Dynamic Survival Bias in Learning from Doubly Censored Signals

Wanyi Chen*

March 31, 2019

A Preliminary Draft

Abstract

This paper introduces a survival bias that emerges when we seek to learn about a random payoff from observing a sequence of doubly censored signals. The sequence of signals can only be observed if it has always been good enough through out its path, namely a signal can only be observed if its precedents sequentially and jointly passes a threshold in each period of time. For example when the signals come from another agent's optimal stopping problem in drawing the random payoff. The second censoring is a truncation in time such that the we only see the more recent signals but not the historical ones. There is a survival bias if the we don't fully consider the unobserved historical signals and the first censoring mechanism. The naive Bayesian updating is also flawed in its commuting property and its sufficient statistics. This bias is dynamic in the sense that the inference about the un-observed signals evolves as the observation continues. This paper characterizes the correct posterior belief about the random payoff that fully considers the censoring process and thus is adjusted for the bias.

Keywords: Optimal stopping; Bayesian Learning; Survival Bias

*I thank Lones Smith, William Sandholm, Marzena Rostek and Daniel Quint for conversations and feedback.

1 Introduction

In many economics scenarios, we need to learn about a random payoff but don't have access to the experimentation, or, the signal generating device. We can only observe the signals from others' experimentations. This type of "second hand learning" always face two kinds of censoring. Fisrtly, the experimentation process which generates signals about the random payoff only continues when the results satisfy certain conditions. In addition, we usually cannot observe the full history of the signal realizations but only the more recent ones.

There are three foco cases that are of this nature. In the first case, the signal generating process is an agent's trials of the random payoff. Typically, the agent faces an optimal stopping problem [1] - he incurs an experimentation cost but gathers a payoff and the information gain determined by the realization of each trial. Specifically, If the reward to the experimenter is the realization of the random payoff each time then it is a classical one-arm bandit problem[2]. The first censoring is done by his stopping rule. For instance, a firm learns about its own capacity through the success and failure of its products (J Miklós-Thal, 2018). The firm may exit the market if it find itself less competent. Consumers also evaluate a firm's capability through its products, but they don't usually consider, or are not aware of the full historical product of the firm. For example the younger generation is not familiar with the once popular iPod by Apple, and walkman by Sony. Their judgement may result in a bias if they do not consider the firm history.

In the second case, the signals come from difference sources. There's a third party that filter the signals. For example, when a consumer considers buying a product online, he sometimes learns about its quality from reading previous buyers' reviews, usually only the several recent ones. Here comes the problem - many online shopping platforms rank products by average review results. This implies that a product is only "visible" to the consumer when it has sufficient good reviews. Given this mechanism, with the same quality, a "visible" product is more likely to have earlier good reviews and later bad reviews. Thus only base one's evaluation on the most recent reviews may lead to underestimating about the the product quality.

In the third case, the signals are from a single agent but the selection is made by a third party. For instance, the Chinese college admission system relies heavily on the

[1] Wald (1947)

[2] As formulated in Gittins (1979)

one-time national college entrance exam. Students from all backgrounds compete in this exam. However students from low-income families are more likely to dropout if they perceive poor performance before they take this exam. Thus, judged from the same final score, students from low income family are more likely to have better historical scores. But this inequality is not taken account of.

Knowing the mechanism of the first censoring, we should updates our belief about the random payoff according to both the observed trial results and the inference of unobserved trial results that are consistent with the on-going experimention. We call this belief the *adjusted posterior*.

The partial information problem may stem from many causes, for instance old historical data are lost, or have limited access. For example when Moody's[3] evaluates a firm, it usually only acquires the most recent two year's finantial statemtents. There are also cases where it's required to only utilize partial data. For example the FDA conducts three phases of clinical trials before a new drug can be marketed. But the information in the physician labeling is only based on the phase 3 data[4]. In this case the consumers may get a biased (conservative) information on the efficacy of the medicine. If the FDA approval policies must be based on observed experiments, then this rule mimics a naive Baysian updating, and thus leads to an excessively conservative drug adoption.

The first goal of this paper is to study the properties of the dynamic survival bias. For simplicity of illustrating this situation, we assume the signals are i.i.d. Bernoulli. The experimenter's trial results are thus a sequence of successes and failures. It's interesting to consider general information structure other than Bernoulli trials and we'll show in the main text that the logic for our main results about the dynamic survival bias applies to more general settings. Yet it's also our goal study the pattern of the information, which degenerates to the path-dependent property in binary signals and delievers tractability.

Essentially, the selection bias exists because the first censoring results in a pattern in the observable data while the second censoring results in a non-representative sample of the full data set. . Since the experimenter will stop experimenting when sufficiently pessimistic about the payoff probability and earlier failures larger drop in the experimenter's incentive to continue experimenting. Therefore, it is more likely that the total

[3]The Moody's KMV EDF RiskCalc v3.1 Model,
https://www.moodys.com/sites/products/ProductAttachments/RiskCalc%203.1%20Whitepaper.pdf

[4]Drug Development and Review Definitions in the Investigational New Drug Application. `https://www.fda.gov/Drugs/DevelopmentApprovalProcess/HowDrugsareDevelopedandApproved/ApprovalApplications/InvestigationalNewDrugINDApplication/ucm176522.htm`

trial history contains early successes and later failures. This logic also makes the pattern of the observed trial sequence important. In the basic Bayesian updating process, the number of successes and failures in the observation is a sufficient statistics. However, to derive the adjusted posterior, the analyst needs to keep track of the ordering of signal realizations. Specially, a signal realizations sequence with earlier failures suggests better inference of pre-observation trial results, and thus implies a higher adjusted posterior. Intuitively, suppose the pre-observation history already contains many failures, then observing earlier failures implies that the experimenter has to endure consecutive failures, which is a less likely situation. This also justifies the dynamics property of the survival bias, as different observation leads to different inference of the unobserved trial results. The ordering of successes and failures can be more important than the number of them. Signal realization sequences with more failures but earlier failures can imply higher posterior. In addition, the pattern of information also impacts whether the survival bias is larger when the survival time is (stochastically) longer. When the observed sequence of trial results contains early successes and large success numbers, the survival bias accumulates as the survival time gets (stochastically) longer. However, this is not necessarily true when the observation contains early failures and large failure numbers, since in this case, we are more certain that the pre-observation is good when it is shorter.

