Dynamic Learning and Market Making in Spread Betting Markets with Informed Bettors

John R. Birge[1] Yifan Feng[1] N. Bora Keskin[2] Adam Schultz[3]

[1]University of Chicago, Booth School of Business, email: {john.birge,yifan.feng}@chicagobooth.edu

[2]Duke University, Fuqua School of Business, email: bora.keskin@duke.edu

[3]Uber Technologies Inc., email: adam.schultz4@gmail.com

We study the profit maximization problem of a market maker in a spread betting market. In this market, the market maker quotes cutoff lines for the outcome of a certain future event as "prices," and bettors bet on whether the event outcome exceeds the cutoff lines. Anonymous bettors with heterogeneous strategic behavior and information levels participate in the market. The market maker has limited information on the event outcome distribution, aiming to extract information from the market (i.e., "learning") while guarding against an informed bettor's strategic manipulation (i.e., "bluff-proofing"). We show that Bayesian policies that ignore bluffing are typically vulnerable to the informed bettor's strategic manipulation, resulting in exceedingly large profit losses for the market maker as well as market inefficiency. We develop and analyze a novel family of policies, called *inertial policies*, that balance the tradeoff between learning and bluff-proofing. We construct a simple instance of this family which (i) enables the market maker to achieve a near-optimal profit loss and (ii) eventually yields market efficiency.

Key words: sequential learning, dynamic pricing, strategic agent, market manipulation, spread betting market, prediction market, market making, sports analytics

History: First version: November 12, 2018. This version: April 14, 2020.

---

1. Introduction

1.1. Background and Overview

Spread betting markets are a prevalent form of prediction market, where market makers quote cutoff lines (a.k.a. "spread lines") for the outcome of an uncertain future event, and participants take sides on whether the outcome will exceed the spread line. As a salient example, in a point-spread market for sports betting, a bookie (the market maker) sets a sequence of point spreads (the spread lines), which can be interpreted as the number of points taken from the favorite side. The players (bettors) then observe and decide whether to wager on either the favorite side winning with a margin larger than the point spread or the favorite not winning with such a margin.[1]

Mostly popular in sports betting, spread betting constitutes a multibillion-dollar industry in the U.S.; see NGISC (1999), Statista (2018) and Schwartz (2018). For example, Schwartz (2018) reports that the amount of wagers placed in 2017 in the Nevada regulated sports betting market was worth around $4.9 billion. Due to the Supreme Court's recent decision to clear the way for states to legalize sports betting (NYTimes 2019, CNN 2018), the size of the U.S. sports betting market is expected to grow considerably in the near future (OE 2017).

---

[1] Other popular examples of spread betting involve the total number of points in a sports game (Moskowitz 2015) as well as the percentage difference of votes in a political election (Wolfers and Zitzewitz 2004).

It is of great value to understand the market making problem in this context, i.e., how to set the spread lines as "prices" on the part of market makers. From a market maker's perspective, mispriced spread lines are costly, since the market makers take the opposite side of every bet offered.[2] The consequence of mispricing is exacerbated by the fact that professional bettors may systematically exploit the mispricing events; see Haralabos "Bob" Voulgaris as a vivid example in the National Basketball Association (NBA) betting market.[3] From a market designer's perspective, an effective spread betting market (as a prediction market) serves a role in information aggregation. More specifically, the spread line should reveal intrinsic characteristics of the event outcome distribution (at least in the long run).

Despite its value, understanding of the aforementioned market making problem is limited. Levitt (2004) provides some guidelines for a *clairvoyant* market maker, i.e., one who has perfect information about both the event outcome distribution and the systematic bias of the public.[4] However, Gandar et al. (1998) suggest that opening line biases exist in general in the NBA betting market, but the lines change relatively frequently in a way to eliminate the opening line biases over time. Such empirical evidence indicates that market makers (sportsbooks) may not necessarily know the "correct" spread lines in the beginning. Rather, Gandar et al. (1998) imply that the spread lines are influenced by the interplay between sportsbooks and the bettors: informed bettors—those who can identify the teams undervalued by the sportsbooks—are both present and influential in this market; while sportsbooks may adjust the spread lines significantly to correct their prediction errors over time.

Motivated by this empirical evidence, this paper aims to deepen the understanding of the above market making problem in the following directions:

1. In the presence of sophisticated bettors, how should a (non-omnipotent) market maker move the spread lines (dynamically) to maximize overall profits?
2. What is the overall cost of the market maker's lack of information?
3. Can spread lines yield market efficiency (i.e., converge to an unbiased predictor of the event outcome) in the long run?

We formulate a dynamic learning problem for the market maker (hereafter referred to as "she"). In particular, we consider an unbiased market where the market maker has a binary prior belief on the correct spread line. The market maker strives to dynamically extract information from the market. For example, too many bets on one side of the spread line may be treated as a signal of mispricing and the market maker can respond to it by moving the spread lines in the opposite direction. We study policies that respond to such

[2] See Levitt (2004) for several notorious examples of bookmakers suffering large losses in history.

[3] Among the top gamblers in this market (or at least, among the ones who have revealed their identities), Voulgaris reportedly "routinely wagered a million dollars in a single day," with his average winning percentage being around 70% (Eden 2013).

[4] In particular, in an unbiased betting market where there is no systematic bias of the public, a clairvoyant market maker should consistently set the spread line at the median of the event outcome distribution, which equalizes the probabilities of bets on both sides.

market signals in a profit-maximizing way and characterize the corresponding spread line dynamics in the market.

Our model incorporates bettors with heterogeneous strategic behaviors and information levels. Specifically, we consider two types of bettors: a population of myopic bettors and an informed bettor (each hereafter referred to as "he"). The informed bettor has superior knowledge about the event outcome distribution, and can bet repeatedly and strategically to maximize his expected profit. On the other hand, myopic bettors do not exhibit the same level of strategic sophistication as the informed bettor. They form idiosyncratic estimates about the event outcome and bet according to their individual estimates in a myopic way. All bets are anonymous, i.e., the market maker can only base her spread lines on the aggregate statistics of bets rather than on each individual bettor's betting history.

To maximize profit, the market maker faces a trade-off between two goals: *learning* and *bluff-proofing*. On one hand, she needs to extract information from the market and incorporate it into her spread lines. We refer to this goal as *learning*. On the other hand, if the market maker adjusts spread lines in a particular way, the informed bettor may strictly prefer to "bluff," i.e., bet counter to his private information to exacerbate the market maker's mispricing. We refer to the market maker's goal to protect herself from bluffing as *bluff-proofing*. A good pricing policy should balance the trade-off between learning and bluff-proofing.[5] For this purpose, we develop a policy that collects information at a judiciously selected rate. We show that our policy (i) achieves near-optimal profit performance for the market maker, and (ii) eventually yields market efficiency by pushing the spread line to the median of the event outcome distribution.

1.2. Summary of Results and Main Contributions

In our analysis, we first study Bayesian policies (BPs)—a popular class of pricing policies in the literature on dynamic pricing with demand learning (see, e.g., Harrison et al. (2012), Chen and Wang (2016) and references therein). Under a BP, the market maker (i) computes her posterior belief about the event outcome distribution using Bayes' rule but ignoring the informed bettor; and (ii) sets the spread line as a time-invariant function of her posterior belief. The BP family contains various well-studied policies, such as the myopic Bayesian policy that uses myopic profit maximization to set the spread line (see, e.g., Harrison et al. 2012, Chen and Wang 2016).

We find that BPs are weak against the informed bettor. To be more precise, the informed bettor could earn a profit that is linearly growing in the number of bets (Proposition 1). In general, the market maker's regret typically grows linearly when the commission rate is small (Theorem 1), where regret is defined as the profit loss compared to the clairvoyant in Levitt (2004). We also show that the poor performance of BPs in our setting is not due to incomplete learning. When the informed bettor is absent in our setting, many

[5] See Huddart et al. (2001) and Back and Baruch (2004) for similar ideas of bluffing in the financial market. Chen and Wang (2016) also consider a trade-off between learning and bluff-proofing in dynamic pricing for a single strategic customer with unknown valuation.

simple BPs (including the myopic Bayesian policy) eventually learn the event outcome distribution and achieve good performance. Specifically, the spread line converges to the optimal one at an exponential rate, leading to a constant regret independent of the number of bets (Theorem 2). We also show the robustness of our results by considering two extensions where the informed bettor has restricted ability to place bets (Theorems 5 and 6).

We develop a policy framework that protects the market maker from strategic manipulation. Our solution, called the *inertial policy* (IP), is similar to BP except that IP moves the spread line at a slower rate, but at the same time makes it more costly for the informed bettor to bluff. IP is based on a different (yet parsimonious) state variable: the difference between bets on both sides of the spread line. This state variable resembles the market maker's log-likelihood ratio process, but it aggregates historical data differently, effectively discounting the statistical power of each single data point. We construct a simple instance of IP that achieves three goals simultaneously. First, the informed bettor never bluffs and bets according to a threshold strategy (Theorem 3). Second, the spread line converges to the optimal one almost surely, although at a sub-exponential rate (Theorem 4). Third, IP achieves a regret that grows logarithmically in the number of bets (Theorem 4). To gain deeper insights on our design choices for IP, we also provide a generalized analysis (Propositions K.1 and K.2). Our analysis implies that even if the informed bettor is absent, it is impossible to improve from logarithmic regret to bounded regret by choosing a different instance of IP under mild regularity conditions (Theorem 7). Our proof techniques for deriving these results are based on the exact analysis of a certain Markov chain we construct; this approach differs from the commonly used arguments in the antecedent dynamic learning literature.

1.3. Literature Review

Our work is related to three streams of literature, which are discussed in detail below.

Dynamic pricing and learning in the presence of strategic customers. The literature on dynamic pricing and demand learning is vast (see, for instance, Araman and Caldentey 2009, Harrison et al. 2012, Keskin and Zeevi 2014, 2016, 2018, Chen et al. 2015, Ciocan and Farias 2014, Ferreira et al. 2017, den Boer and Keskin 2017, 2019, Shin and Zeevi 2017, Ban and Keskin 2018, Keskin and Birge 2019). Within this literature, a handful of studies focus on strategic customers; see Levina et al. (2009), Kanoria and Nazerzadeh (2019), Devanur et al. (2014), Chen and Wang (2016), and Huang et al. (2018). Our methodologies and solution concepts are inherently connected to this literature, the closest work being that of Chen and Wang (2016), which builds on the work of Harrison et al. (2012). Like these authors, we consider a binary prior belief on the event outcome distribution in our setting. Our work further develops that of Chen and Wang (2016) in two ways. First, we consider bettors with heterogeneous strategic behaviors and information levels, whereas Chen and Wang (2016) consider a single strategic customer with unknown valuation. Second, we propose a different, yet simple policy family to defeat the informed bettor. Our policy is deterministic, and therefore it is verifiable whether the market maker deviates from the policy *ex-post*. This property is helpful to reinforce the market maker's policy.

We make two contributions to the literature discussed above. First, our paper expands the boundary of applications in this literature from online advertising and retail to prediction markets. Second, we hope that our ideas behind the construction of IP, as well as the proof techniques in characterizing its performance, will in turn motivate analogous pricing policies in other contexts.

Insider trading in financial markets. Our spread betting market model is akin to the insider trading literature in securities markets (see, e.g., Kyle 1985, Glosten and Milgrom 1985, Lin and Howe 1990, Back 1992, Back and Baruch 2004, Caldentey and Stacchetti 2010, and Ostrovsky 2012 for related work on financial markets) but considers the special organization of the spread betting market.[6]

Mulitple modeling differences distinguish our paper from other studies in this literature. First, a common characterization of the market maker in this literature is a zero-profit condition, i.e., the market maker's expected profit conditional on the current filtration is zero; see Kyle (1985) and Caldentey and Stacchetti (2010) for more details. Our market maker solves a dynamic learning and profit maximization problem. Therefore, a simple optimality equation may not be available.[7] Second, we impose a "wisdom of crowd" condition for myopic bettors, who bet according to idiosyncratic but unbiased signals of the event outcome. That is different from the purely noisy trading condition as in Kyle (1985). Finally, our market model is closely related to those in Glosten and Milgrom (1985) and Back and Baruch (2004). The main difference is that our model is fully sequential, i.e., we do not assume any probabilistic structure on the bet arrival process (such as Poisson arrivals). Thus the market maker cannot rely on detection of abnormalities in the arrival rate to distinguish the informed bettor from myopic ones.

Spread betting markets. Our paper further develops Levitt's model (Levitt 2004) by considering the uninformed market maker's dynamic profit maximization problem. Ultimately, our paper contributes to a better understanding of market efficiency in spread betting markets (see also Sauer 2005, Hausch and Ziemba 2008). A spread betting market is statistically efficient if the spread line is an unbiased estimator of the event outcome (Lacey 1990, Golec and Tamarkin 1991). It is economically efficient if there are no profit-earning betting strategies (Zuber et al. 1985 and Gray and Gray 1997). In general, the efficiency of this market may depend on multiple factors. Among these are spread line dynamics; in particular, market inefficiency tends to vanish over time (Gray and Gray 1997, Gandar et al. 1998). Another factor

[6] A differentiating feature of a betting market is that there is no inherent value in the outcome for the participants. Hence participants trade primarily based on their predictions of the event outcome. Moreover, the typical contract structure in a spread betting market is different from a financial exchange market: in a spread betting market, a typical contract is specified by a spread line and a (usually fixed) commission rate, instead of the bid/ask prices.

[7] Regarding our modeling choice of a profit maximization problem, a common justification for the zero-profit condition is competition among market makers. In the context of the spread betting markets, the presence of competition would probably drive the commission rate down to the market maker's marginal cost of a single bet. In that case, the market maker takes commission rates as given and the only way for a market maker to avoid a systematic loss is to price the spread line at the median, akin to the profit maximization problem in our paper.

is the strategic behavior of bettors; e.g., Golec and Tamarkin (1991) point out that in American football betting, the college market is more efficient than the National Football League (NFL) market, which is consistent with the fact that there are more professional bettors in the college market than in the NFL market. Market efficiency is also influenced by bettors' misconceptions about random events like the "hot hand" (Camerer 1989) and even corruption, as in point-shaving scandals (Wolfers 2006). Regarding spread betting market efficiency, our paper conveys the insight that in an unbiased market, even if the market maker is initially uninformed of the event outcome and there is potential strategic manipulation by informed bettors, the market maker is still able to drive the spread line to the efficient one using an inertial policy.

By showing the effectiveness of the market maker's policy, our paper also sheds light on the question of why the spread betting market usually has a market maker instead of being organized as a pure exchange market. While Bossaerts et al. (2002) discuss this question from a market thickness perspective, our paper suggests that the manipulation of informed bettors could also contribute to this phenomenon.[8] Thus, a market maker with commitment power is needed to mitigate informed bettors' strategic manipulations.

Prediction markets. Our paper also fits into the broad theme of prediction markets and belief aggregation rules. For example, in a formulation where participants bet on the specific outcomes of an event, different wagering mechanisms are considered such as the scoring rules (see, e.g., Hanson 2003, Hanson 2007, Chen and Pennock 2007, Chen and Vaughan 2010, Ban 2018) among others (see, e.g., Agrawal et al. 2011, Freeman et al. 2017, Freeman and Pennock 2018). Besides the noticeable difference between the market organization in our paper and those in the common settings in this literature,[9] a unique feature of our paper is that we formulate an online learning problem for the market maker in an environment where a strategic expert can bet *multiple* times.

2. Problem Formulation

In this section, we build a sequential model for the spread betting market with a mixture of bettors with heterogeneous strategic behaviors and information levels. In this market, we formulate a learning-and-pricing problem for the market maker.

2.1. Universal Notations

Throughout the sequel, we use $\mathbb{R}$, $\mathbb{Z}$, $\mathbb{Z}_+$, $\mathbb{Z}_-$, $\mathbb{N}$ and $\mathbb{Q}$ to denote the sets of real numbers, integers, strictly positive integers, strictly negative integers, natural numbers (including zero), and rational numbers, respectively. For all $x, y \in \mathbb{R}$, we use the following notation: $x \wedge y := \min\{x, y\}$, and $x \vee y := \max\{x, y\}$. In

[8] For example, consider a pure exchange market where a mixture of strategic (but uninformed) traders with a common prior and myopic traders participate. The symmetric equilibrium should be that the strategic traders submit the same bid/ask prices (based on their posterior beliefs) and adjust the prices by learning from the myopic traders. In that case, the strategic traders behave like the market maker in our paper, and the informed trader may harm the strategic traders in the same way that the informed bettor harms the market maker in our paper. In this case, the strategic bettors rationally would not enter and the market would not achieve efficiency.

[9] For example, in scoring rules, participants submit their entire belief distributions over outcomes, while in spread betting, participants only give binary responses to the market maker's spread lines.

particular, $x^+ := x \vee 0$ and $x^- := -(x \wedge 0)$. For a function $f : \mathbb{R} \to \mathbb{R}$, we use f' and f'' to denote the derivative and second derivative functions of f, respectively.

2.2. Spread Betting Market

Organization of the market. We consider a betting market for a specific event happening in the future. We model the event outcome as a continuous random variable $X \in \mathbb{R}$ with cumulative distribution function (c.d.f.) $F(\cdot)$. Anonymous bets are placed sequentially, and we index them by $t \in \mathbb{Z}_+$. (We indistinguishably use the index t to refer to period t or bet t.) For bet t, the market maker first quotes a spread line s_t chosen from a compact interval $\mathcal{S} := [s_L, s_H]$, where $-\infty < s_L < s_H < \infty$. The bettor who bets in period t can then bet on either the event $\{X > s_t\}$ (in which case, we denote the bet as $d_t = +1$, or a "positive bet") or the event $\{X < s_t\}$ (in which case, we denote the bet as $d_t = -1$, or a "negative bet"). The payment to the bettor is made after the event outcome is realized. Letting $c \in (0,1)$ denote the commission rate charged by the market maker, the normalized net payment to the bettor is $1-c$ if the bettor wins (i.e., $(X - s_t)\, d_t > 0$), and -1 if the bettor loses (i.e., $(X - s_t)\, d_t < 0$).

Uncertain event outcome. The event outcome distribution is of either "high type" or "low type." To be precise, in our model, $X = m + \epsilon$, where: $m \in \{m_0, m_1\}$ is the (unknown) median of $F(\cdot)$, ϵ is the noise term with c.d.f. $F_\epsilon(\cdot)$, and $s_L < m_0 < m_1 < s_H$. We propose regularization conditions for $F_\epsilon(\cdot)$ in Assumption 1 below. We introduce *hypothesis* i, H_i, to be the hypothesis that $m = m_i$. We also denote by $F_i(\cdot)$ the event outcome distribution under H_i. We illustrate the event outcome distributions in Figure 1.

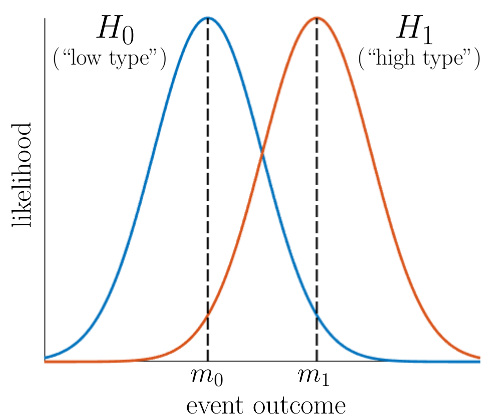

Figure 1 **Illustration of the event outcome distributions.** The bell-shaped curve on the left displays the probability density function of $F_0(\cdot)$ (i.e., the event outcome distribution is of "low type"). The bell-shaped curve on the right displays the probability density function of $F_1(\cdot)$ (i.e., the event outcome distribution is of "high type"). For this graph, $m_0 = 0$, $m_1 = 1$, and $\epsilon \sim \text{Normal}(0,1)$.

Myopic bettors. There is a population of myopic bettors who bet according to idiosyncratic signals of the event. As a whole, they represent the market's (unbiased) knowledge about the event outcome X. In our model, the myopic bettor who bets in period t receives a signal x_t, which is independently drawn from the true event outcome distribution $F(\cdot)$. He bets on the side $X > s_t$ if and only if $x_t > s_t$. That is, his bet $\vartheta_t := \mathbb{I}\{x_t > s_t\} - \mathbb{I}\{x_t \le s_t\}$ follows a binary distribution that equals $+1$ with probability $\bar{F}_i(s_t)$ and -1 with probability $F_i(s_t)$.

Informed bettor. There is an informed bettor who knows the correct median m and bets in a strategic way. For brevity, if the median $m = m_i$, then the informed bettor is referred to as the *type-i informed bettor*. An admissible strategy ξ_i of the type-i informed bettor specifies a (possibly randomized) action $a_t \in \mathcal{A} := \{-1, 0, +1\}$, which can depend on both the market maker's policy π and the underlying hypothesis H_i. Here, actions $a_t = +1$ and $a_t = -1$ correspond to placing the bets $d_t = +1$ and $d_t = -1$, respectively, and $a_t = 0$ corresponds to a "waiting" action. That is, if $a_t = 0$, then bet t is placed by a myopic bettor. The action a_t is adapted to (i) the (public) transaction history $h_{t-1} := (s_1, d_1, \ldots, s_{t-1}, d_{t-1})$, (ii) the informed bettor's action history $A_{t-1} := (a_1, \ldots, a_{t-1})$, and (iii) the most recent spread line s_t. Accordingly, bet t can be expressed as $d_t = \mathbb{I}\{a_t \neq 0\} a_t + \mathbb{I}\{a_t = 0\} \vartheta_t$.

Market maker's policy. An admissible policy for the market maker is any function π that maps the public transaction history h_{t-1} to the spread line $s_t \in \mathcal{S}$. To represent the market maker's knowledge, the mapping π takes neither $\{a_t\}$ nor $\{\vartheta_t\}$ as arguments. This means that the market maker neither knows nor observes whether a bet is placed by the informed bettor or a myopic one. Moreover, her pricing function can depend only on the problem input parameter $\Xi := (c, m_0, m_1, F_\epsilon)$, but neither on the correct median m nor the total number of bets T. In our model, the market maker picks the policy π in the beginning and commits to her policy afterward. To make the commitment credible, we also require that her pricing function can be verified ex-post by the market. For example, this condition holds when the market maker's spread line decision is a deterministic function of the bet history.

Informed bettor's decision problem as a best response strategy. The informed bettor seeks to maximize his total expected profit, given his private information about m, as well as the public knowledge of the market maker's pricing policy π.[10] We introduce some notation to describe the informed bettor's decision making problem. Given the event outcome distribution $F_i(\cdot)$ and the market maker's spread line s, we denote the type-i informed bettor's expected profit from a single positive and negative bet as

$$\begin{cases} j_i^+(s) := (1-c)\mathbb{P}_i(X > s) + (-1)\mathbb{P}_i(X < s) = (c-2)F_i(s) + 1 - c, \\ j_i^-(s) := (1-c)\mathbb{P}_i(X < s) + (-1)\mathbb{P}_i(X > s) = (2-c)F_i(s) - 1, \end{cases} \tag{2.1}$$

respectively, where the probabilities are taken over the event outcome X. Throughout the sequel, we refer to the quantities in (2.1) above as the informed bettor's one-stage profit function. Furthermore, given the market maker's policy π and informed bettor's response strategy ξ, the informed bettor's T-period expected profit function under H_i is

$$V_i^{\pi,\xi}(T) := \mathbb{E}_i^{\pi,\xi}\left[\sum_{t=1}^{T} \left[j_i^+(s_t)\,\mathbb{I}\{a_t = +1\} + j_i^-(s_t)\,\mathbb{I}\{a_t = -1\}\right]\right], \tag{2.2}$$

[10] This assumption implies the market maker is able to credibly commit to a policy. In our case, bettors can verify that the market maker is following the policy, making the commitment assumption reasonable.

where the expectation is taken over the spread lines $\{s_t\}$ and the informed bettor's (possibly randomized) actions $\{a_t\}$, which are specified by the market maker's policy π, informed bettor's response strategy ξ, and the underlying hypothesis H_i. Given π and H_i, the informed bettor aims to find a policy ξ that maximizes his total expected profit, and in case the informed bettor's profit becomes unbounded, the informed bettor chooses a policy that maximizes his long-run average profit per bet. Formally speaking, we say that policy ξ_i^* is a best response strategy if it satisfies the following condition:

$$\xi_i^* \in \begin{cases} \arg\max_\xi \liminf_{T\to\infty} \left\{V_i^{\pi,\xi}(T)\right\} & \text{if } \sup_\xi \liminf_{T\to\infty} \left\{V_i^{\pi,\xi}(T)\right\} < \infty, \\ \arg\max_\xi \liminf_{T\to\infty} \left\{\frac{1}{T} V_i^{\pi,\xi}(T)\right\} & \text{if } \sup_\xi \liminf_{T\to\infty} \left\{V_i^{\pi,\xi}(T)\right\} = \infty. \end{cases} \tag{2.3}$$

We consider an undiscounted formulation because common spread betting markets, including sports betting, typically have frequent bets within short deadlines (Moskowitz 2015).

2.3. Market Maker's Decision Problem

The market maker's goal is to choose a policy π to maximize her T-period profit, given by

$$\sum_{t=1}^{T} \mathbb{E}_i^{\pi,\xi}\Big[\underbrace{\mathbb{I}\{(X-s_t)\,d_t<0\}}_{\text{bettor loses}} - (1-c)\underbrace{\mathbb{I}\{(X-s_t)\,d_t>0\}}_{\text{bettor wins}}\Big],$$

where the expectation is taken over the history h_T generated by strategy profile (π,ξ). To evaluate a given policy π, we use as a performance metric the market maker's *regret*, which is the profit loss of π relative to a clairvoyant market maker who knows the underlying event outcome distribution $F(\cdot)$. The clairvoyant's optimal policy is to consistently set the spread line at the median of $F(\cdot)$. To see this, let us first consider a myopic bettor. Let $r_i(s)$ be the market maker's expected profit from a myopic bettor under hypothesis $i \in \{0,1\}$ when the spread line is s; i.e.,

$$\begin{aligned} r_i(s) &:= \underbrace{\mathbb{P}_i(X>s)}_{\text{prob. of positive bet}} \underbrace{[\mathbb{P}_i(X<s)+(c-1)\mathbb{P}_i(X>s)]}_{\text{market maker's profit}} + \underbrace{\mathbb{P}_i(X<s)}_{\text{prob. of negative bet}} \underbrace{[\mathbb{P}_i(X>s)+(c-1)\mathbb{P}_i(X<s)]}_{\text{market maker's profit}} \\ &= (2c-4)\left(F_i(s)-\tfrac{1}{2}\right)^2 + \tfrac{c}{2}. \end{aligned} \tag{2.4}$$

According to Equation (2.4), the clairvoyant earns the optimal expected profit $\frac{c}{2}$ if and only if $F_i(s) = \frac{1}{2}$. Meanwhile, the same pricing policy drives the informed bettor out of the market because of the commission cost c. That is, if $F_i(s) = \frac{1}{2}$, the informed bettor's profit from betting, $j_i^+(s) = j_i^-(s) = -\frac{c}{2}$, is strictly negative; hence, the informed bettor is incentivized to refrain from betting. For every $i \in \{0,1\}$ and $T \in \mathbb{Z}_+$, the market maker's regret is her T-period profit loss relative to the clairvoyant under hypothesis H_i; that is,

$$\Delta_i^{\pi,\xi}(T) := \frac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_i^{\pi,\xi}\big[\mathbb{I}\{(X-s_t)\,d_t<0\} - (1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\big]. \tag{2.5}$$

We also let $\Delta^\pi(T) := \max\left\{\Delta_0^{\pi,\xi_0^*}(T), \Delta_1^{\pi,\xi_1^*}(T)\right\}$ be the worst-case regret of policy π. Throughout the sequel, we study how $\Delta^\pi(T)$ increases in T under certain classes of policies. Specifically, we use the following Big O notation in our asymptotic performance evaluations.

DEFINITION 1. For every pair of functions $f(\cdot), g(\cdot) : \mathbb{Z}_+ \to \mathbb{R}$, we say that:

- $f(T) = O(g(T))$ if there exists $M < \infty$ and $T_0 \in \mathbb{Z}_+$ such that $|f(T)| \le Mg(T)$ for all $T \ge T_0$;
- $f(T) = \Omega(g(T))$ if there exists $\delta > 0$ and $T_0 \in \mathbb{Z}_+$ such that $f(T) \ge \delta |g(T)|$ for all $T \ge T_0$;
- $f(T) = \Theta(g(T))$ if both $f(T) = O(g(T))$ and $f(T) = \Omega(g(T))$.

2.4. Assumptions and Discussions

ASSUMPTION 1. *The noise distribution $F_\epsilon(\cdot)$ satisfies the following properties:*

(A1:1) $F_\epsilon(\cdot)$ has a continuously differentiable probability density function (p.d.f.) $f_\epsilon(\cdot)$.

(A1:2) $f_\epsilon(\cdot)$ is symmetric around zero; i.e., $f_\epsilon(x) = f_\epsilon(-x)$ for all $x \in \mathbb{R}$.

(A1:3) $f_\epsilon(x) > 0$ for all x satisfying $|x| \le \max\{s_H - m_0, m_1 - s_L\}$.

Statement (A1:1) implies that, for each $i \in \{0,1\}$, $F_i(\cdot)$ has a smooth p.d.f., which we denote as $f_i(\cdot)$. This continuous distribution assumption is a reasonable approximation when the event outcome is continuous (e.g., majority vote percentages) or when the number of possible event outcomes is large (e.g., basketball games). This modeling choice grants us greater analytical tractability to provide insights into the market maker's problem. Statement (A1:2) implies that the noise term ϵ is symmetrically distributed around zero. In particular, ϵ has zero mean. This symmetric distribution assumption is supported by empirical tests on sports betting for American football games (Stern 1991) as well as basketball and baseball games (Stern 1994). The last assumption, Statement (A1:3), has two implications: (i) both $f_0(\cdot)$ and $f_1(\cdot)$ are strictly positive on the feasible region $\mathcal{S} = [s_L, s_H]$, and (ii) $F_1(s_L) > 0$ and $F_0(s_H) < 1$. The first implication ensures that $F_0(\cdot)$ and $F_1(\cdot)$ are separable, while the second implication rules out the degenerate case of instant learning. For a more detailed discussion on Assumption 1, we refer readers to Appendix A.

Informed bettor. Besides superior information about $F(\cdot)$, there are several implicit assumptions about the informed bettor in our model. First, we treat each bet as anonymous. The reason is that the informed bettor may find agencies or proxies to place the bet for him. As a consequence, the market maker cannot differentiate the informed bettor from myopic ones based on their identities. This anonymity accommodates market manipulation in the spirit of Allen and Gale (1992). Second, the informed bettor can repeatedly bet without budget constraints and can also bet an arbitrarily large amount of money before any myopic bettor (while maintaining anonymity). Since bettors make profits at the cost of the market maker, our assumption of a powerful, informed bettor imposes a "stress test" for the market maker's pricing policy. Specifically, the bet arrival model of our paper can be viewed as a combination of an adversarial model (when the informed bettor bets) and a stochastic model (when the myopic bettors bet), where the informed bettor has the power to choose between these two models to maximize his profit. In practice, we believe that such a "stress test" is relevant because the sizes of game-specific sports betting markets are relatively small. In Section 4, we propose a pricing policy, IP, that passes this stress test. That is, it defeats the informed bettor by allowing him to extract at most a constant profit (Lemmas 1 and 2). In Section 5, we also consider more restricted versions of the informed bettor to understand better the robustness of our findings.

In connection to the broader literature on stochastic bandit problems with adversarial opponents, it is worth noting that we show the existence of a pricing policy whose performance does not deteriorate infinitely in the opponent's budget.[11] In part, this is due to the informed bettor's outside option of withdrawing the bets. Recall that the profit benchmark for the market maker is $\frac{cT}{2}$. If there is a bluffing strategy under which the market maker's profit is $\frac{cT}{4}$ and the informed bettor's loss is $\frac{cT}{4}$, the informed bettor would prefer to quit (with a profit of 0) even though he could have caused a regret of $\Omega(T)$ for the market maker.

Market maker's decision problem. There is some debate over whether market makers maximize expected profit, or minimize risk by setting the spread line to balance the wagers on both sides (Paul and Weinbach 2012). In our model, since the market is unbiased, these two spread lines are equal, and there is no ambiguity regarding what the "ideal" spread line is for the market maker. In a biased market, it is still possible to study the market making problem under a profit maximization framework. For example, if the systematic bias is known to the market maker (or can be empirically estimated), we can generalize our current formulation by adding a bias term to the market maker's objective function.

3. Failure (and Success) of Bayesian Policies

This section studies Bayesian policies, a fairly general class of simple and intuitive policies for the market maker. A Bayesian policy (BP) consists of two components: a belief state and a pricing function. Under such a policy, the market maker updates her belief about the (unknown) event outcome distribution, but assuming that there is no informed bettor. To be more precise, we denote b_t as the market maker's posterior probability that $m = m_1$ in period t. The market maker's spread line depends exclusively on her belief state through a time-invariant pricing function $s^{\pi_B}(\cdot)$; see Algorithm 1 in Appendix B for details. We find it convenient to equivalently express the market maker's belief state b_t using the log-likelihood ratio between F_1 and F_0; i.e., $b_{t+1} = \frac{b_1}{b_1+(1-b_1)\exp(-L_t)}$, where

$$L_t = \sum_{\ell=1}^{t} \left[\mathbb{I}\{d_\ell = 1\} \log\left(\frac{\bar{F}_1(s_\ell)}{\bar{F}_0(s_\ell)}\right) + \mathbb{I}\{d_\ell = -1\} \log\left(\frac{F_1(s_\ell)}{F_0(s_\ell)}\right) \right]. \tag{3.1}$$

Note that L_t is a linear aggregation of the betting sequence $\{d_t\}$ with weights $\left\{\log \frac{\bar{F}_1(s_t)}{\bar{F}_0(s_t)}\right\}$ and $\log\left\{\frac{F_1(s_t)}{F_0(s_t)}\right\}$. To avoid pathological cases, we restrict our analysis to the case where both $s^{\pi_B}(0+) := \lim_{b\downarrow 0} s^{\pi_B}(b)$ and $s^{\pi_B}(1-) := \lim_{b\uparrow 1} s^{\pi_B}(b)$ exist; that is, the spread line s_t converges to a certain level as the market maker's belief state b_t converges to 0 or 1. This is a fairly mild condition and useful in our asymptotic analysis below.

[11] Apart from imposing exogenous budget constraints (Lykouris et al. 2018, Jun et al. 2018), there are other ways to restrict the opponent that we do not require in our analysis, such as focusing on oblivious strategies (Slivkins 2019) and restricting the information structure (Jun et al. 2018).

3.1. Performance of Bayesian Policies

In this subsection, we evaluate the performance of Bayesian policies. We find that when the commission rate c is low, the informed bettor is able to earn a constant amount from the market maker per bet on average, and the price does not converge to the median of the event outcome distribution. For ease of notation, we let $\mathfrak{d}_t := |s_t - m_i|$ denote the distance between the spread line s_t and the correct median m_i. In particular, s_t converges to m_i if and only if $\mathfrak{d}_t$ vanishes. Our main finding in this subsection is in Theorem 1 below. The proof of this theorem is in Appendix C.

THEOREM 1. *Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$. Then, for every initial belief $b_1 \in (0,1)$ and sufficiently small $c > 0$, we have the following:*

(T1:1) (non-convergence) For some hypothesis $i \in \{0,1\}$, with strictly positive $\mathbb{P}_i^{\pi_B,\xi_i^*}$-probability, $\mathfrak{d}_t$ does not converge to zero.

(T1:2) (linearly growing regret) $\Delta^{\pi_B}(T) = \Omega(T)$.

Theorem 1 states that, under a BP, the spread line does not converge to the correct median, and the market maker's regret grows linearly in T.

The key step in deriving Theorem 1 is identifying profitable strategies for the informed bettor if the market maker uses a BP. Among all possible cases for $s^{\pi_B}(0+)$ and $s^{\pi_B}(1-)$, the most relevant one is perhaps when the market maker is *asymptotically myopic*, i.e., $s^{\pi_B}(0+) = m_0$ and $s^{\pi_B}(1-) = m_1$. In this case, the market maker learns from the bet history and asymptotically moves the spread line to one of the two possible medians as the market maker becomes almost certain about the event outcome distribution under a BP. (We also study all other cases of BPs in Appendix C.) We find in this case, the informed bettor may earn a constant amount of profit per bet on average, by betting on both sides proportionally. We formalize this finding in Proposition 1 below. We briefly discuss the intuition behind Proposition 1 in Section 3.2 and present its proof in Appendix C.4.

PROPOSITION 1. *Suppose that the market maker uses a Bayesian policy π_B with a pricing function $s^{\pi_B}(\cdot)$ such that $s^{\pi_B}(0+) = m_0$ and $s^{\pi_B}(1-) = m_1$. Then, for every initial belief $b_1 \in (0,1)$, hypothesis $i \in \{0,1\}$, and sufficiently small commission rate c, the type-i informed bettor has a "bluffing" policy ξ_b satisfying the following:*

(P1:1) (belief and spread line dynamics) The posterior belief b_t converges to $(1-i)$ and the spread line s_t converges to m_{1-i} almost surely under $\mathbb{P}_i^{\pi_B,\xi_b}$.

(P1:2) (linearly growing profit of the informed bettor) $V_i^{\pi_B,\xi_b}(T) = \Omega(T)$.

Proposition 1 is the key result in characterizing the performance of BPs. It demonstrates that when the market maker learns from market signals, the informed bettor earns a systematic profit by driving the spread line away from the correct median.

3.2. On the Informed Bettor's Profitable Manipulation Strategy (Proposition 1)

In this subsection, we provide some intuition behind Proposition 1, i.e., how the informed bettor can obtain a linearly growing profit from the market maker when the market maker is learning. In brief, the informed bettor gains from exploiting the constant learning rate of BPs (formally defined as the strictly positive drift of the log-likelihood process). For ease of illustration, let us focus on hypothesis H_1. (The intuition for the analysis under the other hypothesis is the same.)

The informed bettor faces a trade-off between two effects. On one hand, bluffing (i.e., placing a negative bet) means betting on the losing side, which is costly. On the other hand, an honest (i.e. positive) bet pushes the market maker's belief b_t towards 1, which corrects her spread lines in the future. Our question is the following: is there a balance between bluffing and honest betting for the informed bettor in order to confuse the market maker while maintaining overall profitability? The informed bettor can make a profit by betting honestly more often than he bluffs (in the limit). To see why, suppose that $c = 0$, and observe that $j_1^+(s) = -j_1^-(s) > 0$ for every $s < m_1$. That is, the one-stage cost of bluffing is offset by the profit of honest betting. Thus in the limit where $c \to 0$ and s_t converges to some $s_\infty < m_1$, the informed bettor gains a linearly growing profit if the average fraction of honest betting is strictly larger than one half.

Using honest bets more often than bluffing, the informed bettor may still push b_t down to 0 (in the limit). To see that, first suppose that $s_t = m_0$ for all t. Then the probability of a positive bet is $\bar{F}_0(m_0) = \frac{1}{2}$ under H_0 and $\bar{F}_1(m_0) > \frac{1}{2}$ under H_1. With the fraction of positive bets exactly equal to one half, L_t diverges (linearly) in favor of H_0; that is $L_t \to -\infty$.[12] Since L_t is a linear aggregation of the bet sequence, the drift of L_t remains negative if the informed bettor perturbs the fraction of honest (i.e., positive) bets by ε. That is, there exists $\varepsilon > 0$ such that $L_t \to -\infty$ (i.e., $b_t \to 0$), even if the average fraction of honest bets is $\frac{1}{2} + \varepsilon$.

We have thus identified an opportunity for the informed bettor to obtain a linearly growing profit. The informed bettor may first bet negatively consecutively to drive the market maker's posterior belief close to zero. Even though this is costly for the informed bettor, once the market maker's posterior belief is sufficiently close to zero, he gains a strictly positive net profit per average bet by (i) betting honestly with an average ratio of $\frac{1}{2} + \varepsilon$, and (ii) keeping ε sufficiently small, so as to drive the market maker's belief further closer to zero.

3.3. Informed Bettor's Manipulation versus Incomplete Learning

This subsection shows that the impact of the informed bettor on the market maker's profit is distinct from the impact of incomplete learning. We do so by evaluating the performance of BPs in an environment with no informed bettors. Under a mild regularity condition, we find that a BP performs well in our setting when the informed bettor is absent. This finding follows from a separability condition regarding hypotheses H_0 and H_1.

[12] In fact, the average increment of L_t per bet is the negative of the Kullback-Leibler divergence between two Bernoulli random variables with success rates $\bar{F}_0(m_0)$ and $\bar{F}_1(m_0)$, which is strictly negative.

For notational simplicity, we introduce a vacuous policy $\xi_\emptyset$ for the informed bettor, under which his action $a_t = 0$ for all t.[13] We say that a pricing function $s^{\pi_B}(\cdot)$ is *regular*, if $\max\left\{\limsup_{b\downarrow 0} \frac{|s^{\pi_B}(b)-m_0|}{b}, \limsup_{b\uparrow 1} \frac{|s^{\pi_B}(b)-m_1|}{1-b}\right\} < \infty$. This regularity condition is stronger than than saying that $s^{\pi_B}(\cdot)$ is asymptotically myopic, i.e., $s^{\pi_B}(0+) = m_0$ and $s^{\pi_B}(1-) = m_1$, because it also requires a certain speed of convergence. However, this regularity condition is still mild. In fact, it subsumes many intuitive policies as special cases, for example, the myopic Bayesian policy, as well as the linear interpolation pricing function, $s^{\pi_B}(b) = bm_0 + (1-b)m_1$ (see Appendix D.3 for a more detailed discussion on the myopic Bayesian policy).

