda_1)\Delta_\tau}}{r+\lambda_1}\Big) \ge W_2^*e^{-r\Delta_\tau}\frac{\lambda_1-r\rho_1(1-e^{-\lambda_1\Delta_\tau})}{\lambda_1-r\rho_1}.$$

Proposition 1 If (6) holds, there exists a unique equilibrium in which $G_1^* = W_1^*$ and $G_2^* = W_2^*$.

Proof. If (6) holds with strict inequality, it is a dominant strategy for player 1 to exit completely at τ_i . Given this, the pure-strategy equilibrium is clearly the only possible equilibrium. If (6) holds with equality, player 1 is indifferent. If $G_2^* < W_2^*$, however, (7) is violated, implying that $G_1^* = W_1^*$ must hold.

The proposition implies that there exists a non-degenerate set of parameters which support the efficient allocation as an equilibrium outcome: among other things, (6) is satisfied if Δ_{τ} is sufficiently large. This stands in clear contrast to the standard setup where the efficient allocation can never be attained. If τ_1 and τ_2 are sufficiently close to each other, on the other hand, (6) fails to hold, in which case the termination points (T_1, T_2) play a crucial role in determining the equilibrium allocation. Given that (6) does not hold, Lemma 6 implies that there must exist $\overline{G}_2 \in (0, W_2^*)$ such that

$$(1-p_1^*)(1-e^{-r\Delta_{\tau}})\rho_1 = p_1^* \frac{(\lambda_1 - r\rho_1)(1-e^{-(r+\lambda_1)\Delta_{\tau}})}{r+\lambda_1} + \overline{G}_2 e^{-r\Delta_{\tau}}(1-p_1^* + p_1^* e^{-\lambda_1\Delta_{\tau}}).$$
 (7)

When (7) holds, player 1 is indifferent between exiting at τ_1 and at τ_2 (or after).

Proposition 2 Suppose that (6) does not hold. Then, there exists a unique equilibrium in which

1. $G_1^* \in (0, W_1^*)$ and $G_2^* = \overline{G}_2 \in (0, W_2^*)$ such that $T_1(G_1^*) = T_2(\overline{G}_2)$ if $T_1(0) > T_2(\overline{G}_2)$; 2. $G_1^* = 0$ and $G_2^* \in (0, W_2^*)$ such that $T_1(0) = T_2(G_2^*)$ if $T_2(\overline{G}_2) \ge T_1(0)$.

Proof. Suppose that player 2 exits at τ_2 with probability \overline{G}_2 , in which case player 1 is indifferent. If $T_1(0) > T_2(\overline{G}_2)$, then player 1 can pick G_1^* which satisfies $T_1(G_1^*) = T_2(\overline{G}_2)$. If $T_2(\overline{G}_2) \ge T_1(0)$, then G_2^* must be high enough to satisfy $T_1(0) = T_2(G_2^*)$. Given this, player 1 has no incentive to exit at τ_1 , i.e., $G_1^* = 0$.

We can also show that no other equilibrium exists. Observe that if $G_1^* = W_1^*$, then $G_2^* = W_2^*$, which gives player 1 an incentive to deviate and stay in the game. This implies that $G_1^* < W_1^*$. If $G_2^* = W_2^*$, on the other hand, player 1 has a strict incentive to stay in the game, i.e., $G_1^* = 0$. This cannot constitute an equilibrium, however, because

 $T_1(0) > T_2(W_2^*) = \tau_2$. Given that $G_1^* \in [0, W_1^*)$ and $G_2^* \in (0, W_2^*)$, the indifference condition and the termination condition must be satisfied, giving a unique pair (G_1^*, G_2^*) .

Propositions 1 and 2 characterize all equilibria in this setup. Remarkably, player 2 must exit with some positive probability in any equilibrium of this game. The reason for this is that if $G_2^* = 0$, then player 1 has no incentive to stay in the game if he is weak. Given that $G_1^* = W_1^*$, however, player 2 also has no incentive to stay in the game beyond τ_2 if he is weak. Therefore, $G_2^* = 0$ cannot be sustained as an equilibrium, and player 2 must exit with positive probability at τ_2 .