There has been lots of literatures that also consider the learning in stopping problems. For instance, the applications of the Bandit problem in a multi-agent setting. Bolton, Patrick and Harris, Christopher (1999) studies a model in which different agents play bandits with the same payoff distribution and can observe the other's trial results (signals). In this case, information is public good and their focus is the free riding problem. Rosenberg, Dinah and Solan, Eilon and Vieille, Nicolas (2008) studies a similar setting where the agents can only observe the actions of continuing or stopping playing the bandit, but not their trial results. On the other hand, there's also the applications of the Bandit problem in the contract problem framework. For example, Gomes, Renato and Gottlieb, Daniel and Maestri, Lucas (2016) studies the screening and learning where the analyst sees the information but not the action of the experimenter. In all cases, the available information to the observer is either all of the signals, all of the actions, or both. Instead, I assume a natural situation where both action and signals are partially observable and emphasize the learning in a dynamic setting. Another branch of literature my work fits into is the social learning problem. In Smith, Lones and Sørensen, Peter and Tian, Jianrong (2012), they argue that the information herding problem stems from the incomplete learning from a bandit player who forgets previous signals and ignores the information impact of his actions. Applying this idea to my model, I study the optimal Bayesian learning of a partially forgetful experimenter who can only keeps track on recent trial histories.

2 Model Setting

Consider a random payoff P needed to be learned. There's a sequence of i.i.d binary signals $s \in \{0, 1\}$ where $\Pr(s = 1) = P$. We call the realization "1" as the high signal and "0" as the low signal.

Assume that the trials start at period $-T$ and is currently at period Z where T and Z are known positive integers. Denote the signal realization in period t by s_t and the history of signals up to Z by $\langle s_{-T} \ldots s_0 \ldots s_Z \rangle$, or $\langle s_t \rangle_{t=-T}^{Z}$.

A. The First Censoring: There is a function $I(\cdot, \cdot) : \mathbb{Z} \times \mathbb{Z} \to \mathbb{R}$ that maps binary sequences of arbitrary length into $\mathbb{R}$ and we call it the index. Let $n_1(\langle s_t \rangle_{t=-T}^{Z}) = \sum_{t=-T}^{Z} s_t$ and $n_0(\langle s_t \rangle_{t=-T}^{Z}) = T + Z + 1 - \sum_{t=-T}^{Z} s_t$ counts the number of high and low signals in the history $\langle s_t \rangle_{t=-T}^{Z}$. Assume that a signal sequence $\langle s_t \rangle_{t=-T}^{Z}$ is only observable at period t if it passes the index $I(\cdot)$ along its path up to t, namely

$$I(n_1(\langle s_t \rangle_{t=-T}^{Z}), n_0(\langle s_t \rangle_{t=-T}^{Z})) \geq 0$$

for $k = -T, \ldots t$.

Assumption: $I(n_1, n_0)$ weakly increases in n_1 and weakly decreases in n_0. $I(\cdot)$ is publically known. For simplicity, we sometimes abuse the notation and use $I(\langle s_t \rangle_{t=-T}^{Z})$ to denote $I(n_1(\langle s_t \rangle_{t=-T}^{Z}), n_0(\langle s_t \rangle_{t=-T}^{Z}))$.

B. The Second Censoring:

We as an outsider-analyst only observed a *public history* starting from period 0. Denote it as H_Z. $H_Z = \langle s_t \rangle_{t=0}^{Z}$. We call the part that's not seen, $\langle s_t \rangle_{t=-T}^{-1}$ as the *pre-history* and denote it as H_0. When the public history length is not emphaszied, we sometimes omit the subscrpit in H_Z.

For example, there's a product of the product of unknown quality Q. Consumers write binary reviews - good review and bad review. Q is positively related to the probability P that a consumer will give the product good review. Products are ranked by their average review and a consumer usually only reads the first few pages of product reviews.

When the average review falls below certain level, this product becomes invisible. The consumer also only read the first page reviews, which are ranked by dates from the most recent to the most obsolete ones.

The performance measure of a student is V. With probability V he (she) will have good scores in a exam/quiz and with probability $1-V$ he or she has bad exam scores. There's an evaluation in each season to determine whether this student stays in this program.

Apart from perception constraint, there are also some policy constraint that exclude us from knowing or utilizing early trial results. For example FDA

For instance, the signals could come from the experiment of a single agent who faces an optimal stopping problem. Specifically, If the payoff to the experimenter is the trial results each time, then it's a bandit problem and $I(\cdot)$ is the Gittin's index. Or the signals are from different sources where there's a third party that chooses their visibility.

C. The Posterior Beliefs

We first consider as a benchmark the *naive posterior* based solely on the public history H and disregard the pre-history. Let $G(p|H)$ denote its cdf, namely the chance that the success chance $P \leq p$, given H alone.

Assume that the prior belief of P has continuous density $f_0(p)$. The naive posterior is computed from the basic Bayesian rule.

$$G(p|H) = \frac{\int_0^p q^{n_1(H)}(1-q)^{n_0(H)} f_0(q)\mathrm{d}q}{\int_0^1 p^{n_1(H)}(1-p)^{n_0(H)} f_0(p)\mathrm{d}p}$$

Intuitively, the naive posterior could be downwardly biased because it ignores the positive information that the signals pass the first censoring. Although this conclusion is correct, as will be shown formally in Section 3, the logic could be quite fuzzy. The signals are in ex-ante independent, but the first censoring creates an ex-post pattern in the signal sequence. Specifically it's more likely that the history contain early high signals and later low signals, or it would become invisible. Our observation being a later truncation in the signal sequence is a non-representative sample of the full signal sequence, thus creating a bias. This bias stems from the joint effect of the first and second censoring.

To correct this bias, we compose the *adjusted posterior* $F(p|H;I,T)$ that accounts for all possible pre-histories consistent with the index rule I. We first consider the problem with fixed pre-history length time T, for simplicity, let's write the adjusted posterior as $F(p|H;I)$. We explore the properties of the corrected posterior $F(p|H;I)$ and its difference from the naive posterior $G(p|H)$. We formally define the dynamic survival bias as the difference between the naive posterior and the adjusted posterior. We will also explain why this is dynamic in Section 3.

Since the naive posterior only considers the public history while the adjusted posterior considers all possible pre-histories that are consistent with the public history and the filter, the bias is essential the inference about the un-observed pre-history. To characterize the corrected posterior, we first explore the possible pre-hitories. Let $\mathcal{H}_0$ be the set of all pre-histories, namely it's the set of length T strings of 0's and 1's. We use '+" between sequences denote the operation of concatenation (E.g. $\langle 10\rangle + \langle 01\rangle = \langle 1001\rangle$). We say a pre-history is admissible with public history H if it passes the index rule along its path after been cancatenated by the history H. Let $\Gamma(H|I)$ denote the set of admissible pre-histories for public history H with the index I,

$$\Gamma(H;I) = \{\langle s_t\rangle_{t=-T}^{-1} \in \mathcal{H}_0 | I(\langle s_t\rangle_{t=-T}^{k}) \geq 0, \text{ for all } k = -T,\ldots,Z \text{ and } H = \langle s_t\rangle_{t=0}^{Z}\} \tag{1}$$