Theorem 2 below characterizes the performance of a Bayesian policy when the informed bettor is absent. We defer the proof of this theorem to Appendix D.

THEOREM 2. *Suppose that the market maker uses a Bayesian policy π_B with a regular pricing function $s^{\pi_B}(\cdot)$. Then for every initial belief $b_1 \in (0,1)$ and hypothesis $i \in \{0,1\}$, we have the following:*

(T2:1) (convergence of spread lines) $\mathfrak{d}_t$ converges to zero almost surely under $\mathbb{P}_i^{\pi_B,\xi_\emptyset}$.

(T2:2) (exponential convergence) $\mathbb{E}_i^{\pi_B,\xi_\emptyset}[\mathfrak{d}_t] = O\left(e^{-\lambda t}\right)$ for some constant $\lambda > 0$.

(T2:3) (bounded regret) $\Delta_i^{\pi_B,\xi_\emptyset}(T) = O(1)$.

Theorem 2 implies that (statistical) incomplete learning does not happen in our context and many simple Bayesian policies exhibit remarkably good profit performance when there is no informed bettor (for the antecedent work on incomplete learning, see, e.g., McLennan 1984, Harrison et al. 2012, Keskin and Zeevi 2018). Thus, the informed bettor's strategic manipulation, instead of incomplete learning, is the market maker's major challenge in the context of our problem formulation. The main intuition behind Theorem 2 is that in our setting, a BP satisfies a separability condition similar to being a δ-discriminative policy as in Harrison et al. (2012) (see Lemma A.3 in Appendix A for more details).

4. Defeating the Informed Bettor with an Inertial Policy

In this section, we construct a simple policy, called the inertial policy (IP), that allows the market maker to combat the informed bettor. We find that no matter how small the commission rate is, IP is immune to the strategic manipulation of the informed bettor while guaranteeing that the market maker's regret grows at most logarithmically in the number of bets.

4.1. Preliminaries

Definition of the inertial policy. IP employs a univariate state variable that represents the evidence in support of $m = m_1$, as well as a pricing function. The aforementioned state variable is the difference between the number of positive and negative bets before period t, given by

$$Z_t := \sum_{\ell=1}^{t-1} d_\ell. \tag{4.1}$$

[13] Equivalently, we could interpret $\xi_\emptyset$ as the informed bettor's best response strategy if the commission is sufficiently high. See Appendix D.2 for a discussion.

Intuitively, we may interpret the new state variable Z_t as an "unweighted" version of the log-likelihood ratio process L_t in (3.1). On one hand, both Z_t and L_t are linear aggregations of the betting sequence $\{d_t\}$. As a result, an intuitive property that Z_t inherits from L_t is that a high value of Z_t corresponds to strong evidence in support of H_1. On the other hand, however, Z_t differs from L_t in how bet outcomes are aggregated. Given a spread line s_t, L_t gives the weight of $\log \frac{\bar{F}_1(s_t)}{\bar{F}_0(s_t)}$ to a positive bet observation and the weight of $\log \frac{F_1(s_t)}{F_0(s_t)}$ to a negative one. In comparison, Z_t weighs these observations with weights $+1$ and -1. Such difference in weights effectively adjusts the statistical power of each data point observed. As shown below, we construct Z_t so that it accounts for the informed bettor's incentives, while maintaining tractability in both policy implementation and performance evaluation.

Similar to BPs, IP specifies the market maker's spread line through a time-invariant pricing function $\tilde{s}(\cdot)$ of state variable Z_t. That is, given $Z_t = z \in \mathbb{Z}$, the market maker's spread line is $s_t = \tilde{s}(z)$.

Representing IP via residual probabilities. For notational simplicity and interpretability, we use a sequence of numbers called *residual probabilities* to provide an alternative representation of the pricing function $\tilde{s}(\cdot)$. The idea behind the residual probabilities is to represent a spread line s by the quantity $|F_i(s) - \frac{1}{2}|$, which captures how far s is from the median m_i. Letting $\alpha := F_1(m_0)$, we explain in Proposition 2 below how we can represent $\tilde{s}(\cdot)$ by residual probabilities. The proof of this result is in Appendix E.

PROPOSITION 2. *(residual probability) For all $\rho(\cdot): \mathbb{Z}_+ \to \left(0, \frac{1}{2} - \alpha\right)$, there uniquely exist a pricing function $\tilde{s}(\cdot): \mathbb{Z} \to [m_0, m_1]$ and an extension of $\rho(\cdot)$ from $\mathbb{Z}_+$ to $\mathbb{Z}$ such that for all $z \in \mathbb{Z}$,*

$$F_0\big(\tilde{s}(-z)\big) = \tfrac{1}{2} + \rho(z) \text{ and } F_1\big(\tilde{s}(z)\big) = \tfrac{1}{2} - \rho(z). \tag{4.2}$$

The closed-form expression for $\tilde{s}(\cdot)$ is given by:

$$\tilde{s}(z) = \begin{cases} F_1^{-1}\left(\frac{1}{2} - \rho(z)\right) & \text{if } z \in \mathbb{Z}_+, \\ \frac{m_0+m_1}{2} & \text{if } z = 0, \\ F_0^{-1}\left(\frac{1}{2} + \rho(-z)\right) & \text{if } z \in \mathbb{Z}_-. \end{cases} \tag{4.3}$$

The closed-form expression for the extension of $\rho(\cdot)$ is given by (E.1) in Appendix E.

The function $\rho(\cdot)$ quantifies how much the policy incorporates historical information into the next spread line.[14] For example, if $Z_t \in \mathbb{Z}_+$, a small $\rho(Z_t)$ means that $\tilde{s}(Z_t)$ is close to m_1; while if $Z_t \in \mathbb{Z}_-$, a small $\rho(Z_t)$ means that $\tilde{s}(Z_t)$ is close to m_0. We illustrate the correspondence between $\tilde{s}(\cdot)$ and $\rho(\cdot)$ in Figure 2; see also Algorithm 2 in Appendix B for details.

Constructing a candidate for the residual probability sequence. Noting that IP is broadly defined for a generic function $\rho(\cdot)$, we propose a simple candidate for $\rho(\cdot)$. Let $\rho(z) = \frac{1}{r_0 + rz}$ for $z \in \mathbb{Z}_+$. Here, we choose $r_0 := \left[\frac{1}{2} - F_1\left(\frac{m_0+m_1}{2}\right)\right]^{-1}$ so that $\rho(0) = \frac{1}{r_0}$, where $\rho(0)$ is defined in the sense of Proposition 2. In this construction, the only tuning parameter is $r > 0$, which controls the rate of convergence of $\rho(z)$ as $z \uparrow +\infty$. A larger value of r means that $\rho(z)$ converges to 0 faster.

[14] Because $\rho(\cdot)$ is a function of integers, we also refer to $\rho = \{\rho(z), z \in \mathbb{Z}_+\}$ as a *residual probability sequence*.

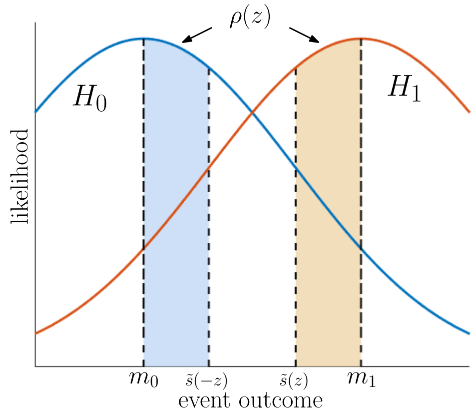

Figure 2 **Illustration of the residual probabilities.** For every $z \in \mathbb{Z}_+$, $\tilde{s}(z)$ and $\tilde{s}(-z)$ are chosen such that each of the two shaded regions has an area equal to $\rho(z)$. For this graph, $m_0 = 0$, $m_1 = 1$, $\epsilon \sim \text{Normal}(0, 0.7)$, $\tilde{s}(-z) = 0.3$, and $\tilde{s}(z) = 0.7$.

Discussion on the inertial policy. Our construction of IP possesses several desirable properties. First, it is simple to implement. To be specific, the update of the state variable Z_t, the evaluation of $\rho(\cdot)$, and the calculation of the spread line s_t all require direct function evaluations only.

Second, the market dynamics under IP are tractable from an analytical point of view. Observe that Z_t is a stationary Markov chain that represents the whole market. To see why, note that IP is a stationary Markov policy that exclusively depends on the state variable Z_t. Therefore, it is sufficient to consider stationary Markov policies for the informed bettor as well. In fact, if we fix the market maker's inertial policy π_I and the informed bettor's policy ξ, then Z_t becomes a birth and death chain under $\mathbb{P}_i^{\pi_I,\xi}$.

Third, IP makes makes manipulation more difficult. Recall from Section 3.2 that we described a simple manipulation strategy where the informed bettor mixes bluffing and honest betting to gain a linearly growing profit from BP. It is straightforward to check that IP guards the market maker from the same manipulation strategy. For example, suppose that H_1 is correct while the spread line is near m_0. The informed bettor may push L_t to $-\infty$ by betting a $\left(\frac{1}{2} + \varepsilon\right)$ fraction of honest (i.e. positive) bets. However, the same strategy only pushes Z_t to ∞, which means that the spread line will be corrected eventually.[15] In Theorem 3, we show that IP guards the market maker against all bluffing behaviors in general.

4.2. Performance of the Inertial Policy

In this subsection, we quantify the performance of IP. We find that under IP, (i) the informed bettor never bluffs and bets under a threshold strategy (Theorem 3), (ii) the uninformed market maker's regret grows at most logarithmically in the number of bets T (Theorem 4), and (iii) the spread line converges to the median of the event outcome distribution with probability one (Theorem 4).

We first characterize the informed bettor's best response policy ξ_i^* as well as his total profit. Recall that the market state is encoded by Z_t defined in (4.1). With a slight abuse of notation, we specify the type-i informed bettor's optimal strategy ξ_i^* by a function of the state $z \in \mathbb{Z}$. We also introduce the value function

[15] Thus, no matter how small the commission rate $c > 0$ is, the informed bettor does not have an incentive to bluff, at least when the spread line is close to either of the medians m_i.

$J^i(\cdot):\mathbb{Z}\to\mathbb{R}\cup\{+\infty\}$ such that $J^i(z)$ is the type-i informed bettor's maximum total expected continuation profit given that the market maker uses IP and the current market state is z. In particular, $J^i(0)$ is the type-i informed bettor's best total profit, because the market starts with state $Z_1=0$. Note that conceptually, $J^i(z)$ is possibly infinite if the market maker's policy is not carefully designed. But IP rules out this possibility as shown in Theorem 3 below. This theorem characterizes the informed bettor's profit and best response to the market maker's inertial policy π_I. We relegate the proof of of Theorem 3 to Appendix G, and discuss our key proof approach in Section 4.3.

THEOREM 3. *There exists $\bar{r}>0$ such that for every policy parameter $r\in(0,\bar{r})$, we have the following:*

(T3:1) (informed bettor's bounded profit) For every hypothesis $i\in\{0,1\}$ and $z\in\mathbb{Z}$, $J^i(z)<+\infty$.

(T3:2) (informed bettor's best response strategy) For every hypothesis $i\in\{0,1\}$, the informed bettor's optimal strategy $\xi_i^*(\cdot)$ is a threshold strategy; i.e., there exists $\bar{z}\in\mathbb{Z}\cup\{-\infty\}$ such that

$$\xi_1^*(z)=\mathbb{I}\{z<\bar{z}\}\ \text{ and }\ \xi_0^*(z)=-\mathbb{I}\{z>-\bar{z}\}\ \text{ for every } z\in\mathbb{Z}. \tag{4.4}$$

The expressions of $\bar{r}$ and $\bar{z}$ depend only on the problem input parameter Ξ and are given in Appendix G.2.

To interpret Theorem 3, let us say that the market state Z_t changes in the "right" direction if it increases under H_1 and decreases under H_0, and in the "wrong" direction otherwise. Theorem 3 means that under IP (with a sufficiently small r), the informed bettor will bet honestly if and only if the market state evolves sufficiently far in the wrong direction. Because the market maker's spread line s_t is a function of the market state Z_t through the pricing function $\tilde{s}(\cdot)$ defined in (4.3), it is equivalent to say that the informed bettor bets if and only if the market maker's spread line is sufficiently close to the wrong median. Otherwise, the informed bettor chooses not to bet on either side because of the transaction cost (in the form of the market maker's commission rate). In either case, the informed bettor does not have an incentive to bluff. To illustrate our inertial policy as well as the informed bettor's best response, we show in Figure 3 a sample path of $\{Z_t\}$ in a numerical example. Under IP, the total profit that the informed bettor gains from the market maker is finite.

Finally, Theorem 4 below characterizes the performance of our inertial policy π_I with the residual probability sequence $\{\rho(z)=\frac{1}{r_0+rz}, z\in\mathbb{Z}_+\}$. We relegate the proof of this result to Section 4.4.

THEOREM 4. *For every commission rate $c\in(0,1)$, hypothesis $i\in\{0,1\}$, and policy parameter $r\in(0,\bar{r})$, we have the following:*

(T4:1) (convergence of spread lines) $Z_t\to\infty$ (resp. $Z_t\to-\infty$) almost surely under $\mathbb{P}_1^{\pi_I,\xi_1^*}$ (resp. $\mathbb{P}_0^{\pi_I,\xi_0^*}$). As a result, $\mathfrak{d}_t$ converges to zero almost surely under $\mathbb{P}_i^{\pi_I,\xi_i^*}$.

(T4:2) (sub-exponential convergence) $\sum_t \mathbb{E}_i^{\pi_I,\xi_i^*}[\mathfrak{d}_t]$ diverges at a rate satisfying $\sum_{t=1}^T \mathbb{E}_i^{\pi_I,\xi_i^*}[\mathfrak{d}_t]=O(\sqrt{T\log T})$.

(T4:3) (logarithmic regret) $\Delta^{\pi_I}(T)=O(\log T)$.

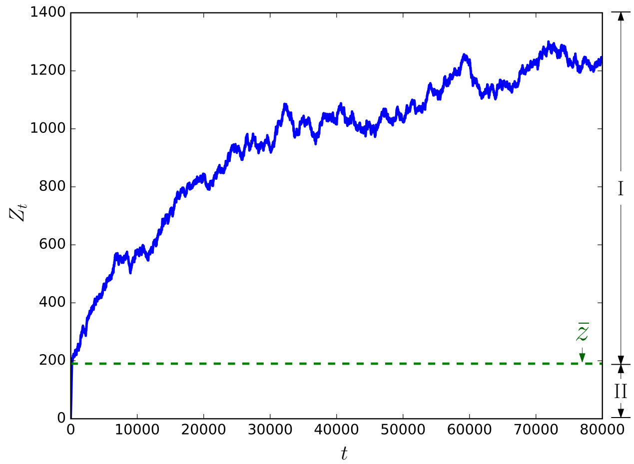

Figure 3 Sample path illustration for the inertial policy under hypothesis H_1. The solid curve displays a sample path of $\{Z_t\}$ whereas the dashed line displays the informed bettor's betting threshold $\bar{z}$. When the market state Z_t is in Region I (i.e., above the dashed line), the informed bettor is inactive and only myopic bettors bet. In comparison, when Z_t is in Region II (i.e., below the dashed line), the informed bettor actively exploits his inside information by betting honestly. In this graph, $m_0 = 0$, $m_1 = 1$, $\epsilon \sim \text{Normal}(0,1)$, $c = 0.1$, and $r = 0.99\bar{r}$, where $\bar{r} = 0.1667$ is calculated as in Appendix G.2.

Theorem 4 implies that IP asymptotically sets the spread line s_t at the correct median. Since the informed bettor's best response policy is of threshold type (as shown in Theorem 3), Theorem 4 also implies that IP eventually drives the informed bettor out of the market with probability one. Consequently, under IP, the market maker's T-period regret is at most in the order of $\log T$. Together with Theorem 3, Theorem 4 gives a comprehensive characterization of IP.

Noticeably, if the commission rate is sufficiently high, the informed bettor's best response strategy is to never bet at all; see Appendix D.2 for more details. As a result, in connection with Theorem 2, our analysis in Theorems 3 and 4 subsumes the environment with no informed bettors as a special case. Corollary 1 below formally states this result, which characterizes the performance of IP with the residual probability sequence $\{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$ when the informed bettor is absent.

COROLLARY 1. *For every hypothesis $i \in \{0,1\}$ and policy parameter $r \in (0, \bar{r})$, we have the following:*

(C1:1) (convergence of spread lines) $Z_t \to \infty$ (resp. $Z_t \to -\infty$) almost surely under $\mathbb{P}_1^{\pi_I,\xi_\emptyset}$ (resp. $\mathbb{P}_0^{\pi_I,\xi_\emptyset}$). As a result, $\mathfrak{d}_t$ converges to zero almost surely under $\mathbb{P}_i^{\pi_I,\xi_\emptyset}$.

(C1:2) (sub-exponential convergence) $\sum_t \mathbb{E}_i^{\pi_I,\xi_\emptyset}[\mathfrak{d}_t]$ diverges at a rate satisfying $\sum_{t=1}^{T} \mathbb{E}_i^{\pi_I,\xi_\emptyset}[\mathfrak{d}_t] = O(\sqrt{T \log T})$.

(C1:3) (logarithmic regret) $\Delta_i^{\pi_I,\xi_\emptyset}(T) = O(\log T)$.

From a managerial standpoint, IP stands in stark contrast to the Bayesian policies in Section 3. On one hand, IP effectively protects the market maker against the informed bettor's strategic manipulation. On the other hand, this protection is at a cost: IP learns (from myopic bettors) more slowly than BPs do. For example, suppose there are no informed bettors. Under hypothesis H_1, the drift (i.e., expected one-stage increment) of Z_t equals $2\rho(z)$, which converges to zero as z grows to infinity, while the drift of L_t is bounded away from zero. Roughly speaking, this means Z_t has a sublinear growth rate, which leads to a sub-exponential convergence rate of s_t and logarithmic regret under IP. In contrast, L_t has a linear growth rate, which leads to an exponential convergence rate of s_t and constant regret under BP. The related results are formally stated in Theorems 2 and 4 as well as Corollary 1. Moreover, we also discuss how to intuitively understand the dynamics of $\{Z_t\}$ in Section 4.4.

4.3. On the Informed Bettor's Optimal Strategy and Profit (Theorem 3)

This subsection summarizes our proof approach for Theorem 3, explaining why the informed bettor never bluffs and gains a finite total profit under IP. In a nutshell, IP pushes the market state Z_t in the right direction. Meanwhile, IP judiciously controls the growth rate of Z_t: the drift of Z_t vanishes as Z_t grows (so that the informed bettor finds it costly to bluff), but slowly (so that the informed bettor cannot gain an infinite profit from simply waiting for mispricing events).

Market state as a birth and death Markov chain. As alluded to earlier, we represent the market by Z_t, which is a birth and death Markov chain. That is, Z_t increases or decreases by one after each bet, and its transition rule is determined by the market participants' stationary Markovian policies. Based on the informed bettor's best response strategy $\xi_i^*(\cdot)$ described in Theorem 3, the transition rule of Z_t can be described by two cases. In the first case, Z_t is sufficiently far in the right direction (i.e., $Z_t \geq \bar{z}$ under hypothesis H_1 and $Z_t \leq -\bar{z}$ under hypothesis H_0), the informed bettor is inactive and only myopic bettors participate. As a result, Z_t increases (resp. decreases) with probability $\frac{1}{2} + \rho(Z_t)$ $\big($resp. $\frac{1}{2} + \rho(-Z_t)\big)$ under hypothesis H_1 (resp. hypothesis H_0). In the second case, our informed bettor actively exploits his inside information by betting honestly. Hence Z_t moves in the right direction with probability one. As a result, the birth and death Markov chain Z_t has a reflecting boundary point $\bar{z}-1$ (resp. $-\bar{z}+1$) under hypothesis H_1 (resp. hypothesis H_0) after a finitely many of steps.

To formally describe the dynamics of Z_t, it is convenient to use the following notation:

$$\mathbb{P}_i^z(\cdot) := \mathbb{P}_i^{\pi_I,\xi_i^*}(\cdot \,|\, Z_1 = z) \text{ and } \mathbb{E}_i^z[\cdot] := \mathbb{E}_i^{\pi_I,\xi_i^*}[\cdot \,|\, Z_1 = z] \text{ for all } i \in \{0,1\} \text{ and } z \in \mathbb{Z}. \tag{4.5}$$

In short, $\mathbb{P}_i^z(\cdot)$ (resp. $\mathbb{E}_i^z[\cdot]$) is a translated version of $\mathbb{P}_i^{\pi_I,\xi_i^*}$ (resp. $\mathbb{E}_i^{\pi_I,\xi_i^*}$), under which Z_t starts with z almost surely. Moreover, in the context of Theorems 3 and 4, it is clear that the market maker implements the inertial policy π_I, and the type-i informed bettor implements the threshold strategy ξ_i^*. Hence we drop the superscript (π_I, ξ_i^*) for notational brevity. In particular, since the market starts with state $Z_1 = 0$, we have $\mathbb{P}_i^0 = \mathbb{P}_i^{\pi_I,\xi_i^*}$. We explicitly write out the transition matrix $\mathcal{P}_{z,\breve{z}}^i := \mathbb{P}_i^z(Z_2 = \breve{z})$ in Equation (4.6) and illustrate the dynamics of Z_t under $\mathbb{P}_i^z$ in Figure 4.

$$\begin{cases} \mathcal{P}_{z,\breve{z}}^0 = 1 & \text{if } z \geq -\bar{z}+1 \text{ and } \breve{z} = z-1, \\ \mathcal{P}_{z,\breve{z}}^0 = \frac{1}{2} + \rho(-z) & \text{if } z \leq -\bar{z} \text{ and } \breve{z} = z-1, \\ \mathcal{P}_{z,\breve{z}}^0 = \frac{1}{2} - \rho(-z) & \text{if } z \leq -\bar{z} \text{ and } \breve{z} = z+1, \end{cases} \quad \text{and} \quad \begin{cases} \mathcal{P}_{z,\breve{z}}^1 = 1 & \text{if } z \leq \bar{z}-1 \text{ and } \breve{z} = z+1, \\ \mathcal{P}_{z,\breve{z}}^1 = \frac{1}{2} + \rho(z) & \text{if } z \geq \bar{z} \text{ and } \breve{z} = z+1, \\ \mathcal{P}_{z,\breve{z}}^1 = \frac{1}{2} - \rho(z) & \text{if } z \geq \bar{z} \text{ and } \breve{z} = z-1. \end{cases} \tag{4.6}$$

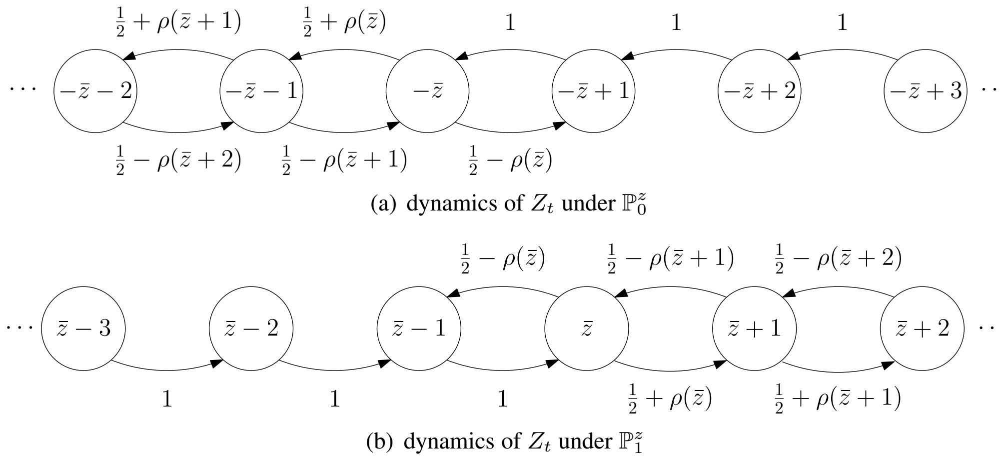

Figure 4 **Illustration of the Markov chain $\{Z_t\}$.** The nodes represent the set of integers, $\mathbb{Z}$, as the state space. The numbers associated with the arrows display the transition probabilities.

Informed bettor's one-stage profit function. We introduce a compact representation of the informed bettor's one-stage profit functions as follows. Invoking (2.1) and (4.2), we use use $j_i^+(z)$ and $j_i^-(z)$ in place of $j_i^+\big(\tilde s(z)\big)$ and $j_i^-\big(\tilde s(z)\big)$, respectively, for shorthand notation throughout this section; that is,

$$\underbrace{j_1^+(z)=j_0^-(-z)}_{\text{honest betting}}=(2-c)\rho(z)-\tfrac{c}{2}\ \text{ and }\ \underbrace{j_1^-(z)=j_0^+(-z)}_{\text{bluffing}}=(c-2)\rho(z)-\tfrac{c}{2}. \tag{4.7}$$

Proof sketch for Theorem 3. In the proof of Theorem 3, we employ a verification argument. We evaluate the informed bettor's profit under the threshold strategy ξ_i^* defined in (4.4), showing that: (i) ξ_i^* is a best response strategy, and (ii) ξ_i^* generates a finite profit for the informed bettor.

Constructing a candidate value function $\bar J^i(\cdot)$. Let us first define a particular function $\bar J^i(\cdot)$ as follows:

$$\bar J^1(z)=\bar J^0(-z)=\begin{cases}0\ \text{ for all } z\in\mathbb{Z} & \text{if } \bar z=-\infty,\\ j_1^+(\bar z-1)\sum_{n=(z-\bar z)^+}^{\infty}\Lambda_n+\sum_{i=1}^{(\bar z-z)^+}j_1^+(\bar z-i)\ \text{ for all } z\in\mathbb{Z} & \text{if } \bar z>-\infty.\end{cases} \tag{4.8}$$

Here, the constants $\{\Lambda_n\}$ depend only on $\bar z$ and $\rho(\cdot)$, and are given by

$$\Lambda_n:=\prod_{k=0}^{n}\frac{\frac{1}{2}-\rho(\bar z+k)}{\frac{1}{2}+\rho(\bar z+k)}>0. \tag{4.9}$$

The intuition behind this construction is as follows. Intuitively, $\bar J^1(\cdot)$ is the informed bettor's continuation profit function under the threshold strategy ξ_1^*. Drawing on the informed bettor's one-stage profit function and the transition rule of Z_t, we expect $\bar J^1(\cdot)$ to satisfy the following recursive relation for $z\in\mathbb{Z}$:

$$\bar J^1(z)=\begin{cases}j_1^+(z)+\bar J^1(z+1) & \text{if } z<\bar z,\\ \bar F_1(\tilde s(z))\,\bar J^1(z+1)+F_1(\tilde s(z))\,\bar J^1(z-1) & \text{if } z\ge\bar z,\end{cases}$$

subject to the boundary condition $\lim_{z\to\infty}\bar J^1(z)=0$. In the meanwhile, by symmetry, $\bar J^0(\cdot)$ should be a "reflected" version of $\bar J^1(\cdot)$; that is, $\bar J^0(z)=\bar J^1(-z)$ for all $z\in\mathbb{Z}$. Thus, the construction in (4.8) can be viewed as a solution to the aforementioned recursive relation.

Key properties of $\bar J^i(\cdot)$. The construction of $\bar J^i(\cdot)$ raises three questions: (i) Is $\bar J^i(\cdot)$ well-defined (i.e., finitely valued)? (ii) Is $\bar J^i(\cdot)$ indeed the informed bettor's continuation profit function under the threshold strategy ξ_i^*? (iii) Does the informed bettor have an incentive to deviate from ξ_i^*? In Lemmas 1-3 below, we give definite answers to all of the three questions. The proofs of these lemmas are in Appendix G.5.

LEMMA 1. *(range of the value function) For all $z \in \mathbb{Z}$, we have $0 \le \bar{J}^1(z) = \bar{J}^0(-z) < \infty$.*

LEMMA 2. *(continuation profit) Let $r \in (0, \bar{r})$. Given hypothesis $i \in \{0,1\}$, the market maker's inertial policy π_I, and the informed bettor's response strategy ξ_i^*, $\bar{J}^i(\cdot)$ is the expected continuation profit function for the type-i informed bettor. That is, for all $z \in \mathbb{Z}$,*

$$\bar{J}^1(z) = \lim_{T\to\infty} \mathbb{E}_1^z \left[\sum_{t=1}^{T} j_1^+(Z_t)\mathbb{I}\{Z_t < \bar{z}\} \right] \text{ and } \bar{J}^0(z) = \lim_{T\to\infty} \mathbb{E}_0^z \left[\sum_{t=1}^{T} j_0^-(Z_t)\mathbb{I}\{Z_t > -\bar{z}\} \right]. \tag{4.10}$$

LEMMA 3. *(Bellman optimality) Let $r \in (0, \bar{r})$. Given hypothesis $i \in \{0,1\}$, the market maker's inertial policy π_I, and current state $z \in \mathbb{Z}$, the type-i informed bettor does not have any incentive to deviate from the action specified by ξ_i^* in (4.4). That is, $\bar{J}^i(\cdot)$ satisfies the following Bellman equation:*

$$\begin{aligned}
\bar{J}^1(z) &= \max\Big\{ \underbrace{j_1^+(z) + \bar{J}^1(z+1)}_{\text{bettor's profit from } a_t=+1}, \underbrace{j_1^-(z) + \bar{J}^1(z-1)}_{\text{bettor's profit from } a_t=-1}, \underbrace{\bar{F}_1(\tilde{s}(z))\bar{J}^1(z+1) + F_1(\tilde{s}(z))\bar{J}^1(z-1)}_{\text{bettor's profit from } a_t=0} \Big\}, \\
\bar{J}^0(z) &= \max\Big\{ \underbrace{j_0^+(z) + \bar{J}^0(z+1)}_{\text{bettor's profit from } a_t=+1}, \underbrace{j_0^-(z) + \bar{J}^0(z-1)}_{\text{bettor's profit from } a_t=-1}, \underbrace{\bar{F}_0(\tilde{s}(z))\bar{J}^0(z+1) + F_0(\tilde{s}(z))\bar{J}^0(z-1)}_{\text{bettor's profit from } a_t=0} \Big\}.
\end{aligned} \tag{4.11}$$

<u>Verification of optimality of $\bar{J}^i(\cdot)$.</u> To summarize, Lemma 2 implies that $\bar{J}^i(\cdot)$ defined in (4.8) is the continuation profit function of the threshold strategy ξ_i^* defined in (4.4), and Lemma 1 ensures that the total profit generated by ξ_i^* is finite. The nonnegativity of $\bar{J}^i(\cdot)$ in Lemma 1 plus the Bellman optimality in Lemma 3 imply that $\bar{J}^i(\cdot)$ is an upper bound of the informed bettor's value function.[16] Lemmas 1-3 jointly establish the optimality of threshold strategy ξ_i^*, as well as the fact that $\bar{J}^i(\cdot) = J^i(\cdot)$. This verifies the optimality of $\bar{J}^i(\cdot)$. The complete proof of Theorem 3 is in Appendix G.

Discussion of proof methodology. Our analysis in Theorem 3 (especially Lemma 2) builds on an exact analysis of $\{Z_t\}$. To achieve tight results, we summarize our key proof step in Lemma 4 below. We relegate its proof to Appendix F.

LEMMA 4. *(key step for performance evaluation) Consider an arbitrary stationary discrete-time Markov chain $\{Y_t,\ t = 1, 2, \ldots\}$ with state space $S \subset \mathbb{R}$ defined on some probability measure space $(\Omega, \mathbb{P})$. Suppose that $u, f : S \to \mathbb{R}$ are functions that satisfy $f(z) = \mathbb{E}[u(Y_2)|Y_1 = z] - u(z)$ for all $z \in S$. Then, $\mathbb{E}f(Y_1) + \mathbb{E}f(Y_2) + \cdots + \mathbb{E}f(Y_t) = \mathbb{E}u(Y_{t+1}) - \mathbb{E}u(Y_1)$ for all t.*

Lemma 4 gives us a guideline on how to evaluate quantities of the form $\mathbb{E}f(Y_1) + \mathbb{E}f(Y_2) + \cdots + \mathbb{E}f(Y_t)$ for any given function $f(\cdot)$. In the first step of this evaluation, we solve for the difference equation $f(z) =$

[16] Because we consider the total profit problem for the informed bettor, the Bellman equation (4.11) in Lemma 3 alone implies neither the optimality of ξ_i^* as the informed bettor's strategy nor the optimality of $\bar{J}^i(\cdot)$ as his value function. In fact, the solution to the Bellman equation (4.11) is not even unique: if any $\tilde{J}^i(\cdot)$ solves the Bellman equation, so does $\tilde{J}^i(\cdot) + c$, where c is an arbitrary constant.

$\mathbb{E}[u(Y_2)|Y_1=z]-u(z)$ to obtain the function $u(\cdot)$. In the second step, we can replace the t-period sum with only two quantities: $\mathbb{E}u(Y_{t+1})$ and $\mathbb{E}u(Y_1)$, which are usually much easier to work with. Provided that solving the difference equation is tractable, this machinery has two main advantages over the commonly used large-deviation based arguments. First, large-deviation based arguments are sometimes not easily available because the "deterministic" part of $\{Y_t\}$ does not overwhelm the "stochastic noise" part. Second, this evaluation is *exact*, and hence tighter results can be obtained from this machinery. Readers familiar with the stochastic calculus literature may view this machinery as a discrete-time analog of Dynkin's formula (Øksendal 2003, Theorem 7.4.1), which is commonly used to estimate various random quantities via solving differential equations (Krylov 2002, Chapter 6.10).

In the context of our paper, we face several technical challenges: the growth of Z_t suffers from vanishing drift (see Section 4.4 for a more detailed discussion), and we need to evaluate the informed bettor's continuation profit *exactly* to verify the Bellman optimality. That is why we use the method in Lemma 4 to overcome our challenges. More specifically, we take $f(x)=j_1^{+}(x)\,\mathbb{I}\{x<\bar{z}\}$ (resp. $f(x)=j_0^{-}(x)\,\mathbb{I}\{x>-\bar{z}\}$) to evaluate the informed bettor's continuation profit function under ξ_1^* (resp. ξ_0^*). The same machinery is also a key step in Proposition 3, where we take other forms of $f(x)$ to show (i) the almost sure convergence of spread lines, (ii) the convergence rate of spread lines, and (iii) the logarithmic growth rate of regret under IP.

4.4. Discussion on the Market Maker's Regret (Theorem 4)

This subsection provides some intuition for why the market maker's regret is $O(\log T)$ under IP (as shown in Theorem 4). This performance guarantee is derived via an exact analysis of the market state $\{Z_t\}$ via Lemma 4. Roughly speaking, the market state Z_t grows in the order of $\sqrt{t}$, with the market maker's one-period regret vanishing in the order of $1/Z_t^2=1/t$, and her T-period regret growing in the order of $\log T$.

Representation of the market maker's regret. The following lemma expresses the market maker's regret, $\Delta^{\pi_I}(T)$, in an additive form. We relegate its proof to Appendix H.

LEMMA 5. *We have*

$$\Delta^{\pi_I}(T)=\sum_{t=1}^{T}\mathbb{E}_1^0[l(Z_t)], \tag{4.12}$$

where the loss function $l(\cdot):\mathbb{Z}\to\mathbb{R}_+$ is given by

$$l(z)=(2-c)\rho(z)\mathbb{I}\{z<\bar{z}\}+(4-2c)\rho^2(z)\mathbb{I}\{z\geq\bar{z}\}. \tag{4.13}$$

The key advantage of Lemma 5 is that through the above additive form, we not only link the regret evaluation problem with the performance evaluation step in Lemma 4, but also have a more parsimonious way of understanding regret through an intuitive characterization (see the discussions below for more details).

Understanding the dynamics of $\{Z_t\}$. In order to make sense of the convergence of s_t and the overall regret, let us give some remarks on the dynamics of $\{Z_t\}$ (both intuitively and rigorously).

Let us first characterize $\{Z_t\}$ in a heuristic fashion to gain intuition. For simplicity, we provide a characterization under hypothesis H_1 (the reasoning for H_0 is the same). Define an auxiliary stochastic process $\{Y_t\}$ such that $Y_t := 1/\rho^2(Z_t)$ for all t. Note that for sufficiently large z (i.e., $z \geq \bar{z} \vee 1$),

$$\mathbb{E}_1^0[Y_{t+1} - Y_t | Z_t = z] = \left[\tfrac{1}{2} + \rho(z)\right](r_0 + rz + 1)^2 + \left[\tfrac{1}{2} - \rho(z)\right](r_0 + rz - 1)^2 - (r_0 + rz)^2 = 5$$

is a constant. Since $Z_t \uparrow \infty$ almost surely (as in Theorem 4), we expect Y_t to grow linearly in t, expressed as $Y_t \sim t$ for brevity. Taking the appropriate transformations, we thus expect that $\rho(Z_t) = 1/\sqrt{Y_t} \sim 1/\sqrt{t}$ and $Z_t \sim 1/\rho(Z_t) \sim \sqrt{t}$. This heuristic characterization of Z_t is related to Theorem 4 in two ways. First, through a first-order Taylor expansion of $\tilde{s}(Z_t)$ as a function of $\rho(Z_t)$, we expect that $\mathfrak{d}_t \approx \rho(Z_t) \sim 1/\sqrt{t}$.[17] Second, in light of Lemma 5, we expect that $\Delta^{\pi_I}(T) = \sum_{t=1}^T \mathbb{E}_1^0[l(Z_t)] \approx \sum_{t=1}^T \rho^2(Z_t) \sim \log T$.

We formalize our intuition on the dynamics of $\{Z_t\}$ (under hypothesis H_1) in Proposition 3 below. We relegate the proof of Proposition 3 to Appendix H.

PROPOSITION 3. *For every commission rate $c \in (0,1)$ and policy parameter $r \in (0, \bar{r})$, we have the following:*

(P3:1) For all sufficiently large $M > 0$, $\sum_t \mathbb{E}_1^0[\mathbb{I}\{Z_t \leq M\}]$ converges.
(P3:2) $\sum_t \mathbb{E}_1^0[\rho(Z_t)]$ diverges at a rate satisfying $\sum_{t=1}^T \mathbb{E}_1^0[\rho(Z_t)] = O\left(\sqrt{T \log T}\right)$.
(P3:3) $\sum_t \mathbb{E}_1^0[l(Z_t)]$ diverges at a rate satisfying $\sum_{t=1}^T \mathbb{E}_1^0[l(Z_t)] = O(\log T)$.

Proposition 3 above provides the main steps for proving Theorem 4. These steps are related to (i) the almost sure convergence of spread lines, (ii) the convergence rate of spread lines, and (iii) the logarithmic growth rate of regret under IP (see the proof in Theorem 4 below). In Section 6, we revisit this proposition in the case where $\rho(\cdot)$ belongs to a more general family of functions (see Proposition K.2 for details).

We also make a technical remark on Proposition 3 (and hence Theorem 4). In formalizing our intuition on the dynamics of Z_t, the main technical barrier is that the growth rates of Y_t and Z_t do not necessarily overwhelm stochastic fluctuations. To see why, observe that Y_t is essentially a linearly growing process, but the (conditional) second moment of the increment of Y_t, formally defined as $\mathbb{E}_1^0[(Y_{t+1} - Y_t)^2 | Z_t = z]$, grows without bound in t. In other words, Z_t is a martingale with bounded increments, but the growth rate of Z_t is sublinear ($Z_t \sim \sqrt{t}$). This barrier makes estimating $\mathbb{E}[f(Z_t)]$ difficult for each fixed t,[18] as such estimation typically relies on large-deviation based arguments. In comparison, Proposition 3 showcases how our method in Lemma 4 estimates the partial sum $\sum_{t=1}^T \mathbb{E}[f(Z_t)]$ directly. We believe that this approach is generally helpful if quantifying $\mathbb{E}[f(Z_t)]$ is more complicated than solving the difference equation in Lemma 4.

[17] Recall that Statement (T4:2) in Theorem 4 expresses that $\sum_{t=1}^T \mathbb{E}_1^0[\mathfrak{d}_t] = O(\sqrt{T \log T})$. This statement is a $\sqrt{\log T}$ factor weaker than the above heuristic characterization, which implies that $\sum_{t=1}^T \mathfrak{d}_t \sim \sqrt{T}$. While a tighter estimate is possible via Lemma 4, we did not pursue it because we only need to show that $\sum_t \mathbb{E}_1^0[\mathfrak{d}_t]$ diverges in order to demonstrate the sub-exponential convergence of spread lines. This also does not affect our main goal of characterizing the regret performance.

[18] In Proposition 3, the choices for $f(x)$ are $\mathbb{I}\{x \leq M\}$ and $l(x)$.

Proof of Theorem 4. Without loss of generality, we focus our analysis on H_1 with the corresponding probability measure $\mathbb{P}_1^0(\cdot) = \mathbb{P}_1^{\pi_I,\xi_1^*}(\cdot)$; the reasoning for H_0 is the same.