4.3 Discussion: experimentation versus private learning

Our model incorporates experimentation and private learning into an otherwise standard war of attrition, meaning that our learning process consists of two distinct elements. The first necessary element is *private learning* where each player's type changes over time when a onetime breakthrough arrives. The second element, which turns out to be equally important, is *experimentation* where each player cannot directly observe his own state of nature. If we take out the first element, it is a model with fixed types which has already been discussed in section 4.1. Here, we examine the role of the second element by assuming that each player *i* can (privately) observe his own state θ_i .

Consider first the problem faced by a player who knows that the state is good. For this player, the belief $p_i(t)$ stays invariably at one for all t, regardless of the history of the game. Under the assumption that $\lambda_i > r\rho_i$, it is a dominant strategy to stay in the game, i.e., a player who knows that the state is good never exits.

Given this strategy, what should a player do if he knows that the state is bad? Since $p_{i,t} = 0$, player *i* exits at *t* if

$$r\rho_i \ge \frac{g_j(t)}{1 - G_j(t)},$$

which is the same condition as in section 4.1. Integrating both sides, the equilibrium strategy is given by

$$G_i(t) = 1 - (1 - G_i(0))e^{-r\rho_j t},$$

until $G_i(t) = 1 - \alpha_i$. We can then follow the same procedure to obtain $T_i(0)$ and $G_i(0)$ and derive the equilibrium.

This argument shows that when the state is directly observable, the model is essentially equivalent to the benchmark case with fixed types, provided that λ_i is large enough to satisfy

Assumption 1. In other words, a player's behavior does not depend on whether a player is currently strong or not, as long as he knows that the state is good. This shows that although the element of private learning is clearly indispensable, it is not enough by itself to alter the equilibrium allocation in a qualitative sense; all of our main results can be obtained only when it is combined with the element of experimentation.

5 Example: one-sided learning

5.1 The asymptotic efficiency and the equilibrium payoff

An interesting special case of our model is where only one player can experiment and learn while the other player's type is fixed over time. The case with one-sided learning illuminates the role of experimentation and how it can affect the equilibrium payoffs in a clear way. We in particular show that slow learning can be a blessing, as it provides a credible commitment device to stay in the game and is hence instrumental in inducing the other player to concede earlier. This stands in sharp contrast to the standard setup in which the stronger player always earns a higher payoff.

Now suppose that $\eta_1 = 0$ and $\eta_2 = 1$, so that player 1 has no margin of experimentation. To assure that this is the only source of heterogeneity, we fix other parameters at the same level: $\alpha_i = \alpha$, $\lambda_i = \lambda$ and $\rho_i = \rho$ for i = 1, 2. Under this specification, the probability that player 1 is strong stays constant at α , leaving no room for experimentation regardless of λ . In contrast, the probability that player 2 is strong at t is $\alpha(1 - e^{-\lambda t})$ which is strictly lower than α at any finite t, i.e., player 2 is strictly less efficient than player 1.

Player 1 knows that if he is weak at time 0, he will always be weak. Since $\tau_1 = 0$, player 1 must decide whether to exit immediately or wait until $\tau_2 := \frac{1}{\lambda} \ln[\frac{(\lambda - r\rho)\alpha}{r\rho(1-\alpha)}]$. Given $p_{1,0} = \alpha \eta_1 = 0$, from Proposition 1, player 1 exits immediately if and only if

$$(1 - e^{-r\tau_2})\rho \ge W_2^* e^{-r\tau_2} = (1 - \alpha + \alpha e^{-\lambda\tau_2})e^{-r\tau_2}.$$

which can be written as

$$\left(1 - \left(\frac{r\rho(1-\alpha)}{(\lambda-r\rho)\alpha}\right)^{\frac{r}{\lambda}}\right)\rho \ge (1-\alpha)\left(\frac{r\rho(1-\alpha)}{(\lambda-r\rho)\alpha}\right)^{\frac{r}{\lambda}}\frac{\lambda}{\lambda-r\rho}.$$
(8)

If (8) is satisfied, player 1, if he is weak, exits immediately in the unique equilibrium. As a consequence, player 1 earns positive payoff if and only if he is strong at time 0. Let $\Pi_i(G_i^*, G_j^*)$ denote player *i*'s expected equilibrium payoff. The equilibrium payoff for player 1 is then given by

$$r\Pi_1(W_1^*, W_2^*) = \alpha.$$

In contrast, player 2 can earn positive payoff from the outset with probability $1 - \alpha$. With the remaining probability, player 1 is strong and stays in the game indefinitely. Given this, player 2 waits until τ_2 and exits once and for all if he cannot achieve a breakthrough by then. The equilibrium payoff for player 2 is hence given by

$$r\Pi_2(W_2^*, W_1^*) = 1 - \alpha + \alpha \left(\alpha \frac{(\lambda - r\rho)(1 - e^{-(r+\lambda)\tau_2})}{r+\lambda} - (1 - \alpha)(1 - e^{-r\tau_2})\rho \right) > 1 - \alpha.$$