Specifically, when $H = \varnothing$,

$$\Gamma(\varnothing|I) = \{\langle s_t\rangle_{t=-T}^{-1} \in \mathcal{H}_0 | I(\langle s_t\rangle_{t=-T}^{k}) \geq 0, \text{ for all } k = -T,\ldots,-1\} \tag{2}$$

For any $H \in \mathcal{H}$, we have $\Gamma_H \subseteq \Gamma_\varnothing$. This property that Γ depends on H is the essential cause for the survival bias to be dynamic, which we will discuss in the next section.

$$F(p|H;I) = \frac{\int_0^p \sum_{H_0\in\Gamma(H;I)} q^{\sigma(H_0+H)}(1-q)^{\phi(H_0+H)} f_0(q)\mathrm{d}q}{\int_0^1 \sum_{H_0\in\Gamma(H;I)} p^{\sigma(H_0+H)}(1-p)^{\phi(H_0+H)} f_0(p)\mathrm{d}p} \tag{3}$$

Generally it's very hard to compute, even for a very simple forms of I. For example, when $I(H) = n_1(H)/(n_1(H)+n_0(H)) - 1/2$, namely the average of signals realizations has to be always greater than $1/2$ for the signal generating process to continue, or the number of high signals has to be greater than the number of low signals along the path.

$$F(p|H;I) = \frac{\int_0^p \sum_{s=\lceil T/2\rceil}^{T} \left(\sum_{m=0}^{T-s} c(m,2,r_1)c(T-s-m,2,r_2)\right) q^s(1-q)^f f_0(q)\mathrm{d}q}{\int_0^1 \sum_{s=\lceil T/2\rceil}^{T} \left(\sum_{m=0}^{T-s} c(m,2,r_1)c(T-s-m,2,r_2)\right) p^s(1-p)^f f_0(p)\mathrm{d}p}$$

where $c(m,\rho,r) = \frac{r}{m\rho+r}\binom{m\rho+r}{m}$ is the mth fuss-catalan number[5], $r_1(s) = |H| - 2\sigma(H) + 2s - T + 1$ and $r_2 = a - b + 1$.

Similarly, when $v = 2/3$.

$$F(p|H;I) = \frac{\int_0^p \sum_{s=\lceil 2T/3\rceil}^{T} \left(\sum_{m=0}^{T-s} c(m,3,r_1)c(T-s-m,3,r_2)\right) z^s(1-z)^f \mathrm{d}\Upsilon(z)}{\int_0^1 \sum_{s=\lceil 2T/3\rceil}^{T} \left(\sum_{m=0}^{T-s} c(m,3,r_1)c(T-s-m,3,r_2)\right) z^s(1-z)^f \mathrm{d}\Upsilon(z)}$$

where $r_1(s) = 2|H| - 3\sigma(H) + \mathrm{mod}(T,3) + 3s - 2(T-1)$, $r_2 = a - b + 2$ and $\delta \leq 0.72$.

One may wonder whether is it worth to run these complicated computations to get the adjusted posterior. Next, we show in which aspects are the naive posterior flawed and why it's critical to employ the adjusted posterior in some cases.

3 Dynamic Survival Bias

Firstly, we explore how does different admisible pre-history affect the adjusted posterior in the likelihood ratio order[6] $\succeq_{\mathrm{lr}}$.

Lemma 1. *The adjusted posterior falls when the admissible pre-history set gets larger: for all public histories H,*

$$\Gamma(H|I) \subseteq \Gamma(H|I') \Rightarrow F(\cdot|H;I) \succeq_{lr} F(\cdot|H;I')$$

Since the first censoring is a quasi-right censoring, when the pre-history set grows larger, it only includes more unfavorable pre-histories and thus make the adjusted posterior lower. A formal proof is given in the appendix.

We can regard the naive posterior as the corrected posterior with an index rule that allows for all prehistories, namely $G(\cdot|H) = F(\cdot|H;I_0)$ where $\Gamma(H|I_0) = \mathcal{H}_0$, the fact that $F(\cdot|H) \succeq_{\mathrm{lr}} G(\cdot|H)$ follows easily from Lemma 1 as $\Gamma(H|I) \subseteq \mathcal{H}_0$ for all I.

Corollary 1. *The adjusted posterior dominates the naive posterior in the likelihood ratio order: for all public histories H,*

$$F(\cdot|H;I) \succeq_{lr} G(\cdot|H)$$

[5] or the dyck path numbers

[6] $F_2 \succeq_{\mathrm{lr}} F_1$ if $F_2'(p)/F_1'(p)$ increases in $p \in [0,1]$. Note that the likelihood ratio order is a strong notion of stochastic dominance. It implies the first order stochastic dominance.

Since the corrected posterior is always higher than the naive poserior, we measure this survival bias by their difference with the L^1 norm:

$$B(H; I) = \int_0^1 |F(p|H; I) - G(p|H)| dp$$

More inportantly, this bias is dynamic as our inference about the pre-history changes over time. We first show this property by an exmaple

Example 1. Consider index $I(n_1, n_0) = n_1 + 1 - n_0$ and pre-history length $T = 2$. Suppose we see the public history $H_1 = \langle 0 \rangle$, then from Figure we can see that $\Gamma(H_2) = \{\langle 11 \rangle, \langle 10 \rangle\}$. Next, let's continue for 1 period of observation and suppose we see another $\langle 0 \rangle$ namley $H_2 = \langle 00 \rangle$, then $\Gamma(H_2) = \{\langle 11 \rangle\}$.

Although this extreme examples exploit the discrete nature of this problem, but te general intuition is that discovering subsequent low signal may exclude more bad pathes in the admissible history, compensating for the additional single low signal.

Another important message from this example is that additional low signals doesn't always reduce our estimate. Suppose our prior is uniform, then the pdfs

$$f(p|\langle 0 \rangle) \propto 2p^2(1-p)^2 + 3p^3(1-p)$$

$$f(p|\langle 00 \rangle) \propto p^3(1-p)^2$$

We have $F(p|\langle 00 \rangle) \succ_2 F(p|\langle 0 \rangle)$ where $\succ_2$ stands for the second order stochastic dominance. This also breaks down the result from naive Bayesian updating which always give higher estimate with more high signals and worse estimate with more low signals. Next, we show in the survival bias, setting, when is this rule broken down.

We say sequence H'_Z *lexicographically* precedes H_Z if $n(H')$ *****. For example, $\langle 011 \rangle \sqsubseteq_1 \langle 101 \rangle \sqsubseteq_1 \langle 110 \rangle$.