Statement (T4:1) in Theorem 4 follows from Statement (P3:1) in Proposition 3. Invoking the Borel-Cantelli lemma, we conclude that $Z_t \le M$ infinitely often with $\mathbb{P}_1^0$-probability zero. As a result, $Z_t \to \infty$ almost surely under $\mathbb{P}_1^0$. Because $\rho(\cdot)$ is asymptotically vanishing (i.e., $\rho(z) \to 0$ as $z \to \infty$), this implies that $\rho(Z_t) \to 0$ and $\mathfrak{d}_t = |s_t - m_1| = |F_1^{-1}(\frac{1}{2} - \rho(Z_t)) - F_1^{-1}(\frac{1}{2})| \to 0$ almost surely under $\mathbb{P}_1^0$.

Statement (T4:2) in Theorem 4 follows from Statement (P3:2) in Proposition 3. Note that under IP,

$$\mathfrak{d}_t = |s_t - m_1| = m_1 - s_t = \frac{1}{f_1\left(F_1^{-1}\left(\frac{1}{2} - \rho(Z_t)\right)\right)}\rho(Z_t) + o(\rho(Z_t))$$

as $\rho(Z_t) \to 0$ and $Z_t \to \infty$. By Assumption 1, the term $1/f_1\left(F_1^{-1}\left(\frac{1}{2} - \rho(Z_t)\right)\right)$ is bounded away from both zero and infinity. Therefore, $\sum_{t=1}^T \mathbb{E}_1^0[\mathfrak{d}_t] = \Theta\left(\sum_{t=1}^T \mathbb{E}_1^0[\rho(Z_t)]\right)$. Invoking Statement (P3:2), $\sum_t \mathbb{E}_1^0[\mathfrak{d}_t]$ diverges, and its growth rate is such that $\sum_{t=1}^T \mathbb{E}_1^0[\mathfrak{d}_t] = O(\sqrt{T \log T})$.

Statement (T4:3) in Theorem 4 follows from Statement (P3:3) in Proposition 3. In fact, $\Delta^{\pi_I}(T) \overset{\text{Lemma 5}}{=} \sum_{t=1}^T \mathbb{E}_1^0[l(Z_t)] \overset{\text{Statement (P3:3)}}{=} O(\log T)$. ■

5. Generalized Analysis of Bayesian Policies

To deepen our understanding of how BPs perform, this section studies BPs under two distinct generalizations of our base model. These generalizations not only demonstrate the robustness of our main results for BPs but also shed light on how BPs transition from "success" into "failure" as the informed bettor's bets become more prominent. Depending on whether the informed bettor can dominate the market (at least temporarily), he needs either $\Theta(T)$ or $o(T)$ betting opportunities to make BPs systematically fail.

5.1. Random Blocking by Myopic Bettors

As a generalization of our base model, assume that informed bettor's each action attempt is randomly "blocked" by myopic bettors with probability q. To be more precise, suppose that in every period t, the following events happen sequentially:

1. The market maker quotes a spread line s_t.
2. The informed bettor chooses an action $a_t \in \{+1, 0, -1\}$.
3. Nature randomly picks whether the informed bettor is blocked by a myopic bettor. We encode that by χ_t, which follows an independent and identically distributed (i.i.d.) Bernoulli sequence with mean q.
   - if $\chi_t = 1$ (i.e., the informed bettor is blocked) or $a_t = 0$ (i.e., the informed bettor chooses to wait), then bet t comes from a myopic bettor. That is, $d_t = \vartheta_t$.
   - if $\chi_t = 0$ (i.e., the informed bettor has a betting opportunity) and $a_t \neq 0$ (i.e., the informed bettor chooses to act), then bet t comes from the informed bettor. That is, $d_t = a_t$.

The random blocking model differs from our base model by introducing random blocking by myopic bettors with probability q. Note that if $q = 0$, this probabilistic blocking model reduces to our base model where the informed bettor is present. If $q = 1$, the blocking model reduces to our base model any without informed bettors. Thus, this generalization bridges two extreme cases in the base model by restricting the informed bettor's ability to bet.

We define the market maker's BP, π_B, and the informed bettor's response strategy, ξ, in the same way as before, except that the realized transaction becomes $d_t = \mathbb{I}\{\chi_t = 0 \text{ and } a_t \neq 0\}a_t + \mathbb{I}\{\chi_t = 1 \text{ or } a_t = 0\}\vartheta_t$.[19] Note that the probabilistic blocking by myopic bettors adds another source of randomness to the model. To accommodate this difference, we denote by $\hat{\mathbb{P}}_i^{\pi_B,\xi}(\cdot)$ the probability measure governing the market statistics $\{(s_t, d_t, a_t, \chi_t)\}$ and by $\hat{\mathbb{E}}_i^{\pi_B,\xi}(\cdot)$ the corresponding expectation operator, given that the market maker's policy is π_B, the informed bettor's response strategy is ξ, and the underlying hypothesis is H_i. Accordingly, we let $\hat{V}_i^{\pi_B,\xi}(T) := \sum_{t=1}^{T} \hat{\mathbb{E}}_i^{\pi_B,\xi}\left[\mathbb{I}\{\chi_t = 0\}\left(j_i^+(s_t)\,\mathbb{I}\{a_t = +1\} + j_i^-(s_t)\,\mathbb{I}\{a_t = -1\}\right)\right]$ be the informed bettor's T-period profit. With the introduction of $\hat{\mathbb{P}}_i(\cdot), \hat{\mathbb{E}}_i[\cdot]$ and $\hat{V}_i^{\pi_B,\xi}(\cdot)$, we define the informed bettor's best response strategy $\hat{\xi}_i^*$ and the market maker's regret $\hat{\Delta}^{\pi_B}(\cdot)$ in the same way as in the base model.

Recall that $\Xi = (c, m_0, m_1, F_\epsilon)$ is the collection of all problem input parameters. Let $\hat{\Xi} := (m_0, m_1, F_\epsilon)$ be the collection of problem input parameters concerning the *distribution* information only, i.e., those in Ξ except the commission rate c. Theorem 5 below summarizes our main results for our model extension where the informed bettor is randomly blocked by myopic bettors. We relegate its proof to Appendix I.

THEOREM 5. *Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$. Then there exist $\underline{q}, \bar{q} \in (0,1)$, which depend only on $\hat{\Xi}$, such that for every initial belief $b_1 \in (0,1)$ and sufficiently small commission rate $c > 0$, we have the following:*

- *(T5:1) (low blocking probability) If $q < \underline{q}$, then for some hypothesis $i \in \{0,1\}$, with strictly positive $\hat{\mathbb{P}}_i^{\pi_B,\xi_i^*}$-probability, $\mathfrak{d}_t$ does not converge to zero. Moreover, $\hat{\Delta}^{\pi_B}(T) = \Omega(T)$.*
- *(T5:2) (high blocking probability) If $q > \bar{q}$ and the pricing function $s^{\pi_B}(\cdot)$ is regular, then for every initial belief $b_1 \in (0,1)$ and hypothesis $i \in \{0,1\}$, $\mathfrak{d}_t$ converges to zero almost surely under $\hat{\mathbb{P}}_i^{\pi_B,\hat{\xi}_i^*}$, at a rate such that $\hat{\mathbb{E}}_i^{\pi_B,\xi_i^*}[\mathfrak{d}_t] = O\left(e^{-\lambda t}\right)$ for some constant $\lambda > 0$. Moreover, $\hat{\Delta}^{\pi_B}(T) = O(1)$.*

Theorem 5 generalizes our analysis of BPs to incorporate random blocking by myopic bettors. It means that all the conclusions about the failure (Theorem 1) and success (Theorem 2) of Bayesian Policies are robust even if we perturb the blocking probability q from $\{0,1\}$ by a constant independent of T.

Theorem 5 also provides us a guidance on the transition of BPs from good to poor performance as the number of the informed bettor's betting opportunities increases. Roughly speaking, BPs display good profit performance even if the informed bettor has $(1-\bar{q})T$ betting opportunities. On the other hand, BPs suffer from a linear regret even if the market maker observes $\underline{q}T$ number of bets from myopic bettors. As a result, the critical number of informed bets that make BPs transition from success to failure is $\Theta(T)$.

[19] Here, a_t could be interpreted as a "virtual" action that may not necessarily be executed because of random blocking.

5.2. Budget-constrained Informed Bettor

As another generalization of our base model, assume that the informed bettor can place at most K bets. In this generalization, the informed bettor's decision problem is the same as the base model except that he can place up to K bets. This setting differs from our base model due to a hard constraint on the total number of bets by the informed bettor. Note that if $K = \infty$, this budget-constrained model reduces to our base model with a informed bettor. If $K = 0$, the budget-constrained model reduces to our base model without any informed bettors. This generalization, like its counterpart in the preceding subsection, connects the two extreme cases in the base model by restricting the informed bettor's ability to bet, but there is a fundamental difference between the two generalizations. Specifically, the former generalization imposes a stochastic restriction, making it difficult for the informed bettor to place many bets without the intervention of myopic bettors. The latter generalization imposes a deterministic restriction, and hence the informed bettor can still perfectly coordinate his bets as long as they are within the budget constraint.

In the budget-constrained model, the expression for the realized transaction in period t becomes $d_t = \mathbb{I}\{\sum_{\ell=1}^{t} |a_\ell| \le K \text{ and } a_t \neq 0\} a_t + \mathbb{I}\{\sum_{\ell=1}^{t} |a_\ell| > K \text{ or } a_t = 0\} \vartheta_t$. To account for this, we denote by $\breve{\mathbb{P}}_i^{\pi_B,\xi}(\cdot)$ the probability measure governing the market statistics $\{(s_t, d_t, a_t)\}$, and by $\breve{\mathbb{E}}_i^{\pi_B,\xi}[\cdot]$ the corresponding expectation operator in the budget-constrained model, given that the market maker's policy is π_B, the informed bettor's response strategy is ξ, and the underlying hypothesis H_i. Furthermore, we let $\breve{V}_i^{\pi_B,\xi}(T;K) := \sum_{t=1}^{T} \breve{\mathbb{E}}_i^{\pi_B,\xi}\big[\big(\mathbb{I}\{\sum_{\ell=1}^{t} |a_\ell| \le K\}\big)\big(j_i^+(s_t)\,\mathbb{I}\{a_t = +1\} + j_i^-(s_t)\,\mathbb{I}\{a_t = -1\}\big)\big]$ be the informed bettor's T-period profit when the informed bettor has K betting opportunities remaining. Given the market maker's Bayesian policy π_B, the informed bettor's adaptive strategy $\breve{\xi}_i^* = \breve{\xi}_i^*(K)$ is a best response strategy if

$$\breve{\xi}_i^* \in \begin{cases} \arg\max_\xi \liminf_{T\to\infty} \big\{\breve{V}_i^{\pi_B,\xi}(T;K)\big\} & \text{if } \sup_\xi \liminf_{T\to\infty} \big\{\breve{V}_i^{\pi_B,\xi}(T;K)\big\} < \infty, \\ \arg\max_\xi \liminf_{T\to\infty} \big\{\frac{1}{T}\breve{V}_i^{\pi_B,\xi}(T;K)\big\} & \text{if } \sup_\xi \liminf_{T\to\infty} \big\{\breve{V}_i^{\pi_B,\xi}(T;K)\big\} = \infty. \end{cases}$$

Note that in general, even if the market maker's BP is a Markov policy with the posterior belief b_t serving as a state variable, $\breve{\xi}_i^*$ may not be Markov with the same state space because $\breve{\xi}_i^*$ can depend on the number of remaining bets of the informed bettor. With the introduction of $\breve{\mathbb{P}}_i^{\pi_B,\xi}(\cdot), \breve{\mathbb{E}}_i^{\pi_B,\xi}[\cdot]$, $\breve{V}_i^{\pi_B,\xi}(\cdot)$ and $\breve{\xi}_i^*$, we define the market maker's regret $\breve{\Delta}^{\pi_B}(T) = \breve{\Delta}^{\pi_B}(T;K)$ in the same way as in the base model. In our asymptotic analysis, we study how $\breve{\Delta}^{\pi_B}(T;K)$ increases as T and K grow.

Theorem 6 below summarizes our main results for our model extension where the informed bettor is budget-constrained. We relegate its proof to Appendix J.

THEOREM 6. *Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$. Then for every initial belief $b_1 \in (0,1)$ and sufficiently small commission rate $c > 0$, $\breve{\Delta}^{\pi_B}(T;K) = \Omega(T \wedge K)$. If in addition, the pricing function $s^{\pi_B}(\cdot)$ is regular, then $\breve{\Delta}^{\pi_B}(T;K) = \Theta(T \wedge K)$.*

Theorem 6 states that the market maker's regret under a BP is in the order of $T \wedge K$ (under certain regularity conditions) under the budget-constrained model. In particular, the regret of a BP is unbounded as long as both T and K grow to infinity, and the regret becomes $\Omega(\log T)$ if $K = \Omega(\log T)$.

5.3. Discussion

Our findings for the generalized models in this section (Theorems 5 and 6) demonstrate the robustness of our main results for BPs (Theorems 1 and 2). In addition, these findings consistently imply that the success of BPs depends on the number of betting opportunities for the informed bettor. Specifically, if the informed bettor is restricted to place up to a "small" number of bets, BPs can still achieve good revenue performance. Otherwise, the market maker should consider a policy from the IP family. Roughly speaking, in the random blocking model, the transition line from success to failure of BPs is when the informed bettor has $\Theta(T)$ betting opportunities. In the budget-constrained model, the same transition line is when the informed bettor has $o(T)$ betting opportunities.

Contrasting both models further reveals how manipulation-proofness of BPs depends on the bet arrival process beyond the volume of bets from the informed bettor. Note that the aforementioned transition line between success and failure of BPs is different in the two generalizations studied in this section. The reason is that in the budget-constrained model, the informed bettor can inject large batch of bets within the budget without the intervention of any myopic bettor. But, in the random blocking model, the informed bettor's bets are randomly mixed with myopic bettors' bets. From a managerial standpoint, the difference between transition lines can be viewed as the net value to the informed bettor of the ability to "flood" the market while still maintaining anonymity.

6. Generalized Analysis of Inertial Policies

This section sheds light on the general designed of IP, especially focusing on why our choice of the residual probability sequence $\rho = \{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$ is a good one. We find that there exists a problem instance such that under mild regularity conditions, it is impossible to improve performance from logarithmic regret to bounded regret by choosing a different type of residual probability sequence.

For intuition, consider a residual probability sequence ρ. If $\rho(z)$ becomes too small as z increases, IP does not push $\{Z_t\}$ in the correct direction sufficiently, and $\{Z_t\}$ behaves like a random walk. To see this, recall that we argued in the preceding section that $\{Z_t\}$ evolves as a birth and death chain, and its drift is essentially proportional to $\rho(Z_t)$ (except for a finite number of states). Thus, if $\rho(z) < \frac{1}{4z}$ in the limit, Z_t does not diverge to infinity and the spread line does not converge to the correct median. Regarding our original choice of ρ, this intuition suggests that for the sake of ensuring convergence, it is undesirable to pick a sequence that vanishes faster. However, if $\rho(z)$ remains too large as z increases, the market maker does not fully exploit the historical data reflected in $\{Z_t\}$. In the end, if $\rho(z) > \frac{1}{4z}$ in the limit, the market maker's T-period regret would be $O\big(\sum_{t=1}^{T} \rho(t)\big)$. For our original choice of ρ, this means that it is also undesirable to pick a sequence that vanishes slower. As a result, the choice of $\rho = \{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$ (with an appropriate value of r) makes the residual probabilities vanish at just the right rate to regularize the dynamics of $\{Z_t\}$ so as to achieve good profit performance.

Let us now formalize the above intuition. First, since we expect $\tilde{s}(z) \to m_1$ as $z \to \infty$ and $\tilde{s}(z) \to m_0$ as $z \to -\infty$, we restrict our attention to the residual probability sequences that are *vanishing*, i.e., $\rho(z) \to 0$ as $z \to \infty$. Let us further consider the following two regimes of residual probability sequences in terms of their behaviors in the limit.

DEFINITION 2. We say that the sequence $\rho = \{\rho(z) \in (0, \frac{1}{2} - \alpha), z \in \mathbb{Z}_+\}$ is *fast vanishing* if $\limsup_{z\to\infty}\{z\rho(z)\} < \frac{1}{4}$, and *slowly vanishing* if $\liminf_{z\to\infty}\{z\rho(z)\} > \frac{1}{4}$ and $\lim_{z\to\infty}\{\rho(z)\} = 0$.

In Definition 2, we compare the function $\rho(\cdot)$ with the critical function, $z \mapsto \frac{1}{4z}$. Specifically, $\{\rho(z), z \in \mathbb{Z}_+\}$ is fast (resp. slowly) vanishing if it vanishes faster (resp. more slowly) than $\frac{1}{4z}$ as $z \to \infty$. The two different cases covers all the possible vanishing residual probability sequences such that $\lim_{z\to\infty}\{z\rho(z)\}$ exists and is not equal to $\frac{1}{4}$. In particular, our original choice, $\rho = \{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$ with a sufficiently small r, belongs to the family of slowly vanishing sequences.

DEFINITION 3. The residual probability sequence $\rho = \{\rho(z), z \in \mathbb{Z}_+\}$ is *regular* if there exists $A \in \mathbb{R} \cup \{\pm\infty\}$ such that $\frac{\rho(z)}{\rho(z+1)} = 1 + \frac{A}{z} + o\left(\frac{1}{z}\right)$ as $z \to \infty$.

The regularity condition above means that the sequence ρ does not alternate excessively in the limit. This is closely related to the Raabe's test of convergence (Bromwich 1908, p. 33) applied to the series $\sum_z \rho(z)$. This condition is satisfied by many well-behaved sequences such as the popular family of "test" sequences, $\{C(a+bz)^{-\mu}, z \in \mathbb{Z}_+\}$ for given $C, a, b, \mu > 0$, which includes our original choice of ρ.

Theorem 7 below studies how inertial policies with generic residual probability sequences perform when the commission rate is sufficiently large. The proof of this result is in Appendix K.

THEOREM 7. *Let π_I be an inertial policy with a regular residual probability sequence $\rho = \{\rho(z) : z \in \mathbb{Z}_+\}$ and the commission rate c be sufficiently large so that $\xi_i^* = \xi_\emptyset$ (i.e., the type-i informed bettor's best response strategy is to never bet) for every hypothesis $i \in \{0, 1\}$. Then, we have the following:*

- *If ρ is fast vanishing, then $\{Z_t\}$ is recurrent.*
- *If ρ is slowly vanishing, then $\Delta^{\pi_I}(T)$ diverges in T at a rate satisfying $\Delta^{\pi_I}(T) = O\left(\sum_{t=1}^T \rho(t)\right)$.*

In either case, $\Delta^{\pi_I}(T)$ is unbounded in T.

Theorem 7 indicates that there exists a problem instance such that, if we use almost any other type of residual probability sequence, then either $\{Z_t\}$ becomes recurrent (and thus the spread line s_t fails to converge) or the regret guarantee becomes weaker than our original result. Consequently, we cannot improve performance from logarithmic regret to bounded regret by choosing a different type of residual probability sequence. This gives a partial characterization of the best achievable regret performance. While we leave the complete characterization as an open question, we conjecture that this is generally the case, i.e., it is not possible to pick any adaptive policy such that the market maker's T-period regret is bounded in T.

It is worth emphasizing that our proof of Theorem 7 is valid in more general setting in which the informed bettor participates in the market and places bets according to a threshold strategy in the following form:

$$\xi_1^{\bar{Z}}(z) = \mathbb{I}\{z < \bar{Z}\} \text{ and } \xi_0^{\bar{Z}}(z) = -\mathbb{I}\{z > -\bar{Z}\} \text{ for every } z \in \mathbb{Z},$$

where $\bar{Z} \in \mathbb{Z} \cup \{-\infty\}$. This is a generalization of the particular threshold policy in Theorem 3. This extra level of generality in the proof of Theorem 7 reveals that the above performance results can be applied to more general problem instances, as long as the informed bettor's best response strategy is of threshold type.

7. Concluding Remarks

Wolfers and Zitzewitz (2006) identify the question of how the market limits manipulation as one of the five open questions about prediction markets. Partially in response to this question, we study a stylized model to analyze the pricing policies of a monopolist market maker operating a spread betting market, who is uninformed of the event outcome distribution. We demonstrate that if the market maker ignores the existence of informed and strategic bettors, an informed bettor can manipulate the market by bluffing and eventually gain an abnormal amount of profit. This (negative) finding still holds even we consider a informed bettor who is restricted in the ability to dominate the market. We propose a policy, called the inertial policy, which eliminates the informed bettor's incentive to bluff, resulting in a regret up to a logarithmic factor of the total number of bets.

There are many possibilities for future work and extensions to our model. For example, one could extend the model so that myopic agents are systematically biased. One approach would be to add a bias term. If the bias term is known to the market maker (or can be empirically estimated), then our model could be extended to incorporate this setting.

Another extension would be to relax the continuity assumption for the event outcome distribution. If the feasible set of spread lines is finite, the spread line cannot be arbitrarily close to the correct median. We conjecture that in this setting, the same inertial policy with proper randomization would be also nearly optimal. One could also extrapolate the insights of this paper to other forms of prediction markets, e.g., the odds market and the index market (see Wolfers and Zitzewitz 2004). The main difference between those organizations and the spread betting market is the payoff structure. We conjecture that the main insights, such as the strategic manipulation of the informed bettor, as well as the rule of thumb to be inertial, carry through to those organizations of prediction markets. Lastly, we could consider relaxing the commitment assumption by considering a perfect Bayesian equilibrium (PBE) between the market maker and the informed bettor. While it is significantly harder to characterize a PBE in a similar setting (see Routledge 1999 for some discussions in the financial market), our Theorem 1 hints that there is no PBE that always leads the spread lines to the correct median. The reason is that in a PBE where the informed bettor eventually quits, the equilibrium condition requires the market maker to essentially use a BP studied in our paper (except that we did not cover the pathological case where $s^{\pi_B}(0+)$ or $s^{\pi_B}(1-)$ does not exist), but a BP is unable to keep the informed bettor out of the market. Hence a market maker with commitment power is needed. Two further extensions of the PBE characterization are (i) multiple market makers and informed bettors; and (ii) the endogenous formation of market makers/informed bettors/commission rates. We leave these extensions for future research.

Acknowledgments

Financial support from Duke University Fuqua School of Business and University of Chicago Booth School of Business is gratefully acknowledged.

References

Agrawal, S., Delage, E., Peters, M., Wang, Z., and Ye, Y. (2011). A unified framework for dynamic prediction market design. *Operations Research*, 59(3):550–568.

Allen, F. and Gale, D. (1992). Stock-price manipulation. *The Review of Financial Studies*, 5(3):503–529.

Araman, V. F. and Caldentey, R. (2009). Dynamic pricing for nonperishable products with demand learning. *Operations Research*, 57(5):1169–1188.

Back, K. (1992). Insider trading in continuous time. *The Review of Financial Studies*, 5(3):387–409.

Back, K. and Baruch, S. (2004). Information in securities markets: Kyle meets Glosten and Milgrom. *Econometrica*, 72(2):433–465.

Ban, A. (2018). Strategy-proof incentives for predictions. In *International Conference on Web and Internet Economics*, pages 51–65. Springer.

Ban, G.-Y. and Keskin, N. B. (2018). Personalized dynamic pricing with machine learning: High dimensional features and heterogeneous elasticity. Available at SSRN: https://ssrn.com/abstract=2972985.

Bossaerts, P., Fine, L., and Ledyard, J. (2002). Inducing liquidity in thin financial markets through combined-value trading mechanisms. *European Economic Review*, 46(9):1671–1695.

Bromwich, T. J. I. (1908). *An Introduction to the Theory of Infinite Series*. Macmillan and Company, Ltd.

Caldentey, R. and Stacchetti, E. (2010). Insider trading with a random deadline. *Econometrica*, 78(1):245–283.

Camerer, C. F. (1989). Does the basketball market believe in the “hot hand”? *The American Economic Review*, 79(5):1257–1261.

Chen, B., Chao, X., and Ahn, H.-S. (2015). Coordinating pricing and inventory replenishment with nonparametric demand learning. Available at SSRN: https://ssrn.com/abstract=2694633.

Chen, X. and Wang, Z. (2016). Bayesian dynamic learning and pricing with strategic customers. Available at SSRN: https://ssrn.com/abstract=2715730.

Chen, Y. and Pennock, D. M. (2007). A utility framework for bounded-loss market makers. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*, pages 49–56. AUAI Press.

Chen, Y. and Vaughan, J. W. (2010). A new understanding of prediction markets via no-regret learning. In *Proceedings of the 2010 ACM Conference on Economics and Computation (EC '10)*, pages 189–198. ACM.

Ciocan, D. F. and Farias, V. F. (2014). Fast demand learning for display advertising revenue management. Working paper, MIT.

CNN (2018). Supreme court lets states legalize sports gambling. Accessed: Oct 30, 2018.

den Boer, A. and Keskin, N. B. (2017). Dynamic pricing with demand learning and reference effects. Available at SSRN: https://ssrn.com/abstract=3092745.

den Boer, A. and Keskin, N. B. (2019). Discontinuous demand functions: Estimation and pricing. Forthcoming in *Management Science*. Available at SSRN: https://ssrn.com/abstract=3003984.

Devanur, N. R., Peres, Y., and Sivan, B. (2014). Perfect Bayesian equilibria in repeated sales. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*, pages 983–1002.

Durrett, R. (2019). *Probability: theory and examples*, volume 49. Cambridge university press.

Eden, S. (2013). Meet the world's top NBA gambler. Accessed: Oct 30, 2018.

Ferreira, K., Simchi-Levi, D., and Wang, H. (2017). Online network revenue management using Thompson sampling. Available at SSRN: https://ssrn.com/abstract=2588730.

Freeman, R. and Pennock, D. M. (2018). An axiomatic view of the parimutuel consensus wagering mechanism. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1936–1938. International Foundation for Autonomous Agents and Multiagent Systems.

Freeman, R., Pennock, D. M., and Wortman Vaughan, J. (2017). The double clinching auction for wagering. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC '17)*, pages 43–60. ACM.

Gandar, J. M., Dare, W. H., Brown, C. R., and Zuber, R. A. (1998). Informed traders and price variations in the betting market for professional basketball games. *The Journal of Finance*, 53(1):385–401.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100.

Golec, J. and Tamarkin, M. (1991). The degree of inefficiency in the football betting market: Statistical tests. *Journal of Financial Economics*, 30(2):311 – 323.

Gray, P. K. and Gray, S. F. (1997). Testing market efficiency: Evidence from the NFL sports betting market. *The Journal of Finance*, 52(4):1725–1737.

Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119.

Hanson, R. (2007). Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15.

Harrison, J. M., Keskin, N. B., and Zeevi, A. (2012). Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586.

Hausch, D. B. and Ziemba, W. T., editors (2008). *Handbook of Sports and Lottery Markets*. Handbooks in Finance. North Holland, 1st edition.

Huang, Z., Liu, J., and Wang, X. (2018). Learning optimal reserve price against non-myopic bidders. *arXiv preprint arXiv:1804.11060*.

Huddart, S., Hughes, J. S., and Levine, C. B. (2001). Public disclosure and dissimulation of insider trades. *Econometrica*, 69(3):665–681.

Jun, K.-S., Li, L., Ma, Y., and Zhu, J. (2018). Adversarial attacks on stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 3640–3649.

Kanoria, Y. and Nazerzadeh, H. (2019). Dynamic reserve prices for repeated auctions: Learning from bids. Available at SSRN: https://ssrn.com/abstract=2444495.

Keskin, N. B. and Birge, J. R. (2019). Dynamic selling mechanisms for product differentiation and learning. *Operations Research*, 67(4):1069–1089.

Keskin, N. B. and Zeevi, A. (2014). Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations Research*, 62(5):1142–1167.

Keskin, N. B. and Zeevi, A. (2016). Chasing demand: Learning and earning in a changing environment. *Mathematics of Operations Research*, 42(2):277–307.

Keskin, N. B. and Zeevi, A. (2018). On incomplete learning and certainty-equivalence control. *Operations Research*, 66(4):1136–1167.

Krylov, N. V. (2002). *Introduction to the Theory of Random Processes*, volume 43. American Mathematical Soc.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Lacey, N. J. (1990). An estimation of market efficiency in the nfl point spread betting market. *Applied Economics*, 22(1):117–129.

Levina, T., Levin, Y., McGill, J., and Nediak, M. (2009). Dynamic pricing with online learning and strategic consumers: an application of the aggregating algorithm. *Operations Research*, 57(2):327–341.

Levitt, S. (2004). Why are gambling markets organised so differently from financial markets? *The Economic Journal*, 114:223–246.

Lin, J.-C. and Howe, J. S. (1990). Insider trading in the OTC market. *The Journal of Finance*, 45(4):1273–1284.

Lykouris, T., Mirrokni, V., and Paes Leme, R. (2018). Stochastic bandits robust to adversarial corruptions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 114–122.

McLennan, A. (1984). Price dispersion and incomplete learning in the long run. *Journal of Economic Dynamics and Control*, 7(3):331–347.

Moskowitz, T. J. (2015). Asset pricing and sports betting. Available at SSRN: https://ssrn.com/abstract=2635517.

NGISC (1999). National gambling impact study commission final report. Accessed: Oct 30, 2018.

NYTimes (2019). New jersey embraces sports gambling, and a billion-dollar business is born. Accessed: Feb 18, 2019.

OE (2017). Economic impact of legalized sports betting. Technical report, Oxford Economics, Wayne, PA. Accessed: Oct 30, 2018.

Øksendal, B. (2003). *Stochastic Differential Equations*. Springer.

Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647.

Paul, R. J. and Weinbach, A. P. (2012). Does sportsbook.com set pointspreads to maximize profits? *The Journal of Prediction Markets*, 1(3):209–218.

Routledge, B. R. (1999). Adaptive learning in financial markets. *The Review of Financial Studies*, 12(5):1165–1202.

Rudin, W. (1976). *Principles of Mathematical Analysis*. McGraw-Hill.

Sauer, R. D. (2005). The state of research on markets for sports betting and suggested future directions. *Journal of Economics and Finance*, 29(3):416–426.

Schwartz, D. G. (2018). Nevada sports betting totals: 1984-2017. Accessed: Oct 30, 2018.

Shin, D. and Zeevi, A. (2017). Dynamic pricing and learning with online product reviews. Working Paper, Columbia University.

Slivkins, A. (2019). Introduction to multi-armed bandits. *Foundations and Trends in Machine Learning*, 12(1-2):1–286.

Statista (2018). Sports betting - statistics & facts. Accessed: Oct 30, 2018.

Stern, H. (1991). On the probability of winning a football game. *The American Statistician*, 45(3):179–183.

Stern, H. S. (1994). A brownian motion model for the progress of sports scores. *Journal of the American Statistical Association*, 89(427):1128–1134.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations Research*, 26(2):305–321.

Wolfers, J. (2006). Point shaving: Corruption in NCAA basketball. *The American Economic Review*, 96(2):279–283.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126.

Wolfers, J. and Zitzewitz, E. (2006). Five open questions about prediction markets. Technical report, National Bureau of Economic Research.

Zuber, R. A., Gandar, J. M., and Bowers, B. D. (1985). Beating the spread: Testing the efficiency of the gambling market for national football league games. *Journal of Political Economy*, 93(4):800–806.

Appendix A: Facts Related to Problem Inputs (Assumption 1)

In this section, we explore immediate consequences of Assumption 1, centering around the symmetry and separability of certain problem inputs.

A.1. Symmetry

Much of our analysis throughout the paper is based on the ideas that we only need to analyze the system under one of the two hypotheses, and the case for the other hypothesis follows by symmetry. The following lemmas reveals the problem structure that allows us to enjoy such simplification.

LEMMA A.1. *(symmetry of probability distribution) For every $x \in \mathbb{R}$, $f_0(x) = f_1(m_0 + m_1 - x)$ and $F_0(x) + F_1(m_0 + m_1 - x) = 1$.*

Proof. Note that for every $i \in \{0,1\}$, $f_i(\cdot)$ is the p.d.f. for $m_i + \epsilon$, where ϵ has a p.d.f. $f_\epsilon(\cdot)$. As a result, for every $x \in \mathbb{R}$ and $i \in \{0,1\}$, $f_i(x) = f_\epsilon(x - m_i)$ and $F_i(x) = F_\epsilon(x - m_i)$. Due to (A1:2), which states that f_ϵ is symmetric around zero, we have $f_\epsilon(x) = f_\epsilon(-x)$. Hence, $f_0(x) = f_\epsilon(x - m_0) = f_\epsilon(m_0 - x) = f_1(m_0 + m_1 - x)$, thereby establishing the first statement in the lemma.

To prove the second statement, we observe that for all $x \in \mathbb{R}$, $F_\epsilon(x) = \int_{-\infty}^{x} f_\epsilon(t)dt = 1 - \int_{x}^{\infty} f_\epsilon(t)dt = 1 - \int_{-\infty}^{-x} f_\epsilon(-t)dt = 1 - \int_{-\infty}^{-x} f_\epsilon(t)dt = 1 - F_\epsilon(-x)$. Thus, $F_\epsilon(x) + F_\epsilon(-x) = 1$, and $F_0(x) = F_\epsilon(x - m_0) = 1 - F_\epsilon(m_0 - x) = 1 - F_1(m_1 + m_0 - x)$. Therefore, we have the second statement in the lemma. ■

LEMMA A.2. *(symmetry of profit function) For every $s \in [m_0, m_1]$, $j_1^+(s) = j_0^-(m_0 + m_1 - s)$ and $j_1^-(s) = j_0^+(m_0 + m_1 - s)$.*

Proof. We observe that $j_1^+(s) \overset{(2.1)}{=} (c-2)F_1(s) + 1 - c \overset{\text{Lem. A.1}}{=} (c-2)[1 - F_0(m_0 + m_1 - s)] + 1 - c = (2-c)F_0(m_0 + m_1 - s) - 1 \overset{(2.1)}{=} j_0^-(m_0 + m_1 - s)$. Similarly, note that $j_1^-(s) \overset{(2.1)}{=} (2-c)F_1(s) - 1 \overset{\text{Lem. A.1}}{=} (2-c)[1 - F_0(m_0 + m_1 - s)] - 1 = (2-c)F_0(m_0 + m_1 - s) + 1 - c \overset{(2.1)}{=} j_0^+(m_0 + m_1 - s)$. ■

A.2. Separability

The following lemma is concerned with the separability of the hypotheses, which is a similar condition to the δ-discriminative condition in Harrison et al. (2012).

LEMMA A.3. *(separability) There exists $\bar{\delta} > 0$ such that for all $s \in \mathcal{S}$, (i) $F_0(s) - F_1(s) \geq \bar{\delta}$, (ii) $\log \frac{F_0(s)}{F_1(s)} \geq \bar{\delta}$, and (iii) $\log \frac{\bar{F}_1(s)}{\bar{F}_0(s)} \geq \bar{\delta}$.*

Proof. Due to (A1:1), $F_0(s) - F_1(s)$, $\log \frac{F_0(s)}{F_1(s)}$, and $\log \frac{\bar{F}_1(s)}{\bar{F}_0(s)}$ are all continuous functions of s on the compact set $\mathcal{S} = [s_L, s_H]$ and hence obtain minimal values in $\mathcal{S}$. Moreover, by (A1:3), $F_0(s) > F_1(s) > 0$ for all $s \in \mathcal{S}$. Thus, for all $s \in \mathcal{S}$, $F_0(s) - F_1(s) > 0$, $\log \frac{F_0(s)}{F_1(s)} > 0$, and $\log \frac{\bar{F}_1(s)}{\bar{F}_0(s)} > 0$. In particular, the minimal values of all three functions are strictly positive. Based on this, we complete the proof by picking $\bar{\delta} > 0$ such that $\min\left\{ \min_{s\in\mathcal{S}} \{F_0(s) - F_1(s)\}, \min_{s\in\mathcal{S}} \log \frac{F_0(s)}{F_1(s)}, \min_{s\in\mathcal{S}} \log \frac{\bar{F}_1(s)}{\bar{F}_0(s)} \right\} > \bar{\delta} > 0$. ■

Appendix B: Summary of Algorithms

Algorithm 1: Bayesian policy (BP)

Data: initial belief $b_1 \in (0,1)$, pricing function $s^{\pi_B}: (0,1) \to \mathcal{S}$.
Result: the spread line s_t for each bet t.

```
t ← 1;
while t ≤ T do
    s_t ← s^{π_B}(b_t);
    observe bet d_t ∈ {−1, +1};
    if d_t = +1 then
        b_{t+1} ← b_t F̄_1(s_t) / (b_t F̄_1(s_t) + (1−b_t) F̄_0(s_t));
    else
        b_{t+1} ← b_t F_1(s_t) / (b_t F_1(s_t) + (1−b_t) F_0(s_t));
    end
    t ← t + 1;
end
```

Algorithm 2: Inertial Policy (IP)

Data: the residual probability sequence $\rho(\cdot): \mathbb{Z}_+ \to (0, \frac{1}{2} - \alpha)$.
Result: the spread line s_t for each bet t.

```
t ← 1, Z_1 ← 0;
while t ≤ T do
    s_t ← s̃(Z_t) according to (4.3);
    observe bet d_t ∈ {−1, +1};
    if d_t = +1 then
        Z_{t+1} ← Z_t + 1;
    else
        Z_{t+1} ← Z_t − 1;
    end
    t ← t + 1;
end
```

Appendix C: On the Failure of Bayesian Policies (Theorem 1)

This section provides the details for the proof of Theorem 1.

C.1. Roadmap

We first aim to identify profitable strategies for the informed bettor when the market maker uses BPs. Our search for such strategies depends only on the values of $s^{\pi_B}(0+)$ and $s^{\pi_B}(1-)$, i.e., the limiting spread lines as the posterior beliefs converge to $\{0,1\}$. There are two cases of the values of $s^{\pi_B}(0+)$ and $s^{\pi_B}(1-)$, each corresponding to a profitable strategy for the informed bettor.

- Case 1 (profitable manipulation): $s^{\pi_B}(0+) \le m_0$ and $s^{\pi_B}(1-) \ge m_1$. In this case, the informed bettor's honest bets a clear correcting power on the market maker's spread lines. We construct a policy for the informed bettor such that he can still gain a linear profit by mixing bluffing bets and honest bets in a certain manner. We formally state our result regarding Case 1 in Proposition C.1 below, which is a direct generalization of Proposition 1. We present the proofs of both propositions in Appendix C.4.

  PROPOSITION C.1. *(bluffing) Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$ such that $s^{\pi_B}(0+) \le m_0$ and $s^{\pi_B}(1-) \ge m_1$. Then, there exists $\bar{c}_0 = \bar{c}_0(\hat{\Xi}) \in (0,1)$ such that for every initial belief $b_1 \in (0,1)$, hypothesis $i \in \{0,1\}$, and commission rate $c \le \bar{c}_0$, the type-i informed bettor has a "bluffing" policy ξ_b satisfying the following:*
  1. *(belief and spread line dynamics) The posterior belief b_t converges to $(1-i)$ and the spread line s_t converges to a limit $s_\infty \ne m_i$ almost surely under $\mathbb{P}_i^{\pi_B,\xi_b}$.*
  2. *(linearly growing profit of the informed bettor) $V_i^{\pi_B,\xi_b}(T) = \Omega(T)$.*

- Case 2 (profitable honest betting): $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$. In this case, the informed bettor's honest bets do not have a sufficiently strong correcting power on the market maker's spread lines.

That is, even if a certain type of informed bettor honestly bet all the time, the spread line does not converge to the correct median. Thus the informed bettor can earn a linear profit by simply betting honestly all the time. We formally state our result regarding Case 2 in Proposition C.2 below, and present the proof of this proposition in Appendix C.5.