It is immediate to see that player 2, who is strictly inefficient, can earn a higher payoff if $0.5 \ge \alpha$. The question is then if (8) can be satisfied for some $\alpha \in (0, 0.5]$. The following statement shows that this is indeed possible when experimentation is less costly, provided that the players are relatively impatient.

Proposition 3 For any $\alpha \in (0, 0.5]$, if $r \ge \lambda > 0$, there exists $\overline{\rho}(\alpha, \frac{r}{\lambda}) \in (0, \frac{\alpha\lambda}{r})$ such that there exists an equilibrium in which the strictly inefficient player earns a higher payoff for $\rho < \overline{\rho}(\alpha, \frac{r}{\lambda})$.

Proof. A necessary condition for this statement is $\tau_2 > 0$ which can be written as $\alpha > \frac{r\rho}{\lambda}$. To prove this, therefore, it suffices to show that there exists some $\alpha \in (0, 0.5]$ which satisfies both $\alpha > \frac{r\rho}{\lambda}$ and (8). To this end, fix $\rho = \frac{\alpha\lambda}{(1+x)r}$ for some given $\alpha \in (0, 0.5]$, so that $\alpha > \frac{r\rho}{\lambda} = \frac{\alpha}{1+x}$ holds for any x > 0. Then, (8) can be written as

$$\alpha \left(1 + (1 + \frac{x}{1-\alpha})^{-\frac{r}{\lambda}}\right) \frac{(1 + \frac{x}{1-\alpha})^{1+\frac{r}{\lambda}}}{(1+x)^2} \ge \frac{\lambda}{r},$$

which holds for a sufficiently large x for any $\alpha \in (0, 0.5]$ if $\frac{r}{\lambda} \ge 1$. It is clear from this that the threshold value can be obtained as a function of α and $\frac{r}{\lambda}$.

The impact of a decrease in ρ is not immediately clear, as it yields two opposing effects: on one hand, it lowers the cost of staying in the game, rendering player 1 less willing to exit immediately; on the other hand, it can also lower the cost of experimentation for player 2, allowing him to extend τ_2 . Intuitively, the latter effect dominates the former when player 1 is less patient and the benefit of waiting (until τ_2) is smaller. The proposition suggests that for any $r \ge \lambda > 0$, we can find ρ that is low enough to let player 1 concede immediately.

There is another way to interpret this result. Consider an alternative specification in which $\eta_1 = \eta_2 = 1$ but $\lambda_1 > \lambda_2 = \lambda > 0$. The case with one-sided learning can then be seen as the limit case where $\lambda_1 \to \infty$. In this sense, Proposition 3 implies that slow learning can be a blessing when outlasting the opponent is the major concern. With a high λ_i , the belief $p_i(t)$ goes down and the option value of experimentation dissipates rather quickly; in

fact, as $\lambda_i \to \infty$, $p_i(t) \to 0$ for any arbitrarily small t > 0, thereby approximating a situation where $\eta_i = 0$. In contrast, if λ_i is small, the belief $p_i(t)$ goes down only gradually, and the option value of experimentation is hence kept relatively high over time. Due to this effect, slow learning can actually be beneficial, as it induces the other player to concede early even if he is strictly more efficient in the asymptotic sense.

5.2 Implications

Standards competitions: One application which fits our framework particularly well is standards competitions, mainly on two accounts. First, standards competitions are a typical example of war of attrition in that the market can accommodate only a few standards in the end largely due to network externalities. Second, it is also important to note that many of those competitions have historically taken place in innovation-intensive industries. To analyse a standards competition as a war of attrition, it seem imperative to take into account the possibility of innovation and technological progress.