Intuitively, starting from any state (σ, ϕ), a following trial results sequence that contains earlier failures always reaches states with lower index than a sequence that contains later failures. This can be seen in the exemplary table of Gittins indexes in Figures 1. Each of the grids represents a state. As the experiment goes on, states $(\sigma_t(H_0 + H), \phi_t(H_0 + H))$, $t = -T, \ldots, -1, 0, 1, \ldots$ forms a path in this table. A success leads the path to the right and a failure leads the path down. Thus earlier failures make the experimenter more likely to stop. Therefore, observing earlier failures, while the

Changes in the Possible Total Histories with Different Public History

(a)

Number of Successes

Number of Failures	0	1	2	3	4	5	6	7
0	0.50	0.67	0.75	0.80	0.83	0.86	0.88	0.89
1	0.33	0.50	0.60	0.67	0.71	0.75	0.78	0.80
2	0.25	0.40	0.50	0.57	0.63	0.67	0.70	0.73
3	0.20	0.33	0.43	0.50	0.56	0.60	0.64	0.67
4	0.17	0.29	0.38	0.44	0.50	0.55	0.58	0.62
5	0.14	0.25	0.33	0.40	0.45	0.50	0.54	0.57
6	0.13	0.22	0.30	0.36	0.42	0.46	0.50	0.53
7	0.11	0.20	0.27	0.33	0.38	0.43	0.47	0.50

(b)

Number of Successes

Number of Failures	0	1	2	3	4	5	6	7
0	0.50	0.67	0.75	0.80	0.83	0.86	0.88	0.89
1	0.33	0.50	0.60	0.67	0.71	0.75	0.78	0.80
2	0.25	0.40	0.50	0.57	0.63	0.67	0.70	0.73
3	0.20	0.33	0.43	0.50	0.56	0.60	0.64	0.67
4	0.17	0.29	0.38	0.44	0.50	0.55	0.58	0.62
5	0.14	0.25	0.33	0.40	0.45	0.50	0.54	0.57
6	0.13	0.22	0.30	0.36	0.42	0.46	0.50	0.53
7	0.11	0.20	0.27	0.33	0.38	0.43	0.47	0.50

Figure 1:
These tables show the Gittins Indexes for each combination of success and failure numbers. The blue areas in (a) are the ending states corresponding to all possible total histories of the experimenter from whose trials the public history is $\langle 10 \rangle$. The blue areas in (b) are the ending states corresponding to all possible total histories of the experimenter from whose trials the public history is $\langle 01 \rangle$. The red area highlights the states that are possible with public history $\langle 10 \rangle$ but not with $\langle 01 \rangle$.

experimenter continues experimenting suggests that the pre-observation history must be good enough for the experimenter to endure the early failures.

Figures 1 shows the ending states $(\sigma_{|H|}(H_0 + H), \phi_{|H|}(H_0 + H))$ of all possible total histories $H_0 + H \in \Gamma_H$ for an experimenter with uniform prior, discount factor $\delta = 0$, and cost $v = 0.5$, where the public history are $H_1 = \langle 10 \rangle$ and $H_2 = \langle 01 \rangle$ respectively. Note that with $\langle 01 \rangle$, the pre-histories with the worst posteriors are eliminated. Essentially, while the smaller is the continuation set Γ_H, the higher is the adjusted posterior.

4 Dynamic Feature

Corollary 2. *The adjusted posterior rises when H is permuted earlier in the lexicographic order:*

$$H' \sqsubseteq H \Rightarrow F(\cdot|H') \succeq_{lr} F(\cdot|H)$$

Previously, we showed that the naive Bayesian updating results in a downward bias and breaks its monotone property in the number of high/low signals. Now we've shown that it is also flawed in treating the number of successes and failures in the history as a sufficient statistic. To compute the adjusted posterior, we need to keep track of the ordering of the public signals, namely the adjusted posterior is path-dependent.

Eg1. $T = 2$, $I = a + 2 - 2b$
$F(\cdot|\langle 0\rangle) \succeq_{\mathrm{lr}} F(\cdot|\langle 10\rangle)$
$f(\cdot|\langle 0\rangle) \sim p^2(1-p) = p^3(1-p) + p^2(1-p)^2$
$f(\cdot|\langle 10\rangle) \sim p^3(1-p) + 2p^2(1-p)^2$

Corollary 3. *The adjusted posterior rises, or the bias is larger when the index is more strict: for all public histories H.*

$$I' \leq I \Rightarrow F(\cdot|H; I') \succeq_{lr} F(\cdot|H; I)$$

In addition, assume the analyst knows the pre-history length is 2**.** The only possible pre-histories are $\langle 10\rangle$ and $\langle 11\rangle$. The density $F'(p|\varnothing, \text{go})$ is thus proportional to $p^2 + p(1-p) = p$, the likelihood of the event that the experimenter is still going conditional on the realization of P being p. This dominates the uniform distribution in the likelihood ratio order as the likelihood ratio p is an increasing function of p.

Moreover, this bias is dynamic since the posterior on the pre-history changes with the public history, let's consider the setting in the previous example again. Now instead of a null public history, let's assume the public history is $\langle 1\rangle$, then the possible pre-histories are the same as before, $\langle 10\rangle$ and $\langle 11\rangle$. However if the public history is $\langle 0\rangle$, the only possbile history is $\langle 11\rangle$, resulting in a posterior in period 0 that's proportional to p^2. We will address this later in section.

Now we have shown this comparative statics results regarding public signals with the same length. Considering public signals with different length is more complecated since we cannot fix both the starting time and ending time. To avoid the complication, we instead consider two special cases: 1. Fixing the pre-history length T, continue one period of observation; 2. Fixing the ending time, trace back one period of observation.

Lemma 2. *(a)discovering a subsequent high signal increases the adjusted posterior:*

$$F(\cdot|H + \langle 1\rangle; I) \succeq_{lr} F(\cdot|H; I)$$

This result is obvious since discovering a subsequent high signal does not change the inference of the pre-history, namely $\Gamma(H|I) = \Gamma(H + 1|I)$, the additional good signal only raises the adjusted posterior. However the opposite does not necessarily hold for discovering a subsequent low signal. Example 1 is a counter example.

The adjusted Posterior Rises in the Pre-history Length

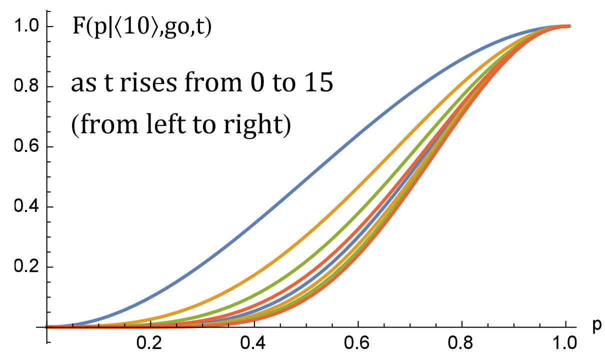

Figure 2:
This figrues plots $F(p|\langle 10\rangle, \mathrm{go}; t)$ with various t when the experimenter is myopic, has uniform prior and outside option valua 0.5.

Claim2. Now since time is not fixed, we resume the full expression of F that includes the parameter T.