PROPOSITION C.2. *(honest betting) Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$ such that $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$. Then there exists $\bar{c}_1 = \bar{c}_1(\pi_B, \hat{\Xi})$ such that for some hypothesis $i \in \{0,1\}$, every initial belief $b_1 \in (0,1)$, and every commission rate $c \leq \bar{c}_1$, the type-i informed bettor has an "honest" policy ξ_h satisfying the following:*

1. *(belief and spread line dynamics) The posterior belief b_t converges to i and the spread line s_t converges to a limit $s_\infty \neq m_i$ almost surely under $\mathbb{P}_i^{\pi_B,\xi_b}$.*
2. *(linearly growing profit of the informed bettor) $V_i^{\pi_B,\xi_h}(T) = \Omega(T)$.*

C.2. Proof of Theorem 1

Fix any policy π_B with pricing function $s^{\pi_B}(\cdot)$. Pick $\bar{c}_0 \in (0,1)$ as in Proposition C.1, and hypothesis i and $\bar{c}_1$ as in Proposition C.2. Let $\bar{c} := \min\{\bar{c}_0, \bar{c}_1\} \in (0,1)$, and pick any $c \in (0, \bar{c})$. We first claim that the market maker's regret is linear, i.e., $\liminf_{T\to\infty}\{\frac{1}{T}\Delta^{\pi_B}(T)\} > 0$. Because $0 < c < \bar{c}$, we deduce from Propositions C.1 and C.2 that the type-i informed bettor has a feasible policy ξ_i such that $\liminf_{T\to\infty}\{\frac{1}{T}V_i^{\pi_B,\xi_i}(T)\} > 0$. The type-$i$ bettor's best response policy, ξ_i^*, maximizes his long-run average profit. Hence, $\liminf_{T\to\infty}\{\frac{1}{T}V_i^{\pi_B,\xi_i^*}(T)\} \geq \liminf_{T\to\infty}\{\frac{1}{T}V_i^{\pi_B,\xi_i}(T)\} > 0$. From the market maker's point of view, her regret is at least the informed bettor's profit, which leads to a linear regret. In fact, we can decompose the bets into two groups: the first group comes from myopic bettors, and the second from the informed bettor:

$$\begin{aligned}
\Delta_i^{\pi_B,\xi_i^*}(T) &\overset{(2.5)}{=} \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\left[\mathbb{I}\{(X-s_t)d_t < 0\} - (1-c)\mathbb{I}\{(X-s_t)d_t > 0\}\right] \\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\mathbb{I}\{a_t = 0\}\left[\mathbb{I}\{(X-s_t)d_t < 0\} - (1-c)\mathbb{I}\{(X-s_t)d_t > 0\}\right] \\
&\quad - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\mathbb{I}\{a_t \neq 0\}\left[\mathbb{I}\{(X-s_t)d_t < 0\} - (1-c)\mathbb{I}\{(X-s_t)d_t > 0\}\right] \\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\mathbb{I}\{a_t = 0\}\left[\mathbb{I}\{(X-s_t)\vartheta_t < 0\} - (1-c)\mathbb{I}\{(X-s_t)\vartheta_t > 0\}\right] \\
&\quad - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\mathbb{I}\{a_t = +1\}\left[\mathbb{I}\{X < s_t\} - (1-c)\mathbb{I}\{X > s_t\}\right] \\
&\quad - \sum_{t=1}^{T} \mathbb{E}_i^{\pi_B,\xi_i^*}\mathbb{I}\{a_t = -1\}\left[\mathbb{I}\{X > s_t\} - (1-c)\mathbb{I}\{X < s_t\}\right]
\end{aligned}$$

$$
\begin{aligned}
&= \underbrace{\mathbb{E}_i^{\pi_B,\xi_i^*} \sum_{t=1}^{T} \left[\tfrac{c}{2} - \mathbb{I}\{a_t = 0\} r_i(s_t)\right]}_{\overset{(2.4)}{\geq} 0} + \underbrace{\mathbb{E}_i^{\pi_B,\xi_i^*} \sum_{t=1}^{T} \left[\mathbb{I}\{a_t = +1\} j_i^+(s_t) + \mathbb{I}\{a_t = -1\} j_i^-(s_t)\right]}_{\overset{(2.2)}{=} V_i^{\pi_B,\xi_i^*}(T)} \\
&\geq V_i^{\pi_B,\xi_i^*}(T).
\end{aligned}
$$

In conclusion, the market maker's regret is asymptotically linear in T:

$$
\liminf_{T\to\infty} \left\{\tfrac{1}{T}\Delta^{\pi_B}(T)\right\} \geq \liminf_{T\to\infty} \left\{\tfrac{1}{T}\Delta_i^{\pi_B,\xi_i^*}(T)\right\} \geq \liminf_{T\to\infty} \left\{\tfrac{1}{T} V_i^{\pi_B,\xi_i^*}(T)\right\} > 0.
$$

We next claim that for some $i \in \{0,1\}$, with positive $\mathbb{P}_i^{\pi_B,\xi_i^*}$-probability, s_t does not converge to m_i. Suppose towards a contradiction that for all $i \in \{0,1\}$, s_t converges to m_i almost surely. It implies that the informed bettor, who makes a linear profit, bets finite times (i.e., $\sum_{t=1}^{\infty} \mathbb{I}\{a_t \neq 0\} < \infty$) almost surely. Thus,

$$
0 < \liminf_{T\to\infty} \{\tfrac{1}{T} V_i^{\pi_B,\xi_i^*}(T)\} \overset{(a)}{\leq} \lim_{T\to\infty} \left\{\tfrac{1}{T}\mathbb{E}_i^{\pi_B,\xi_i^*}\left[\textstyle\sum_{t=1}^{T} \mathbb{I}\{a_t \neq 0\}\right]\right\} \overset{(b)}{=} \mathbb{E}_i^{\pi_B,\xi_i^*}\left[\lim_{T\to\infty} \tfrac{1}{T}\textstyle\sum_{t=1}^{T} \mathbb{I}\{a_t \neq 0\}\right] = 0,
$$

where (a) follows because the informed bettor's profit per bet is at most 1, and (b) follows from the bounded convergence theorem. We have thus arrived at a contradiction. ■

C.3. Main Proof Idea: One-stage Analysis

The above proof is based on what we call a "one-stage analysis." That is, if the informed bettor places one (possibly randomized) bet, we characterize the expected impact on the market maker's belief as well as the informed bettor's profit from this single bet. Two functions are of interest throughout the proof. First, define

$$
D(b,\mathfrak{p}) := (1-\mathfrak{p}) \log\left(\frac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right) + \mathfrak{p} \log\left(\frac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right) \tag{C.1}
$$

as the expected increment (i.e., drift) of the market maker's log-likelihood process L_t after a single bet if (i) the current belief state is b and (ii) the informed bettor bets positively with probability $\mathfrak{p}$ and negatively with probability $1-\mathfrak{p}$. In (C.1), the expectation is taken over the randomized action of the informed bettor. The informed bettor misleads the market maker if $D(b,\mathfrak{p}) < 0$ under H_1 and $D(b,\mathfrak{p}) > 0$ under H_0. Second, let

$$
R_i(b,\mathfrak{p}) := (1-\mathfrak{p}) j_i^-(s^\pi(b)) + \mathfrak{p} j_i^+(s^\pi(b)) \tag{C.2}
$$

be the informed bettor's expected profit from a single bet under H_i if (i) the current belief state is b and (ii) the informed bettor bets positively with probability $\mathfrak{p}$ and negatively with probability $1-\mathfrak{p}$. In (C.2), the expectation is taken over the randomized action of the informed bettor and the final realization of the event outcome X. The type-i informed bettor makes a profit in expectation if $R_i(b,\mathfrak{p}) > 0$.

The following result demonstrates how we utilize the aforementioned one-stage analysis in our proofs. It builds on standard large-deviation based arguments. With a slight abuse of notation, we let $\{\mathfrak{p}(b)\}$ denote the informed bettor's following *behavioral strategy*: he randomly (and independently) chooses to bet positively with probability $\mathfrak{p}(b)$ and negatively with probability $1-\mathfrak{p}(b)$ given the belief state b.

LEMMA C.1. *Let $i \in \{0,1\}$. Suppose that the market maker uses a Bayesian policy π_B and the type-i informed bettor's policy ξ is given by the behavioral strategy $\{\mathfrak{p}(b)\}$. Then, we have the following:*

1. *If there exists $\delta > 0$ such that $\mathbb{E}_i^{\pi_B,\xi}[L_{t+1} - L_t | b_t = b] = D(b, \mathfrak{p}(b)) < -\delta$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$, then (i) $\mathbb{E}_i^{\pi_B,\xi}[b_t] = O(e^{-\lambda t})$ for some $\lambda > 0$, and (ii) $b_t \to 0$, $L_t \to -\infty$, $s_t \to s^\pi\ (0+)$ almost surely. If in addition, there exists $\bar{b} \in (0,1)$ such that $R_i(b, \mathfrak{p}(b)) > \delta$ for all $b \in (0, \bar{b}]$, then $V_i^{\pi_B,\xi}(T) = \Omega(T)$.*
2. *If there exists $\delta > 0$ such that $\mathbb{E}_i^{\pi_B,\xi}[L_{t+1} - L_t | b_t = b] = D(b, \mathfrak{p}(b)) > \delta$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$, then (i) $\mathbb{E}_i^{\pi_B,\xi}[1 - b_t] = O(e^{-\lambda t})$ for some $\lambda > 0$, and (ii) $b_t \to 1$, $L_t \to \infty$, $s_t \to s^\pi\ (1-)$ almost surely. If in addition, there exists $\bar{b} \in (0,1)$ such that $R_i(b, \mathfrak{p}(b)) > \delta$ for all $b \in [\bar{b}, 1)$, then $V_i^{\pi_B,\xi}(T) = \Omega(T)$.*

In our proofs, we employ our one-stage analysis in different settings, and combine it with Lemma C.1 to obtain desired results. Specifically, we use the one-stage analysis in Lemma C.2, in Step 1 in the proof of Proposition C.2 in Appendix C.5, and in Step 1 in the proof of Theorem 2 in Appendix D.1. We combine these instances of the one-stage analysis with Lemma C.1 to prove Propositions C.1 and C.2 as well as Theorem 2.

Proof of Lemma C.1. Without loss of generality, suppose that $D(b, \mathfrak{p}(b)) < -\delta$ for all $b \in (0,1)$. The proof for the other case follows by repeating the same arguments verbatim. We complete the proof in two steps.

Step 1: concentration and convergence for $\{L_t\}$. We claim that there exists $\varepsilon > 0$ such that $\mathbb{P}_i^{\pi,\xi}(L_t \geq -\frac{\delta t}{2}) \leq \exp(-\varepsilon t)$ for $t \in \mathbb{Z}_+$, in which case the following hold: (i) $\mathbb{E}_i^{\pi_B,\xi}[b_t] = O(e^{-\lambda t})$ for some $\lambda > 0$; (ii) $L_t \to -\infty$ almost surely; (iii) $b_t \to 0$ almost surely; and (iv) $s_t \to s^\pi\ (0+)$ almost surely.

Observe that the market maker's pricing policy π_B is Markovian with respect to the belief state b_t. The informed bettor's behavioral strategy $\{\mathfrak{p}(b)\}$ is also Markovian. As a result, b_t is a Markov chain under $\mathbb{P}_i^{\pi_B,\xi}$, and so is the market maker's log-likelihood ratio process L_t.

We apply Doob's decomposition to the process L_t. That is, we define $\mathfrak{A}_t := \sum_{\ell=1}^t \mathbb{E}_i^{\pi_B,\xi}[L_\ell - L_{\ell-1} | L_{\ell-1}]$ and $\mathfrak{M}_t := \sum_{\ell=1}^t \left(L_\ell - \mathbb{E}_i^{\pi_B,\xi}[L_\ell | L_{\ell-1}]\right)$ so that $L_t = \mathfrak{A}_t + \mathfrak{M}_t$ for all t. We may interpret $\mathfrak{A}_t$ as the "drift" of L_t and $\mathfrak{M}_t$ as its "noise." Because $D(b, \mathfrak{p}(b)) < -\delta$ for all $b \in (0,1)$, $\mathfrak{A}_t < -\delta t$ almost surely. Moreover, $\mathfrak{M}_t$ is a martingale with bounded increments: note that

$$
\begin{aligned}
|\mathfrak{M}_\ell - \mathfrak{M}_{\ell-1}| &\leq \left| L_\ell - \mathbb{E}_i^{\pi_B,\xi}[L_\ell | L_{\ell-1}] - L_{\ell-1} + \mathbb{E}_i^{\pi_B,\xi}[L_{\ell-1} | L_{\ell-2}] \right| \\
&= \left| L_\ell - \mathbb{E}_i^{\pi_B,\xi}[L_\ell - L_{\ell-1} | L_{\ell-1}] - 2L_{\ell-1} + \mathbb{E}_i^{\pi_B,\xi}[L_{\ell-1} - L_{\ell-2} | L_{\ell-2}] + L_{\ell-2} \right| \\
&\leq |L_\ell - L_{\ell-1}| + |L_{\ell-1} - L_{\ell-2}| + \mathbb{E}_i^{\pi_B,\xi}[|L_\ell - L_{\ell-1}| | L_{\ell-1}] + \mathbb{E}_i^{\pi_B,\xi}[|L_{\ell-1} - L_{\ell-2}| | L_{\ell-2}] \\
&\overset{(a)}{\leq} 4M,
\end{aligned}
$$

where (a) follows by defining the constant $M := \max\left\{\sup_{s \in \mathcal{S}} \log\left(\frac{F_0(s)}{F_1(s)}\right), \sup_{s \in \mathcal{S}} \log\left(\frac{\bar{F}_1(s)}{\bar{F}_0(s)}\right)\right\} \leq \max\left\{\log \frac{1}{F_0(s_L)}, \log \frac{1}{F_1(s_H)}\right\}$, which is finite due to Assumption (A1:3). As a result, L_t is a Markov chain with a non-vanishing drift and bounded increments.

Finally, we deduce the following for all $t \in \mathbb{Z}_+$:

$$\begin{aligned}
&\mathbb{P}_i^{\pi_B,\xi}(L_t \geq -\tfrac{\delta t}{2}) \\
&\quad = \mathbb{P}_i^{\pi_B,\xi}\left(\mathfrak{A}_t + \mathfrak{M}_t \geq -\tfrac{\delta t}{2}\right) \\
&\quad \leq \mathbb{P}_i^{\pi_B,\xi}\left(\mathfrak{M}_t \geq \tfrac{\delta t}{2}\right) && [\mathfrak{A}_t < -\delta t \text{ almost surely}] \\
&\quad \leq \exp\left(-\tfrac{t^2}{2t(64M^2/\delta^2)}\right) && [\text{by Azuma-Hoeffding inequality; } \left|\tfrac{2\mathfrak{M}_t}{\delta} - \tfrac{2\mathfrak{M}_{t-1}}{\delta}\right| \leq \tfrac{8M}{\delta}\ ] \\
&\quad = \exp\left(-\tfrac{t}{128M^2/\delta^2}\right) \\
&\quad = \exp(-\varepsilon t) && [\varepsilon := \tfrac{\delta^2}{128M^2} > 0]
\end{aligned}$$

The four facts mentioned in the beginning of this step, namely, (i) $\mathbb{E}_i^{\pi_B,\xi}[b_t] = O(e^{-\lambda t})$ for some $\lambda > 0$; (ii) $L_t \to -\infty$ almost surely; (iii) $b_t \to 0$ almost surely; and (iv) $s_t \to s^{\pi}(0+)$ almost surely, follow from the above analysis in a standard manner; see, e.g., the proof of Proposition 7 in Harrison et al. (2012).

Step 2: profit evaluation. Suppose that $R_i(b, \mathfrak{p}(b)) > \delta$ for all $b \in (0, \bar{b}]$. We claim that $\liminf_{T\to\infty}\{\frac{1}{T}V_i^{\pi_B,\xi}\} > 0$. Pick $\bar{L} := \log\left(\frac{\bar{b}}{1-\bar{b}}\right) - \log\left(\frac{b_1}{1-b_1}\right)$ so that $L_t \geq \bar{L}$ if and only if $b_t \geq \bar{b}$. Note that $\bar{L}$ is finite because $b_1 \in (0,1)$ and $\bar{b} \in (0,1)$. Now, let us evaluate the type-i informed bettor's payoff:

$$\begin{aligned}
&V_i^{\pi_B,\xi}(T) \\
&= \sum_{t=1}^{T}\left(\mathbb{E}_i^{\pi_B,\xi}[R_i(b_t,\mathfrak{p}(b_t))]\mathbb{I}\{b_t \leq \bar{b}\} + \mathbb{E}_i^{\pi_B,\xi}[R_i(b_t,\mathfrak{p}(b_t))]\mathbb{I}\{b_t > \bar{b}\}\right) \\
&\geq \sum_{t=1}^{T}\left(\mathbb{E}_i^{\pi_B,\xi}[\delta\,\mathbb{I}\{b_t \leq \bar{b}\}] + \mathbb{E}_i^{\pi_B,\xi}[(-1)\mathbb{I}\{b_t > \bar{b}\}]\right) && [R_i(b,\mathfrak{p}(b)) > \delta\ \forall b \in (0,\bar{b}]] \\
&= \delta T - (1+\delta)\sum_{t=1}^{T}\mathbb{P}_i^{\pi_B,\xi}(b_t > \bar{b}) \\
&\geq \delta T - (1+\delta)\sum_{t=1}^{\infty}\mathbb{P}_i^{\pi_B,\xi}(b_t > \bar{b}) \\
&\geq \delta T - (1+\delta)\sum_{t=1}^{\infty}\mathbb{P}_i^{\pi_B,\xi}(L_t \geq \bar{L})
\end{aligned}$$

We deduce from Step 1 that $(1+\delta)\sum_{t=1}^{\infty}\mathbb{P}_i^{\pi_B,\xi}(L_t \geq \bar{L})$ is a finite constant that is independent of T. Therefore, $\liminf_{T\to\infty}\{\frac{1}{T}V_i^{\pi_B,\xi}\} \geq \delta > 0$. ■

C.4. Profitable Manipulation (Proofs of Propositions 1 and C.1)

Let π_B be the market maker's Bayesian policy satisfying $s^{\pi_B}(0+) \leq m_0$ and $s^{\pi_B}(1-) \geq m_1$. The existence of $s^{\pi_B}(0+)$ and $s^{\pi_B}(1-)$ are guaranteed by the definition of a Bayesian policy in our setting. Roughly speaking, the proofs of Propositions 1 and C.1 rely on the construction of a strategy for the informed bettor that randomizes between bluffing and honest betting. Under such a strategy, the informed bettor keeps misleading the market maker while making profits. To formalize this idea, we recall that $\Xi = (c, m_0, m_1, F_\epsilon)$ is the collection of problem input parameters and $\alpha = F_1(m_0)$, and state the following auxiliary result.

LEMMA C.2. *(one-stage analysis for manipulation) There exist $\bar{c}_0, \mathfrak{p}_0 \in (0,1)$, which depend only on α, such that for all $c \in (0, \bar{c}_0]$, there exist $\bar{b} = \bar{b}(\Xi, \mathfrak{p}_0) \in (0,1)$ and $\delta = \delta(\Xi, \mathfrak{p}_0) > 0$ satisfying the following:*

1. *(global manipulability) For all $b \in (0,1)$, $D(b,0) < -\delta$ and $D(b,1) > \delta$.*
2. *(local profitable manipulation; type-1) For all $b \in (0, \bar{b}]$, $D(b, \mathfrak{p}_0) < -\delta$ and $R_1(b, \mathfrak{p}_0) > \delta$.*
3. *(local profitable manipulation; type-0) For all $b \in [1-\bar{b}, 1)$, $D(b, 1-\mathfrak{p}_0) > \delta$ and $R_0(b, 1-\mathfrak{p}_0) > \delta$.*

Proof of Proposition C.1. We focus on the type-1 bettor, as the proof for the type-0 bettor follows from the same arguments. Let $b_1 \in (0,1)$, and choose $\bar{c}_0, c, \mathfrak{p}_0, \bar{b}, \delta$ as in Lemma C.2. In particular, $\bar{c}_0$ depends only on α (and hence on $\hat{\Xi}$). Consider the following (behavioral) betting strategy ξ_b for the type-1 informed bettor: under ξ_b, $\mathfrak{p}(b) = \mathfrak{p}_0 \mathbb{I}\{b \le \bar{b}\}$. That is, the probability that he bets positively is $\mathfrak{p}_0$ if $b_t \le \bar{b}$ and 0 otherwise. Lemma C.2 implies that $\mathbb{E}_i^{\pi_B, \xi_b}[L_{t+1} - L_t | b_t = b] = D(b, \mathfrak{p}(b)) = D(b, \mathfrak{p}_0)\, \mathbb{I}\{b \le \bar{b}\} + D(b, 0)\, \mathbb{I}\{b > \bar{b}\} < -\delta$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$, and $R_1(b, \mathfrak{p}(b)) > \delta$ for all $b \in (0, \bar{b}]$. Hence, we are in the first case in the statement of Lemma C.1. As a result, $b_t \to 0$, $L_t \to -\infty$, and $s_t \to s^\pi(0+)$ almost surely, as well as $V_1^{\pi_B, \xi}(T) = \Omega(T)$. ■

Proof of Proposition 1. Proposition 1 is a special case of Proposition C.1, because we focus on the case where $s^\pi(0+) = m_0$ and $s^\pi(1-) = m_1$ in Proposition 1 while we consider all possible cases satisfying $s^\pi(0+) \le m_0$ and $s^\pi(1-) \ge m_1$ in Proposition C.1. ■

Proof of Lemma C.2. We complete the proof in four steps.

Step 1. We claim that there exists $\delta_1 = \delta_1(m_0, m_1, F_\epsilon)$ such that (i) $D(b,0) < -\delta_1$ and (ii) $D(b,1) > \delta_1$ for all $b \in (0,1)$. By Lemma A.3, there exists $\bar{\delta} > 0$ such that $D(b,0) = \log\left(\frac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right) \le \sup_{s \in \mathcal{S}} \log\left(\frac{F_1(s)}{F_0(s)}\right) = -\inf_{s \in \mathcal{S}} \log\left(\frac{F_0(s)}{F_1(s)}\right) \le -\bar{\delta}$ for all $b \in [0,1]$. Similarly, $D(b,1) = \log\left(\frac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right) \ge \inf_{s \in \mathcal{S}} \log\left(\frac{\bar{F}_1(s)}{\bar{F}_0(s)}\right) \ge \bar{\delta}$. To prove our claim in this step, we choose $\delta_1 = \frac{\bar{\delta}}{2}$.

Step 2. We claim that $\bar{c}_0, \mathfrak{p}_0 \in (0,1)$, which depend only on α, satisfying the following for all $c \in (0, \bar{c}_0]$:

- $D(0+, \mathfrak{p}_0) < 0$ and $R_1(0+, \mathfrak{p}_0) > 0$;
- $D(1-, 1-\mathfrak{p}_0) > 0$ and $R_0(1-, 1-\mathfrak{p}_0) > 0$.

In other words, the type-1 (resp. type-0) informed bettor can enjoy profitable manipulation when the market maker's belief is close to 0 (resp. 1). To prove this claim, let us first introduce some constants. Define $\kappa := \frac{-\log 2\alpha}{\log 2(1-\alpha)}$, $\bar{c}_0 := \frac{(\kappa-1)(1-2\alpha)}{2(\kappa-1)(1-\alpha)+1}$, and $\hat{\kappa} := \frac{(\bar{c}_0-2)\alpha+1}{(\bar{c}_0-2)\alpha+1-\bar{c}_0}$. By definition, all of the three constants depend only on α. The following auxiliary result below summarizes the relationship among these constants.

LEMMA C.3. *(ranges of and relations between $\kappa, \hat{\kappa}$ and $\bar{c}_0$) We have $\bar{c}_0 \in (0,1)$. Moreover, for all $c \in (0, \bar{c}_0]$, $\kappa > \hat{\kappa} \ge \frac{(c-2)\alpha+1}{(c-2)\alpha+1-c} > 1$.*

In light of the result above, we choose $\mathfrak{p}_0$ so that $\kappa > \frac{\mathfrak{p}_0}{1-\mathfrak{p}_0} > \hat{\kappa}$. For example, we can choose $\mathfrak{p}_0$ as the solution to the equation $\frac{\mathfrak{p}}{1-\mathfrak{p}} = \frac{\kappa+\hat{\kappa}}{2}$. Such a construction is valid because $\kappa > \hat{\kappa}$ and the mapping $\mathfrak{p} \mapsto \frac{\mathfrak{p}}{1-\mathfrak{p}}$ maps $(0,1)$ onto $(0,\infty)$. Since κ and $\hat{\kappa}$ depend only on α, so does $\mathfrak{p}_0$. Intuitively, we may interpret $\mathfrak{p}_0$ as the probability of honest betting (instead of bluffing), which means positive betting for the type-1 informed bettor and negative betting for the type-0 informed bettor. Similarly, we may interpret $\frac{\mathfrak{p}_0}{1-\mathfrak{p}_0}$ as the probability ratio of honest betting over bluffing. The constants κ and $\hat{\kappa}$ are respectively upper and lower benchmarks for this ratio: if bluffing is sufficiently frequent (i.e., $\frac{\mathfrak{p}_0}{1-\mathfrak{p}_0} < \kappa$) then manipulation happens, and if honest betting is sufficiently frequent (i.e., $\frac{\mathfrak{p}_0}{1-\mathfrak{p}_0} > \hat{\kappa}$) then manipulation is profitable for the informed bettor. A more detailed derivation is presented below. The fact that there is a strict gap between κ and $\hat{\kappa}$ is a key construct in this proof, ensuring that the strategy ξ_b achieves profitable manipulation.

Second, we verify that $D(0+, \mathfrak{p}_0) < 0$ and $D(1-, 1-\mathfrak{p}_0) > 0$. In other words, as the informed bettor bluffs with high probability (i.e., the ratio of honest betting over bluffing less than κ), he misleads the market maker. To see this, observe that

$$
\begin{aligned}
&D(0+, \mathfrak{p}_0) \\
&\quad = (1-\mathfrak{p}_0)\log\left(\tfrac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right) + \mathfrak{p}_0 \log\left(\tfrac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right) \\
&\quad < \tfrac{1}{1+\kappa}\log\left(\tfrac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right) + \tfrac{\kappa}{1+\kappa}\log\left(\tfrac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right) && [\mathfrak{p}_0 < \tfrac{\kappa}{\kappa+1}] \\
&\quad = H\big(F_1(s^\pi(0+))\big) - H\big(F_0(s^\pi(0+))\big) && [H(x) := \tfrac{1}{\kappa+1}\log(x) + \tfrac{\kappa}{\kappa+1}\log(1-x)] \\
&\quad \le H(x_0) - H\big(F_0(s^\pi(0+))\big) && [x_0 := F_1(s^\pi(0+)) \le \alpha] \\
&\quad \overset{(a)}{=} \big(H(x_0) - H(\tfrac{1}{2})\big) \vee 0 \\
&\quad \overset{(b)}{\le} \big(H(\alpha) - H(\tfrac{1}{2})\big) \vee 0 \overset{(c)}{=} 0 \vee 0 = 0.
\end{aligned}
$$

To derive part (a) above, we use the following two facts: (i) $H(\cdot)$ increases in the region $(0, \frac{1}{\kappa+1})$ and decreases in the region $(\frac{1}{\kappa+1}, 1)$, and hence is a quasi-concave function; and (ii) $x_0 = F_1(s^\pi(0+)) \le F_0(s^\pi(0+)) \le F_0(m_0) = \frac{1}{2}$. These two facts imply that $H\big(F_0(s^\pi(0+))\big) \ge H(x_0) \wedge H(\frac{1}{2})$. Rearranging terms, we deduce that $H(x_0) - H\big(F_0(s^\pi(0+))\big) \le \big(H(x_0) - H(\frac{1}{2})\big) \vee 0$. For part (b), we use two facts as well. First, $H(\alpha) - H(\frac{1}{2}) = \frac{1}{\kappa+1}\log(2\alpha) + \frac{\kappa}{\kappa+1}\log(2(1-\alpha)) = 0$, implying that $H(\cdot)$ is increasing in the region $(0, \alpha)$. Second, $x_0 = F_1(s^\pi(0+)) \le F_1(m_0) = \alpha$. Thus, $H(x_0) \le H(\alpha)$. Part (c) follows because $H(\alpha) = H(\frac{1}{2})$. Similarly,

$$
\begin{aligned}
&D(1-, 1-\mathfrak{p}_0) \\
&\quad = \mathfrak{p}_0\log\left(\tfrac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right) + (1-\mathfrak{p}_0)\log\left(\tfrac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right) \\
&\quad > \tfrac{\kappa}{1+\kappa}\log\left(\tfrac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right) + \tfrac{1}{1+\kappa}\log\left(\tfrac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right) && [\mathfrak{p}_0 < \tfrac{\kappa}{\kappa+1}] \\
&\quad = H\big(\bar{F}_1(s^\pi(1-))\big) - H\big(\bar{F}_0(s^\pi(1-))\big) \\
&\quad \ge \big(H(\tfrac{1}{2}) \vee H(y_0)\big) - H(y_0) && [y_0 := \bar{F}_0(s^\pi(1-)) \le \alpha] \\
&\quad = \big(H(\tfrac{1}{2}) - H(y_0)\big) \vee 0 \\
&\quad \ge \big(H(\tfrac{1}{2}) - H(\alpha)\big) \vee 0 = 0.
\end{aligned}
$$

Third, we verify that for all $c \in (0, \bar{c}_0]$, $R_1(0+, \mathfrak{p}_0) > 0$ and $R_1(1-, 1-\mathfrak{p}_0) > 0$. In other words, If the informed bettor bets honestly with high probability (i.e., the ratio of honest betting over bluffing higher than $\hat{\kappa}$), he makes profits from every bet in expectation. To see this point, note that

$$\begin{aligned}
R_1(0+, \mathfrak{p}_0) &= (1-\mathfrak{p}_0) j_1^-(s^\pi(0+)) + \mathfrak{p}_0 j_1^+(s^\pi(0+)) \\
&= (1-\mathfrak{p}_0)\left[(2-c)F_1(s^\pi(0+)) - 1\right] + \mathfrak{p}_0\left[(c-2)F_1(s^\pi(0+)) + 1 - c\right] \qquad \text{[by (2.1)]} \\
&\overset{(d)}{\geq} (1-\mathfrak{p}_0)\left[(2-c)F_1(m_0) - 1\right] + \mathfrak{p}_0\left[(c-2)F_1(m_0) + 1 - c\right] \\
&= (1-\mathfrak{p}_0)\left[(2-c)\alpha - 1\right] + \mathfrak{p}_0\left[(c-2)\alpha + 1 - c\right] \overset{(e)}{>} 0,
\end{aligned}$$

where (d) follows because $\frac{\partial}{\partial F_1(s^\pi(0+))} R_1(0+, \mathfrak{p}_0) = (2-c)(1-2\mathfrak{p}_0) < 0$ as $\frac{\mathfrak{p}_0}{1-\mathfrak{p}_0} > \hat{\kappa} > 1$ by construction, and (e) follows because $\frac{\mathfrak{p}_0}{1-\mathfrak{p}_0} > \hat{\kappa} = \frac{(\bar{c}_0-2)\alpha+1}{(\bar{c}_0-2)\alpha+1-\bar{c}_0} \geq \frac{(c-2)\alpha+1}{(c-2)\alpha+1-c}$ (see Lemma C.3). Similarly,

$$\begin{aligned}
R_0(1-, 1-\mathfrak{p}_0) &= \mathfrak{p}_0 j_0^-(s^\pi(1-)) + (1-\mathfrak{p}_0) j_0^+(s^\pi(1-)) \\
&= \mathfrak{p}_0\left[(2-c)F_0(s^\pi(1-)) - 1\right] + (1-\mathfrak{p}_0)\left[(c-2)F_0(s^\pi(1-)) + 1 - c\right] \qquad \text{[by (2.1)]} \\
&\geq \mathfrak{p}_0\left[(2-c)F_0(m_1) - 1\right] + (1-\mathfrak{p}_0)\left[(c-2)F_0(m_1) + 1 - c\right] \\
&= \mathfrak{p}_0\left[(2-c)(1-\alpha) - 1\right] + (1-\mathfrak{p}_0)\left[(c-2)(1-\alpha) + 1 - c\right] \\
&= \mathfrak{p}_0\left[(c-2)\alpha + 1 - c\right] + (1-\mathfrak{p}_0)\left[(2-c)\alpha - 1\right] > 0.
\end{aligned}$$

The preceding derivations confirm that $D(0+, \mathfrak{p}_0)$, $D(1-, 1-\mathfrak{p}_0)$, $R_1(0+, \mathfrak{p}_0)$, and $R_0(1-, 1-\mathfrak{p}_0)$ are all well-defined as $s^\pi(0+)$ and $s^\pi(1-)$ exist.

Step 3. By Step 2, there exists $\bar{b}, \delta_2, \delta_3, \delta_4, \delta_5 > 0$, all of which depend only on $\mathfrak{p}_0$ and Ξ, such that

- *(local profitable manipulation; type-*1*)* $D(b, \mathfrak{p}_0) < -\delta_2$ and $R_1(b, \mathfrak{p}_0) > \delta_3$ for all $b \in (0, \bar{b}]$;
- *(local profitable manipulation; type-*0*)* $D(b, 1-\mathfrak{p}_0) > \delta_4$ and $R_0(b, 1-\mathfrak{p}_0) > \delta_5$ for all $b \in (1-\bar{b}, 1]$.

The existence is guaranteed by the local continuity of $D(b, \mathfrak{p}_0)$ and $R_1(b, \mathfrak{p}_0)$ with respect to b at $0+$, as well as that of $D(b, 1-\mathfrak{p}_0)$ and $R_0(b, 1-\mathfrak{p}_0)$ with respect to b at $1-$.

Step 4. Based on Steps 1 and 3, we complete the proof by choosing $\delta := \min\{\delta_1, \delta_2, \delta_3, \delta_4, \delta_5\}$. ∎

Proof of Lemma C.3. To prove that $\bar{c}_0 > 0$, it suffices to verify that $\kappa > 1$. Note that $\alpha \in (0, \frac{1}{2})$, and $\log\big(2(1-\alpha)\big) > 0 > \log 2\alpha$. Moreover, due to Jensen's inequality, $\log 2(1-\alpha) + \log 2\alpha < 2\log\frac{2(1-\alpha)+2\alpha}{2} = 0$. Thus, $\kappa = \frac{-\log 2\alpha}{\log 2(1-\alpha)} > 1$, and $\bar{c}_0 > 0$. To see why $\bar{c}_0 < 1$, note that $\bar{c}_0 = \frac{(\kappa-1)(1-2\alpha)}{2(\kappa-1)(1-\alpha)+1} < \frac{(\kappa-1)(1-2\alpha)}{(\kappa-1)(1-\alpha)+1} < \frac{(\kappa-1)(1-\alpha)}{(\kappa-1)(1-\alpha)+1} < 1$. We have thus verified that $\bar{c}_0 \in (0,1)$. Now, we choose $c \in (0, \bar{c}_0]$ to verify that $\kappa > \hat{\kappa} \geq \frac{(c-2)\alpha+1}{(c-2)\alpha+1-c} > 1$. To see why $\frac{(c-2)\alpha+1}{(c-2)\alpha+1-c} > 1$, note that $c \leq \bar{c}_0 < \frac{1-2\alpha}{1-\alpha} \Rightarrow (1-c) + (c-2)\alpha > 0$. Moreover, $\big[(c-2)\alpha+1\big] - \big[(c-2)\alpha+1-c\big] = c > 0$. Hence, $\frac{(c-2)\alpha+1}{(c-2)\alpha+1-c} > 1$. To see why $\hat{\kappa} \geq \frac{(c-2)\alpha+1}{(c-2)\alpha+1-c}$, observe that the function $c \mapsto \frac{(c-2)\alpha+1}{(c-2)\alpha+1-c}$ increases in c because $\alpha < 1$. As a result, $c \leq \bar{c}_0$, which implies that $\frac{(c-2)\alpha+1}{(c-2)\alpha+1-c} \leq \frac{(\bar{c}_0-2)\alpha+1}{(\bar{c}_0-2)\alpha+1-\bar{c}_0} = \hat{\kappa}$. Finally, let us verify that $\kappa > \hat{\kappa}$. Note that $\bar{c}_0 = \frac{(\kappa-1)(1-2\alpha)}{2(\kappa-1)(1-\alpha)+1} < \frac{(\kappa-1)(1-2\alpha)}{(\kappa-1)(1-\alpha)+1}$. By rearranging terms, we see that

$$(\kappa-1)(1-2\alpha) > [(1-\alpha)(\kappa-1)+1]\bar{c}_0 \implies (1-2\alpha)\kappa + 2\alpha - 1 > (1-\alpha)\bar{c}_0\kappa + \alpha\bar{c}_0$$

$$\implies [(1-\bar{c}_0)+(\bar{c}_0-2)\alpha]\kappa > 1+(\bar{c}_0-2)\alpha$$
$$\implies \kappa > \tfrac{1+(\bar{c}_0-2)\alpha}{(1-\bar{c}_0)+(\bar{c}_0-2)\alpha} = \hat{\kappa}.$$

We have thus completed the proof. ■

C.5. Profitable Honest Betting (Proof of Proposition C.2)

Proof of Proposition C.2. Choose

$$\bar{c}_1 := \min\left\{\max\left\{\tfrac{2F_0(s^{\pi_B}(0+))-1}{2F_0(s^{\pi_B}(0+))}, \tfrac{1-2F_1(s^{\pi_B}(1-))}{2\bar{F}_1(s^{\pi_B}(1-))}\right\}, \tfrac{1}{2}\right\}. \tag{C.3}$$

By construction, $\bar{c}_1 \in (0,1)$, because either $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$. Let $c \in (0, \bar{c}_1]$.

First, suppose that $s^{\pi_B}(0+) > m_0$. We define the type-0 informed bettor's "honest" strategy ξ_h as always betting negatively, i.e., under ξ_h, $\mathfrak{p}(b) = 0$ for all $b \in (0,1)$. We claim that there exists $\delta, \bar{b} > 0$ such that $\mathbb{E}_0^{\pi_B,\xi_h}[L_{t+1} - L_t | b_t = b] = D(b,0) < -\delta$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$, and $R_0(b,\mathfrak{p}(b)) = R_0(b,0) > 0$ for all $b \in (0,\bar{b}]$. To that end, we deduce from Lemma A.3 that there exists $\bar{\delta} > 0$ such that $D(b,0) = \log \frac{F_1(s^{\pi}(b))}{F_0(s^{\pi}(b))} \le -\bar{\delta}$. Note that $F_0(s^{\pi_B}(0+)) > F_0(m_0) = \frac{1}{2}$, and $0 < c < \frac{2F_0(s^{\pi_B}(0+))-1}{F_0(s^{\pi_B}(0+))}$. In that case, $R_0(0+,0) = j_0^-(s^{\pi_B}(0+)) = (2-c)F_0(s^{\pi_B}(0+)) - 1 = [2F_0(s^{\pi_B}(0+)) - 1] - cF_0(s^{\pi_B}(0+)) > 0$. Hence, there exist $\varepsilon > 0$ and $\bar{b}$ (independent of T) such that $R_0(b,0) \ge \varepsilon$ for all $b \in (0,\bar{b}]$. Choosing $\delta = \min\{\bar{\delta}, \varepsilon\}$, we deduce from Lemma C.1 that $b_t \to 0$ and $s_t \to s^{\pi_B}(0+)$ almost surely, and that $V_0^{\pi_B,\xi_h}(T) = \Omega(T)$.

In the case where $s^{\pi_B}(1-) < m_1$, our analysis is similar. In fact, type-1 informed bettor's honest policy ξ_h is specified as always betting positively. Notice that $0 < c < \frac{1-2F_1(s^{\pi_B}(1-))}{\bar{F}_1(s^{\pi_B}(1-))}$; thus, $R_1(0+,0) = j_1^+(s^{\pi_B}(1-)) = (c-2)F_1(s^{\pi_B}(1-)) + 1 - c = [1 - 2F_1(s^{\pi_B}(1-))] - c\bar{F}_1\big(s^{\pi_B}(1-)\big) > 0$. The rest of the proof for this case follows by repeating the above arguments. ■

Appendix D: On the Success of Bayesian Policies (Theorem 2)

This section provides the details for the proof of Theorem 2, as well as additional discussions on the myopic Bayesian policy (MBP) as a special case of BP. In what follows, we use the following little-o notation: for all functions f, g defined in a neighborhood around zero, we say that $f(x) = o(g(x))$ if $\lim_{x\to 0} \frac{f(x)}{g(x)} = 0$.

D.1. Proof of Theorem 2

Let $s^{\pi_B}(\cdot)$ be a regular pricing function and $b_1 \in (0,1)$. We assume without loss of generality that $i = 0$ (the analysis that follows can be repeated verbatim for the case where $i = 1$). Throughout this proof, we also denote $\mathbb{P}_0^{\pi_B,\xi_\emptyset}(\cdot)$ and $\mathbb{E}_0^{\pi_B,\xi_\emptyset}[\cdot]$ as $\mathbb{P}_0(\cdot)$ and $\mathbb{E}_0[\cdot]$ for brevity. We complete the proof in three steps.

Step 1. We claim that there exists $\delta > 0$ such that $\mathbb{E}_0[L_{t+1} - L_t | b_t = b] < -\delta$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$. To prove this claim, we deduce from Lemma A.3 that π_B is a $\bar{\delta}$-discriminative policy (in the sense of Harrison et al. 2012) for some $\bar{\delta} > 0$. By Lemma A.2 in Harrison et al. (2012), $\mathbb{E}_0[L_{t+1} - L_t | b_t = b] < -2\bar{\delta}^2$ for all $b \in (0,1)$ and $t \in \mathbb{Z}_+$. As a result, Lemma C.1 implies that $\mathbb{E}_0[b_t] = O(e^{-\lambda t})$ for some $\lambda > 0$, and $s_t \to s^{\pi}(0+)$ almost surely.