Concerning standards competitions, a puzzle that has attracted much attention in the literature is that we often observe instances where what appears to be an inferior technology survives and emerges as the defacto standard. One of the most famous anecdotes of this sort is perhaps the video format war between Betamax and VHS. In this new market for home VCR, Sony – the creator of Betamax – was the predecessor and quickly captured the entire market share by the end of 1975 before VHS was introduced in the following year. Sony at the time also possessed a more matured technology due to its prior involvement in the U-matic system – one of the first cassette-based tape formats. In fact, many recognized Betamax as the superior technology to VHS, with its better resolution, superior sound, more stable image, faster response to operating keys, and so on. Despite all these advantages, however, Betamax somehow faded out of popularity, forcing SONY to retreat from the format effectively by the end of the 1980s.

Many explanations have been suggested to account for why Betamax lost in the war,⁶ but our analysis sheds new light on this issue from a previously unexplored channel, by providing a mechanism through which a technologically less matured standard can outlast other possibly more matured predecessors in this type of competition. Although a lot remains to be seen at this point,⁷ we argue that the margin of experimentation, which has largely been ignored,

⁶The literature generally focuses on SONY's pricing and marketing strategies. Many also argue that Betamax's shorter recording time was a crucial deciding factor. See, e.g., Casumano et al. [4] and Liebowitz and Margolis [18].

⁷One important caveat is that a war of attrition is typically characterized by a mixed-strategy equilibrium, so almost any outcome, when it is seen in isolation, is consistent with theory. It is also emphasized in the literature that network externalities may give rise to the multiplicity of equilibrium, with some more efficient

can be a powerful force that can dictate the outcome of a standards competition, thereby offering some empirical implications.

Technology adoption: From a broader perspective, our analysis yields an important implication for technology adoption where there are choices over how to develop a particular standard. To put this idea in context, suppose that there are two available technologies, a (more matured) fast technology λ_F and a (less matured) slow technology λ_S where $\lambda_F > \lambda_S$. Suppose further that a player can choose which technology to adopt at time 0 (to simplify the argument, suppose that only one of the players has this choice). Our analysis then implies that there are cases where it is optimal to deliberately choose the slow technology λ^S , which takes more time to develop, even when the two technologies are perfectly equivalent in every other dimension.

Of course, in reality, technology adoption often involves a tradeoff between short-run and long-run gains. Although we consider an extreme case where the two technologies differ only in the learning rate for the sake of illustration, it is more realistic to assume that some technologies take more time to develop but are more efficient in the long run. One way to capture this possibility is to consider technologies that differ also in the asymptotic efficiency, i.e., a fast technology (α_F, λ_F) and a slow technology (α_S, λ_S) where $\alpha_S > \alpha_F$. This type of tradeoff is often present in the choice between exploitation (of old ideas) and exploration (of new ideas): it is extensively discussed in the innovation literature that exploration inherently entails more uncertainty and is more time-consuming but potentially brings about a bigger benefit (March [19]). Our analysis then implies that there arises an additional value of exploration in particular, and long-run innovation strategies in general, when a firm faces a war-of-attrition type of competition.

Industry dynamics: Consider yet another variant of our model where one player, say player 1, enters the game earlier at time $-\tilde{t}$, and the other player, player 2, joins later at time 0. Suppose that the two players are homogeneous in every aspect, except for the timing of entry. Then, having been in existence for a duration \tilde{t} of time, the probability that player 1 is weak at time 0 is $1 - \alpha + \alpha \eta e^{-\lambda \tilde{t}}$. Letting $\eta_1 := \eta e^{-\lambda \tilde{t}}$ and $\eta_2 = \eta$, therefore, the game is equivalent to the one in which the players differ in η_i . Taken at face value, our model thus predicts that younger firms, with less matured technologies, tend to outperform older counterparts, thereby generating and accelerating product life cycle. Although this hypothesis is hard to test empirically due to the lack of information about innovation and also to the difficulty in

than others, and an inefficient equilibrium can at least in theory emerge due to the lack of coordination. Even in this case, one still needs an explanation for why concerned parties coordinate on the wrong side. Our framework building on a war of attrition can provide a way.

controlling for the asymptotic efficiency,⁸ our analysis suggests a testable implication linking firm age and survival.

6 Conclusion

In this paper, we provide a simple extension of a standard war of attrition to incorporate the possibility that players may innovate through experimentation. We find that this extension yields some qualitative impacts on the strategic nature of the problem. Most notably, we show that the option value of experimentation provides a credible commitment device which affects the equilibrium allocation in a non-trivial way. As a direct consequence of this effect, there arises an equilibrium in which the strictly less efficient player, but with larger room for experimentation, can earn a higher payoff in equilibrium. Our analysis provide a plausible mechanism through which an inferior technology survives and becomes the de facto standard in a standards competition, and more broadly offers implications for technology adoption and industry dynamics.