Lemma 3. *$F(\cdot|H; I, T-1) \succeq_{lr} F(\cdot|\langle 0\rangle + H; I, T)$ and $F(\cdot|\langle 1\rangle + H; I, T-1) \succeq_{lr} F(\cdot|H; I, T)$, namely discovering a precedent high signal will increase the adjusted posterior, similarly, discovering a precedent low signal will decrease the adjusted posterior.*

The difference in discovering subsequent and precedent signals implies the importance of timing of the observation. It's intuitive to say that the adjusted posterior should rise if the pre-history length increases as the longer is the pre-history length, the signals pass more cencoring stages. However this intuition is only correct when the public history is favorable enough.

Proposition 1. *The adjusted posterior rises in the pre-history length: when $H \sqsupseteq \bar{H}$ or $H \sqsupseteq \bar{H} + \langle 1 \ldots\rangle$ for some $\bar{H}$,*

$$T' > T \Rightarrow F(\cdot|H; T') \succeq_{lr} F(\cdot|H; T)$$

When the public history is favorable enough, the survival bias accumulates as the survival time grows longer. We show the magnitude of the dynamic survival bias by plotting $F(p|H, \mathrm{go}, t)$ varying the certain pre-history length t in Figure 2.

However, note that this property doesn't necessarily hold for public histories containing early failures or too many large failures compared to successes. In this case, we are certain about pre-history successes when the history length is short and this certainty is lost when the pre-history is long. For instance, if we see 10 failures and we know there's only one un-observed trial before that, we are pretty sure the un-observed trial is a success. However when we know there are 100 previous trials, we can hardly make judgements about them.

5 Summary and Discussion

This paper studies the survival bias when there are two types of truncations of a"panel data" (or a group of sequential signals). On the cross-section dimension, sample data are only observable when they jointly pass certain filter, other than pass this filter individually as in usual survival bias problems. On the time dimension, the truncation is due to the observer's perception constraint. The first truncation results in a pattern in the observable data while the second truncation results in a non-representative sample of the full data set. A survival bias emerges and this bias is dynamic in the sense that a different observation affects the inference of the un-observed pre-history. And the inference is path-dependent.

6 Appendix

Specify of the Gittins Index and the adjusted Posterior

Let Ω_H denote the set of states $(c_1(H), c_0(H))$ corresponding to the histories in the continuation set Γ_H, namely

$$\Omega_H = \{(n_1(H_0), n_0(H_0))|H_0 \in \Gamma_H\}$$

Assume state (n_1, n_0) is the n_1, n_0th vertex in a 2-dimensional regular grid. Then Ω is a quasi-upper-triangular region in the square grid graph of (i, j) as $I(i, j)$ rises in i and falls in j. We can partition it into $\Omega = \cup_{i=b(T)}^{T}\{(i, T-i)\}$ where $b(T+1) = b(T)+1$ or $b(T+1) = b(T)$. Let $N(i, j|\Omega)$ be the number of staircase walks whose pathes are entirely inside Ω from $(0, 0)$ to (i, j). Note that $N(i, j|\Omega)$ is the number of admissible pre-histories with success number i and failure number j. $N(i, j|\Omega) = 0$ when $(i, j) \notin \Omega$, Note that the state transition can only take the form $(i, j) \to (i+1, j)$ and $(i, j) \to (i, j+1)$.

$$N(i, j|\Omega) = N(i-1, j|\Omega) + N(i, j-1|\Omega)$$

when $(i,j)\in\Omega$ and $(i-1,j)\in\Omega$.

Incorparating this, we can write (**??**) and (**??**) as

$$\begin{aligned} F(p|H;I) &= \frac{\int_0^p \sum_{(i,j)\in\Omega_H(I)} z^{i+\sigma(H)}(1-z)^{j+\phi(H)}\mathrm{d}z}{\int_0^1 \sum_{(i,j)\in\Omega_H(I)} z^{i+\sigma(H)}(1-z)^{j+\phi(H)}\mathrm{d}z} \\ &= \frac{\int_0^p \sum_{T=0}^{\infty}\sum_{i=b(T)}^{T} N(i-\sigma(\mathrm{H}), j-\phi(\mathrm{H})|\Omega_H) z^i(1-z)^{T-j}\mathrm{d}z}{\int_0^1 \sum_{T=0}^{\infty}\sum_{i=b(T)}^{T} N(i-\sigma(\mathrm{H}), j-\phi(\mathrm{H})|\Omega_H) z^i(1-z)^{T-j}\mathrm{d}z} \end{aligned} \tag{4}$$

Proofs of the Main Results

Lemma 1, $F(p|H;I)$ falls falls in $\Gamma(H;I)$

Proof of Lemma 1:

Step 2. From Claim 2, we prove $F(p|H;I)\succeq_{\mathrm{lr}} F(p|H;I')$ by showing that for $\forall j>i$,

$$N(j,T-i|\Omega)/N(i,T-i|\Omega)\geq \binom{T}{j}\Big/\binom{T}{i} \tag{5}$$

Obviously when $0\leq i<i'\leq b(T)$, the desired inequality holds as $N(i,T-i|\Omega)=0$.

Note that

$$N(i,j|\mathbb{N}^2)=\binom{i+j}{i} \text{ and } \frac{N(i',T-i'|\mathbb{N}^2)}{N(i,T-i|\mathbb{N}^2)}=\binom{T}{i'}\Big/\binom{T}{i}$$

We show (6) by showing that $N(j,T-j|\Omega)/N(i,T-i|\Omega)$ falls in the "size" of Ω:

Claim 1. If $\Omega'=\Omega\cup\{(a,b)\}$ for $\forall c\in\mathbb{N}$ and $i<i'$, we have

$$N(i',c-i'|\Omega')N(i,c-i|\Omega)\leq N(i,c-i|\Omega')N(i',c-i'|\Omega) \tag{6}$$

For non-triviality, assume $b\geq 1$ and (a,b) is on the lower boundary of Ω.

Firstly, as $N(a,b|\Omega)=0<N(a,b|\Omega')$ and $N(a-k,b+k|\Omega)=N(a-k,b+k|\Omega')$, for all $k\neq 0$, we have (6) for all $c=a+b$.

Next, since for arbitrary non-negative numbers A,B,C,D,E,F satisfying $AD\leq BC$ and $CF<DE$, we have $(A+C)(D+F)\leq(B+D)(C+E)$, we have (6) for $c=a+b+1$ by Claim **??**.

Finally, by induction, (6) holds for $c=T$ as desired.