Step 2. We claim that $\mathfrak{d}_t \to 0$ almost surely, and $\mathbb{E}_0[\mathfrak{d}_t] = O\big(e^{-\lambda t}\big)$, thus verifying Statements (T2:1) and (T2:2) in Theorem 2. To that end, choose $\mu_0 > 0$ such that $\mathbb{E}_0[b_t] \le \mu_0 e^{-\lambda t}$ for all $t \in \mathbb{Z}_+$. Since $s^{\pi_B}(\cdot)$ is regular, $\limsup_{b\downarrow 0} \frac{|s^{\pi_B}(b)-m_0|}{b} < \infty$. That is, there exists $C > 0$ such that $|s^{\pi_B}(b) - m_0| < bC$ for all $b \in (0,\delta)$. Hence, $\mathfrak{d}_t = |s^{\pi_B}(b_t) - m_0| \to 0$ $\mathbb{P}_0$-almost surely. Moreover,

$$\begin{aligned}\mathbb{E}_0[\mathfrak{d}_t] &= \underbrace{\mathbb{E}_0[\mathfrak{d}_t \big| b_t \ge \delta]}_{\le s_H - s_L} \underbrace{\mathbb{P}_0(b_t \ge \delta)}_{\le \frac{\mathbb{E}_0[b_t]}{\delta}} + \underbrace{\mathbb{E}_0[\mathfrak{d}_t | b_t \le \delta]}_{\le \mathbb{E}_0[b_t]C} \underbrace{\mathbb{P}_0(b_t \le \delta)}_{\le 1} \\ &\le \tfrac{(s_H - s_L)\mathbb{E}_0[b_t]}{\delta} + \mathbb{E}_0[b_t]C = \big[\tfrac{s_H - s_L}{\delta} + C\big]\mathbb{E}_0[b_t] \le \big[\tfrac{s_H - s_L}{\delta} + C\big]\mu_0 e^{-\lambda t}.\end{aligned}$$

Therefore, $\mathbb{E}_0[\mathfrak{d}_t] = O\big(e^{-\lambda t}\big)$.

Step 3. We claim that $\Delta_0^{\pi_B,\xi_\emptyset}(T) = O(1)$, verifying Statement (T2:3) in Theorem 2. For this purpose, observe that $\frac{c}{2} - r_0(s^\pi(b)) = o(b)$ and $\frac{c}{2} - r_1(s^\pi(b)) = o(1-b)$. To see this, recall from (2.4) that $r_i(s) = (2c-4)\left(F_i(s) - \frac{1}{2}\right)^2 + \frac{c}{2}$ for $i \in \{0,1\}$. Hence, $r_i'(s) = (2c-4)(2F_i(s)-1)f_i(s)$. In particular, $r_0'(m_0) = r_1'(m_1) = 0$. By Taylor's theorem, we have $r_0(s) = \frac{c}{2} - o(s - m_0)$ and $r_1(s) = \frac{c}{2} - o(m_1 - s)$. Because the BP in question is regular, $\max\left\{\limsup_{b\downarrow 0} \frac{|s^{\pi_B}(b)-m_0|}{b}, \limsup_{b\uparrow 1} \frac{|s^{\pi_B}(b)-m_1|}{1-b}\right\} < \infty$. Thus, $r_0(s^\pi(b)) - \frac{c}{2} = o\big(s^{\pi_B}(b) - m_0\big) = o(b)$ and $r_1(s^\pi(b)) - \frac{c}{2} = o\big(m_1 - s^{\pi_B}(b)\big) = o(1-b)$. Based on this, and repeating the arguments in Step 2, we deduce that there exists $\mu_2 > 0$ such that $\mathbb{E}_0[\frac{c}{2} - r_0(s^\pi(b_t))] \le \mu_2 e^{-\lambda t}$. Consequently, $\Delta_0^{\pi_B,\xi_\emptyset}(T) = \frac{cT}{2} - \mathbb{E}_i\big[\sum_{t=1}^T \mathbb{I}\{(X - s_t)d_t < 0\} - (1-c)\mathbb{I}\{(X-s_t)d_t > 0\}\big] = \sum_{t=1}^T \mathbb{E}_0\left[\frac{c}{2} - r_0(s^\pi(b_t))\right] \le \sum_{t=1}^T \mu_2 e^{-\lambda t} < \sum_{t=1}^\infty \mu_2 e^{-\lambda t} = O(1)$. ■

D.2. Discussion on an Equivalent Interpretation of the Absence of the Informed Bettor

The informed bettor's vacuous strategy $\xi_\emptyset$ is equivalent to his best response strategy when the commission rate c is sufficiently high. We formalize this observation in the result below.

LEMMA D.1. *For every policy π of the market maker, hypothesis $i \in \{0,1\}$, and sufficiently large c that depends only on $F_0(\cdot)$ and $F_1(\cdot)$, $\xi_i^* = \xi_\emptyset$.*

Proof of Lemma D.1. Let $\tilde{\alpha} := \min\{F_1(s_L), 1 - F_0(s_H)\}$, which is strictly positive by Assumption 1. Let $c > \frac{1-2\tilde{\alpha}}{1-\tilde{\alpha}}$ so that $\tilde{\alpha} < \frac{1-c}{2-c}$ (this choice of c is feasible because $\frac{1-2\tilde{\alpha}}{1-\tilde{\alpha}} < 1$). For all $s \in \mathcal{S}$,

$$\begin{aligned}j_1^+(s) &= (c-2)F_1(s) + 1 - c \le (c-2)\tilde{\alpha} + 1 - c < 0;\\ j_0^+(s) &= (c-2)F_0(s) + 1 - c \le (c-2)F_1(s) + 1 - c < 0;\\ j_0^-(s) &= (2-c)F_0(s) - 1 \le (2-c)(1-\tilde{\alpha}) - 1 = (1-c) - \tilde{\alpha}(2-c) < 0;\\ j_1^-(s) &= (2-c)F_1(s) - 1 \le (2-c)F_0(s) - 1 < 0.\end{aligned}$$

That is, the informed bettor finds it (strictly) better off not to bet at all, regardless of the market maker's spread line s. In that case, it is easy to verify that the informed bettor's best response strategy is to quit the market, regardless the market maker's policy. ■

D.3. Discussion on the Myopic Bayesian Policy (MBP)

In this section, we briefly discuss the MBP. The main purpose of this discussion is to connect to the previous dynamic pricing and learning literature (in particular, Harrison et al. 2012 and Chen and Wang 2016), and demonstrate that the MBP satisfies the additional regularity condition in Theorem 2. For all $b \in [0,1]$, we denote the market maker's myopic expected profit function by

$$r_b(s) := br_1(s) + (1-b)r_0(s) = (2c-4)\left[b(F_1(s)-\tfrac{1}{2})^2 + (1-b)(F_0(s)-\tfrac{1}{2})^2\right] + \tfrac{c}{2}. \tag{D.1}$$

The MBP uses the following policy function:

$$s^{\pi_B}(b) := s^\dagger(b) = \sup \arg\max_{s\in\mathbb{R}} \{r_b(s)\}. \tag{D.2}$$

The supremum operator in Equation (D.2) is introduced to ensure that $s^\dagger(b)$ is uniquely defined. In the following series of results, we establish the key properties of the MBP.

LEMMA D.2. *For* $i \in \{0,1\}$, $\arg\max_{s\in\mathbb{R}} r_i(s) = m_i$.

Proof. Observe that $r_0(s) = (2c-4)\left(F_0(s)-\frac{1}{2}\right)^2 + \frac{c}{2}$ and $r_1(s) = (2c-4)\left(F_1(s)-\frac{1}{2}\right)^2 + \frac{c}{2}$. The statement follows by noticing that $F_0(\cdot)$ and $F_1(\cdot)$ have unique medians m_0 and m_1, respectively, due to (A1:3). ■

LEMMA D.3. *For* $b \in [0,1]$, $s^\dagger(b) \in [m_0, m_1]$. *Moreover, for* $b \in [0,1]$, $s^\dagger(b)$ *strictly increases in* b.

Proof. To prove the first statement, note that for $i \in \{0,1\}$, $r_i(\cdot)$ has a unique maximizer, m_i. Thus, it suffices to consider $b \in (0,1)$. Observe that $r_b'(s) = b(2c-4)f_1(s)\big(2F_1(s)-1\big) + (1-b)(2c-4)f_0(s)\big(2F_0(s)-1\big)$. Hence, for all $s \geq m_1$, $r_b'(s) \leq b(2c-4)f_1(m_1)\big(2F_1(m_1)-1\big) + (1-b)(2c-4)f_0(m_1)\big(2F_0(m_1)-1\big)$. By (A1:3), $F_0(m_1) > F_0(x) > F_0(m_0) = \frac{1}{2} = F_1(m_1) > F_1(x) > F_1(m_0)$ for all $x \in (m_0, m_1)$. As a result, $r_b'(s) \leq 0$ and $r_b(s) \leq s_b(m_1)$ for all $s \geq m_1$. Moreover, $r_b'(m_1) = (1-b)(2c-4)f_0(m_1)\big(2F_0(m_1)-1\big) < 0$, because $b<1$, $c<1$, $f_0(m_1) > 0$ by (A1:3), and $2F_0(m_1) - 1 > 0$. Thus, $r_b(s) < s_b(m_1)$ for all $s > m_1$. Consequently, by a similar argument, $r_b(s) < s_b(m_0)$ for all $s < m_0$.

To prove the second statement, we deduce from the first statement that it suffices to consider the case where $s \in [m_0, m_1]$. Note that, for all $b \in [0,1]$ and $s \in [m_0, m_1]$,

$$\begin{aligned}\frac{\partial^2 r_b(s)}{\partial b \partial s} &= \frac{\partial}{\partial b}[r_b'(s)] = \frac{\partial}{\partial b}\big[b(2c-4)f_1(s)(2F_1(s)-1) + (1-b)(2c-4)f_0(s)(2F_0(s)-1)\big] \\ &= \underbrace{(2c-4)}_{<0}\Big[\underbrace{f_1(s)(2F_1(s)-1) + f_0(s)(1-2F_0(s))}_{<0}\Big] > 0,\end{aligned}$$

where the strict inequality is due to (A1:3). Thus, $r_b(s)$ is a strictly supermodular function of (b,s). Consequently, we deduce from Topkis's theorem (Topkis 1978) that $s^\dagger(b)$ is non-decreasing in b. To see why $s^\dagger(b)$ is strictly increasing in b, note that $s^\dagger(b)$ satisfies the first order condition $r_b'(s) = 0$, which is equivalent to

$$b = \frac{f_0(s)(2F_0(s)-1)}{f_0(s)(2F_0(s)-1) + f_1(s)(1-2F_1(s))} =: \mathcal{G}(s). \tag{D.3}$$

For all $s \in [m_0, m_1]$, the denominator of $\mathcal{G}(s)$ is strictly positive and hence well-defined. Suppose that there exist $b_x, b_y \in [0,1]$ such that $s^\dagger(b_x) = s^\dagger(b_y)$. Then, the first order condition states that $b_x = \mathcal{G}(s^\dagger(b_x)) = \mathcal{G}(s^\dagger(b_y)) = b_y$. Hence, $s^\dagger(b)$ must be strictly increasing in b. ■

PROPOSITION D.1. *There exist $C_0, C_1 > 0$ such that $s^\dagger(b) = m_0 + C_0 b + o(b) = m_1 + C_1(b-1) + o(1-b)$. That is, $s^\dagger(b)$ is a regular pricing function.*

Proof. Recall that $s^\dagger(\cdot)$ satisfies the first order condition in Equation (D.3). By Lemma D.2, m_0 and m_1 are the unique solutions to $\mathcal{G}(s) = 0$ and $\mathcal{G}(s) = 1$, respectively. Moreover, it is straightforward verify that

$$\mathcal{G}'(m_0) = \frac{2f_0^2(m_0)}{f_1(m_0)(1-2F_1(m_0))} > 0 \text{ and } \mathcal{G}'(m_1) = \frac{2f_1^2(m_1)}{f_0(m_1)(2F_0(m_1)-1)} > 0.$$

Note that the strict positivity is guaranteed by (A1:3). Define $C_0 := \frac{1}{\mathcal{G}'(m_0)} > 0$ and $C_1 := \frac{1}{\mathcal{G}'(m_1)} > 0$. By the inverse function theorem (Rudin 1976, Theorem 9.24), $\mathcal{G}^{-1}(b)$ is uniquely defined in $[0,\varepsilon) \cup (1-\varepsilon, 1]$ for some $\varepsilon > 0$, with $(\mathcal{G}^{-1})'(i) = C_i$ for $i \in \{0,1\}$. Thus, by Taylor's theorem, for $b \in [0,\varepsilon) \cup (1-\varepsilon, 1]$, $s^\dagger(b) = \mathcal{G}^{-1}(b) = m_0 + C_0 b + o(b) = m_1 + C_1(b-1) + o(1-b)$. For $b \in [\varepsilon, 1-\varepsilon]$, $s^\dagger(b)$ is uniquely defined, and the statement of this lemma trivially holds. ■

Appendix E: Residual Probability Representation of Inertial Policies (Proposition 2)

In this section, we provide the details for the proof of Proposition 2, as well as additional discussions on the extension of the function $\rho(\cdot)$ from $\mathbb{Z}_+$ to $\mathbb{Z}$.

E.1. Proof of Proposition 2

We first extend $\rho(\cdot)$ from $\mathbb{Z}_+$ to $\mathbb{Z}$ as follows:

$$\rho(z) = \begin{cases} \frac{1}{2} - F_1\left(\frac{m_0+m_1}{2}\right) & \text{if } z = 0, \\ \frac{1}{2} - F_1 \circ F_0^{-1}\left(\frac{1}{2} + \rho(-z)\right) & \text{if } z \in \mathbb{Z}_-. \end{cases} \tag{E.1}$$

In particular, we find it useful to combine (E.1) with our specific construction of the residual probability sequence $\{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$, and write out the extended version of $\rho(\cdot)$ as the following:

$$\rho(z) = \begin{cases} \frac{1}{r_0+rz} & \text{if } z \in \mathbb{N}, \\ \frac{1}{2} - F_1 \circ F_0^{-1}\left(\frac{1}{2} + \frac{1}{r_0 - rz}\right) & \text{if } z \in \mathbb{Z}_-. \end{cases} \tag{E.2}$$

We complete the rest of proof in three steps.

Step 1. We claim that both $\tilde{s}(\cdot)$ in (4.3) and the extension of $\rho(\cdot)$ in (E.1) are well-defined (we prove this statement to verify that the inverse functions in (4.3) and (E.1) both exist). Because $\rho(z) \in (0, \frac{1}{2} - \alpha)$ for all $z \in \mathbb{Z}_+$, it suffices to verify that (i) $F_0^{-1}(\cdot)$ exists in $(\frac{1}{2}, 1-\alpha)$, and (ii) $F_1^{-1}(\cdot)$ exists in $(\alpha, \frac{1}{2})$. Moreover, as $F_\epsilon(\cdot) = F_i(\cdot + m_i)$ for $i \in \{0,1\}$, it is sufficient to show that $F_\epsilon^{-1}(\cdot)$ exists in $(\alpha, 1-\alpha)$. Note that $F_\epsilon(m_0 - m_1) = F_1(m_0) = \alpha$, and $F_\epsilon(m_1 - m_0) = F_0(m_1) = 1 - F_1(m_0) = 1 - \alpha$, where the second equality follows from Lemma A.1. Since $|m_1 - m_0| = m_1 - m_0 \le s_H - m_0$, we deduce from (A1:3) that F_ϵ is strictly increasing in the interval $[m_0 - m_1, m_1 - m_0]$. This means that $F_\epsilon^{-1}(\cdot)$ exists and strictly increases in $(\alpha, 1-\alpha)$.

Step 2. We now claim that the construction of $\tilde{s}(\cdot)$ and the extension of $\rho(\cdot)$ satisfy (4.2). First, let $z \in \mathbb{Z}_+$. Then, $F_0(\tilde{s}(-z)) \overset{(4.3)}{=} F_0\left(F_0^{-1}\left(\frac{1}{2} + \rho(z)\right)\right) = \frac{1}{2} + \rho(z)$. Furthermore, $F_1(\tilde{s}(z)) \overset{(4.3)}{=} F_1\left(F_1^{-1}\left(\frac{1}{2} - \rho(z)\right)\right) =$

$\frac{1}{2}-\rho(z)$. Now, let $z=0$. In this case, $F_0(\tilde{s}(0)) \overset{(4.3)}{=} F_0\left(\frac{m_0+m_1}{2}\right) \overset{\text{Lem. A.1}}{=} 1-F_1\left(\frac{m_0+m_1}{2}\right) \overset{(E.1)}{=} \frac{1}{2}+\rho(0)$. Moreover, $F_1(\tilde{s}(0)) \overset{(4.3)}{=} F_1\left(\frac{m_0+m_1}{2}\right) \overset{(E.1)}{=} \frac{1}{2}-\rho(0)$. Finally, let $z\in\mathbb{Z}_-$. To analyze this case, we use the following result, which states that the pricing function $\tilde{s}(\cdot)$ in (4.3) is symmetric around the point $\left(0,\frac{m_0+m_1}{2}\right)$.

LEMMA E.1. *For all $z\in\mathbb{Z}$, $\tilde{s}(z)+\tilde{s}(-z)=m_0+m_1$.*

Thus, $F_0(\tilde{s}(-z)) \overset{\text{Lem. A.1 \& E.1}}{=} 1-F_1(\tilde{s}(z)) \overset{(4.3)}{=} 1-F_1\circ F_0^{-1}\left(\frac{1}{2}+\rho(-z)\right) \overset{(E.1)}{=} \frac{1}{2}+\rho(z)$. In addition, $F_1\big(\tilde{s}(z)\big) \overset{(4.3)}{=} F_1\left(F_0^{-1}\left(\frac{1}{2}+\rho(-z)\right)\right) \overset{(E.1)}{=} \frac{1}{2}-\rho(z)$.

Step 3. Lastly, we claim that given $\{\rho(z), z\in\mathbb{Z}_+\}$, the construction of $\tilde{s}(\cdot)$ and the extension of $\rho(\cdot)$ that satisfy (4.2) are both unique. The proof of uniqueness also provides us some intuition for the choices of $\tilde{s}(\cdot)$ and $\rho(\cdot)$. Note that both $F_0(\cdot)$ and $F_1(\cdot)$ are strictly increasing by Step 1. Thus, we first uniquely determine the values of $\tilde{s}(\cdot)$ from (4.2). In fact, given the residual probability sequence $\{\rho(z), z\in\mathbb{Z}_+\}$,

- $\{\tilde{s}(z), z\in\mathbb{Z}_+\}$ is uniquely determined by the relationship $F_1\big(\tilde{s}(z)\big)=\frac{1}{2}-\rho(z)$ for all $z\in\mathbb{Z}_+$ (this corresponds to the zone where the betting sequence is in favor of H_1, and hence $\tilde{s}(\cdot)$ is closer to m_1),
- $\{\tilde{s}(z), z\in\mathbb{Z}_-\}$ is uniquely determined by the relationship $F_0\big(\tilde{s}(-z)\big)=\frac{1}{2}-\rho(z)$ for all $z\in\mathbb{Z}_+$ (this corresponds to the zone where the betting sequence is in favor of H_0, and hence $\tilde{s}(\cdot)$ is closer to m_0),
- $\tilde{s}(0)$ is uniquely determined by the relationships $F_1\big(\tilde{s}(0)\big)=\frac{1}{2}-\rho(0)$ and $F_0\big(\tilde{s}(0)\big)=\frac{1}{2}-\rho(0)$, which imply that $F_1\big(\tilde{s}(0)\big)+F_0\big(\tilde{s}(0)\big)=1$ (this corresponds to the zone where the betting sequence is in favor of neither hypothesis, and hence $\tilde{s}(0)=\frac{m_0+m_1}{2}$).

Because the values of $\tilde{s}(\cdot)$ are uniquely determined, the value of $\rho(z)$ for every $z\in\mathbb{Z}_-\cup\{0\}$ is uniquely determined by (4.2). ■

Proof of Lemma E.1. Let $z\in\mathbb{Z}$. If $z=0$, then we deduce from (4.3) that $\tilde{s}(0)+\tilde{s}(0)=m_0+m_1$, and the claim holds. On the other hand, if $z\in\mathbb{Z}_+$, then we note that $F_0\left(m_0+m_1-F_1^{-1}\left(\frac{1}{2}-\rho(z)\right)\right) \overset{\text{Lem. A.1}}{=} 1-F_1\circ F_1^{-1}\left(\frac{1}{2}-\rho(z)\right)=\frac{1}{2}+\rho(z)$. In this case, by Step 1 of the proof of Proposition 2, the inverse of F_0 is well-defined. Thus, $F_0^{-1}\left(\frac{1}{2}+\rho(z)\right)=m_0+m_1-F_1^{-1}\left(\frac{1}{2}-\rho(z)\right)$. As a result, $\tilde{s}(z)+\tilde{s}(-z) \overset{(4.3)}{=} F_1^{-1}\left(\frac{1}{2}-\rho(z)\right)+F_0^{-1}\left(\frac{1}{2}+\rho(z)\right)=m_0+m_1$. ■

E.2. Discussions

Let us now explore general properties of the extended residual probability sequence $\{\rho(z), z\in\mathbb{Z}\}$ in (E.1).

LEMMA E.2. *(upper bound) $\rho(z)<\frac{1}{2}-\alpha$ for all $z\in\mathbb{Z}$.*

Proof. By definition, $\rho(z)<\frac{1}{2}-\alpha$ for all $z\in\mathbb{Z}_+$. Note that $\rho(0)=\frac{1}{2}-F_1\left(\frac{m_0+m_1}{2}\right)<\frac{1}{2}-F_1(m_0)=\frac{1}{2}-\alpha$ by the (strict) monotonicity of $F_1(\cdot)$. For all $z\in\mathbb{Z}_-$, $\rho(z)=\frac{1}{2}-F_1\circ F_0^{-1}\left(\frac{1}{2}+\rho(-z)\right)<\frac{1}{2}-F_1\circ F_0^{-1}\left(\frac{1}{2}\right)=\frac{1}{2}-\alpha$ by the (strict) monotonicity of $F_1(\cdot)$ and $F_0(\cdot)$. ■

LEMMA E.3. *(lower bound) $\rho(z)>0$ for all $z\in\mathbb{Z}$. Moreover, if $\sup\{\rho(z):z\in\mathbb{Z}_+\}<\frac{1}{2}-\alpha$, then $\inf\{\rho(z):z\le M\}>0$ for all $M\in\mathbb{Z}$.*

Proof. By definition, $\rho(z) > 0$ for all $z \in \mathbb{Z}_+$. Moreover, $\rho(0) = \frac{1}{2} - F_1\left(\frac{m_0+m_1}{2}\right) > \frac{1}{2} - F_1(m_1) = 0$ by the (strict) monotonicity of $F_1(\cdot)$. Choose $\delta \geq 0$ so that $\rho(z) \leq \frac{1}{2} - \alpha - \delta$ for all $z \in \mathbb{Z}_+$. For all $z \in \mathbb{Z}_-$, $\rho(z) = \frac{1}{2} - F_1 \circ F_0^{-1}\left(\frac{1}{2} + \rho(-z)\right) \geq \frac{1}{2} - F_1 \circ F_0^{-1}(1-\alpha-\delta) \geq 0$, where the last inequality is strict if $\delta > 0$ by the (strict) monotonicity of $F_1(\cdot)$ and $F_0(\cdot)$. If $\delta > 0$, $\rho(z) \geq \min\{\frac{1}{2} - F_1 \circ F_0^{-1}(1-\alpha-\delta), \rho(1), \ldots, \rho(M)\} > 0$ for all $z \leq M$. ■

Appendix F: Key Proof Steps for the Results in Section 4

Proof of Lemma 4. Since Y_t is time-homogeneous, we have $f(z) = \mathbb{E}[u(Y_2)|Y_1 = z] - u(z) = \mathbb{E}[u(Y_{t+1})|Y_t = z] - u(z)$ for all z, t. Consequently,

$$\begin{aligned} f(Y_1)+f(Y_2)+\cdots+f(Y_t) &= \mathbb{E}[u(Y_2|Y_1)] - u(Y_1) + \mathbb{E}[u(Y_3|Y_2)] - u(Y_2) + \cdots + \mathbb{E}[u(Y_{t+1}|Y_t)] - u(Y_t) \\ &= \underbrace{\sum_{i=1}^{t-1}\Big(\mathbb{E}[u(Y_{i+1})|Y_i] - u(Y_{i+1})\Big)}_{\mathcal{M}_t} + \mathbb{E}[u(Y_{t+1})|Y_t] - u(Y_1) \end{aligned}$$

for all t, where $\mathcal{M}_t$ is a martingale with respect to the σ-algebra $\mathcal{F}_t = \sigma(Y_1, \cdots, Y_t)$ because

$$\mathbb{E}[\mathcal{M}_{t+1}|\mathcal{F}_t] \overset{(a)}{=} \mathbb{E}\left[\mathcal{M}_t + \mathbb{E}[u(Y_{t+1})|Y_t] - u(Y_{t+1})|\mathcal{F}_t\right] \overset{(b)}{=} \mathcal{M}_t + \mathbb{E}[u(Y_{t+1})|Y_t] - \mathbb{E}[u(Y_{t+1})|\mathcal{F}_t] \overset{(c)}{=} 0.$$

In the preceding chain of equalities, (a) follows from the definition of $\mathcal{M}_t$, (b) from the fact that both $\mathcal{M}_t$ and $\mathbb{E}[u(Y_{t+1})|Y_t]$ are $\mathcal{F}_t$-measurable, and (c) from the Markov property of the Markov chain Y_t. Thus,

$$\mathbb{E}[f(Y_1)] + \mathbb{E}[f(Y_2)] + \cdots + \mathbb{E}[f(Y_t)] = \mathbb{E}\mathcal{M}_t + \mathbb{E}\mathbb{E}[u(Y_{t+1})|Y_t] - \mathbb{E}u(Y_1) = \mathbb{E}u(Y_{t+1}) - \mathbb{E}u(Y_1). \blacksquare$$

The following lemma is another key proof step for the results in Section 4.

LEMMA F.1. *For all constants $\delta \in \mathbb{R}, \hat{z} \in \mathbb{Z}$, and any two sequences $x(\cdot), p(\cdot) : \{\hat{z}, \hat{z}+1, \ldots\} \to \mathbb{R}$, consider the difference equation (F.1) below:*

$$\begin{cases} y(\hat{z}-1) = 0, \ y(\hat{z}) = \delta, \\ p(z)y(z+1) + \bar{p}(z)y(z-1) - y(z) = x(z) \ \textit{for all } z \geq \hat{z}, \end{cases} \tag{F.1}$$

where $\bar{p}(z) := 1 - p(z)$. If $p(\cdot) \notin \{0,1\}$, the difference equation (F.1) above admits the following solution $y_\delta^{\hat{z}}(\cdot) : \{\hat{z}-1, \hat{z}, \ldots\} \to \mathbb{R}$:

$$y_\delta^{\hat{z}}(z) = \begin{cases} 0 & \textit{if } z = \hat{z}-1, \\ \left(1 + \sum_{n=\hat{z}}^{z-1} \prod_{m=\hat{z}}^{n} \frac{\bar{p}(m)}{p(m)}\right)\delta + \sum_{n=\hat{z}}^{z-1}\sum_{k=\hat{z}}^{n}\left(\prod_{m=k}^{n}\frac{\bar{p}(m)}{p(m)}\right)\frac{x(k)}{\bar{p}(k)} & \textit{if } z \geq \hat{z}. \end{cases} \tag{F.2}$$

In the notation above, we use the convention that $\sum_{k=n}^{n-1}(\cdot) := 0$, and $\prod_{k=n}^{n-1}(\cdot) := 1$.

Proof. Let $\delta \in \mathbb{R}$, $\hat{z} \in \mathbb{Z}$, and $x(\cdot)$ be a function from $\{\hat{z}, \hat{z}+1, \ldots\}$ to $\mathbb{R}$. By construction, $y_\delta^{\hat{z}}(\cdot)$ satisfies the boundary conditions $y_\delta^{\hat{z}}(\hat{z}-1) = 0$ and $y_\delta^{\hat{z}}(\hat{z}) = \delta$. To verify the inductive relation, we first evaluate the term $y_\delta^{\hat{z}}(z+1) - y_\delta^{\hat{z}}(z)$ for $z \geq \hat{z}-1$:

$$y_\delta^{\hat{z}}(z+1) - y_\delta^{\hat{z}}(z) = \begin{cases} \delta & \text{if } z = \hat{z}-1, \\ \delta \prod_{m=\hat{z}}^{z} \frac{\bar{p}(m)}{p(m)} + \sum_{k=\hat{z}}^{z}\left(\prod_{m=k}^{z}\frac{\bar{p}(m)}{p(m)}\right)\frac{x(k)}{\bar{p}(k)} & \text{if } z \geq \hat{z}. \end{cases}$$

Next, we evaluate the term $p(z)y_\delta^{\hat{z}}(z+1)+\bar{p}(z)y_\delta^{\hat{z}}(z-1)-y_\delta^{\hat{z}}(z)$: for all $z\geq\hat{z}$,

$$
\begin{aligned}
&p(z)y_\delta^{\hat{z}}(z+1)+\bar{p}(z)y_\delta^{\hat{z}}(z-1)-y_\delta^{\hat{z}}(z)\\
&=p(z)[y_\delta^{\hat{z}}(z+1)-y_\delta^{\hat{z}}(z)]-\bar{p}(z)[y_\delta^{\hat{z}}(z)-y_\delta^{\hat{z}}(z-1)]\\
&=\delta\left(\prod_{m=\hat{z}}^{z-1}\frac{\bar{p}(m)}{p(m)}\right)\left(p(z)\frac{\bar{p}(z)}{p(z)}-\bar{p}(z)\right)+p(m)\frac{\bar{p}(m)}{p(m)}\frac{x(z)}{\bar{p}(z)}+\left[\sum_{k=\hat{z}}^{z-1}\left(\prod_{m=k}^{z-1}\frac{\bar{p}(m)}{p(m)}\right)\frac{x(k)}{\bar{p}(k)}\right]\left(p(z)\frac{\bar{p}(z)}{p(z)}-\bar{p}(z)\right)\\
&=0+x(z)+0=x(z).\ \blacksquare
\end{aligned}
$$

Appendix G: The Informed Bettor's Best Response to IP (Theorem 3)

In this section, we provide the details for the proof of Theorem 3.

G.1. Summary of Intuition

Lemmas 1 and 3 correspond to two separate mechanisms through which the informed bettor's profit may be unbounded. The first mechanism is that the threshold strategy ξ_i^* itself generates an infinite amount of profit for the informed bettor. In the context of the threshold strategy ξ_i^* defined in (4.4), this happens when the market state Z_t behaves so noisily that the event of severe mispricing (i.e., the spread line being sufficiently far away from the true median) occurs infinitely often. In such cases, the market maker is effectively not collecting information from the market. Our inertial policy guards against this mechanism by preventing the learning rate (formally defined as the drift of $\{Z_t\}$) from vanishing too fast, in which case $\{Z_t\}$ diverges in the right direction and the spread line s_t converges to the correct median almost surely.

The second mechanism is that the informed bettor may have an incentive to deviate from ξ_i^*. To mathematically verify that IP guards against this mechanism, it suffices to only algebraically verify the Bellman equation (4.11). But to see it intuitively, let us separately discuss the following two cases, each corresponding to a different type of deviation:

- (*case 1*) The informed bettor may deviate from ξ_i^* by bluffing. IP prevents this type of deviation because under IP, a pair of positive-negative bets have no net impact on the market state Z_t. Thus, the informed bettor can only push the market state Z_t to the wrong direction by bluffing more often than honest betting. But as discussed in Section 3.2, he also needs to bet honestly more often than bluffing to gain a positive net profit. Based on these two contradicting facts, we reach the conclusion that the informed bettor does not have an incentive to bluff.
- (*case 2*) The informed bettor may also deviate from ξ_i^* by not following the threshold structure of betting honestly versus waiting. IP induces the informed bettor to bet according to a threshold strategy, because we can marginally change the action from betting ($a_t=+1$) to waiting ($a_t=0$) at every state $z\in\mathbb{Z}$, and evaluate the difference in his continuation profits. It turns out this difference function only crosses zero once, leading to a threshold structure.[20]

[20] Algebraically, this is ultimately reduced to verifying that the function $j_1^+(z-1)\left[(\frac{1}{2}-\rho(z))+(\frac{1}{2}-\rho(z))^2+\cdots\right]-j_1^+(z)=j_1^+(z-1)\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}-j_1^+(z)$ crosses zero only once.

G.2. Preliminaries

First, we provide the explicit expressions for $\bar{z}$ and $\bar{r}$. Let

$$\bar{z} := \inf\left\{z : \tfrac{\rho(z-1)}{2\rho(z)} - \rho(z) - \rho(z-1) > \tfrac{1}{2} - \tfrac{c}{2-c}\right\} \tag{G.1}$$

be the threshold used in the informed bettor's optimal policy ξ_i^* in (4.4). We follow the common convention that $\bar{z} = \infty$ if the set inside the infimum in (G.1) is empty, and $\bar{z} = -\infty$ if this set is unbounded from below. By Lemma G.2 below, $\bar{z}$ is finite if and only if $j_1^+(-\infty) > 0$ (i.e., the type-1 informed bettor finds it profitable to act at some point in time). Otherwise, $\bar{z} = -\infty$, which corresponds to the informed bettor's policy of never betting according to (4.4). Given strictly positive constants $\bar{r}_0, \bar{r}_1$ that depend only on $m_0, m_1, F_\epsilon(\cdot)$, and c, let

$$\bar{r} = \min\{2, \bar{r}_0, \bar{r}_1\} \tag{G.2}$$

be the upper bound of r for Theorem 3 to hold. The closed form expressions for $\bar{r}_0$ and $\bar{r}_1$ are as follows:

$$\bar{r}_0 := \tfrac{c r_0}{(2-c)\zeta_0}, \quad \text{where } \zeta_0 := \max\left\{1, \max_{s\in[m_0,m_1]} \tfrac{f_1(s)}{f_0(s)}\right\}, \tag{G.3}$$

and

$$\bar{r}_1 := \sup\{r : \zeta_1 r + 2r_0 > 0\} = \begin{cases} \frac{2r_0}{-\zeta_1} & \text{if } \zeta_1 < 0, \\ \infty & \text{if } \zeta_1 \geq 0, \end{cases} \quad \text{where } \zeta_1 := \min_{s\in[m_0,m_1]} \left\{\left(\tfrac{f_1'(s)}{f_1(s)} - \tfrac{f_0'(s)}{f_0(s)}\right)\tfrac{1}{f_0(s)}\right\}. \tag{G.4}$$

REMARK 1. Theorem 3 requires $\bar{r}$ to be strictly positive. It is straightforward to verify the strict positivity of $\bar{r}$. By (A1:1) and (A1:3), $f_i(\cdot)$ is continuous and strictly positive in the interval $[m_0, m_1]$. Hence, ζ_0 is strictly positive and finite (which implies that $\bar{r}_0 > 0$), and ζ_1 is finite (which implies that $\bar{r}_1 > 0$).

G.3. Auxiliary Lemmas

We employ the following auxiliary lemmas to prove Theorem 3, deferring their proofs to Appendix G.6. The first auxiliary lemma summarizes the properties of $\rho(\cdot)$ in (E.2).

LEMMA G.1. *The (extended) residual probability sequence $\{\rho(z), z\in\mathbb{Z}\}$ in (E.2) satisfies the following:*

(LG.1-1) The natural further extension of $\rho(z)$ from domain $\mathbb{Z}$ to domain $\mathbb{R}$, defined as

$$\rho(x) = \begin{cases} \frac{1}{r_0 + rx} & \text{if } x \geq 0, \\ \frac{1}{2} - F_1 \circ F_0^{-1}\left(\frac{1}{2} + \frac{1}{r_0 - rx}\right) & \text{if } x < 0, \end{cases} \tag{G.5}$$

is a continuous and strictly decreasing function. Moreover, $\rho(\cdot)$ is twice differentiable in $\mathbb{R}\setminus\{0\}$.

(LG.1-2) For all $z\in\mathbb{Z}$, $\frac{1}{2} - \alpha = \rho(-\infty) > \rho(z) > \rho(\infty) = 0$.

(LG.1-3) For all $r\in(0,4)$ and $z_0\in\mathbb{Z}$, $\sum_{n=0}^{\infty}\prod_{k=0}^{n} \frac{\frac{1}{2}-\rho(z_0+k)}{\frac{1}{2}+\rho(z_0+k)} < \infty$.

(LG.1-4) Suppose that $0 < r < \bar{r}_0$. Then, for every integer $z\in\mathbb{Z}$, $\frac{\rho(z+1)}{\rho(z)} > 1 - \frac{c}{2-c}$. Therefore,

(a) $\rho(z) - \rho(z+1) < \frac{c}{2-c}$,

(b) $\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)} < \frac{(2-c)\rho(z+1)+\frac{c}{2}}{(2-c)\rho(z)+\frac{c}{2}}$.

In Lemma G.1 above, (LG.1-1)-(LG.1-2) are (intuitive) regularity conditions for $\rho(\cdot)$. Property (LG.1-3) is closely related to the convergence of $\sum_{n=1}^{\infty}\Lambda_n$, which ensures that Z_t diverges to infinity with probability one. The last property, (LG.1-4), ensures that $\rho(z+1)$ is "close enough" to $\rho(z)$. This "inertia" property eliminates the informed bettor's incentive to bluff; see the proof of Lemma 3 for further details. The following lemma summarizes the properties of the threshold $\bar{z}$, which is defined in (G.1).

LEMMA G.2. *Suppose that $0<r<\min\{2,\bar{r}_1\}$. Then, $\bar{z}$ possesses the following properties:*

(LG.2-1) (finiteness) If $j_1^+(-\infty)>0$, then $\bar{z}$ is finite. Otherwise, $\bar{z}=-\infty$.

(LG.2-2) (single-crossing property) $j_1^+(z)\left[\frac{1}{2}+\rho(z)\right]<j_1^+(z-1)\left[\frac{1}{2}-\rho(z)\right]$ if and only if $z\geq\bar{z}$.

(LG.2-3) (profitable action) $j_1^+(z)>0$ for all $z<\bar{z}$.

The following lemma characterizes the summation of the probabilities that Z_t hits the region $(-\infty,M]$ up to period T under the probability measure $\mathbb{P}_1^z$.

LEMMA G.3. *For all $r\in(0,\bar{r}),\bar{z}\in\mathbb{Z}\cup\{-\infty\}$, and $M\in\mathbb{Z}$ satisfying $M>\bar{z}-2$, there exists an increasing function $\tilde{u}:\{z\in\mathbb{Z}:z>\bar{z}-2\}\to\mathbb{R}$ such that*

$$\sum_{t=1}^{T}\mathbb{E}_1^z\mathbb{I}\{Z_t\leq M\}=\mathbb{E}_1^z\tilde{u}(Z_{T+1})-\tilde{u}(z)\ \text{ for all } z\in\mathbb{Z} \text{ satisfying } z>\bar{z}-2 \text{ and } T\in\mathbb{Z}_+. \tag{G.6}$$

The closed-form expression for $\tilde{u}(\cdot)$ is as follows:

$$\tilde{u}(z)=\begin{cases}\left(1+\sum\limits_{n=z+1}^{M}\prod\limits_{m=n}^{M}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right)\tilde{\beta}+\sum\limits_{n=z+1}^{M}\sum\limits_{k=n}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod\limits_{m=n}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} & \text{if } \bar{z}-2<z\leq M,\\ 0 & \text{if } z=M+1,\\ \left(1+\sum\limits_{n=M+2}^{z-1}\prod\limits_{m=M+2}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\beta & \text{if } z\geq M+2,\end{cases} \tag{G.7}$$

where $\beta>0$ and $\tilde{\beta}<0$ are finite constants given by:

$$\begin{cases}\tilde{\beta}=-\prod_{m=\bar{z}}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}-\sum_{k=\bar{z}}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)},\\ \beta=-\frac{\frac{1}{2}-\rho(M+1)}{\frac{1}{2}+\rho(M+1)}\tilde{\beta}.\end{cases} \tag{G.8}$$

G.4. Main Body of the Proof of Theorem 3

For the market maker, fix an arbitrary inertial policy π_I with tuning parameter $r\in(0,\bar{r})$, where $\bar{r}$ is as in (G.2). We make two assumptions without loss of generality. First, we assume that $i=1$, because the analysis for the case where $i=0$ can be repeated verbatim. Second, we assume that $j_1^+(-\infty)>0$ (which implies that $\bar{z}>-\infty$ by (LG.2-1)) because otherwise the type-1 informed bettor never finds it profitable to act and the theorem's statement holds trivially.