References

- [1] Dilip Abreu and Faruk Gul. Bargaining and reputation. *Econometrica*, 68:85–117, 2000.
- [2] Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99:116–131, 1989.
- [3] Alessandro Barbarino and Boyan Jovanovic. Shakeouts and market crashes. International Economic Review, 48:385–420, 2007.
- [4] Michael A. Casumano, Yiorgos Mylonadis, and Richard S. Rosenbloom. Strategic maneuvering and mass-market dynamics: The triumph of vhs over beta. *Business History Review*, 66:51–94, 1992.
- [5] Christophe Chamley and Douglas Gale. Information revelation and strategic delay in a model of investment. *Econometrica*, 62:1065–1085, 1994.
- [6] Chia-Hui Chen and Junichiro Ishida. Hierarchical experimentation. ISER Discussion Paper No. 949, 2015.

⁸A more specific empirical challenge is that the timing of entry is most likely to be endogenous: some firms enter early because of their technological superiority (Klepper and Simons [17]), which might be translated into a higher α_i in our setup; as such, we might expect some correlation between the timing of entry and the asymptotic efficiency.

- [7] Jay P. Choi. Irreversible choice of uncertain technologies with network externalities. *RAND Journal of Economics*, 25:382–401, 1994.
- [8] Jean-Paul Decamps and Thomas Mariotti. Investment timing and learning externalities. Journal of Economic Theory, 118:80–102, 2004.
- [9] Joseph Ferrell and Garth Saloner. Standadization, compatibility, and innovation. RAND Journal of Economics, 16:70–83, 1985.
- [10] Joseph Ferrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncementa, and predation. *American Economic Review*, 76:940–955, 1986.
- [11] Drew Fudenberg and Jean Tirole. A theory of exit in duopoly. *Econometrica*, 54:943–960, 1986.
- [12] Shinsuke Kambe. Bargaining with imperfect commitment. Games and Economic Behavior, 28:217–237, 1999.
- [13] Michael Katz and Carl Shapiro. Network externalities, competition and compatibility. American Economic Review, 75:424–440, 1985.
- [14] Michael Katz and Carl Shapiro. Technology adoption in the presence of network externalities. Journal of Political Economy, 94:822–841, 1986.
- [15] Michael Katz and Carl Shapiro. Systems competition and network effects. Journal of Economic Perspectives, 8:93–115, 1994.
- [16] Kyungmin Kim and Frances Z. Xu Lee. Information acquisition in a war of attrition. American Economic Journal: Microeconomics, 6:37–78, 2014.
- [17] Steven Klepper and Kenneth L. Simons. Industry shakeouts and technological change. International Journal of Industrial Organization, 23:23–43, 2005.
- [18] Stan J. Liebowitz and Stephen E. Margolis. Path dependence, lock-in, and history. Journal of Law, Economics, and Organization, 11:205–226, 1995.
- [19] James March. Exploration and exploitation in organizational learning. Organizational Science, 2:71–87, 1991.
- [20] Pauli Murto and Juuso Välimäki. Learning and information aggregation in an exit game. *Review of Economic Studies*, 78:1426–1461, 2011.

- [21] Janusz A. Ordover and Ariel Rubinstein. A sequential concession game with asymmetric information. Quarterly Journal of Economics, 101:879–888, 1986.
- [22] Rafael Rob. Learning and capacity expansion under demand uncertainty. Review of Economic Studies, 58:655–675, 1991.
- [23] Rafael Rob and Nikolaos Vettas. Foreign direct investment and exports with growing demand. *Review of Economic Studies*, 70:629–648, 2003.
- [24] Victor Stango. The economics of standards wars. Review of Network Economics, 3:1–19, 2004.
- [25] Nikolaos Vettas. Demand and supply in new markets: Diffusion with bilateral learning. RAND Journal of Economics, 29:215–233, 1998.

Appendix

Proof of Lemma 2. We first establish that G_i must be strictly increasing and continuous, which implies that the players must be held indifferent. We can show that G_i must be increasing by standard argument. Suppose otherwise, i.e., there exists an interval [t', t''] such that $G_i(t)$ is constant for $t \in [t', t'']$ and increases for t > t''. Then, it is strictly better for player j to exit at t rather at any point in an interval shortly after t''. However, if $g_j(t) = 0$ for $t \in [t'', t'' + \varepsilon]$, it is strictly better for player i to exit at t'' rather than in this interval, which is a contradiction. This means that G_i must be strictly increasing for $t > \tau$ until until no weak player remains in the game.