Claim 2. Suppose a sequence of differentiable probability density functions $f_1, f_2 \dots f_n$ satisfy $f_j'/f_j > f_i'/f_i$ for $\forall i < j$. Let $g^k = \sum_{i=1}^n a_i^k f_i / \sum_{i=1}^n a_i^k$, $k = 1, 2$. When $a_j^1/a_i^1 > a_j^2/a_i^2$ for all $j > i$, we have $g^2 \succeq_{\text{lr}} g^1$. This follows easily from the fact that

$$\left(\frac{g^2}{g^1}\right)' = \frac{\sum_{i=1}^n a_i^1 f_i' \sum_{j=1}^n a_j^2 f_j - \sum_{i=1}^n a_i^1 f_i \sum_{j=1}^n a_j^2 f_j'}{(\sum_{j=1}^n a_j^2 f_i)^2}$$

Step 1.

Proof of Proposition 2

when $H' \geq H$, $\Omega_T(H|I) \supseteq \Omega_T(H'|I)$ for any I and T

Proof of Proposition 4

when $I' \geq I$, $\Omega_T(H|I) \supseteq \Omega_T(H|I')$ for any H and T

Proof of Proposition ??:

Step 1. For simplicity we show the case where the analyst has a prior with continuous density function $f^0(z)$. The arguments can easily be applied to general cases.

When $H \succeq_1 \bar{H}$ or $H \succeq_1 \bar{H} + \langle 1 \dots \rangle$ for some $\bar{H} \in B(\Gamma_H)$, every sequence in $\Gamma_\varnothing$ is a possible pre-history, namely $\Gamma_H = \{H_0 + H | \ H_0 \in \Gamma_\varnothing\}$. Thus we only need to show $F(\cdot|\varnothing, \text{go}; T+1) \sqsupseteq_1 F(\cdot|\varnothing, \text{go}; T)$

Following Step 1 of the proof for Lemma 1, the pdf of the adjusted posterior is

$$\begin{aligned} f(p|\varepsilon, \text{go}; T) &= \sum_{i=b(T)}^{T} N(i, T-i|\Omega_\varnothing) p^i (1-p)^{T-i} f^0(p)/A & (7) \\ &= \sum_{i=b(T)}^{T} N(i, T-i|\Omega_\varnothing) p^i (1-p)^{T-i} (p+1-p) f^0(p)/A & (8) \\ &= \sum_{i=b(T)}^{T+1} \left(N(i, T-i|\Omega_\varnothing) + N(i-1, T-i+1|\Omega_\varnothing)\right) p^i (1-p)^{T+1-i} f^0(p)/A \end{aligned}$$

On the other hand,

$$f(p|\varepsilon, \text{go}; T+1) = \sum_{i=b(T+1)}^{T+1} N(i, T+1-i|\Omega_\varnothing) p^i (1-p)^{T+1-i} f^0(p)/A$$

From the quasi-upper-triangular shape of Ω, either $\bar{i}(T+1) = \bar{i}(T) + 1$ or $\bar{i}(T+1) = \bar{i}(T)$.

Step 2. Recall (**??**) that $N(i, T+1-i|\Omega) = N(i, T-i|\Omega) + N(i-1, T+1-i|\Omega)$. When $\bar{i}(T+1) = \bar{i}(T)$, we have $f(p|\varepsilon, \text{go}; T+1) = f(p|\varepsilon, \text{go}; T)$ for any p and thus $F(\varepsilon, \text{go}; T+1) = F(\varepsilon, \text{go}; T)$. When $b(T+1) = b(T)+1$, by Claim 2, $F(\cdot|\varnothing, \text{go}; T+1) \sqsupseteq_1 F(\cdot|\varnothing, \text{go}; T)$. Thus when $H \succeq_1 \bar{H}$ or $H \succeq_1 \bar{H} + \langle 1 \ldots \rangle$ for some $\bar{H} \in B(\Gamma_H)$, the adjusted posterior $F(\cdot|\varnothing, \text{go}; T)$ rises in T in the likelihood ratio order. Still by Claim 2, $F(\cdot|\varnothing, \text{go}; \gamma)$ rises in γ.

Formally, when $H \succeq_1 \bar{H}$ or $H \succeq_1 \bar{H} + \langle 1 \ldots \rangle$ for some $\bar{H} \in B(\Gamma_H)$, we can express (8) alternatively as

$$f(p|\varepsilon, \text{go}; T) = \sum_{(i,j)\in\Omega_\varnothing} N(i, j|\Omega_\varnothing) p^i (1-p)^j f^0(p)/A$$

However when the condition does not hold,

$$f(p|\varepsilon, \text{go}; T) = \sum_{(i,j)\in\Omega_H} N(i, j|\Omega_\varnothing) p^i (1-p)^j f^0(p)/A$$

Note the difference in the set for summation. The arguments in Step 2 fails.

Proof of Proposition ??:

We first show this result fixing the pre-history length, namely

$$H' \sqsupseteq_1 H \Rightarrow F(\cdot|H', \text{go}; T) \succeq_2 F(\cdot|H, \text{go}; T)$$

From the semi-upper-triangular shape of Ω, obviously $\Omega_H \subset \Omega_{H'} \cup B(\Gamma_{H'})$. Let Ω_H^T denote the set of states corresponding to the histories in the continuation set with length T, then $\Omega_H \cup \{(i,j)|i+j = T\}$. Either $\Omega_H^T = \Omega_{H'}^T$ or $\Omega_{H'}^T = \Omega_{H'}^T \backslash \{(b(T), T-b(T))\}$. In the second case, $\Omega_{H'}^T$ can be derived from Ω_H^T by exclude the state with the most failures.

From (4),

$$F(p|H, \text{go}; T) = \frac{\int_0^p \sum_{i=b(T)}^T z^{i+\sigma(H)}(1-z)^{T-i+\phi(H)}\gamma^T \text{d}z}{\int_0^1 \sum_{i=b(T)}^T z^{i+\sigma(H)}(1-z)^{T-i+\phi(H)}\gamma^T \text{d}z}$$

and

$$F(p|H', \text{go}; T) = \frac{\int_0^p \sum_{i=\tilde{b}(T)}^T z^{i+\sigma(H)}(1-z)^{T-i+\phi(H)}\gamma^T \text{d}z}{\int_0^1 \sum_{i=\tilde{b}(T)}^T z^{i+\sigma(H)}(1-z)^{T-i+\phi(H)}\gamma^T \text{d}z}$$

where $\tilde{b}(T) = b(T)$ or $b(T)+1$.

By Claim 2, $F(p|H, \text{go}; T) \succeq_{\text{lr}} F(p|H', \text{go}; T)$ for any T.