Given the function $\bar{J}^1(\cdot)$ defined in (4.8), we have the following for all $T\in\mathbb{Z}_+$ and any admissible response policy ξ_1 of type-1 informed bettor:

$$0\leq\mathbb{E}_1^{\pi_I,\xi_1}[\bar{J}^1(Z_{T+1})] \qquad \text{[by Lemma 1]}$$

$$
\begin{aligned}
&= \mathbb{E}_1^{\pi_I,\xi_1}[\bar{J}^1(Z_1)] + \sum_{t=1}^{T}\Big[\mathbb{E}_1^{\pi_I,\xi_1}[\bar{J}^1(Z_{t+1})] - \mathbb{E}_1^{\pi_I,\xi_1}[\bar{J}^1(Z_t)]\Big] \\
&= \bar{J}^1(0) + \mathbb{E}_1^{\pi_I,\xi_1}\sum_{t=1}^{T}\Big[\mathbb{I}\{a_t=+1\}\big(\underbrace{\bar{J}^1(Z_t+1)-\bar{J}^1(Z_t)}_{\overset{\text{Lem. 3}}{\le} -j_1^+(Z_t)}\big) + \mathbb{I}\{a_t=-1\}\big(\underbrace{\bar{J}^1(Z_t-1)-\bar{J}^1(Z_t)}_{\overset{\text{Lem. 3}}{\le} -j_1^-(Z_t)}\big) \\
&\qquad + \mathbb{I}\{a_t=0\}\Big(\underbrace{\big[\tfrac{1}{2}+\rho(Z_t)\big]\bar{J}^1(Z_t+1)+\big[\tfrac{1}{2}-\rho(Z_t)\big]\bar{J}^1(Z_t-1)-\bar{J}^1(Z_t)}_{\overset{\text{Lem. 3}}{\le} 0}\Big)\Big] \\
&\le \bar{J}^1(0) + \mathbb{E}_1^{\pi_I,\xi_1}\sum_{t=1}^{T}\Big[\mathbb{I}\{a_t=+1\}[-j_1^+(Z_t)] + \mathbb{I}\{a_t=-1\}[-j_1^-(Z_t)]\Big] \\
&= \bar{J}^1(0) - V_1^{\pi_I,\xi_1}(T). \qquad\qquad [\text{by } (2.2)]
\end{aligned}
$$

As a result, under the strategy profile (π_I, ξ_1), the informed bettor's continuation profit function is bounded above by the constant $\bar{J}^1(0)$, which is independent of T. That is, $V_1^{\pi_I,\xi_1}(T) \le \bar{J}^1(0)$. Taking the limit infimum over T on the term on both sides and invoking Lemma 2, we reach the following chain of relations:

$$
\liminf_{T\to\infty} V_1^{\pi_I,\xi_1}(T) \le \bar{J}^1(0) \overset{\text{Lem. 2}}{=} \lim_{T\to\infty}\mathbb{E}_1^0\left[\sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar{z}-1\}\right] \overset{(2.2)}{=} \lim_{T\to\infty} V_1^{\pi_I,\xi_1^*}(T).
$$

Thus, $\xi_1^* \in \arg\max_\xi \liminf_{T\to\infty} V_i^{\pi_I,\xi}(T)$. By Lemma 1, we deduce that the informed bettor's total profit is finite: $\sup_\xi \liminf_{T\to\infty} V_i^{\pi_I,\xi}(T) = \lim_{T\to\infty} V_1^{\pi_I,\xi_1^*}(T) = \bar{J}^1(0) \overset{\text{Lem. 1}}{<} \infty$. Consequently, ξ_1^* is the type-1 informed bettor's best response strategy in the sense of (2.3). Moreover, $\bar{J}^1(\cdot) = J^1(\cdot)$ is the optimal value function. ■

G.5. Proofs of Lemma 1-3

Proof of Lemma 1. Recall from (G.2) that $\bar{r} \le 2 < 4$. Hence, when $r \in (0, \bar{r})$,

$$
\sum_{n=0}^{\infty}\Lambda_n \overset{(4.9)}{=} \sum_{n=0}^{\infty}\prod_{k=0}^{n}\frac{\frac{1}{2}-\rho(\bar{z}+k)}{\frac{1}{2}+\rho(\bar{z}+k)} \overset{\text{(LG.1-3)}}{<} \infty.
$$

Recalling (4.8), we note that the finiteness and nonnegativity of $\bar{J}^i(\cdot)$ follows from the convergence of $\sum_{n=0}^{\infty}\Lambda_n$ as well as the nonnegativity of $j_1^+(z)$ for all $z < \bar{z}$ because of (LG.2-3). ■

Proof of Lemma 2. To verify that (4.10) holds, let us assume that $\bar{z} > -\infty$ without loss of generality. Otherwise, both sides of the equation are zero trivially. We complete the proof in three steps.

Step 1. We claim that if $\bar{J}^1(\cdot)$ satisfies (4.10), where the limit term on the right-hand side exists, then $\bar{J}^0(\cdot)$ also satisfies (4.10), and the limit term on the right-hand side exists. In other words, we may assume that $i=1$ without loss of generality. To prove this claim, note that for all $z, \breve{z} \in \mathbb{Z}$, $\mathcal{P}^1_{z,\breve{z}} = \mathcal{P}^0_{-z,-\breve{z}}$. Therefore, $\{Z_t\}$ under $\mathbb{P}_1^z$ has the same law as $\{-Z_t\}$ under $\mathbb{P}_0^{-z}$. Formally speaking, given any $T \in \mathbb{Z}_+$, measurable function $f : \mathbb{R}^T \to \mathbb{R}$, and $z \in \mathbb{Z}$, we have

$$
\mathbb{E}_1^z f(Z_1, Z_2, \ldots, Z_T) = \mathbb{E}_0^{-z} f(-Z_1, -Z_2, \ldots, -Z_T). \tag{G.9}
$$

As a result, if $\bar{J}^1(\cdot)$ satisfies (4.10), we have

$$\begin{aligned}
&\bar{J}^0(z) = \bar{J}^1(-z) && \text{[by (4.8)]}\\
&= \lim_{T\to\infty} \mathbb{E}_1^{-z}\left[\sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar{z}-1\}\right] && [\bar{J}^1(\cdot) \text{ satisfies (4.10) where the limit exists by assumption}]\\
&= \lim_{T\to\infty} \mathbb{E}_0^{z}\left[\sum_{t=1}^{T} j_1^+(-Z_t)\,\mathbb{I}\{-Z_t \le \bar{z}-1\}\right] && \text{[by (G.9)]}\\
&= \lim_{T\to\infty} \mathbb{E}_0^{z}\left[\sum_{t=1}^{T} j_0^-(Z_t)\,\mathbb{I}\{Z_t \ge 1-\bar{z}\}\right] && \text{[by (4.7)]}
\end{aligned}$$

<u>Step 2.</u> We claim that for all $z \in \mathbb{Z}$, there exists a bounded function $u_0(\cdot)$ such that for all $T \ge \bar{z} - z$,

$$\mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar{z}-1\} = \begin{cases} j_1^+(\bar{z}-1)\,[\mathbb{E}_1^z u_0(Z_{T+1}) - u_0(z)] & \text{if } z \ge \bar{z},\\ \sum_{i=1}^{\bar{z}-z} j_1^+(\bar{z}-i) + j_1^+(\bar{z}-1)\mathbb{E}_1^{\bar{z}} u_0(Z_{T+1-\bar{z}+z}) & \text{if } z \le \bar{z}-1. \end{cases} \tag{G.10}$$

Here, $u_0(\cdot)$ is as follows:

$$u_0(z) = \begin{cases} -1 & \text{if } z = \bar{z}-1,\\ 0 & \text{if } z = \bar{z},\\ \sum_{n=0}^{z-\bar{z}-1} \Lambda_n & \text{if } z \ge \bar{z}+1. \end{cases} \tag{G.11}$$

Let us first consider the case where $z \ge \bar{z}$. Starting with initial value z, the Markov chain Z_t is restrained to the region $[\bar{z}-1, \infty)$. Consider a special case of Lemma G.3 where $\bar{z} > -\infty$ and $M = \bar{z}-1$. In that special case, the corresponding $\tilde{u}(\cdot)$ function simplifies to:

$$\begin{aligned}
\tilde{u}(z) &= \begin{cases} -1 & \text{if } z = \bar{z}-1,\\ 0 & \text{if } z = \bar{z},\\ \left(1 + \sum_{n=\bar{z}+1}^{z-1} \prod_{m=\bar{z}+1}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right) \frac{\frac{1}{2}-\rho(\bar{z})}{\frac{1}{2}+\rho(\bar{z})} & \text{if } z \ge \bar{z}+1, \end{cases} && \text{[by (G.7)]}\\
&= \begin{cases} -1 & \text{if } z = \bar{z}-1,\\ 0 & \text{if } z = \bar{z},\\ \sum_{n=\bar{z}}^{z-1} \prod_{m=\bar{z}}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} & \text{if } z \ge \bar{z}+1, \end{cases}\\
&= \begin{cases} -1 & \text{if } z = \bar{z}-1,\\ 0 & \text{if } z = \bar{z},\\ \sum_{n=0}^{z-\bar{z}-1} \Lambda_n & \text{if } z \ge \bar{z}+1, \end{cases} && \text{[by (4.9)]}\\
&= u_0(z). && \text{[by (G.11)]}
\end{aligned}$$

Thus, $u_0(\cdot)$ is a special case of the function $\tilde{u}(\cdot)$ in Lemma G.3. By the conclusion of Lemma G.3, $u_0(\cdot)$ is an increasing function such that $\sum_{t=1}^{T} \mathbb{E}_1^z \mathbb{I}\{Z_t \le \bar{z}-1\} = \mathbb{E}_1^z \tilde{u}(Z_{T+1}) - \tilde{u}(z) = \mathbb{E}_1^z u_0(Z_{T+1}) - u_0(z)$ for all $z \ge \bar{z}-1$. Moreover, since $\sum_n \Lambda_n$ is a convergent series (see the derivations in the proof of Lemma 1), the function $u_0(\cdot)$ is bounded. Finally, $\mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar{z}-1\} = \mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(\bar{z}-1)\,\mathbb{I}\{Z_t \le \bar{z}-1\} = j_1^+(\bar{z}-1) \sum_{t=1}^{T} \mathbb{E}_1^z \mathbb{I}\{Z_t \le \bar{z}-1\} = j_1^+(\bar{z}-1)[\mathbb{E}_1^z u_0(Z_{T+1}) - u_0(z)]$.

Now, let us consider the case where $z < \bar{z}$. Note that Z_t increases with certainty until it hits $[\bar{z}, \infty)$. Thus,

$$\mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar{z}-1\}$$

$$
\begin{aligned}
&= \mathbb{E}_1^z \sum_{t=1}^{\bar z - z} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar z - 1\} + \mathbb{E}_1^z \sum_{t=\bar z - z+1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar z - 1\} \\
&= \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + \mathbb{E}_1^z \sum_{t=\bar z - z+1}^{T} j_1^+(\bar z - 1)\,\mathbb{I}\{Z_t \le \bar z - 1\} \\
&= \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + \mathbb{E}_1^{\bar z} \sum_{t=1}^{T-\bar z+z} j_1^+(\bar z - 1)\,\mathbb{I}\{Z_t \le \bar z - 1\} && \text{[by the Markov property of } Z_t] \\
&\overset{(a)}{=} \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + \mathbb{E}_1^{\bar z} \sum_{t=1}^{T-\bar z+z} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar z - 1\} \\
&= \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + j_1^+(\bar z - 1)\mathbb{E}_1^{\bar z} u_0(Z_{T-\bar z+z+1}). && \text{[by the analysis for } z \ge \bar z;\ u_0(\bar z)=0]
\end{aligned}
$$

In the derivations above, (a) follows since Z_t increases by one with certainty as soon as it hits $(-\infty, \bar z - 1]$. Combining our results for the cases where $z \ge \bar z$ and $z < \bar z$, we finish this step.

Step 3. We evaluate the term $\bar J^1(z)$ for all $z \in \mathbb{Z}$:

$$
\begin{aligned}
&\bar J^1(z) \\
&= \begin{cases} j_1^+(\bar z - 1)\sum_{n=z-\bar z}^{\infty} \Lambda_n & \text{if } \bar z \le z, \\ \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + j_1^+(\bar z - 1)\sum_{n=0}^{\infty}\Lambda_n & \text{if } z \le \bar z - 1, \end{cases} && \text{[by (4.8)]} \\
&= \begin{cases} j_1^+(\bar z - 1)\left[u_0(\infty) - u_0(z)\right] & \text{if } \bar z \le z, \\ \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + j_1^+(\bar z - 1)u_0(\infty) & \text{if } z \le \bar z - 1, \end{cases} && \text{[by (G.11)]} \\
&\overset{(b)}{=} \lim_{T\to\infty} \begin{cases} j_1^+(\bar z - 1)\left[\mathbb{E}_1^z u_0(Z_{T+1}) - u_0(z)\right] & \text{if } \bar z \le z, \\ \sum_{i=1}^{\bar z - z} j_1^+(\bar z - i) + j_1^+(\bar z - 1)\mathbb{E}_1^{\bar z} u_0(Z_{T+1-\bar z+z}) & \text{if } z \le \bar z - 1, \end{cases} \\
&= \lim_{T\to\infty} \mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar z - 1\}, && \text{[by (G.10)]}
\end{aligned}
$$

where (b) follows from the bounded convergence theorem (Rudin 1976, Theorem 11.32) and the following two facts: (i) $Z_t \uparrow \infty$ almost surely (by Statement (T4:1) in Theorem 4, which is proven independently), and (ii) $u_0(\cdot)$ in (G.11) is a bounded function (see Step 2). As a result, $\lim_{T\to\infty} \mathbb{E}_1^z \sum_{t=1}^{T} j_1^+(Z_t)\,\mathbb{I}\{Z_t \le \bar z - 1\}$ exists and equals $\bar J^1(z)$. Combining Steps 1-3, we complete the proof. ■

Proof of Lemma 3. We first assume that $\bar z > -\infty$ without loss of generality. To see this, suppose that $\bar z = -\infty$. By (LG.2-1), $j^+(-\infty) \le 0$. As a result, we have $j_1^+(z) \le 0$ and $j_1^-(z) \le 0$ for all $z \in \mathbb{Z}$, and the Bellman equation (4.11) holds trivially. We complete the rest of the proof in four steps.

Step 1. We claim that we can assume that $i = 1$ without loss of generality. To be more accurate, $\bar J^0(\cdot)$ satisfies the Bellman equation (4.11) if $\bar J^1(\cdot)$ does. Note that if $\bar J^1(\cdot)$ satisfies (4.11), then we have the following for all $z \in \mathbb{Z}$:

$$
\bar J^0(z) = \bar J^1(-z) \qquad \text{[by (4.8)]}
$$

$$
\begin{aligned}
&= \max\Big\{ \underbrace{j_1^+(-z)}_{\overset{(a)}{=} j_0^-(z)} + \underbrace{\bar{J}^1(-z+1)}_{\overset{(b)}{=} \bar{J}^0(z-1)}, \quad \underbrace{j_1^-(-z)}_{\overset{(a)}{=} j_0^+(z)} + \underbrace{\bar{J}^1(-z-1)}_{\overset{(b)}{=} \bar{J}^0(z+1)}, \\
&\qquad \underbrace{\bar{F}_1(\tilde{s}(-z))}_{\overset{(c)}{=} F_0(\tilde{s}(z))} \underbrace{\bar{J}^1(-z+1)}_{\overset{(b)}{=} \bar{J}^0(z-1)} + \underbrace{F_1(\tilde{s}(-z))}_{\overset{(c)}{=} \bar{F}_0(\tilde{s}(z))} \underbrace{\bar{J}^1(-z-1)}_{\overset{(b)}{=} \bar{J}^0(z+1)} \Big\} && \text{[by (4.11)]} \\
&= \max\left\{ j_0^-(z) + \bar{J}^0(z-1),\, j_0^+(z) + \bar{J}^0(z+1),\, F_0(\tilde{s}(z))\, \bar{J}^0(z-1) + \bar{F}_0(\tilde{s}(z))\, \bar{J}^0(z+1) \right\}
\end{aligned}
$$

where (a), (b), and (c) follow from (4.7), (4.8), and (4.2) respectively. Therefore, $\bar{J}^0(\cdot)$ satisfies the Bellman equation (4.11) as well.

Step 2. We claim that the following equation holds:

$$
\bar{J}^1(z) = \begin{cases} j_1^+(z) + \bar{J}^1(z+1) & \text{if } z \le \bar{z}-1, \\ \bar{F}_1(\tilde{s}(z))\, \bar{J}^1(z+1) + F_1(\tilde{s}(z))\, \bar{J}^1(z-1) & \text{if } z \ge \bar{z}. \end{cases} \tag{G.12}
$$

The intuition for (G.12) is that the function $\bar{J}^1(\cdot)$ satisfies the above recursive relation if the type-1 bettor follows the threshold strategy ξ_1^*. This step is a direct consequence of Lemma 2, the transition rule of Z_t, and the Markov property of Z_t.

Step 3. We build on Step 2, and claim that for all $z \in \mathbb{Z}$, the following holds:

$$
\bar{J}^1(z) = \max\left\{ j_1^+(z) + \bar{J}^1(z+1),\, \bar{F}_1(\tilde{s}(z))\, \bar{J}^1(z+1) + F_1(\tilde{s}(z))\, \bar{J}^1(z-1) \right\}. \tag{G.13}
$$

Equation (G.13) can be interpreted as a "weakened" version of the Bellman equation for the decision problem where bluffing (or $a_t = -1$) is not a feasible action for the type-1 bettor. Invoking (G.12), we deduce that it suffices to verify that

$$
\underbrace{j_1^+(z) + \bar{J}^1(z+1) - \left[\bar{F}_1(\tilde{s}(z))\, \bar{J}^1(z+1) + F_1(\tilde{s}(z))\, \bar{J}^1(z-1)\right]}_{(*)} \ge 0 \iff z \le \bar{z}-1.
$$

Let us evaluate the term $(*)$ above. For all $z \in \mathbb{Z}$,

$$
\begin{aligned}
(*) &= j_1^+(z) + \bar{J}^1(z+1) - \left[\bar{F}_1(\tilde{s}(z))\, \bar{J}^1(z+1) + F_1(\tilde{s}(z))\, \bar{J}^1(z-1)\right] \\
&= j_1^+(z) + \bar{J}^1(z+1) - \left[\tfrac{1}{2} + \rho(z)\right] \bar{J}^1(z+1) - \left[\tfrac{1}{2} - \rho(z)\right] \bar{J}^1(z-1) && \text{[by (4.2)]} \\
&= j_1^+(z) + \left[\tfrac{1}{2} - \rho(z)\right] \left[\bar{J}^1(z+1) - \bar{J}^1(z-1)\right] && \text{[rearranging terms]} \\
&= \begin{cases} j_1^+(z) + \left[\tfrac{1}{2} - \rho(z)\right] \left[j_1^+(\bar{z}-1) \sum_{n=z+1-\bar{z}}^{\infty} \Lambda_n - j_1^+(\bar{z}-1) \sum_{n=z-1-\bar{z}}^{\infty} \Lambda_n \right] & \text{if } z \ge \bar{z}+1, \\ j_1^+(z) + \left[\tfrac{1}{2} - \rho(z)\right] \left[j_1^+(\bar{z}-1) \sum_{n=1}^{\infty} \Lambda_n - j_1^+(\bar{z}-1) \sum_{n=0}^{\infty} \Lambda_n - j_1^+(\bar{z}-1) \right] & \text{if } z = \bar{z}, \\ j_1^+(z) + \left[\tfrac{1}{2} - \rho(z)\right] \left[\sum_{i=1}^{\bar{z}-z-1} j_1^+(z-i) - \sum_{i=1}^{\bar{z}-z+1} j_1^+(z-i) \right] & \text{if } z \le \bar{z}-1, \end{cases} && \text{[by (4.8)]} \\
&= \begin{cases} j_1^+(z) - j_1^+(\bar{z}-1) \Lambda_{z-\bar{z}-1} \frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)} & \text{if } z \ge \bar{z}+1, \\ j_1^+(z) - j_1^+(\bar{z}-1) \frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)} & \text{if } z = \bar{z}, \\ j_1^+(z) - \left[\tfrac{1}{2} - \rho(z)\right] \left[j_1^+(z-1) + j_1^+(z) \right] & \text{if } z \le \bar{z}-1, \end{cases} && \text{[by (4.9)]} \\
&= \begin{cases} j_1^+(z) - j_1^+(\bar{z}-1) \left(\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}\right) \left(\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)}\right) \cdots \left(\frac{\frac{1}{2}-\rho(\bar{z})}{\frac{1}{2}+\rho(\bar{z})}\right) & \text{if } z \ge \bar{z}, \\ \left[\tfrac{1}{2} + \rho(z)\right] j_1^+(z) - \left[\tfrac{1}{2} - \rho(z)\right] j_1^+(z-1) & \text{if } z \le \bar{z}-1. \end{cases} && \text{[by (4.9)]}
\end{aligned}
$$

For all $z \geq \bar{z}$, we have the following due to (LG.2-2):

$$\begin{aligned} j_1^+(z) &< j_1^+(z-1)\left(\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}\right) \\ &< j_1^+(z-2)\left(\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}\right)\left(\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)}\right) \\ &\vdots \\ &< j_1^+(\bar{z}-1)\left(\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}\right)\left(\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)}\right)\cdots\left(\frac{\frac{1}{2}-\rho(\bar{z})}{\frac{1}{2}+\rho(\bar{z})}\right). \end{aligned}$$

Thus, for all $z \geq \bar{z}$, the term $(*)$ is less than 0. On the other hand, for all $z \leq \bar{z}-1$, $j_1^+(z) \geq j_1^+(z-1)\,\frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}$ by (LG.2-2); thus the term $(*)$ is greater than or equal to 0. In either case, (G.13) holds.

Step 4. We claim that the negative bet (i.e., bluffing) is a dominated action for type-1 bettor. Thus, the type-1 bettor never needs to bluff. By the Bellman equation (4.11), it suffices to verify the following:

$$\underbrace{\bar{J}^1(z) - \left[j_1^-(z) + \bar{J}^1(z-1)\right]}_{(**)} > 0. \tag{G.14}$$

Let us first evaluate the term $(**)$ when $z \leq \bar{z}$:

$$\begin{aligned} (**) &= \left[\bar{J}^1(z) - \bar{J}^1(z-1)\right] - j_1^-(z) \\ &= -j_1^+(z-1) - j_1^-(z) && [\text{by } (4.8);\ z \leq \bar{z}\ ] \\ &= -(2-c)\rho(z-1) + \tfrac{c}{2} - (c-2)\rho(z) + \tfrac{c}{2} && [\text{by } (4.7)] \\ &= c - (2-c)[\rho(z-1) - \rho(z)] > 0, \end{aligned}$$

where the inequality follows because by the lemma's hypothesis, $r < \bar{r}$, and hence by (LG.1-4), $\rho(z-1) - \rho(z) < \frac{c}{2-c}$. We now evaluate the term $(**)$ when $z \geq \bar{z}+1$:

$$\begin{aligned} (**) &= \left[\bar{J}^1(z) - \bar{J}^1(z-1)\right] - j_1^-(z) \\ &= -j_1^-(z) - j_1^+(\bar{z}-1)\Lambda_{z-1-\bar{z}} && [\text{by } (4.8);\ z \geq \bar{z}+1] \\ &= (2-c)\rho(z) + \tfrac{c}{2} - j_1^+(\bar{z}-1)\Lambda_{z-1-\bar{z}} && [\text{by } (4.7)] \\ &\overset{(d)}{>} \left[(2-c)\rho(z-1) + \tfrac{c}{2}\right]\left(\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)}\right) - j_1^+(\bar{z}-1)\Lambda_{z-1-\bar{z}} \\ &\vdots \\ &\overset{(d)}{>} \left[(2-c)\rho(\bar{z}) + \tfrac{c}{2}\right]\underbrace{\left(\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)}\right)\cdots\left(\frac{\frac{1}{2}-\rho(\bar{z})}{\frac{1}{2}+\rho(\bar{z})}\right)}_{\overset{(4.9)}{=}\Lambda_{z-1-\bar{z}}} - j_1^+(\bar{z}-1)\Lambda_{z-1-\bar{z}} \\ &= \left[(2-c)\rho(\bar{z}) + \tfrac{c}{2}\right]\Lambda_{z-1-\bar{z}} - j_1^+(\bar{z}-1)\Lambda_{z-1-\bar{z}} \\ &= \left[(2-c)\rho(\bar{z}) + \tfrac{c}{2} - (2-c)\rho(\bar{z}-1) + \tfrac{c}{2}\right]\Lambda_{z-1-\bar{z}} && [\text{by } (4.7)] \\ &= \left[c - (2-c)\big(\rho(\bar{z}-1) - \rho(\bar{z})\big)\right]\Lambda_{z-1-\bar{z}} \overset{(e)}{>} 0, \end{aligned}$$

where (d) follows because by the lemma's hypothesis, $r < \bar{r}$, and thus $\frac{\frac{1}{2}-\rho(z-1)}{\frac{1}{2}+\rho(z-1)} < \frac{(2-c)\rho(z)+\frac{c}{2}}{(2-c)\rho(z-1)+\frac{c}{2}}$ (see (LG.1-4)); and (e) follows because $\rho(\bar{z}-1) - \rho(\bar{z}) < \frac{c}{2-c}$ (see (LG.1-4)). Combining our findings in the cases where $z \le \bar{z}$ and $z \ge \bar{z}+1$, we conclude that the term $(**)$ is greater than 0. Therefore, based on all of the conclusions from Steps 1-4, we have the desired result. ■

G.6. Proofs of Auxiliary Lemmas

Proof of Lemma G.1. We prove each part separately. To prove (LG.1-1), note that (G.5) can be viewed as an extension of the function $\rho(\cdot)$ from the domain $\mathbb{Z}$ to domain $\mathbb{R}$, because it is consistent with (E.2) on the integer domain $\mathbb{Z}$. By construction, $\rho(\cdot)$ is piecewise continuous in the region $(-\infty,0)$ and in the region $(0,\infty)$. To see that $\rho(\cdot)$ is twice differentiable in $\mathbb{R}\setminus\{0\}$, note that $\rho(\cdot)$ is twice differentiable in $(0,\infty)$. By (A1:1) and (A1:3), $F_1\circ F_0^{-1}(\cdot)$ is a twice differentiable function in $(-\infty,0)$. This implies that $\rho(x)$ is twice differentiable in $(0,\infty)$. To verify that $\rho(\cdot)$ is continuous at 0 and hence in $\mathbb{R}$, observe that $\rho(0+)=\rho(0)$ and that $\rho(0)=\frac{1}{r_0}=\frac{1}{2}-F_1\left(\frac{m_0+m_1}{2}\right)\overset{\text{Lem. A.1}}{=}F_0\left(\frac{m_0+m_1}{2}\right)-\frac{1}{2}$. As a result,

$$\rho(0-)=\tfrac{1}{2}-F_1\circ F_0^{-1}\left(\tfrac{1}{2}+\rho(0+)\right)=\tfrac{1}{2}-F_1\circ F_0^{-1}\left(\tfrac{1}{2}+\rho(0)\right)=\tfrac{1}{2}-F_1\left(\tfrac{m_0+m_1}{2}\right)=\rho(0).$$

Lastly, to show that $\rho(x)$ is strictly decreasing in x, note that for all $x\in\mathbb{R}\setminus\{0\}$,

$$\rho'(x)=\begin{cases}-\frac{r}{(r_0+rx)^2} & \text{if } x>0,\\ \frac{f_1\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}{f_0\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}\left(\frac{-r}{(r_0-rx)^2}\right) & \text{if } x<0,\end{cases}$$

is strictly negative. The evaluation of $\rho'(x)$ when $x<0$ is based on the inverse function theorem (IFT) (Rudin 1976, Theorem 9.24). Invoking the mean value theorem (MVT) (Rudin 1976, Theorem 5.9), we conclude that $\rho(x)$ strictly decreases in x.

To prove (LG.1-2), observe that

$$\rho(-\infty)=\tfrac{1}{2}-F_1\circ F_0^{-1}\left(\tfrac{1}{2}+\rho(\infty)\right)=\tfrac{1}{2}-F_1\circ F_0^{-1}\left(\tfrac{1}{2}\right)=\tfrac{1}{2}-F_1(m_0)=\tfrac{1}{2}-\alpha,$$

and $\rho(\infty)=\lim_{x\uparrow\infty}\frac{1}{r_0+rx}=0$. Thus, (LG.1-2) holds by the (strict) monotonicity of $\rho(\cdot)$.

To prove (LG.1-3), note that since $\rho(z)=\frac{1}{r_0+rz}$ for $z\in\mathbb{Z}_+$ and $r\in(0,4)$, $\liminf_{n\to\infty}\{n\rho(n)\}>\frac{1}{4}$. The rest follows from Lemma L.1.

To prove (LG.1-4), observe that for all $z\in\mathbb{Z}$,

$$\begin{aligned}\log\left(\frac{\rho(z)}{\rho(z+1)}\right)&=-\left[\log\rho(x)\right]'=\frac{-\rho'(x)}{\rho(x)} && [\text{for some } x\in(z,z+1)\subset\mathbb{R}\setminus\{0\}, \text{ by MVT}]\\ &=\begin{cases}\frac{r}{r_0+rx} & \text{if } x>0,\\ \frac{f_1\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}{f_0\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}\left(\frac{r}{(r_0-rx)^2}\right)\frac{1}{\rho(x)} & \text{if } x<0,\end{cases} && [\text{by IFT}]\\ &<\begin{cases}\frac{r}{r_0} & \text{if } x>0,\\ \frac{f_1\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}{f_0\left[F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)\right]}\left(\frac{r}{r_0^2}\right)r_0 & \text{if } x<0,\end{cases} && \left[x<0\Longrightarrow\rho(x)>\rho(0)=\tfrac{1}{r_0}\right]\\ &\le\max\left\{1,\max_{s\in[m_0,m_1]}\frac{f_1(s)}{f_0(s)}\right\}\frac{r}{r_0}=\frac{\zeta_0 r}{r_0}. && \left[\text{by (G.3); } F_0^{-1}\left(\tfrac{1}{2}+\rho(-x)\right)\in[m_0,m_1]\right]\end{aligned}$$

As a result, for all r such that $0<r<\bar{r}_0$,

$$\begin{aligned}\frac{\rho(z+1)}{\rho(z)}=\exp\left(\log\frac{\rho(z+1)}{\rho(z)}\right)\ge\exp\left(-\frac{\zeta_0 r}{r_0}\right)&\ge1-\frac{\zeta_0 r}{r_0} && [e^x\ge1+x\ \forall x\in\mathbb{R}]\\ &>1-\frac{\zeta_0}{r_0}\frac{cr_0}{(2-c)\zeta_0}=1-\frac{c}{2-c}. && \left[r<\bar{r}_0=\frac{cr_0}{(2-c)\zeta_0}\right]\end{aligned}$$

To see (LG.1-4)(a), note that for all $z \in \mathbb{Z}$,

$$\rho(z) - \rho(z+1) = \rho(z)\left(1 - \frac{\rho(z+1)}{\rho(z)}\right) < \rho(z)\frac{c}{2-c} \overset{(a)}{<} \frac{c}{2-c},$$

where (a) holds because $\rho(z) < 1$. To see (LG.1-4)(b), observe that that since $\frac{\rho(z+1)}{\rho(z)} > 1 - \frac{c}{2-c} = \frac{2-2c}{2-c}$,

$$\frac{(2-c)\rho(z+1)+\frac{c}{2}}{(2-c)\rho(z)+\frac{c}{2}} > \frac{(2-2c)\rho(z)+\frac{c}{2}}{(2-c)\rho(z)+\frac{c}{2}} = \frac{c\left[\frac{1}{2}-\rho(z)\right]+(2-c)\rho(z)}{c\left[\frac{1}{2}+\rho(z)\right]+(2-2c)\rho(z)} \overset{(b)}{>} \frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)},$$

where (b) holds because $\frac{2-c}{2-2c} > 1 > \frac{\frac{1}{2}-\rho(z)}{\frac{1}{2}+\rho(z)}$. ■

Proof of Lemma G.2. In light of (G.1), let

$$\phi(z) := \frac{\rho(z-1)}{2\rho(z)} - \rho(z) - \rho(z-1), \tag{G.15}$$

so that $\bar{z} = \inf\{z : \phi(z) > \frac{1}{2} - \frac{c}{2-c}\}$. We divide the proof into four steps.

Step 1. We claim that $\phi(z)$ strictly increases in z. That is to say, for all $z \in \mathbb{Z}$, $\phi(z) > \phi(z-1)$. Let us first consider the case where $z \in \mathbb{Z}_+$:

$$\begin{aligned}
\phi(z) &= \frac{\rho(z-1)}{2\rho(z)} - \rho(z) - \rho(z-1) && [\text{by (G.15)}]\\
&= \frac{r_0+rz}{2(r_0+rz-r)} - \frac{1}{r_0+rz} - \frac{1}{r_0+rz-r} && [\rho(z) = \tfrac{1}{r_0+rz}\ \forall z \in \mathbb{Z}_+]\\
&= \frac{x}{2(x-r)} - \frac{1}{x} - \frac{1}{x-r} && [x := r_0 + rz \geq r_0 + r]\\
&= \frac{1}{2} + \left(\frac{r}{2} - 1\right)\frac{1}{x-r} - \frac{1}{x}, && [r < 2]
\end{aligned}$$

which strictly increases in x. Next, let us consider the case where $z \in \mathbb{N}_-$; i.e., z is a negative natural number (including zero). In light of (G.5), let us consider the continuous extension of $\phi(\cdot)$ to $(-\infty, 0]$, which is continuous in $(-\infty, 0]$ and differentiable in $(-\infty, 0)$. For all $z \in \mathbb{N}_-$,

$$\begin{aligned}
\phi(z) - \phi(z-1) &= \phi'(\tilde{x}) && [\text{for some } z-1 < \tilde{x} < z, \text{ by MVT}]\\
&= \frac{\rho'(\tilde{x}-1)\rho(\tilde{x}) - \rho'(\tilde{x})\rho(\tilde{x}-1)}{2\rho^2(\tilde{x})} - \rho'(\tilde{x}) - \rho'(\tilde{x}-1) &&\\
&\geq \frac{\rho'(\tilde{x}-1)/\rho(\tilde{x}-1) - \rho'(\tilde{x})/\rho(\tilde{x})}{2\rho(\tilde{x})/\rho(\tilde{x}-1)} && [\rho'(x) < 0\ \forall x < 0]\\
&= \frac{\tilde{\rho}'(\tilde{x}) - \tilde{\rho}'(\tilde{x}-1)}{2\rho(\tilde{x})/\rho(\tilde{x}-1)} && [\tilde{\rho}(x) := \log(-\rho'(x)) \in \mathcal{C}^1 \text{ on } (-\infty, 0)\ ]\\
&= \frac{\tilde{\rho}'(\breve{x})}{2\rho(\tilde{x})/\rho(\tilde{x}-1)}, && [\text{for some } \tilde{x}-1 < \breve{x} < \tilde{x}, \text{ by MVT}]
\end{aligned}$$

where MVT stands for the mean value theorem (Rudin 1976, Theorem 5.9). Thus, to show that $\phi(z-1) - \phi(z) > 0$ for all $z \in \mathbb{N}_-$, it suffices to show that $\tilde{\rho}'(x) > 0$ for all $x < 0$. Let us compute $\tilde{\rho}(x) = \log(-\rho'(x))$:

$$\begin{aligned}
\tilde{\rho}(x) &= \log\left(-\frac{d}{dx}\left[\frac{1}{2} - F_1 \circ F_0^{-1}\left(\frac{1}{2} + \rho(-x)\right)\right]\right) && [\text{by (G.5)}]\\
&= \log\left(\frac{f_1 \circ F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)}{f_0 \circ F_0^{-1}\left(\frac{1}{2}+\rho(-x)\right)}\left(-\rho'(-x)\right)\right) && [\text{by IFT}]\\
&= \log\left(f_1 \circ g(x)\right) - \log\left(f_0 \circ g(x)\right) + \log\left(-\rho'(-x)\right), && [g(x) := F_0^{-1}(\tfrac{1}{2} + \rho(-x))]
\end{aligned}$$

where IFT stands for the inverse function theorem (Rudin 1976, Theorem 9.24). Let us now evaluate $\tilde{\rho}'(x)$:

$$\tilde{\rho}'(x) = \frac{(f_1' \circ g(x))g'(x)}{f_1 \circ g(x)} - \frac{(f_0' \circ g(x))g'(x)}{f_0 \circ g(x)} + \frac{\rho''(-x)}{-\rho'(-x)}$$

$$
\begin{aligned}
&= \left[\frac{f_1'\circ g(x)}{f_1\circ g(x)} - \frac{f_0'\circ g(x)}{f_0\circ g(x)}\right] g'(x) + \frac{\rho''(-x)}{-\rho'(-x)} \\
&= \underbrace{\left[\frac{f_1'\circ g(x)}{f_1\circ g(x)} - \frac{f_0'\circ g(x)}{f_0\circ g(x)}\right]\left(\frac{1}{f_0\circ g(x)}\right)}_{\geq \zeta_1\text{; see (G.4)}} \underbrace{\left[-\rho'(-x)\right]}_{>0} + \frac{\rho''(-x)}{-\rho'(-x)} \\
&\geq \zeta_1\left[-\rho'(-x)\right] + \frac{\rho''(-x)}{-\rho'(-x)} \\
&= \frac{\zeta_1 r}{(r_0-rx)^2} + \frac{2}{(r_0-rx)} && \left[\rho(-x) = \tfrac{1}{r_0-rx}\right] \\
&= \frac{2r_0+\zeta_1 r - 2rx}{(r_0-rx)^2} \overset{(a)}{>} 0,
\end{aligned}
$$

where (a) holds because (i) $x<0$ and (ii) $\zeta_1 r + 2r_0 > 0$, which is implied by $0<r<\bar r_1$ and (G.4). Therefore, $\phi(z)-\phi(z-1)>0$ for all $z\in\mathbb{Z}$.

Step 2. We claim that (LG.2-1) holds. By Step 1, $\phi(z)$ is strictly increasing in z. Thus, $\bar z<\infty$ if and only if $\phi(\infty) > \frac{1}{2}-\frac{c}{2-c}$. Similarly, $\bar z>-\infty$ if and only if $\phi(-\infty)<\frac{1}{2}-\frac{c}{2-c}$. Let us evaluate $\phi(\infty)$ and $\phi(-\infty)$:

$$
\begin{aligned}
\phi(\infty) &= \lim_{z\uparrow\infty}\left(\frac{\rho(z-1)}{2\rho(z)} - \rho(z) - \rho(z-1)\right) \overset{(b)}{=} \frac{1}{2} - 0 - 0 = \frac{1}{2}, \\
\phi(-\infty) &= \lim_{z\downarrow-\infty}\left(\frac{\rho(z-1)}{2\rho(z)} - \rho(z) - \rho(z-1)\right) \overset{\text{(LG.1-2)}}{=} \frac{1}{2} - 2\left(\frac{1}{2}-\alpha\right) = -\frac{1}{2}+2\alpha.
\end{aligned}
$$

In the derivations above, (b) holds because $\rho(z)=\frac{1}{r_0+rz}$ for every $z\in\mathbb{Z}_+$. Since $c\in(0,1)$, $\frac{1}{2}-\frac{c}{2-c}<\frac{1}{2}=\phi(\infty)$. This implies that $\bar z<\infty$. Moreover,

$$
\begin{aligned}
\phi(-\infty) - \left(\frac{1}{2}-\frac{c}{2-c}\right) &= -\frac{1}{2}+2\alpha-\frac{1}{2}+\frac{c}{2-c} && \left[\phi(-\infty) = -\tfrac{1}{2}+2\alpha\right] \\
&= -2\rho(-\infty) + \frac{c}{2-c} && \left[\rho(-\infty)=\tfrac{1}{2}-\alpha\right] \\
&= \frac{2}{c-2}\left((2-c)\rho(-\infty) - \frac{c}{2}\right) = \frac{2j^+(-\infty)}{c-2}
\end{aligned}
$$

Consequently, $\bar z>-\infty$ if and only if $j_1^+(-\infty)>0$.

Step 3. We claim that (LG.2-2) holds. Note that for all $z\in\mathbb{Z}$,

$$
\begin{aligned}
& z\geq \bar z \\
\iff & \phi(z) \geq \frac{1}{2}-\frac{c}{2-c} && \left[\bar z = \inf\{z:\phi(z)>\tfrac{1}{2}-\tfrac{c}{2-c}\} \text{ and } \phi(z) \text{ increases in } z\right] \\
\iff & \frac{\rho(z-1)}{\rho(z)} - \rho(z) - \rho(z-1) > \frac{1}{2}-\frac{c}{2-c} && [\text{by the definition of } \phi(z)] \\
\iff & \frac{2-c}{2}\rho(z-1) - (2-c)\rho^2(z) - (2-c)\rho(z)\rho(z-1) - \frac{c}{4} > \frac{2-c}{2}\rho(z) - c\rho(z) - \frac{c}{4} \\
\iff & \frac{2-c}{2}\rho(z-1) - (2-c)\rho(z-1)\rho(z) - \frac{c}{4} + \frac{c}{2}\rho(z) > \frac{2-c}{2}\rho(z) - \frac{c}{4} + (2-c)\rho^2(z) - \frac{c}{2}\rho(z) \\
\iff & \left[(2-c)\rho(z-1)-\frac{c}{2}\right]\left[\frac{1}{2}-\rho(z)\right] > \left[(2-c)\rho(z)-\frac{c}{2}\right]\left[\frac{1}{2}+\rho(z)\right] \\
\iff & j_1^+(z-1)\left[\frac{1}{2}-\rho(z)\right] > j_1^+(z)\left[\frac{1}{2}+\rho(z)\right].
\end{aligned}
$$

Step 4. We claim that (LG.2-3) holds. Without loss of generality, assume that $\bar z>-\infty$ (otherwise, the statement trivially holds). By (LG.2-1), we know that $j^+(-\infty)>0$. Let $\tilde z := \inf\{z: j_1^+(z)\leq 0\}\in\mathbb{Z}$. By definition, we have $j_1^+(\tilde z-1)>0$. Moreover, by (LG.1-2), $\rho(\tilde z)\in(0,\frac{1}{2})$. Thus, $j_1^+(\tilde z-1)\left[\frac{1}{2}-\rho(\tilde z)\right] > 0 \geq j_1^+(\tilde z)\left[\frac{1}{2}+\rho(\tilde z)\right]$. Furthermore, in light of (LG.2-2), $\bar z\leq\tilde z$. Therefore, $j_1^+(\bar z-1)\geq j_1^+(\tilde z-1)>0$. ■

Proof of Lemma G.3. By Lemma H.1, Lemma G.3 reduces to a special case of Lemma L.2 where $\bar{z} = \bar{Z}$ and the residual probability sequence is $\{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$. ■

Appendix H: Performance Analysis of IP in Theorem 4

In this section, we provide the proof details regarding the supporting results for Theorem 4.