To show that G_i must be continuous, suppose that G_i makes a discrete jump at some $t > \tau_2$. Then, the other player strictly prefers to wait and exit at t + dt rather than at any time shortly before t, which contradicts that G_i is strictly increasing.

These results suggest that in any continuation equilibrium after τ_2 , a weak player must be held indifferent, i.e.,

$$\frac{g_i(t)}{1 - G_i(t)} = r\rho_j - p_j(t)\lambda_j$$

Integrating both sides yields

$$\ln[1 - G_i(\tau_2)] - \ln[1 - G_i(t)] = r\rho_j(t - \tau_2) + \ln Q_j(t) - \ln Q_j(\tau_2).$$

Alternatively, this can be written as

$$G_i(t) = 1 - e^{-r\rho_j(t-\tau_2)} (1 - G_i(\tau_2)) \frac{Q_j(\tau_2)}{Q_j(t)},$$

which shows the necessary optimality condition that must be satisfied for $t > \tau_2$.

Proof of Lemma 4. Since G_i is a distribution function, it must be weakly increasing. Note also that $W_i(t)$ is strictly decreasing by definition, given that a player can only switch from weak to strong, but not the other way around. Moreover, for any $G_i^* \in [0, W_i^*]$,

$$\lim_{t \to \infty} G_i(t) = 1 > \lim_{t \to \infty} W_i(t).$$

there must exist a unique $T_i(G_i^*) < \infty$.

To show that T_i is decreasing in G_i^* , observe first that

$$\frac{\partial G_i(T_i)}{\partial T_i}\frac{\partial T_i}{\partial G_i^*} + \frac{\partial G_i(T_i)}{\partial G_i^*} = \frac{\partial W_i(T_i)}{\partial T_i}\frac{\partial T_i}{\partial G_i^*} + \frac{\partial W_i(T_i)}{\partial G_i^*}$$

Since $\frac{\partial G_i(T_i)}{\partial T_i} - \frac{\partial W_i(T_i)}{\partial T_i} > 0$, we need to show

$$\frac{\partial T_i}{\partial G_i^*} < 0 \ \Leftrightarrow \ \frac{\partial G_i(T_i)}{\partial G_i^*} > \frac{\partial W_i(T_i)}{\partial G_i^*},$$

which can be written as

$$\phi_j(T_i) > p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) - \int_{\tau_2}^{T_i} \Phi_j(s) (1 - e^{-\lambda(T_i - s)}) ds.$$
(9)

Note that $W_i(T_i) = G_i(T_i)$ by definition, i.e.,

$$1 - (1 - G_i^*)\phi_j(T_i) = Q_i(t) + G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)}) + (1 - G_i^*) \int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds,$$

which can be written as

$$\int_{\tau_2}^{T_i} \Phi_j(s) p_i(s) (1 - e^{-\lambda_i(T_i - s)}) ds = \frac{1 - Q_i(t) - G_i^* p_i^* (1 - e^{-\lambda_i(T_i - \tau_i)})}{1 - G_i^*} - \phi_j(T_j).$$

Plugging this into (9) yields

$$\frac{1 - Q_i(t) - G_i^* p_i^* (1 - e^{-\lambda_i (T_i - \tau_i)})}{1 - G_i^*} > p_i^* (1 - e^{-\lambda_i (T_i - \tau_2)}),$$

which is simplified to

$$1 - Q_i(T_i) > p_i^* (1 - e^{-\lambda_i(T_i - \tau_2)}).$$

This condition can be written as

$$\alpha_i(1-\eta_i e^{-\lambda_i T_i}) > \frac{\alpha_i \eta_i e^{-\lambda_i T_i}}{1-\alpha_i + \alpha_i \eta_i e^{-\lambda_i T_i}} (1-e^{-\lambda_i (T_i-\tau_2)}),$$

which holds because $\alpha_i \geq \frac{\alpha_i \eta_i e^{-\lambda_i T_i}}{1-\alpha_i + \alpha_i \eta_i e^{-\lambda_i T_i}}$ and $1 - \eta_i e^{-\lambda_i T_i} > 1 - e^{-\lambda_i (T_i - \tau_2)}$.