Next, we consider states in the set $\Omega_H^T \backslash \Omega_{H'}^T$. From similar arguments as in the proof of Proposition **??**(b), the cdf function

$$\tilde{F}(p) = \frac{\int_0^p \sum_{H_0 \in \Omega_H^T \backslash \Omega_{H'}^T} z^{\sigma(H_0)+\sigma(H)}(1-z)^{\phi(H_0)+\phi(H)}\gamma^{\sigma(H_0)+\phi(H_0)}\mathrm{d}z}{\int_0^1 \sum_{H_0 \in \Omega_H^T \backslash \Omega_{H'}^T} z^{\sigma(H_0)+\sigma(H)}(1-z)^{\phi(H_0)+\phi(H)}\gamma^{\sigma(H_0)+\phi(H_0)}\mathrm{d}}$$

single crosses $G(p|H)$ at some point $\tilde{p}$ while $E[p|\tilde{F}] \leq E[p|G(\cdot|H)]$. We have

$$\int_{\tilde{p}}^1 \left(\tilde{F}(z) - G(z|H)\right)\mathrm{d}z \geq \int_0^{\tilde{p}} \left(G(z|H) - \tilde{F}(z)\right)\mathrm{d}z$$

Since $F(\cdot|H, \text{go}) \succeq_{\text{lr}} G(\cdot|H)$, thus $F(\cdot|H, \text{go}) \succeq_1 G(\cdot|H)$. Therefore there exists some $\hat{p} < \tilde{p}$ such that $\int_{\hat{p}}^1 \left(\tilde{F}(z) - F(z|H, \text{go})\right)\mathrm{d}z \geq \int_0^{\hat{p}} \left(F(z|H, \text{go}) - \tilde{F}(z)\right)\mathrm{d}z$, namely $F(\cdot|H, \text{go}) \succeq_2 \tilde{F}(\cdot)$. Since $F(p|H, \text{go})$ can be regarded as the average from $F(p|H', \text{go})$ and $\tilde{F}(p)$, we must have $F(\cdot|H', \text{go}) \succeq_2 F(\cdot|H', \text{go})$.

Proof of Proposition ??:

(a) When $H' = H + \langle 1 \rangle$, it is easy since $\Gamma_H = \Gamma_{H'}$. For the other cases, we first show this result fixing the total history length, namely

$$F(\cdot|H', \text{go}; T) \succeq_{\text{lr}} F(\cdot|H, \text{go}; T+1) \text{ if } H' = \langle 1 \rangle + H \text{ or } H' = H + \langle 1 \rangle$$

and

$$F(\cdot|H', \text{go}; T+1) \succeq_{\text{lr}} F(\cdot|H, \text{go}; T) \text{ if } H = \langle 0 \rangle + H' \text{ or } H = H' + \langle 0 \rangle$$

From Claim 2, we will get the desired result by showing that the weight of the good state is higher(lower) when we concatenate 1 (0) to a public history, namely

$$N(a+1, b|\Omega)/N(a, b|\Omega) < N(a, b+1|\Omega)/N(a-1, b+1|\Omega)$$

and

$$N(a, b+1|\Omega)/N(a, b|\Omega) > N(a, b+2|\Omega)/N(a-1, b+1|\Omega)$$

By Claim **??**, in both cases it suffices to show that for $\forall a, b$ and Ω,

$$N(a, b|\Omega)^2 \leq N(a-1, b+1|\Omega)N(a+1, b-2|\Omega) \tag{9}$$

Step 1, Let $n_1, n_2, \cdots n_k$ be a sequence of number such that $n_i n_j \geq n_{i-t} n_{j+t}$ for $\forall j \geq i$ and $t \leq \min\{i, k-j\}$. compose $m_1, m_2 \cdots m_{k-1}$ with some number K such that $m_i = n_i + n_{i+1}$ for $i > K$ and $m_i = 0$ for $i \leq K$. now we show $m_i m_j \geq m_{i-t} m_{j+t}$ for $\forall j \geq i$ and $t \leq \min\{i, k-1-j\}$.

If $i \leq K$, then $i - t \leq K$ and thus $m_i m_j \geq m_{i-t} m_{j+t}$ for $\forall j \geq i$ and $t \leq \min\{i, k-j\}$;

If $i > K$, then for $\forall j \geq i$ and $t \leq \min\{i, k-j\}$,

$$\begin{aligned} m_i m_j &= (n_i + n_{i+1})(n_j + n_{j+1}) = n_i n_j + n_{i+1} n_j + n_i n_{j+1} + n_{i+1} n_{j+1} \\ &\geq n_{i-t} n_{j+t} + n_{i-t} n_{j+t+1} + n_{i-t+1} n_{j+t} + n_{i+1-t} n_{j+1+t} \\ &= (n_{i-t} + n_{i-t+1})(n_{j+t} + n_{j+t+1}) \geq m_{i-t} m_{j+t} \end{aligned}$$

Step 2, for arbitrary a, b, Define

$$\Omega' = \bigcup_{T \leq a+b+1} \bigcup_{t \leq T} \{(a-t, b+t)\}$$

If $(i, j) \in \Omega$, Let $\tilde{N}_{i,j} = N(i, j|\Omega)$ and $N_{i,j} = N_{i-1,j} + N_{i,j-1}$; if $(i, j) \in \Omega' \backslash \Omega$, let $\tilde{N}_{i,j} = 0$.

By Step 1, let $n_i = \tilde{N}_{-b-1+i, b+1-i}$ and $k = a + b + 1$, since $N_{00} = 1$ and $N_{ij} = 0$ for $i \neq 0$ and $j \neq 0$, we have $n_i n_j \geq n_{i-t} n_{j+t}$ for $\forall j \geq i$ and $t \leq \min\{i, k-j\}$. By induction, (9) holds for $\forall a, b$ and Ω.

(b) From proposition **??**(b), we will ge the desired result by showing that the weight of states corresponding to longer histories are higher when the public history contain more failures or less successes. By the expression (**??**), this means for any (i, j), $(i', j') \in \Omega_H^B$ where $i \leq i'$, $j \leq j'$.

$$\frac{N(i' + 1 - \sigma(H), j' - \phi(H)|\Omega)}{N(i + 1 - \sigma(H), j - \phi(H)|\Omega)} > \frac{N(i' - \sigma(H), j' - \phi(H)|\Omega)}{N(i - \sigma(H), j - \phi(H)|\Omega)} \tag{10}$$

and

$$\frac{N(i' - \sigma(H), j' - \phi(H) + 1|\Omega)}{N(i - \sigma(H), j - \phi(H) + 1|\Omega)} < \frac{N(i' - \sigma(H), j' - \phi(H)|\Omega)}{N(i - \sigma(H), j - \phi(H)|\Omega)} \tag{11}$$

Step 1. We first show the case when the experimenter is impatient, the cost $v = k/(k+1)$ for some positive integer k and the prior is Beta$(k \cdot n, n)$ for any positive integer n. In these cases, the continuation set of states $\Omega_0 = \bigcup_{t=1,2,3\ldots} \{(i, j)|i \geq k \cdot j\}$ and $\Omega_H = \bigcup_{t=1,2,3\ldots} \{(i, j)|i + \sigma(H) \geq k(j + \phi(H))\}$. By Aval, Jean-Christophe (2008). $N(i, j|\Omega_0)$ can be expressed as the jth fuss-catalan number with parameter $\rho = k + 1$

and $r = i - k \cdot j + 1$, namely

$$N(i,j|\Omega_0) = c(j,k+1,i) = \frac{i-k\cdot j+1}{j+i+1}\binom{j+i+1}{j} \tag{12}$$