Proof of Lemma 5. Recall that $\Delta^{\pi_I}(T) := \max\left\{\Delta_0^{\pi_I,\xi_0^*}(T), \Delta_1^{\pi_I,\xi_1^*}(T)\right\}$, where $\Delta_i^{\pi_I,\xi_i^*}(T)$ is as in (2.5). Thus, it suffices to evaluate $\Delta_1^{\pi_I,\xi_1^*}(T)$ and $\Delta_0^{\pi_I,\xi_0^*}(T)$. We start by evaluating $\Delta_1^{\pi_I,\xi_1^*}(T)$:

$$
\begin{aligned}
&\Delta_1^{\pi_I,\xi_1^*}(T)\\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_1^{\pi_I,\xi_1^*}\left[\mathbb{I}\{(X-s_t)d_t<0\} - (1-c)\mathbb{I}\{(X-s_t)d_t>0\}\right] && \text{[by (2.5)]}\\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_1^{0}\left[\mathbb{I}\{a_t=0\}r_1(s_t) + \mathbb{I}\{a_t=+1\}(-j_1^+(s_t)) + \mathbb{I}\{a_t=-1\}(-j_1^-(s_t))\right] && \text{[by (2.4) \& (4.5)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_1^{0}\Big[\mathbb{I}\{Z_t \ge \bar{z}\}(4-2c)\underbrace{\left(F_1(\tilde{s}(Z_t)) - \tfrac{1}{2}\right)^2}_{=\rho^2(Z_t)} + \mathbb{I}\{Z_t<\bar{z}\}\underbrace{\left(\tfrac{c}{2} + j_1^+(\tilde{s}(Z_t))\right)}_{=(2-c)\rho(Z_t)}\Big] && \text{[by Thm. 3 \& (2.4)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_1^{0}\left[\mathbb{I}\{Z_t \ge \bar{z}\}(4-2c)\rho^2(Z_t) + \mathbb{I}\{Z_t<\bar{z}\}(2-c)\rho(Z_t)\right] && \text{[by (4.2) \& (4.7)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_1^{0} l(Z_t) = \sum_{t=1}^{T} \mathbb{E}_1^{\pi_I,\xi_1^*}[l(Z_t)].
\end{aligned}
$$

To evaluate $\Delta_0^{\pi_I,\xi_0^*}(T)$, let us leverage the symmetry relation between the cases where $i=0$ and $i=1$:

$$
\begin{aligned}
&\Delta_0^{\pi_I,\xi_0^*}(T)\\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_0^{\pi_I,\xi_0^*}\left[\mathbb{I}\{(X-s_t)d_t<0\} - (1-c)\mathbb{I}\{(X-s_t)d_t>0\}\right] && \text{[by (2.5)]}\\
&= \tfrac{cT}{2} - \sum_{t=1}^{T} \mathbb{E}_0^{0}\left[\mathbb{I}\{a_t=0\}r_0(s_t) + \mathbb{I}\{a_t=+1\}(-j_0^+(s_t)) + \mathbb{I}\{a_t=-1\}(-j_0^-(s_t))\right] && \text{[by (2.4) \& (4.5)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_0^{0}\Big[\mathbb{I}\{Z_t \le -\bar{z}\}(4-2c)\underbrace{\left(F_0(\tilde{s}(Z_t)) - \tfrac{1}{2}\right)^2}_{=\rho^2(-Z_t)} + \mathbb{I}\{Z_t>-\bar{z}\}\underbrace{\left(\tfrac{c}{2} + j_0^-(\tilde{s}(Z_t))\right)}_{=(2-c)\rho(-Z_t)}\Big] && \text{[by Thm. 3 \& (2.4)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_0^{0}\left[\mathbb{I}\{Z_t \le -\bar{z}\}(4-2c)\rho^2(-Z_t) + \mathbb{I}\{Z_t>-\bar{z}\}(2-c)\rho(-Z_t)\right] && \text{[by (4.2) \& (4.7)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_1^{0}\left[\mathbb{I}\{Z_t \ge \bar{z}\}(4-2c)\rho^2(Z_t) + \mathbb{I}\{Z_t<\bar{z}\}(2-c)\rho(Z_t)\right] && \text{[by (G.9)]}\\
&= \sum_{t=1}^{T} \mathbb{E}_1^{0} l(Z_t) = \sum_{t=1}^{T} \mathbb{E}_1^{\pi_I,\xi_1^*}[l(Z_t)].
\end{aligned}
$$

In the derivations above, we adopt the notational convention that when $\bar{z} = -\infty$, (4.13) reduces to $l(z) = (4-2c)\rho^2(z)$ for all $z \in \mathbb{Z}$. As a consequence of the above identities, we have the desired result. ■

Proof of Proposition 3. Let $r \in (0, \bar{r})$, where $\bar{r}$ is as in (G.2). In addition, let π_I be an inertial policy with the residual probability sequence $\rho = \{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$, and extend ρ to $\mathbb{Z}$ according to (E.2). We first the following result, the proof of which is at the end of this section.

LEMMA H.1. *If $r \in (0, \bar{r})$, then the sequence $\rho = \{\rho(z) = \frac{1}{r_0 + rz}, z \in \mathbb{Z}_+\}$ is regular and slow vanishing.*

Because ρ is regular and slowly vanishing (by Lemma H.1), we choose $\bar{Z} := \bar{z}$ in the context of Proposition K.2 and deduce that Statements (PK.2:1) and (PK.2:3) imply Statements (P3:1) and (P3:3), respectively.

The last step is to verify Statement (P3:2). By Lemmas E.2 and E.3, $\rho(z) \in (0, \frac{1}{2} - \alpha)$ for all $z \in \mathbb{Z}$. Thus, we deduce from (4.13) that $(4 - 2c)\rho^2(z) \le l(z)$ for all $z \in \mathbb{Z}$. By Statement (P3:3), $\sum_{t=1}^{T} \mathbb{E}_1^0[\rho^2(Z_t)] \le \frac{1}{4-2c} \sum_{t=1}^{T} \mathbb{E}_1^0[l(Z_t)] = O(\log T)$. Hence, $\sum_{t=1}^{T} \mathbb{E}_1^0[\rho(Z_t)] \le \sqrt{T \sum_{t=1}^{T} \left(\mathbb{E}_1^0[\rho(Z_t)]\right)^2} \le \sqrt{T \sum_{t=1}^{T} \mathbb{E}_1^0[\rho^2(Z_t)]} = O(\sqrt{T} \log T)$. Moreover, $\sum_{t=1}^{T} \mathbb{E}_1^0[\rho(Z_t)] \ge \sum_{t=1}^{T} \mathbb{E}_1^0[\rho^2(Z_t)] \to \infty$ as $T \to \infty$. ∎

Proof of Lemma H.1. Observe that (i) $\lim_{z\to\infty} \{z\rho(z)\} = \lim_{z\to\infty} \left\{\frac{z}{r_0 + rz}\right\} = \frac{1}{r} > \frac{1}{4}$ because $r < \bar{r} < 4$, and (ii) $\lim_{z\to\infty} \{\rho(z)\} = \lim_{z\to\infty} \left\{\frac{1}{r_0 + rz}\right\} = 0$. Therefore, by Definition 2, ρ is slowly vanishing. Now, note that $\lim_{z\to\infty} \left[\left(\frac{\rho(z)}{\rho(z+1)} - 1\right) z\right] = \lim_{z\to\infty} \left[\left(\frac{r_0 + rz + r}{r_0 + rz} - 1\right) z\right] = 1$. As a result, by Definition 3, ρ is regular. ∎

Appendix I: Analysis of the Random Blocking Model (Theorem 5)

This section provides the details for the proof of Theorem 5, which generalizes Theorems 1 and 2 to accommodate random blocking by myopic bettors. Let us restate Theorem 5 by breaking it into two separate results, the first generalizing Theorem 1 and the second generalizing Theorem 2.

The result below, which restates Statement (T5:1) in Theorem 5, generalizes Theorem 1 by incorporating myopic bettors' random blocking. We present its proof in Appendix I.2.

THEOREM I.1. *(low blocking probability) Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$. Then there exists $\underline{q} = \underline{q}(\hat{\Xi}) \in (0,1)$ such that for all $q \le \underline{q}$, $b_1 \in (0,1)$, and sufficiently small $c > 0$, we have the following:*

- *(TI.1:1) (non-convergence of spread line) For some $i \in \{0,1\}$, with strictly positive $\hat{\mathbb{P}}_i^{\pi_B, \hat{\xi}_i^*}$-probability, $\mathfrak{d}_t$ does not converge to zero.*
- *(TI.1:2) (linearly growing regret) $\hat{\Delta}^{\pi_B}(T) = \Omega(T)$.*

When $q = 0$, Theorem I.1 reduces to Theorem 1. This means that all of the conclusions in Theorem 1 hold even if we perturb the random blocking probability by a small constant (independent of T).

The following result, which restates Statement (T5:2) in Theorem 5, generalizes Theorem 2 by allowing for random blocking by myopic bettors. Its proof is in Appendix I.3.

THEOREM I.2. *(high blocking probability) Suppose that the market maker uses a Bayesian policy π_B with a regular pricing function $s^{\pi_B}(\cdot)$. Then there exists $\bar{q} = \bar{q}(\hat{\Xi}) \in (0,1)$ such that for all $q \ge \bar{q}$, $b_1 \in (0,1)$ and $i \in \{0,1\}$, we have:*

(TI.2:1) (convergence of spread lines) $\mathfrak{d}_t$ *converges to zero almost surely under* $\hat{\mathbb{P}}_i^{\pi_B,\hat{\xi}_i^*}$.

(TI.2:2) (rate of convergence) $\hat{\mathbb{E}}_i^{\pi_B,\hat{\xi}_i^*}[\mathfrak{d}_t] = O\left(e^{-\lambda t}\right)$ *for some constant* $\lambda > 0$.

(TI.2:3) (bounded regret) $\hat{\Delta}^{\pi_B}(T) = O(1)$.

When $q = 1$, Theorem I.2 reduces to Theorem 2, meaning that the conclusions in Theorem 2 continue to hold even if the random blocking probability is perturbed by a small constant (independent of T).

I.1. Main Proof Idea: the One-stage Analysis Under Random Blocking

Our proof roadmaps for Theorems I.1 and I.2 are similar to those for Theorems 1 and 2. What differentiates the generalized proofs from their original versions is extending the functions $D(\cdot,\cdot)$ and $R_i(\cdot,\cdot)$ introduced in Appendix C.3 to incorporate random blocking by myopic bettors. Formally, define

$$\hat{D}_i(b,\mathfrak{p}) := \left[(1-q)(1-\mathfrak{p}) + qF_i\left(s^\pi(b)\right)\right]\log\left(\frac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right) + \left[(1-q)\mathfrak{p} + q\bar{F}_i\left(s^\pi(b)\right)\right]\log\left(\frac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)$$

as the expected increment of the market maker's log-likelihood process L_t after a single bet under H_i if (i) the current belief state is b, and (ii) when his bet is not blocked, the informed bettor bets positively with probability $\mathfrak{p}$ and negatively with probability $1-\mathfrak{p}$. Here, the expectation is taken over the random blocking by myopic bettors, the randomized strategy of the informed bettor, and the random behavior of myopic bettors. The informed bettor misleads the market maker if $\hat{D}_1(b,\mathfrak{p}) < 0$ or $\hat{D}_0(b,\mathfrak{p}) > 0$. Now, define

$$\hat{R}_i(b,\mathfrak{p}) := (1-q)\left[(1-\mathfrak{p})j_i^-(s^\pi(b)) + \mathfrak{p}j_i^+(s^\pi(b))\right]$$

to be the informed bettor's expected profit from a single bet under H_i if (i) the current belief state is b, and (ii) when his bet is not blocked, the informed bettor bets positively with probability $\mathfrak{p}$ and negatively with probability $1-\mathfrak{p}$. Here, the expectation is taken over the random blocking by myopic bettors, the randomized strategy of the informed bettor, and the final realization of the event outcome X. The informed bettor makes a profit in expectation if $\hat{R}_i(b,\mathfrak{p}) > 0$.

Summary of key steps. In the proofs of Theorems I.1 and I.2, we utilize our one-stage analysis in the same way we utilize it in the proofs of Theorems 1 and 2. By incorporating the additional randomness from blocking, we immediately obtain the following extended version of Lemma C.1.

LEMMA I.1. *Let* $i \in \{0,1\}$. *Suppose that the market maker uses a Bayesian policy* π_B *and the type-*i *informed bettor's policy* ξ *is given by the behavioral strategy* $\{\mathfrak{p}(b)\}$. *Then, we have the following:*

1. *If there exists* $\delta > 0$ *such that* $\hat{\mathbb{E}}_i^{\pi_B,\xi}[L_{t+1} - L_t | b_t = b] = \hat{D}_i(b,\mathfrak{p}(b)) < -\delta$ *for all* $b \in (0,1)$ *and* $t \in \mathbb{Z}_+$, *then (i)* $\hat{\mathbb{E}}_i^{\pi_B,\xi}[b_t] = O(e^{-\lambda t})$ *for some* $\lambda > 0$, *and (ii)* $b_t \to 0$, $L_t \to -\infty$, $s_t \to s^\pi(0+)$ *almost surely. If in addition, there exists* $\bar{b} \in (0,1)$ *such that* $\hat{R}_i(b,\mathfrak{p}(b)) > \delta$ *for all* $b \in (0,\bar{b}]$, *then* $\hat{V}_i^{\pi_B,\xi}(T) = \Omega(T)$.
2. *If there exists* $\delta > 0$ *such that* $\hat{\mathbb{E}}_i^{\pi_B,\xi}[L_{t+1} - L_t | b_t = b] = \hat{D}_i(b,\mathfrak{p}(b)) > \delta$ *for all* $b \in (0,1)$ *and* $t \in \mathbb{Z}_+$, *then (i)* $\hat{\mathbb{E}}_i^{\pi_B,\xi}[1-b_t] = O(e^{-\lambda t})$ *for some* $\lambda > 0$, *and (ii)* $b_t \to 1$, $L_t \to \infty$, $s_t \to s^\pi(1-)$ *almost surely. If in addition, there exists* $\bar{b} \in (0,1)$ *such that* $\hat{R}_i(b,\mathfrak{p}(b)) > \delta$ *for all* $b \in [\bar{b},1)$, *then* $\hat{V}_i^{\pi_B,\xi}(T) = \Omega(T)$.

Next, we provide below a summary of the results pertaining to the one-stage analysis under random blocking. Because of the new source of randomness, the one-stage analysis is a bit more involved. These results on the generalized one-stage analysis (i.e., Lemmas I.2, I.3, and I.4 below) are what makes the proofs of Theorems I.1 and I.2 different from those of Theorems 1 and 2.

The following result applied to the scenario where $s^{\pi_B}(0+) \le m_0$, $s^{\pi_B}(1-) \ge m_1$, and the blocking probability q is sufficiently small. Recall that $\Xi = (c, m_0, m_1, F_\epsilon)$ is the collection of problem input parameters, and $\hat{\Xi} := (m_0, m_1, F_\epsilon)$ is the collection of problem input parameters concerning distribution information only, i.e., those except the commission rate c.

LEMMA I.2. *(one-stage analysis for manipulation under random blocking) Suppose that the market maker uses a Bayesian policy π_B such that $s^{\pi_B}(0+) \le m_0$ and $s^{\pi_B}(1-) \ge m_1$. Then, there exist $\bar{c}_0, \underline{q}_0, \mathfrak{p}_0 \in (0,1)$, which depend only on $\hat{\Xi}$, such that for all $c \le \bar{c}_0$ and $q \le \underline{q}_0$, there exist $\bar{b} = \bar{b}(\Xi, q) \in (0,1)$ and $\delta = \delta(\Xi, q) > 0$ satisfying the following:*

1. *(global manipulability) For all $b \in (0,1)$, $\hat{D}_1(b,0) < -\delta$ and $\hat{D}_0(b,1) > \delta$.*
2. *(local profitable manipulation; type-1) For all $b \in (0, \bar{b}]$, $\hat{D}_1(b, \mathfrak{p}_0) < -\delta$ and $\hat{R}_1(b, \mathfrak{p}_0) > \delta$.*
3. *(local profitable manipulation; type-0) For all $b \in [1-\bar{b}, 1)$, $\hat{D}_0(b, 1-\mathfrak{p}_0) > \delta$ and $\hat{R}_0(b, 1-\mathfrak{p}_0) > \delta$.*

Observing that Lemma I.2 reduces to Lemma C.2 when $q = 0$, we note that the conclusions in Lemma C.2 hold even if we perturb the random blocking probability by a small constant (independent of T).

The following result applies to the case where $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$, and the blocking probability q is sufficiently small.

LEMMA I.3. *(one-stage analysis for honest betting under random blocking) Suppose that the market maker uses a Bayesian policy π_B such that $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$. Then, there exists $\bar{c}_1 = \bar{c}_1(\hat{\Xi}, \pi_B) \in (0,1)$ such that for all $c \le \bar{c}_1$ and $q \in (0,1)$, there exist $\bar{b} = \bar{b}(\Xi, q, \pi_B) \in (0,1)$ and $\delta = \delta(\Xi, q, \pi_B) > 0$ satisfying the following:*

1. *(correcting power) For all $b \in (0,1)$, $\hat{D}_1(b,1) > \delta$ and $\hat{D}_0(b,0) < -\delta$.*
2. *(local profitable honest betting) Either of the following is true:*
   - *(type-1) for all $b \in [1-\bar{b}, 1)$, $\hat{R}_1(b,1) > \delta$.*
   - *(type-0) for all $b \in (0, \bar{b}]$, $\hat{R}_0(b,0) > \delta$.*

Note that when $q = 0$, Lemma I.3 reduces to the analysis in the proof of Proposition C.2, where the informed bettor honestly bet all the time.

The following result applies to the case where the blocking probability q is sufficiently large.

LEMMA I.4. *(one-stage analysis for high blocking probability) There exists $\bar{q} \in (0,1)$, which depends only on $\hat{\Xi}$, such that for all $\mathfrak{p} \in (0,1)$, $b \in (0,1)$ and $q \ge \bar{q}$, $\hat{D}_1(b, \mathfrak{p}) > 0$ and $\hat{D}_0(b, \mathfrak{p}) < 0$.*

When $q = 0$, Lemma I.4 reduces to the analysis in the proof of Theorem 2, where the informed bettor does not bet at all.

I.2. The Low Blocking Probability Case (Theorem I.1)

Following the same roadmap as in the proof of Theorem 1 (Appendix C.1), we aim to identify profitable strategies for the informed bettor if the market maker uses BPs in presence of random blocking by myopic bettors. There are two cases regarding the values of $s^{\pi_B}(0+)$ and $s^{\pi_B}(1-)$, each corresponding to a profitable strategy for the informed bettor.

The result below generalizes Proposition C.1 by accommodating random blocking by myopic bettors.

PROPOSITION I.1. *(bluffing under random blocking) Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$ such that $s^{\pi}(0+) \leq m_0$ and $s^{\pi_B}(1-) \geq m_1$. Then there exists $\underline{q}_0 = \underline{q}_0(\hat{\Xi}) \in (0,1)$ such that for every blocking probability $q \leq \underline{q}_0$, initial belief $b_1 \in (0,1)$, hypothesis $i \in \{0,1\}$ and sufficiently small commission rate c, the type-i informed bettor has a "bluffing" policy ξ_b satisfying the following:*

- *(PI.1:1) (belief and spread line dynamics) The posterior belief b_t converges to $(1-i)$ and the spread line s_t converges to m_{1-i} almost surely under $\hat{\mathbb{P}}_i^{\pi_B,\xi_b}$;*
- *(PI.1:2) (linearly growing profit of the informed bettor) $\hat{V}_i^{\pi_B,\xi_b}(T) = \Omega(T)$.*

Noting that Proposition I.1 reduces to Proposition C.1 when $q = 0$, we deduce that the conclusions in Proposition C.1 remain to be valid even if the random blocking probability is perturbed by a small constant (independent of T).

Proof of Proposition I.1. After making the appropriate notational changes to incorporate the new source of randomness due to probabilistic blocking (i.e., replacing $\mathbb{P}_i$, $\mathbb{E}_i$, ξ_i^*, V, Δ, D, and R_i with $\hat{\mathbb{P}}_i$, $\hat{\mathbb{E}}_i$, $\hat{\xi}_i^*$, $\hat{V}$, $\hat{\Delta}$, $\hat{D}_i$, and $\hat{R}_i$, respectively), this proposition follows from Lemmas I.1 and I.2 in the same way that Proposition C.1 follows from Lemmas C.1 and C.2. ∎

The result below generalizes Proposition C.2 by accommodating random blocking by myopic bettors.

PROPOSITION I.2. *(honest betting under random blocking) Suppose that the market maker uses a Bayesian policy π_B with pricing function $s^{\pi_B}(\cdot)$ such that $s^{\pi_B}(0+) > m_0$ or $s^{\pi_B}(1-) < m_1$. Then for some hypothesis $i \in \{0,1\}$, every blocking probability $q \in (0,1)$, initial belief $b_1 \in (0,1)$, and sufficiently small commission rate c, the type-i informed bettor has an "honest" policy ξ_h satisfying the following:*

- *(PI.2:1) (belief and spread line dynamics) With $\hat{\mathbb{P}}_i^{\pi_B,\xi_h}$-probability 1, posterior belief b_t converges to the truth, i. The spread line s_t converges to a limit s_∞, but $s_\infty \neq m_i$;*
- *(PI.2:2) (linearly growing profit of the informed bettor) $\hat{V}_i^{\pi_B,\xi_h}(T) = \Omega(T)$.*

Proposition I.2 means that all the conclusions in Proposition C.2 (which is a special case of Proposition I.2 when $q = 0$) hold even if we perturb the random blocking probability arbitrarily within the range $(0,1)$.

Proof of Proposition I.2. This proposition follows from repeating the proof of Proposition C.2 with the following changes to incorporate probabilistic blocking: (i) replacing the notations $\mathbb{P}_i$, $\mathbb{E}_i$, ξ_i^*, V, Δ, D, and

R_i with $\hat{\mathbb{P}}_i$, $\hat{\mathbb{E}}_i$, $\hat{\xi}_i^*$, $\hat{V}$, $\hat{\Delta}$, $\hat{D}_i$, and $\hat{R}_i$, respectively; (ii) replacing Step 1 in the proof of Proposition C.2 with Lemma I.3 (this step corresponds to the one-stage analysis); and (iii) replacing Lemma C.1 with Lemma I.1 (this step corresponds to showing how the one-stage analysis leads to the final result). ■

Proof of Theorem I.1. The logical deduction from Propositions C.1 and C.2 to Theorem 1 is the same as from Propositions I.1 and I.2 to Theorem I.1. In other words, let $\underline{q} := \underline{q}_0$ in Proposition I.1 and fix $q \in (0, \underline{q})$. The rest of the proof follows from repeating the arguments in that of Theorem 1, with the following changes to incorporate random blocking by myopic bettors: (i) replacing $\mathbb{P}_i$, $\mathbb{E}_i$, ξ_i^*, V, Δ, $\mathbb{I}\{a_t = 0\}$, and $\mathbb{I}\{a_t \neq 0\}$ with $\hat{\mathbb{P}}_i$, $\hat{\mathbb{E}}_i$, $\hat{\xi}_i^*$, $\hat{V}$, $\hat{\Delta}$, $\mathbb{I}\{a_t = 0 \text{ or } \chi_t = 1\}$, and $\mathbb{I}\{a_t \neq 0 \text{ and } \chi_t = 0\}$, respectively; and (ii) replacing Propositions C.1 and C.2 with Propositions I.1 and I.2, respectively. ■

I.3. The High Blocking Probability Case (Theorem I.2)

Proof of Theorem I.2. This theorem follows from repeating the proof of Theorem 2 with the following changes to incorporate probabilistic blocking: (i) replacing the notations $\mathbb{P}_i$, $\mathbb{E}_i$, ξ_i^*, V, Δ, D, and R_i with $\hat{\mathbb{P}}_i$, $\hat{\mathbb{E}}_i$, $\hat{\xi}_i^*$, $\hat{V}$, $\hat{\Delta}$, $\hat{D}_i$, and $\hat{R}_i$, respectively; and (ii) replacing Step 1 in the proof of Theorem 2 with Lemmas I.1 and I.4. ■

I.4. Proofs of Auxiliary Lemmas

Proof of Lemma I.2. This proof is similar to the proof of Lemma C.2. We complete the proof in four steps.

Step 1. We claim that there exists $\delta_1 = \delta_1(\hat{\Xi}) > 0$ and $\varepsilon_q = \varepsilon_q(\hat{\Xi}) > 0$ such that (i) $\hat{D}_1(b, 0) < -\delta_1$ and (ii) $\hat{D}_0(b, 1) > \delta_1$ for all $b \in (0, 1)$ and $q \leq \varepsilon_q$. By Lemma A.3 and Assumption (A1:3), there exist $\bar{\delta}, M > 0$, which depend only on $\hat{\Xi}$, such that $-M \leq \log\left(\frac{F_1(s)}{F_0(s)}\right) \leq -\bar{\delta}$ and $\bar{\delta} \leq \log\left(\frac{\bar{F}_1(s)}{\bar{F}_0(s)}\right) \leq M$ for all $s \in \mathcal{S}$. Thus,

$$\begin{aligned}\hat{D}_1(b, 0) &= [(1-q) + qF_1(s^\pi(b))]\log\left(\frac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right) + q\bar{F}_1(s^\pi(b))\log\left(\frac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\\ &\leq (1-q)(-\bar{\delta}) + qM = -\bar{\delta} + (M + \bar{\delta})q,\end{aligned}$$

and

$$\begin{aligned}\hat{D}_0(b, 1) &= qF_0(s^\pi(b))\log\left(\frac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right) + \left[(1-q) + q\bar{F}_0(s^\pi(b))\right]\log\left(\frac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\\ &\geq q(-M) + (1-q)(\bar{\delta}) = \bar{\delta} - (\bar{\delta} + M)q.\end{aligned}$$

To prove our claim in this step, we choose $\delta_1 = \frac{\bar{\delta}}{2}$ and $\varepsilon_q = \frac{\bar{\delta}}{4(\bar{\delta}+M)}$, both of which depend only on $\hat{\Xi}$.

Step 2. We claim that exists $\bar{c}_0, \hat{\varepsilon}_q, \mathfrak{p}_0 \in (0, 1)$, which depend only on $\hat{\Xi}$, such that for all $c \leq \bar{c}_0$ and $q \leq \hat{\varepsilon}_q$,

- $\hat{D}_1(0+, \mathfrak{p}_0) < 0$ and $\hat{R}_1(0+, \mathfrak{p}_0) > 0$;
- $\hat{D}_0(1-, 1-\mathfrak{p}_0) > 0$ and $\hat{R}_0(1-, 1-\mathfrak{p}_0) > 0$.

That is, the type-1 (resp. type-0) informed bettor can make a profitable manipulation when the blocking probability is low, the commission rate is low, and the market maker's belief is close to 0 (resp. 1). To prove this claim, recall the following three quantities introduced in Step 2 of the proof of Lemma C.2 (where random blocking was not present):

$$\kappa = \frac{-\log 2\alpha}{\log 2(1-\alpha)} > 1,\ \hat{\kappa} = \frac{(\bar{c}_0 - 2)(1-\alpha) + 1}{(\bar{c}_0 - 2)(1-\alpha) + 1 - \bar{c}_0},\ \text{and } \bar{c}_0 = \frac{(\kappa-1)(1-2\alpha)}{2(\kappa-1)(1-\alpha)+1} \in (0, 1).$$

In that proof, $\mathfrak{p}_0$ was constructed such that it depends only on $\alpha = F_1(m_0)$ and hence only on $\hat{\Xi}$, and moreover, it satisfies the following inequalities:

$$(1-\mathfrak{p}_0)\log\left(\frac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right)+\mathfrak{p}_0\log\left(\frac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right)<0,$$
$$\mathfrak{p}_0\log\left(\frac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right)+(1-\mathfrak{p}_0)\log\left(\frac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right)>0.$$

In light of this, we choose $\mathfrak{p}_0$ in the same way, and uniquely construct $\hat{\varepsilon}_1,\hat{\varepsilon}_2>0$ that satisfy the following:

$$(1-\hat{\varepsilon}_1)\left[(1-\mathfrak{p}_0)\log\left(\frac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right)+\mathfrak{p}_0\log\left(\frac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right)\right]+\hat{\varepsilon}_1M=0,$$
$$(1-\hat{\varepsilon}_2)\left[\mathfrak{p}_0\log\left(\frac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right)+(1-\mathfrak{p}_0)\log\left(\frac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right)\right]-\hat{\varepsilon}_2M=0.$$

One can verify that both $\hat{\varepsilon}_1$ and $\hat{\varepsilon}_2$ depend only on $\hat{\Xi}$. Now, if $q\le\hat{\varepsilon}_q:=\min\{\frac{\hat{\varepsilon}_1}{2},\frac{\hat{\varepsilon}_2}{2},\frac{1}{2}\}$, then

$$\begin{aligned}
\hat{D}_1(0+,\mathfrak{p}_0)&=[(1-q)(1-\mathfrak{p}_0)+qF_1(s^\pi(0+))]\log\left(\frac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right)\\
&\quad+\left[(1-q)\mathfrak{p}_0+q\bar{F}_1(s^\pi(0+))\right]\log\left(\frac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right)\\
&\le(1-q)\left[(1-\mathfrak{p}_0)\log\left(\frac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right)+\mathfrak{p}_0\log\left(\frac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right)\right]+qM\\
&<(1-\hat{\varepsilon}_1)\left[(1-\mathfrak{p}_0)\log\left(\frac{F_1(s^\pi(0+))}{F_0(s^\pi(0+))}\right)+\mathfrak{p}_0\log\left(\frac{\bar{F}_1(s^\pi(0+))}{\bar{F}_0(s^\pi(0+))}\right)\right]+\hat{\varepsilon}_1M=0,
\end{aligned}$$

and

$$\begin{aligned}
\hat{D}_0(b,1-\mathfrak{p}_0)&=[(1-q)\mathfrak{p}_0+qF_0(s^\pi(1-))]\log\left(\frac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right)\\
&\quad+\left[(1-q)(1-\mathfrak{p}_0)+q\bar{F}_0(s^\pi(1-))\right]\log\left(\frac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right)\\
&\ge(1-q)\left[\mathfrak{p}_0\log\left(\frac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right)+(1-\mathfrak{p}_0)\log\left(\frac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right)\right]-qM\\
&>(1-\hat{\varepsilon}_2)\left[\mathfrak{p}_0\log\left(\frac{F_1(s^\pi(1-))}{F_0(s^\pi(1-))}\right)+(1-\mathfrak{p}_0)\log\left(\frac{\bar{F}_1(s^\pi(1-))}{\bar{F}_0(s^\pi(1-))}\right)\right]-\hat{\varepsilon}_2M=0.
\end{aligned}$$

Moreover, for all $q\le\hat{\varepsilon}_q$ and $c\le\bar{c}_0$,

$$\hat{R}_1(0+,\mathfrak{p}_0)=(1-q)[(1-\mathfrak{p}_0)j_1^-(s^\pi(0+))+\mathfrak{p}_0j_1^+(s^\pi(0+))]\overset{(a)}{>}0,$$
$$\hat{R}_0(1-,1-\mathfrak{p}_0)=(1-q)[\mathfrak{p}_0j_0^-(s^\pi(1-))+(1-\mathfrak{p}_0)j_0^+(s^\pi(1-))]\overset{(b)}{>}0.$$

In the derivations above, (a) and (b) follow from Step 2 in the proof in Lemma C.2.

Step 3. Based on Step 2, there exist $\bar{b},\delta_2,\delta_3,\delta_4,\delta_5>0$, all of which depend only on q and Ξ, such that

- *(local profitable manipulation; type-*1*)* $\hat{D}_1(b,\mathfrak{p}_0)<-\delta_2$ and $\hat{R}_1(b,\mathfrak{p}_0)>\delta_3$ for all $b\in(0,\bar{b}]$;
- *(local profitable manipulation; type-*0*)* $\hat{D}_1(b,1-\mathfrak{p}_0)>\delta_4$ and $\hat{R}_0(b,1-\mathfrak{p}_0)>\delta_5$ for all $b\in(1-\bar{b},1]$.

The existence is guaranteed by the local continuity of $\hat{D}_1(b,\mathfrak{p}_0)$ and $\hat{R}_1(b,\mathfrak{p}_0)$ with respect to b at point $0+$; as well as that of $\hat{D}_0(b,1-\mathfrak{p}_0)$ and $\hat{R}_0(b,1-\mathfrak{p}_0)$ with respect to b at point $1-$.

Step 4. In light of Steps 1 and 3, we complete the proof by taking $\underline{q}_0=\min\{\varepsilon_q,\hat{\varepsilon}_q\}$ and $\delta:=\min\{\delta_1,\delta_2,\delta_3,\delta_4,\delta_5\}$. Note that $\underline{q}_0$ depends only on $\hat{\Xi}$, and δ depends only on (Ξ,q). ■

Proof of Lemma I.3. Choose $\bar{\delta}>0$ as in Lemma A.3 so that $\log\left(\frac{F_0(s)}{F_1(s)}\right)\geq\bar{\delta}$ and $\log\left(\frac{\bar{F}_1(s)}{\bar{F}_0(s)}\right)\geq\bar{\delta}$ for all $s\in\mathcal{S}$. Let $q\in(0,1)$. Note that for all $b\in(0,1)$,

$$\begin{aligned}\hat{D}_1(b,1)&=qF_1\left(s^\pi(b)\right)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\left[(1-q)+q\bar{F}_1\left(s^\pi(b)\right)\right]\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\\&=(1-q)\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)+q\left[F_1\left(s^\pi(b)\right)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\bar{F}_1\left(s^\pi(b)\right)\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\right]\\&\geq(1-q)\bar{\delta}>0.\end{aligned}$$

Similarly,

$$\begin{aligned}\hat{D}_0(b,0)&=\left[(1-q)+qF_0\left(s^\pi(b)\right)\right]\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+q\bar{F}_0\left(s^\pi(b)\right)\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\\&=(1-q)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+q\left[F_0\left(s^\pi(b)\right)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\bar{F}_0\left(s^\pi(b)\right)\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\right]\\&\leq-(1-q)\bar{\delta}<0.\end{aligned}$$

As in the proof of Proposition C.2, choose

$$\bar{c}_1:=\min\left\{\max\left\{\tfrac{2F_0(s^{\pi_B}(0+))-1}{2F_0(s^{\pi_B}(0+))},\tfrac{1-2F_1(s^{\pi_B}(1-))}{2F_1(s^{\pi_B}(1-))}\right\},\tfrac{1}{2}\right\}\in(0,1),$$

and let $c\in(0,\bar{c}_1]$. If $s^{\pi_B}(0+)>m_0$, then $\hat{R}_0(0+,0)=(1-q)j_0^-(s^{\pi_B}(0+))>0$. If $s^{\pi_B}(1-)<m_1$, then $\hat{R}_1(1-,1)=(1-q)j_1^+(s^{\pi_B}(1-))>0$. Since $\hat{R}_0(b,\mathfrak{p})$ is continuous in b at $0+$ and $\hat{R}_1(b,\mathfrak{p})$ is continuous in b at $1-$, there exist $\bar{b},\hat{\delta}>0$, which depend only on Ξ,q,π_B, such that either of the following is true:

- (type-1) for all $b\in[1-\bar{b},1)$, $\hat{R}_1(b,1)>\hat{\delta}$ (this happens if $s^{\pi_B}(1-)<m_1$).
- (type-0) for all $b\in(0,\bar{b}]$, $\hat{R}_0(b,0)>\hat{\delta}$ (this happens if $s^{\pi_B}(0+)>m_0$).

The proof is finished by choosing $\delta:=\min\{\frac{(1-q)\bar{\delta}}{2},\hat{\delta}\}$. ■

Proof of Lemma I.4. By Lemma A.3 and Assumption (A1:3), there exist $\varepsilon,\bar{\delta},M>0$, which depend only on $\hat{\Xi}$, such that $-M\leq\log\left(\frac{F_1(s)}{F_0(s)}\right)\leq-\bar{\delta}$, $\bar{\delta}\leq\log\left(\frac{\bar{F}_1(s)}{\bar{F}_0(s)}\right)\leq M$, $F_0(s)-F_1(s)\geq\bar{\delta}$, $\varepsilon\leq F_0(s)$, and $F_1(s)\leq 1-\varepsilon$ for all $s\in\mathcal{S}$. Let $d_{KL}(x,y):=x\log\left(\frac{x}{y}\right)+(1-x)\log\left(\frac{1-x}{1-y}\right)$ be the Kullback-Leibler divergence between two Bernoulli random variables with success probabilities x and y. This function is continuous and strictly positive on the closed polytope $\mathcal{P}:=\{(x,y):y-x\geq\bar{\delta},\varepsilon\leq x\leq 1-\varepsilon,y\leq 1-\varepsilon\}$ and hence $\inf_{(x,y)\in\mathcal{P}}\{d_{KL}(x,y)\}>0$. Observe that

$$\begin{aligned}\hat{D}_1(b,\mathfrak{p})&=\left[(1-q)(1-\mathfrak{p})+qF_1\left(s^\pi(b)\right)\right]\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\left[(1-q)\mathfrak{p}+q\bar{F}_1\left(s^\pi(b)\right)\right]\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\\&\geq q\left[F_1\left(s^\pi(b)\right)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\bar{F}_1\left(s^\pi(b)\right)\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)\right]-(1-q)M\\&=q\left[d_{KL}(F_1\left(s^\pi(b)\right),F_0\left(s^\pi(b)\right))\right]-(1-q)M\\&=q\left[\inf_{(x,y)\in\mathcal{P}}\{d_{KL}(x,y)\}\right]-(1-q)M\\&=q\left[\inf_{(x,y)\in\mathcal{P}}\{d_{KL}(x,y)\}+M\right]-M.\end{aligned}$$

Similarly,

$$\hat{D}_0(b,\mathfrak{p})=\left[(1-q)(1-\mathfrak{p})+qF_0\left(s^\pi(b)\right)\right]\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\left[(1-q)\mathfrak{p}+q\bar{F}_0\left(s^\pi(b)\right)\right]\log\left(\tfrac{\bar{F}_1(s^\pi(b))}{\bar{F}_0(s^\pi(b))}\right)$$

$$
\begin{aligned}
&\le q\left[F_0\left(s^\pi(b)\right)\log\left(\tfrac{F_1(s^\pi(b))}{F_0(s^\pi(b))}\right)+\bar F_0\left(s^\pi(b)\right)\log\left(\tfrac{\bar F_1(s^\pi(b))}{\bar F_0(s^\pi(b))}\right)\right]+(1-q)M\\
&=-q\left[F_0\left(s^\pi(b)\right)\log\left(\tfrac{F_0(s^\pi(b))}{F_1(s^\pi(b))}\right)+\bar F_0\left(s^\pi(b)\right)\log\left(\tfrac{\bar F_0(s^\pi(b))}{\bar F_1(s^\pi(b))}\right)\right]+(1-q)M\\
&=-q\left[d_{KL}(\bar F_0\left(s^\pi(b)\right),\bar F_1\left(s^\pi(b)\right))\right]+(1-q)M\\
&\le -q\Big[\inf_{(x,y)\in\mathcal{P}}\{d_{KL}(x,y)\}+M\Big]+M.
\end{aligned}
$$

We complete the proof by letting $\bar q := \frac{M}{M+\inf_{(x,y)\in\mathcal{P}}\{d_{KL}(x,y)\}}\in(0,1)$. ■

Appendix J: Analysis of the Budget-constrained Model (Theorem 6)

Proof of Theorem 6. Let π_B be the market maker's Bayesian policy, and $b_1\in(0,1)$ be her initial belief. We complete the proof in three steps.

Step 1. We claim that for a sufficiently small commission rate $c>0$, $\breve\Delta^{\pi_B}(T;K)=\Omega(T\wedge K)$. To prove this claim, observe that in the proofs of Propositions C.1 and C.2, we construct strategies for the informed bettor such that the informed bettor bets every time and the bets do not depend on the time horizon. Thus, using the arguments in the proofs of Proposition C.1 and C.2, we deduce that for some hypothesis $i\in\{0,1\}$ and sufficiently small commission rate c, the type-i informed bettor has a feasible strategy $\breve\xi_i$ such that $\breve V^{\pi_B,\breve\xi_i}(T;K)=\Omega(T\wedge K)$. Because the informed bettor makes profits at the expense of the market maker, we deduce from the arguments in the proof of Theorem 1 that $\breve\Delta^{\pi_B}(T;K)=\Omega(T\wedge K)$.