Now we show (10) for any $i' = i+a$, $j' = j+1$ and $i \geq k \cdot j$. where[7] $a \leq k$. The left hand side of (10) is

$$\frac{N(i+1,j|\Omega_0)}{N(i,j|\Omega_0)} = \frac{j+i+2}{i+2} \cdot \frac{i-k\cdot j+2}{i-k\cdot j+1} \cdot \frac{j+i+1}{j+i+2} \tag{13}$$

The right hand side of (10) is

$$\frac{N(i+1+k,j+1,|\Omega_0)}{N(i+k,j+1|\Omega_0)} = \frac{j+i+a+3}{i+a+2} \cdot \frac{i+a-k-k\cdot j+3}{i+a-k-k\cdot j+2} \cdot \frac{j+i+a+2}{j+i+a+3} \tag{14}$$

Obviously the third term in (13) is smaller than (14). Then let's consider the second term. Since $i \geq k \cdot j \geq a \cdot j$, $(j+i+2)/(i+2) < (a+1)/a$ and thus

$$\frac{j+i+2}{i+2} < \frac{j+i+2+a+1}{i+2+a} = \frac{j+i+a+3}{i+a+2}$$

Similarly the first term in (13) is smaller than (14). Thus (10) holds.

Step 2. In the general case, there exists some $0 < k_1 < k_2$, so that the continuation set

$$\underset{t=1,2,3...}{\cup}\{(i,j)|i \geq k_2 \cdot j\} \subset \Omega_0 \subset \underset{t=1,2,3...}{\cup}\{(i,j)|i \geq k_1 \cdot j\}$$

Fixing j, when i is sufficiently large, the four terms in (10) approaches (12) for some equal $k \in (K_1, K_2)$. Since tha ratio of (13) to (14) is strictly less than 1, (10) holds. To be more specific, there exists $\bar{i}(i,j)$ such that (10) holds when $i > \bar{i}(i,j)$. Let $\bar{\phi} = \sup\{\arg\min_{\phi}(\bar{i}(i,j), j+\phi) \notin \Omega_0\}$ and then when $\phi(H) > \bar{\phi}$, for each $(i,j) \in \Omega_H$, $i > \bar{i}(i,j)$ and (10) holds.

Proof of Proposition ??:

We first show that $\Omega(\Upsilon,\delta,v) \subseteq \Omega(\Upsilon',\delta',v')$ if $\Upsilon' \succeq_1 \Upsilon$, $\delta' \geq \delta_2$ and $v' \leq v$. Since $\Omega(\Upsilon,\delta,v) = \{(i,j) : I(i,j|\Upsilon,\delta,) \geq v\}$, thus the fall of v and increase of Υ and δ

[7] For the special example we consider in this step we only need to show the case $a = k$, but we show the general case with $a \leq k$ for step 2 where the boundary set is not so regular.

expands Ω. We only need to show that the $I(i,j|\Upsilon',\delta') \geq I(s,f|\Upsilon,\delta)$ for$\forall i,j$. Recall the specification of the Gittins index from (**??**), we expand the notation of the value function to $W(\gamma|\sigma,\phi;\Upsilon,\delta)$. since W rises in δ for all γ, the fixed point $I(\sigma,\phi|\Upsilon,\delta)$ also rises in δ by Milgrom, Paul and Shannon, Chris (1994).

On the other hand, when $\Upsilon' \sqsupseteq_1 \Upsilon$, we have $\mu(i,j|\Upsilon') \geq \mu(i,j|\Upsilon)$. Thus $W(\gamma|\sigma,\phi;\Upsilon',\delta) \geq W(\gamma|\sigma,\phi;\Upsilon,\delta)$ for all γ, then $I(\sigma,\phi|\Upsilon',\delta) \geq I(\sigma,\phi|\Upsilon,\delta)$ for $\forall\sigma,\phi$ and δ. The rest of the proof follows from Claim 2 and the same argument in the proof of Proposition **??**.

References

[1] Aval, J.-C. (2008): Multivariate FussCatalan numbers, Discrete Mathematics, 308(20), 4660-4669

[2] Banks, J. S., and R. K. Sundaram (1992): A class of bandit problems yielding myopic optimal strategies, Journal of Applied Probability, 29(3), 625-632.

[3] Bellman, R. (1956): A problem in the sequential design of experiments, Sankhy a: The Indian Journal of Statistics (1933-1960), 16(3/4), 221-229.

[4] Bolton, P., and C. Harris (1999): Strategic experimentation, Econometrica, 67(2), 349-374.

[5] Bubeck, S., N. Cesa-Bianchi, et al. (2012): Regret analysis of stochastic and non-stochastic multi-armed bandit problems, Foundations and Trends R in Machine Learning, 5(1), 11-22.

[6] Choi, M., and L. Smith (2017): Ordinal aggregation results via Karlin's variation diminishing property, Journal of Economic Theory, 168, 1-11.

[7] Gittins, J., K. Glazebrook, and R. Weber (2011): Multi-armed bandit allocation indices. John Wiley & Sons.

[8] Gittins, J. C. (1979): Bandit processes and dynamic allocation indices, Journal of the Royal Statistical Society. Series B (Methodological), 148-177.

[9] Gomes, R., D. Gottlieb, and L. Maestri (2016): Experimentation and project selection: Screening and learning, Games and Economic Behavior, 96, 145-169.

[10] Huelsenbeck, J. P., and K. A. Crandall (1997): Phylogeny estimation and hypothesis testing using maximum likelihood, Annual Review of Ecology and System- atics, 28(1), 437-466.

[11] Kelly, F., et al. (1981): Multi-armed bandits with discount factor near one: The Bernoulli case, The Annals of Statistics, 9(5), 987-1001.

[12] Milgrom, P., and C. Shannon (1994): Monotone comparative statics, Econometrica: Journal of the Econometric Society, pp. 157-180.

[13] Moscarini, G., and L. Smith (2001): The Optimal Level of Experimentation, Econometrica, 69(6), 1629-1644.

[14] Rosenberg, D., E. Solan, and N. Vieille (2007): Social Learning in One-Arm Bandit Problems, Econometrica, 75(6), 1591-1611.

[15] Smith, L., P. Sørensen, and J. Tian (2012): Informational herding, optimal experimentation, and contrarianism, Revised and resubmitted for Review of Economic Studies.

[16] Wald, A. (1947): Foundations of a general theory of sequential decision functions, Econometrica, Journal of the Econometric Society, pp. 279-313.