Step 2. We claim that $\breve\Delta^{\pi_B}(T;K)=O(T)$. The intuition for this is that the market maker's regret is at most a constant per bet. Formally, note that $\breve\Delta^{\pi_B}(T;K)=\max\{\Delta_0^{\pi_B,\breve\xi_0^*}(T),\Delta_1^{\pi_B,\breve\xi_1^*}(T)\}$ and for $i\in\{0,1\}$,

$$
\begin{aligned}
\Delta_i^{\pi_B,\breve\xi_i^*}(T)&=\tfrac{cT}{2}-\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\mathbb{I}\{(X-s_t)\,d_t<0\}-(1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\\
&=\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\tfrac{c}{2}-\mathbb{I}\{(X-s_t)\,d_t<0\}+(1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\\
&\le\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\tfrac{c}{2}+(1-c)\right]=(1-\tfrac{c}{2})T.
\end{aligned}
$$

Step 3. We claim that if the pricing function $s^{\pi_B}(\cdot)$ is regular, then $\breve\Delta^{\pi_B}(T;K)=O(K)$. Without loss of generality, suppose that K is sublinear in T, that is, $\limsup_{T\to\infty}\{\frac{K}{T}\}=0$. Choose $i\in\{0,1\}$, and observe that:

$$
\begin{aligned}
&\Delta_i^{\pi_B,\breve\xi_i^*}(T)\\
&=\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\tfrac{c}{2}-\mathbb{I}\{(X-s_t)\,d_t<0\}+(1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\\
&=\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\tfrac{c}{2}-\mathbb{I}\{(X-s_t)\,d_t<0\}+(1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\mathbb{I}\left\{a_t=0\text{ or }\sum_{\ell=1}^t|a_\ell|>K\right\}\\
&\quad+\sum_{t=1}^T\breve{\mathbb{E}}_i^{\pi_B,\breve\xi_i^*}\left[\tfrac{c}{2}-\mathbb{I}\{(X-s_t)\,d_t<0\}+(1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\mathbb{I}\left\{a_t\neq0\text{ and }\sum_{\ell=1}^t|a_\ell|\le K\right\}
\end{aligned}
$$

$$
\begin{aligned}
&= \sum_{t=1}^{T} \breve{\mathbb{E}}_i^{\pi_B,\breve{\xi}_i^*}\left[\tfrac{c}{2} - r_i(s_t)\right]\mathbb{I}\left\{a_t = 0 \text{ or } \sum_{\ell=1}^{t}|a_\ell| > K\right\} \\
&\quad + \sum_{t=1}^{T} \breve{\mathbb{E}}_i^{\pi_B,\breve{\xi}_i^*}\left[\tfrac{c}{2} - \mathbb{I}\{(X-s_t)\,d_t<0\} + (1-c)\mathbb{I}\{(X-s_t)\,d_t>0\}\right]\mathbb{I}\{a_t \neq 0\} \\
&\leq \sum_{t=1}^{T} \breve{\mathbb{E}}_i^{\pi_B,\breve{\xi}_i^*}\left[\tfrac{c}{2} - r_i(s_t)\right] + (1-\tfrac{c}{2})K.
\end{aligned}
$$

Because K is sublinear in T, L_t diverges in a linear speed and s_t converges to m_i in a linear speed. Hence, the first term above is independent of T and at most linear in K. As a result, $\breve{\Delta}_i^{\pi_B,\breve{\xi}_i^*}(T) = O(K)$. ■

Appendix K: On the Lower Bound of Regret (Theorem 7)

We prove Theorem 7 in a more general setting when the informed bettor participates in the market with a threshold strategy.

K.1. Description of the Setting

The $\bar{Z}$-threshold response strategy. To shed light on how the residual probability sequence should be selected, let us first consider the following family of threshold strategies for the informed bettor.

DEFINITION 4. ($\bar{Z}$-threshold strategy) Given $\bar{Z} \in \mathbb{Z} \cup \{-\infty\}$, the type-$i$ informed bettor's $\bar{Z}$-threshold strategy $\xi_i^{\bar{Z}}$ is such that $\xi_1^{\bar{Z}}(z) = \mathbb{I}\{z < \bar{Z}\}$ and $\xi_0^{\bar{Z}}(z) = -\mathbb{I}\{z > -\bar{Z}\}$.

Note that the $\bar{Z}$-threshold strategy might not necessarily be the informed bettor's best response strategy in general. In fact, although the definition of IP focuses on Markov policies with state variable Z_t, the general characterization of informed bettor's best response strategy for an arbitrary IP is quite complex. However, the analysis of generic $\bar{Z}$-threshold strategies helps us understand how to make a suitable choice for the function $\rho(\cdot)$. In particular, the informed bettor's best response strategy is indeed of threshold type when $\rho(z) = \frac{1}{r_0 + rz}$ with an appropriate value of r (see Theorem 3), where the threshold $\bar{z}$ maximizes the informed bettor's profit. Since we evaluate the profit performance of IP under generic choices for the threshold $\bar{Z}$ and the function $\rho(\cdot)$, our analysis in this section generalizes Theorem 3 and 4. In this sense, a well-performing IP should at least have reasonably good performance against the informed bettor's $\bar{Z}$-threshold strategies.

To account for the generic choice of $\bar{Z}$ instead of $\bar{z}$, let us introduce $\tilde{\mathbb{P}}_i^z(\cdot) := \mathbb{P}_i^{\pi_I,\xi_i^{\bar{Z}}}(\cdot \mid Z_1 = z)$, $\tilde{\mathbb{E}}_i^z[\cdot] := \mathbb{E}_i^{\pi_I,\xi_i^{\bar{Z}}}[\cdot \mid Z_1 = z]$, and $\tilde{l}(z) := (2-c)\rho(z)\mathbb{I}\{z < \bar{Z}\} + (4-2c)\rho^2(z)\mathbb{I}\{z \geq \bar{Z}\}$ for the shorthand notations in this section. We also denote by $\tilde{\Delta}^{\pi_I}(T)$ the market maker's regret under IP against the informed bettor's $\bar{Z}$-threshold strategies.

K.2. Key Intermediate Results

Our main results in this section characterize the dynamics of $\{Z_t\}$ under the market maker's IP and the informed bettor's $\bar{Z}$-threshold strategy. For simplicity, we consider the case where $i = 1$ without loss of generality. The analysis for the case where $i = 0$ follows from the symmetry relations between $\tilde{\mathbb{P}}_0^z$ and $\tilde{\mathbb{P}}_1^z$ akin to (G.9). The next result characterizes the convergence behavior of Z_t when $\rho(\cdot)$ is fast vanishing.

PROPOSITION K.1. *Let $\rho = \{\rho(z), z \in \mathbb{Z}_+\}$ be a fast vanishing residual probability sequence, $m > \bar{Z} - 2$, and $z \in \mathbb{Z}$. Then, $\{m\}$ is recurrent under $\tilde{\mathbb{P}}_1^z$.*

Proposition K.1 means that when $\rho(z) < \frac{1}{4z}$ in the limit, Z_t does not diverge to infinity. Note that this is an undesirable property because it implies that the spread line s_t does not converge to the correct median. The main reason behind Proposition K.1 is that if $\rho(z)$ becomes too small (or $\tilde{s}(z)$ gets close to the median) as $z \to \infty$, the process Z_t behaves like a simple symmetric random walk on $\mathbb{Z}$, which is recurrent. From a technical point of view, this result is closely related to the (standard) characterization of recurrence of a birth and death chain with a reflecting boundary point (see, e.g., Durrett 2019).

In contrast to Proposition K.1, Proposition K.2 below characterizes the dynamics of the market when $\rho(\cdot)$ is slowly vanishing.

PROPOSITION K.2. *Let $\rho = \{\rho(z), z \in \mathbb{Z}_+\}$ be a slowly vanishing residual probability sequence. Then, we have the following:*

(PK.2:1) For all sufficiently large $M > 0$, $\sum_t \tilde{\mathbb{E}}_1^0[\mathbb{I}\{Z_t \le M\}]$ converges.
(PK.2:2) $Z_t \to \infty$ almost surely under $\tilde{\mathbb{P}}_1^0$. As a result, $\mathfrak{d}_t$ converges to zero almost surely under $\tilde{\mathbb{P}}_1^0$.
(PK.2:3) If ρ is regular, then $\sum_t \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)]$ diverges at a rate satisfying $\sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] = O\big(\sum_{t=1}^T \rho(t)\big)$.
(PK.2:4) If ρ is regular, then $\tilde{\Delta}^{\pi_I}(T)$ diverges in T at a rate satisfying $\tilde{\Delta}^{\pi_I}(T) = O\big(\sum_{t=1}^T \rho(t)\big)$.

Proposition K.2 means that when $\rho(z)$ vanishes more slowly than $\frac{1}{4z}$ as $z \to \infty$, the following happens: (i) the spread line converges to the correct median, and (ii) under certain regularity conditions, the market maker's regret (loss) against the $\bar{Z}$-threshold strategy grows in the order of $O\big(\sum_{t=1}^T \rho(t)\big)$.

The above proposition extends the convergence and regret analysis in Proposition 3 and Theorem 4 in the following way: Statements (PK.2:1) and (PK.2:3) generalize Statements (P3:1) and (P3:3) in Proposition 3, respectively. As a consequence, Statements (PK.2:2) and (PK.2:4) generalize Statements (T4:1) and (T4:3) in Theorem 4, respectively. In all of the generalizations, we study a larger family of $\rho(\cdot)$ rather than the particular choice of $\rho(z) = \frac{1}{r_0 + rz}$ with $r \in (0, \bar{r})$, and we study a generic $\bar{Z}$-threshold strategy rather than the particular choice of ξ_1^*.

K.3. Proof of Theorem 7

Consider an IP for the market maker with some residual probability sequence $\rho = \{\rho(z) : \mathbb{Z}_+ \to \frac{1}{2} - \alpha\}$. Let c be sufficiently large as in Lemma D.1 so that $\xi_i^* = \xi_\emptyset$. By construction and using the arguments in the proof of Lemma 5, we deduce that $\Delta^{\pi_I}(T) = \max\{\Delta_0^{\pi_I, \xi_\emptyset}(T), \Delta_1^{\pi_I, \xi_\emptyset}(T)\} = \max\{\tilde{\Delta}_0^{\pi_I, \xi_\emptyset}(T), \tilde{\Delta}_1^{\pi_I, \xi_\emptyset}(T)\} = \tilde{\Delta}^{\pi_I}(T) = \sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)]$, where the last equality follows from (L.5). Based on this, the rest of the proof is analyzes the term $\sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)]$. If ρ is fast vanishing (and not necessarily regular), then the market state Z_t is recurrent; see Proposition K.1. In particular, the state $\{0\}$ is recurrent. As a result, $\lim_{T\to\infty} \sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] \ge \lim_{T\to\infty} \sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\mathbb{I}\{Z_t = 0\}]\tilde{l}(0) = \infty$. If ρ is slowly vanishing and regular, then $\sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)]$ diverges in T at a rate of $\sum_{t=1}^T \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] = O\big(\sum_{t=1}^T \rho(t)\big)$; see Proposition K.2. ■

Appendix L: Convergence and Regret Analysis for IP in Propositions K.1 and K.2

In this section, we provide the details for the proofs of Proposition K.1 and K.2.

L.1. Preliminaries for the Convergence Analysis

We now introduce some auxiliary lemmas, the proofs of which are deferred to Appendix L.4.

LEMMA L.1. *Let $\{x_n, n \in \mathbb{Z}_+\}$ be a sequence satisfying $x_n \in (0, \frac{1}{2})$ for all $n \in \mathbb{Z}_+$, and let $y_n := \prod_{m=1}^{n} \frac{\frac{1}{2}-x_m}{\frac{1}{2}+x_m}$ for all $n \in \mathbb{Z}_+$. Then, we have the following:*

- *If $\liminf_{n\to\infty}\{nx_n\} > \frac{1}{4}$, then $\sum_n y_n$ converges.*
- *If $\limsup_{n\to\infty}\{nx_n\} < \frac{1}{4}$, then $\sum_n y_n$ diverges.*

The following definition provides a general condition for a residual probability sequence to be either fast vanishing or slowly vanishing.

DEFINITION 5. The residual probability sequence $\{\rho(z), z \in \mathbb{Z}_+\}$ is *strictly upper bounded* if $\sup\{\rho(z) : z \in \mathbb{Z}_+\} < \frac{1}{2} - \alpha$.

Definition 5 is equivalent to saying that, as Z_t diverges to infinity (i.e., there is very strong evidence in support of one hypothesis), the spread line does not converge to the median under the *opposite* hypothesis. In particular, if ρ is either fast vanishing or slowly vanishing, then $\lim_{z\to\infty}\{\rho(z)\} = 0$, and hence, ρ is strictly upper bounded.

The next result uses the machinery in Lemma 4 to characterize the quantity $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \mathbb{I}\{Z_t \le M\}$ for sufficiently large M, which is interpreted as the expected staying time of Z_t away from ∞. It is a generalization of Lemma G.3 under the informed bettor's generic $\bar{Z}$-threshold strategy and the market maker's generic residual probability sequence $\{\rho(z), z \in \mathbb{Z}_+\}$.

LEMMA L.2. *Suppose that the residual probability sequence $\{\rho(z) \in (0, \frac{1}{2} - \alpha), z \in \mathbb{Z}_+\}$ is strictly upper bounded. For all $\bar{Z} \in \mathbb{Z} \cup \{-\infty\}$ and $M \in \mathbb{Z}$ satisfying $M > \bar{Z} - 2$, there exists an increasing function $\tilde{u} : \{z \in \mathbb{Z} : z > \bar{Z} - 2\} \to \mathbb{R}$ such that*

$$\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \mathbb{I}\{Z_t \le M\} = \tilde{\mathbb{E}}_1^z \tilde{u}(Z_{T+1}) - \tilde{u}(z) \text{ for all } z \in \mathbb{Z} \text{ satisfying } z > \bar{Z} - 2 \text{ and } T \in \mathbb{Z}_+. \tag{L.1}$$

The closed-form expression for $\tilde{u}(\cdot)$ is as follows:

$$\tilde{u}(z) = \begin{cases} \left(1 + \sum_{n=z+1}^{M} \prod_{m=n}^{M} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right)\tilde{\beta} + \sum_{n=z+1}^{M}\sum_{k=n}^{M} \frac{1}{\frac{1}{2}+\rho(k)} \prod_{m=n}^{k} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} & \text{if } \bar{Z} - 2 < z \le M, \\ 0 & \text{if } z = M+1, \\ \left(1 + \sum_{n=M+2}^{z-1} \prod_{m=M+2}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\beta & \text{if } z \ge M+2, \end{cases} \tag{L.2}$$

where $\beta > 0$ and $\tilde{\beta} < 0$ are finite constants given by

$$\begin{cases} \tilde{\beta} = -\prod_{m=\bar{Z}}^{M} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} - \sum_{k=\bar{Z}}^{M} \frac{1}{\frac{1}{2}-\rho(k)} \prod_{m=k}^{M} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}, \\ \beta = -\frac{\frac{1}{2}-\rho(M+1)}{\frac{1}{2}+\rho(M+1)}\tilde{\beta}. \end{cases} \tag{L.3}$$

L.2. Preliminaries for the Regret Analysis

As mentioned earlier, we generalize the function $l(\cdot)$ to allow for a generic threshold $\bar{Z}$ as follows:

$$\tilde{l}(z) = (2-c)\rho(z)\mathbb{I}\{z < \bar{Z}\} + (4-2c)\rho^2(z)\mathbb{I}\{z \geq \bar{Z}\}. \tag{L.4}$$

The difference between $l(\cdot)$ and $\tilde{l}(\cdot)$ is that the former is defined for the informed bettor's optimal threshold $\bar{z}$ while the latter is defined for a generic threshold $\bar{Z}$. We introduce this new notation so that a generalized version of (4.12) holds, i.e.,

$$\tilde{\Delta}^{\pi_I}(T) = \sum_{t=1}^{T} \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)]. \tag{L.5}$$

Next, we introduce some auxiliary lemmas, whose proofs are in Appendix L.5.

LEMMA L.3. *Let $\{x_n, n \in \mathbb{Z}_+\}$ be a sequence satisfying $x_n \in (0, \frac{1}{2})$ for all $n \in \mathbb{Z}_+$, and let $a_n := \sum_{k=1}^{n} \frac{x_k^2}{\frac{1}{2}-x_k} \prod_{m=k}^{n} \frac{\frac{1}{2}-x_m}{\frac{1}{2}+x_m}$ for all $n \in \mathbb{Z}_+$. Then, $\sum_n a_n$ diverges. If in addition, $\liminf_{n\to\infty}\{nx_n\} > \frac{1}{4}$, $\lim_{n\to\infty}\left\{\left(\frac{x_n}{x_{n+1}} - 1\right)n\right\} \in [0,\infty]$ exists, and $\lim_{n\to\infty}\{x_n\} \in [0, \frac{1}{2}]$ exists, then $a_n = O(x_n)$ as $n \to \infty$.*

The result below uses the machinery in Lemma 4 to characterize the quantity $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \tilde{l}(Z_t)$ when $\bar{Z} > -\infty$, which is closely related to the market maker's regret, $\tilde{\Delta}^{\pi_I}(T)$.

LEMMA L.4. *Let the residual probability sequence $\{\rho(z), z \in \mathbb{Z}_+\}$ be slowly vanishing and regular, and $\bar{Z} > -\infty$. Then, there exists a function $v : \{\bar{Z}-1, \bar{Z}, \dots\} \to \mathbb{R}$ that satisfies the following:*

1. *For all $z \in \{\bar{Z}-1, \bar{Z}, \dots\}$ and $T \in \mathbb{Z}_+$,*

$$\frac{1}{4-2c} \sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \tilde{l}(Z_t) = \tilde{\mathbb{E}}_1^z v(Z_{T+1}) - v(z). \tag{L.6}$$

2. *$v(z)$ is increasing in z.*
3. *$v(z) \uparrow \infty$ as $z \uparrow \infty$ with a rate satisfying $v(z) = O\left(\sum_{k=1}^{z} \rho(k)\right)$.*

The closed-form expression of $v(\cdot)$ is given by

$$v(z) = \begin{cases} 0 & \text{if } z = \bar{Z}-1, \\ \left(1 + \sum_{n=\bar{Z}}^{z-1} \prod_{m=\bar{Z}}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right) \frac{\rho(\bar{Z}-1)}{2} + \sum_{n=\bar{Z}}^{z-1} \sum_{k=\bar{Z}}^{n} \frac{\rho^2(k)}{\frac{1}{2}-\rho(k)} \prod_{m=k}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} & \text{if } z \geq \bar{Z}. \end{cases} \tag{L.7}$$

The following result uses the construction in Lemma 4 to characterize the quantity $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \tilde{l}(Z_t)$ when $\bar{Z} = -\infty$, in order to analyze $\tilde{\Delta}^{\pi_I}(T)$.

LEMMA L.5. *Let the residual probability sequence $\{\rho(z), z \in \mathbb{Z}_+\}$ be slowly vanishing and regular, and $\bar{Z} = -\infty$. Then, there exists function $\tilde{v} : \mathbb{Z} \to \mathbb{R}_+$ that satisfies the following:*

1. *For all $z \in \mathbb{Z}$ and $T \in \mathbb{Z}_+$:*

$$\frac{1}{4-2c} \sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \tilde{l}(Z_t) = \tilde{\mathbb{E}}_1^z \tilde{v}(Z_{T+1}) - \tilde{v}(z). \tag{L.8}$$

2. *$\tilde{v}(z)$ is increasing in z.*
3. *$\tilde{v}(z) \uparrow \infty$ as $z \uparrow \infty$ at a rate satisfying $\tilde{v}(z) = O\left(\sum_{k=1}^{z} \rho(k)\right)$.*

The closed-form expression of $\tilde{v}(\cdot)$ *is given by*

$$\tilde{v}(z) = \begin{cases} \left(1 + \sum_{n=z+1}^{-1} \prod_{m=n}^{-1} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right) \tilde{B} + \sum_{n=z+1}^{-1} \sum_{k=n}^{-1} \frac{\rho^2(k)}{\frac{1}{2}+\rho(k)} \prod_{m=n}^{k} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} & \text{if } z \leq -1, \\ 0 & \text{if } z = 0, \\ \left(1 + \sum_{n=1}^{z-1} \prod_{m=1}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right) B + \sum_{n=1}^{z-1} \sum_{k=1}^{n} \frac{\rho^2(k)}{\frac{1}{2}-\rho(k)} \prod_{m=k}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} & \text{if } z \geq 1. \end{cases} \tag{L.9}$$

where B *and* $\tilde{B}$ *are finite constants.*

L.3. Key Steps of the Convergence and Regret Analysis

Proof of Proposition K.1. Fix a residual probability sequence ρ that is fast vanishing. First, observe that if $\bar{Z} > -\infty$, then Z_t is a standard birth and death chain with a reflecting boundary. Then, $\{\bar{Z}-1\}$ is recurrent (Durrett 2019, Exercise 5.3.4). Because $\{\bar{Z}-1, \bar{Z}, \ldots\}$ is an irreducible class, any state in that class is recurrent. Second, if $\bar{Z} = -\infty$, then we may treat $\{Z_t\}$ as a concatenation of two sub-processes: the first one is $\{Z_t\}$ restricted to $\mathbb{Z}_+ \cup \{0\}$, and the second is $\{Z_t\}$ restricted to $\mathbb{Z}_- \cup \{0\}$. We show that either sub-process visits the state $\{0\}$ infinitely often. That is, $\{0\}$ is recurrent. The first sub-process is a standard birth and death chain with a reflecting boundary $\{0\}$. Let $\tau_0 := \inf\{t > 1 : Z_t = 0\}$ be the hitting time of of state $\{0\}$. Then, $\tilde{\mathbb{P}}_1^0(\tau_0 < \infty \,|\, Z_2 > 0) = 1$ (Durrett 2019). The second sub-process is a "reflected" version of a standard birth and death chain. Since the residual probability sequence $\{\rho(z), z \in \mathbb{Z}_+\}$ is fast vanishing, it is strictly upper bounded (see Definition 5). Now consider the extended residual probability sequence $\{\rho(z), z \in \mathbb{Z}\}$ in the sense of (E.1). By Lemma E.3, $\{\rho(z), z \in \mathbb{Z}_-\}$ is bounded away from 0. Choose $\varepsilon > 0$ such that $\rho(z) \geq \varepsilon > 0$ for all $z \leq \mathbb{Z}_-$. Then, for all $z \in \mathbb{Z}_-$, $\frac{1}{2} + \rho(z) \geq \frac{1}{2} + \varepsilon > \frac{1}{2}$. Hence, $\tilde{\mathbb{P}}_1^0(\tau_0 < \infty \,|\, Z_2 < 0) = 1$. Combining our findings, we have $\tilde{\mathbb{P}}_1^0(\tau_0 < \infty) = 1$; i.e., $\{0\}$ is recurrent. Because Z_t is irreducible when $\bar{Z} = -\infty$, all states in $\{z \in \mathbb{Z} : z > \bar{Z} - 2\}$ are recurrent. ■

Proof of Statements (PK.2:1) and (PK.2:2) in Proposition K.2. Fix a residual probability sequence ρ that is slowly vanishing. Without loss of generality, let us assume that the initial value z is greater than $\bar{Z} - 2$, and Z_t is a Markov chain restrained in $\{z \in \mathbb{Z} : z > \bar{Z} - 2\}$; otherwise, Z_t increases in a deterministic manner under $\tilde{\mathbb{P}}_1^z$ until Z_t hits the region $\{z \in \mathbb{Z} : z > \bar{Z} - 2\}$.

To prove Statement (PK.2:1), we choose $M > \bar{Z} - 2$ and verify that $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \mathbb{I}\{Z_t \leq M\}$ is bounded in T. Because ρ is slowly vanishing, $\liminf_{z\to\infty}\{z\rho(z)\} > \frac{1}{4}$ and $\lim_{z\to\infty}\{\rho(z)\} = 0$ (see Definition 2). Moreover, ρ is also strictly upper bounded (see Definition 5). Invoking Lemma L.2, we have $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^z \mathbb{I}\{Z_t \leq M\} = \tilde{\mathbb{E}}_1^z \tilde{u}(Z_{T+1}) - \tilde{u}(z)$, where $\tilde{u}(\cdot)$ is as in (L.2). Thus, it suffices to show that $\tilde{u}(\cdot)$ is bounded from above. Recall from Lemma L.2 that $\tilde{u}(\cdot)$ is increasing, which implies that it is sufficient to prove that $\lim_{z\to\infty}\{\tilde{u}(z)\} < \infty$. By (L.2), it is equivalent to $\sum_n \prod_{m=1}^{n} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}$ being convergent. This follows from Lemma L.1 as ρ is slowly vanishing.

Statement (PK.2:2) follows from Statement (PK.2:1). By the Borel-Cantelli lemma, we deduce that $\tilde{\mathbb{P}}_1^0\{Z_t \leq M \text{ infinitely often}\} = 0$. As a result, $Z_t \to \infty$ almost surely under $\tilde{\mathbb{P}}_1^0$. Since ρ is vanishing, this implies that $\rho(Z_t) \to 0$ and $\mathfrak{d}_t = |s_t - m_1| = |F_1^{-1}(\frac{1}{2} - \rho(Z_t)) - F_1^{-1}(\frac{1}{2})| \to 0$ almost surely under $\tilde{\mathbb{P}}_1^0$. ■

Proof of Statements (PK.2:3) and (PK.2:4) in Proposition K.2. Fix a residual probability sequence ρ that is both slowly vanishing and regular. For this proof, it suffices to verify Statement (PK.2:3) because Statement (PK.2:4) follows from Statement (PK.2:3) and (L.5). We complete the proof in two steps.

Step 1. We claim that $\sum_{t=1}^{T} \tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] = O\big(\sum_{t=1}^{T} \rho(t)\big)$. Observe that, for all $T \geq \bar{Z} \vee 1$,

$$\frac{1}{4-2c}\sum_{t=1}^{T}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] \overset{(a)}{=} \begin{cases} \tilde{\mathbb{E}}_1^0 \tilde{v}(Z_{T+1}) - \tilde{v}(0) & \text{if } \bar{Z} = -\infty, \\ \tilde{\mathbb{E}}_1^0 v(Z_{T+1}) - v(0) & \text{if } -\infty < \bar{Z} \leq 1, \\ \tilde{\mathbb{E}}_1^{\bar{Z}-1} v(Z_{T+2-\bar{Z}}) - v(\bar{Z}-1) + \sum_{z=0}^{\bar{Z}-2} \frac{\rho(z)}{2} & \text{if } 2 \leq \bar{Z}, \end{cases}$$
$$\overset{(b)}{\leq} \begin{cases} \tilde{v}(T) - \tilde{v}(0) & \text{if } \bar{Z} = -\infty, \\ v(T) - v(0) & \text{if } -\infty < \bar{Z} \leq 1, \\ v(Z_{T+1-\bar{Z}}) - v(\bar{Z}-1) + \sum_{z=0}^{\bar{Z}-2} \frac{\rho(z)}{2} & \text{if } 2 \leq \bar{Z}, \end{cases}$$
$$\overset{(c)}{=} O\left(\sum_{t=1}^{T}\rho(t)\right).$$

In the derivations above, part (a) follows from invoking (L.6) and (L.8) for the cases where $\bar{Z} = -\infty$ and $\bar{Z} > -\infty$, respectively. Moreover, the third piece of (a) holds because when $\bar{Z} \geq 2$, Z_t increments by one with certainty (i.e., $Z_t = t-1$) until Z_t hits $\bar{Z}-1$; consequently,

$$\begin{aligned}
\sum_{t=1}^{T}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] &= \sum_{t=1}^{\bar{Z}-1}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] + \sum_{t=\bar{Z}}^{T+1}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] \\
&= \sum_{t=1}^{\bar{Z}-1}\tilde{l}(t-1) + \sum_{t=\bar{Z}}^{T}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] \\
&= \sum_{t=1}^{\bar{Z}-1}\tilde{l}(t-1) + \sum_{t=1}^{T-\bar{Z}+1}\tilde{\mathbb{E}}_1^{\bar{Z}-1}[\tilde{l}(Z_t)] && \text{[by the Markov property of } Z_t] \\
&\overset{(L.4)}{=} \sum_{z=0}^{\bar{Z}-2}(2-c)\rho(z) + \sum_{t=1}^{T-\bar{Z}+1}\tilde{\mathbb{E}}_1^{\bar{Z}-1}[\tilde{l}(Z_t)] \\
&\overset{(L.6)}{=} \sum_{z=0}^{\bar{Z}-2}(2-c)\rho(z) + (4-2c)\big[\tilde{\mathbb{E}}_1^{\bar{Z}-1}v(Z_{T+2-\bar{Z}}) - v(\bar{Z}-1)\big].
\end{aligned}$$

Part (b) holds because (i) for all $z \in \mathbb{Z}$ and $t \in \mathbb{Z}_+$, $Z_{t+1} \leq t+z$ almost surely under $\tilde{\mathbb{P}}_1^z$; and (ii) by Lemmas L.4 and L.5, $v(\cdot)$ and $\tilde{v}(\cdot)$ are increasing functions. Part (c) holds since $v(T)$ and $\tilde{v}(T)$ are $O\big(\sum_{t=1}^{T}\rho(t)\big)$.

Step 2. We claim that $\sum_{t=1}^{T}\tilde{\mathbb{E}}_1^0[\tilde{l}(Z_t)] \to \infty$ as $T \to \infty$. To prove this claim, it suffices to show that $\tilde{\mathbb{E}}_1^0[v(Z_T)] \to \infty$ and $\tilde{\mathbb{E}}_1^0[\tilde{v}(Z_T)] \to \infty$ as $T \to \infty$; see part (a) in Step 1. Because ρ is slowly vanishing, $Z_t \to \infty$ almost surely under $\tilde{\mathbb{P}}_1^0$; see Statement (PK.2:2). By Lemma L.4, $v(Z_T)$ is a nonnegative random variable such that $v(Z_T) \to \infty$ almost surely as $T \to \infty$ under $\tilde{\mathbb{P}}_1^0$. By Markov's inequality, $\tilde{\mathbb{E}}_1^0[v(Z_T)] \geq x\tilde{\mathbb{P}}_1^0(v(Z_T) \geq x)$ for all $x > 0$. Taking T to ∞ and then x to ∞, we deduce that $\tilde{\mathbb{E}}_1^0[v(Z_T)] \to \infty$ as $T \to \infty$. Repeating the same arguments for $\tilde{\mathbb{E}}_1^0[\tilde{v}(Z_T)]$ and invoking Lemma L.5, we obtain the desired result. ∎

L.4. Proofs of Auxiliary Lemmas for Convergence Analysis

Proof of Lemma L.1. Observe that $\frac{y_n}{y_{n+1}} = \frac{\frac{1}{2}+x_{n+1}}{\frac{1}{2}-x_{n+1}} = 1 + \frac{2x_{n+1}}{\frac{1}{2}-x_{n+1}}$. Thus, we have the following:

- If $\liminf_{n\to\infty}\{nx_n\} > \frac{1}{4}$, then $\liminf_{n\to\infty}\left\{\left(\frac{y_n}{y_{n+1}} - 1\right)n\right\} = \liminf_{n\to\infty}\left\{\frac{2x_{n+1}}{\frac{1}{2}-x_{n+1}}n\right\} \geq \liminf_{n\to\infty}\{4x_{n+1}n\} > 1$.
- If $\limsup_{n\to\infty}\{nx_n\} < \frac{1}{4}$, then $\limsup_{n\to\infty}\left\{\left(\frac{y_n}{y_{n+1}} - 1\right)n\right\} = \limsup_{n\to\infty}\left\{\frac{2x_{n+1}}{\frac{1}{2}-x_{n+1}}n\right\} = \limsup_{n\to\infty}\{4x_{n+1}n\} < 1$.

Invoking Raabe's test of convergence of series with positive terms (Bromwich 1908, p. 33), we complete the proof. ■

Proof of Lemma L.2. With the initial point $z\in\mathbb{Z}$ satisfying $z>\bar{Z}-2$, the Markov chain Z_t is contained in the region $\{z\in\mathbb{Z}: z>\bar{Z}-2\}$. In particular, if $\bar{Z}$ is finite, the state $\bar{Z}-1$ is a reflecting boundary. Invoking Lemma 4 with $S=\{z\in\mathbb{Z}: z>\bar{Z}-2\}$, it is sufficient to solve the difference equation $\tilde{\mathbb{E}}_1^z\tilde{u}(Z_2)-\tilde{u}(z) = \mathbb{I}\{z\leq M\}$ for all $z\in S$. The rest of the proof confirms that (L.2) presents such a solution with the desired monotonicity property. We complete the remainder of the proof in three steps.

Step 1. We verify that both β and $\tilde{\beta}$ are finite, and hence the function $\tilde{u}(\cdot)$ in (L.2) is finitely valued. Note that it suffices to check that $\tilde{\beta}$ is finite. As ρ is strictly upper bounded (see Definition 5), we deduce from Lemma E.3 that there exists $\varepsilon>0$ satisfying $\rho(z)\geq\varepsilon>0$ for all $z\leq M$. Thus,

$$|\tilde{\beta}| = \prod_{m=\bar{Z}}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} + \sum_{k=\bar{Z}}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} = \underbrace{\prod_{m=\bar{Z}}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}}_{\leq 1} + \sum_{k=\bar{Z}}^{M}\underbrace{\frac{1}{\frac{1}{2}+\rho(k)}}_{\leq\frac{1}{\frac{1}{2}+\varepsilon}}\underbrace{\prod_{m=k+1}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}}_{\leq\left(\frac{\frac{1}{2}-\varepsilon}{\frac{1}{2}+\varepsilon}\right)^{M-k}}$$

$$\stackrel{(a)}{\leq} 1 + \sum_{k=-\infty}^{M}\frac{1}{\frac{1}{2}+\varepsilon}\left(\frac{\frac{1}{2}-\varepsilon}{\frac{1}{2}+\varepsilon}\right)^{M-k} < \infty,$$

where (a) follows because $\rho(z)\geq\varepsilon$ for all $z\leq M$, which implies that $\prod_{m=k+1}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\leq\left(\frac{\frac{1}{2}-\varepsilon}{\frac{1}{2}+\varepsilon}\right)^{M-k}$ and $\frac{1}{\frac{1}{2}+\rho(k)}\leq\frac{1}{\frac{1}{2}+\varepsilon}$ for all $k\leq M$.

Step 2. We verify that $\tilde{u}(\cdot)$ is increasing. That is, $\tilde{u}(z-1)-\tilde{u}(z)<0$ for all $z\in\mathbb{Z}$ satisfying $z\geq\bar{Z}$. In particular, when $\bar{Z}>-\infty$, $\tilde{u}(\bar{Z}-1)-\tilde{u}(\bar{Z})=-1$. Since $\tilde{\beta}<0$ and $\beta>0$, we can directly verify from the expression in (L.2) that $\tilde{u}(M)=\tilde{\beta}<0=\tilde{u}(M+1)<\beta=\tilde{u}(M+2)\leq\tilde{u}(M+3)\leq\tilde{u}(M+4)\leq\cdots$ Therefore, it suffices to show that $\tilde{u}(z-1)-\tilde{u}(z)<0$ for all z satisfying $\bar{Z}-1<z\leq M$. Note that

$$\begin{aligned}
\tilde{u}(z-1)-\tilde{u}(z) &= \tilde{\beta}\prod_{m=z}^{M}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} + \sum_{k=z}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=z}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\\
&\stackrel{(L.3)}{=} \left(-\prod_{m=\bar{Z}}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} - \sum_{k=\bar{Z}}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\prod_{m=z}^{M}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} + \sum_{k=z}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=z}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\\
&\stackrel{\bar{Z}\leq z}{\leq} \left(-\prod_{m=\bar{Z}}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} - \sum_{k=z}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{M}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\prod_{m=z}^{M}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} + \sum_{k=z}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=z}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\\
&= -\prod_{m=\bar{Z}}^{z-1}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} - \sum_{k=z}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=z}^{k-1}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} + \sum_{k=z}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=z}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}
\end{aligned}$$

$$= -\prod_{m=\bar{Z}}^{z-1} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} < 0.$$

The inequality above is tight when $\bar{Z} > -\infty$ and $z = \bar{Z}$.That is, when $\bar{Z} > -\infty$, $\tilde{u}(\bar{Z}-1) - \tilde{u}(\bar{Z}) = -1$.

Step 3. As explained above, by Lemma 4, it suffices to show that for all $z > \mathbb{Z} - 2$ satisfying $z \geq \bar{Z}$, $\tilde{\mathbb{E}}_1^z \tilde{u}(Z_2) - \tilde{u}(z) = \mathbb{I}\{z \leq M\}$. Let us first look at the left-hand side of the equation:

$$\tilde{\mathbb{E}}_1^z \tilde{u}(Z_2) - \tilde{u}(z) = \begin{cases} \left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1) + \left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1) - \tilde{u}(z) & \text{if } z \geq \bar{Z}, \\ \tilde{u}(z+1) - \tilde{u}(z) & \text{if } z < \bar{Z}. \end{cases}$$

Let us then look at the right-hand side of the equation:

$$\mathbb{I}\{z \leq M\} = \begin{cases} 1 & \text{if } z \leq M, \\ 0 & \text{if } z > M. \end{cases}$$

Note that we only consider the values of M that exceed $\bar{Z} - 2$. We study four cases for the value of z: $z \geq M+2$, $z = M+1$, $\bar{Z} - 1 < z \leq M$, and $z = \bar{Z} - 1$ (the last case is valid only when $\bar{Z} > -\infty$).

1. When $z \geq M+2$, we show that $\left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1) + \left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1) - \tilde{u}(z) = 0$. By Lemma F.1, $\tilde{u}(\cdot)$ restricted to the domain $[M+1, \infty)$ solves the difference equation (F.1) with $\hat{z} = M+2$, $\delta = \beta$, $p(z) = \frac{1}{2} + \rho(z)$, and $x(z) = 0$ for all $z \geq M+1$. That is, $\tilde{u}(\cdot)$ satisfies the following:
$$\begin{cases} \tilde{u}(M+1) = 0,\ \tilde{u}(M+2) = \beta \\ \left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1) + \left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1) - \tilde{u}(z) = 0 \ \text{ for all } z \geq M+2. \end{cases}$$
Thus, for all $z \geq M+2$, $\left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1) + \left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1) - \tilde{u}(z) = 0$.
2. When $z = M+1$, the definitions of β and $\tilde{\beta}$ imply that $\left[\frac{1}{2}+\rho(M+1)\right]\tilde{u}(M+2) + \left[\frac{1}{2}-\rho(M+1)\right]\tilde{u}(M) - \tilde{u}(M+1) = \left[\frac{1}{2}+\rho(M+1)\right]\beta + \left[\frac{1}{2}-\rho(M+1)\right]\tilde{\beta} = 0$.
3. When $\bar{Z} - 1 < z \leq M$, we show that $\left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1) + \left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1) - \tilde{u}(z) = 1$. To that end, we introduce an auxiliary sequence $\{y(z), z \in \mathbb{N}\}$ such that
$$y(z) = \begin{cases} 0 & \text{if } z = 0, \\ \left(1 + \sum_{n=1}^{z-1}\prod_{m=1}^{n} \frac{\frac{1}{2}+\rho(M+1-m)}{\frac{1}{2}-\rho(M+1-m)}\right)\tilde{\beta} + \sum_{n=1}^{z-1}\sum_{k=1}^{n} \frac{1}{\frac{1}{2}+\rho(M+1-k)} \prod_{m=n}^{k} \frac{\frac{1}{2}+\rho(M+1-m)}{\frac{1}{2}-\rho(M+1-m)} & \text{if } z \geq 1. \end{cases}$$
The above construction of $y(\cdot)$ has two desirable properties. First, $y(\cdot)$ is a "reflected" version of $\tilde{u}(\cdot)$. Specifically, for all $z \in \mathbb{Z}$ satisfying $\bar{Z} - 1 < z \leq M$,
$$\begin{aligned} y(M+1-z) &= \left(1 + \sum_{n=1}^{M-z}\prod_{m=1}^{n} \frac{\frac{1}{2}+\rho(M+1-m)}{\frac{1}{2}-\rho(M+1-m)}\right)\tilde{\beta} + \sum_{n=1}^{M-z}\sum_{k=1}^{n} \frac{1}{\frac{1}{2}+\rho(M+1-k)} \prod_{m=n}^{k} \frac{\frac{1}{2}+\rho(M+1-m)}{\frac{1}{2}-\rho(M+1-m)} \\ &\stackrel{(a)}{=} \left(1 + \sum_{\check{n}=z+1}^{M}\prod_{\check{m}=\check{n}}^{M} \frac{\frac{1}{2}+\rho(\check{m})}{\frac{1}{2}-\rho(\check{m})}\right)\tilde{\beta} + \sum_{\check{n}=z+1}^{M}\sum_{\check{k}=\check{n}}^{M} \frac{1}{\frac{1}{2}+\rho(\check{k})} \prod_{\check{m}=\check{n}}^{\check{k}} \frac{\frac{1}{2}+\rho(\check{m})}{\frac{1}{2}-\rho(\check{m})} = \tilde{u}(z), \end{aligned}$$
where (a) holds by the following change of variables: $\check{n} = M+1-n$, $\check{m} = M+1-m$, and $\check{k} = M+1-k$. Second, we deduce from Lemma F.1 that $y(\cdot)$ solves the difference equation (F.1) with

$\hat{z}=1, \delta=\tilde{\beta}, p(z)=\frac{1}{2}-\rho(M+1-z)$. That is,

$$\begin{cases} y(0)=0,\, y(1)=\tilde{\beta}, \\ \left[\frac{1}{2}-\rho(M+1-z)\right] y(z+1)+\left[\frac{1}{2}+\rho(M+1-z)\right] y(z-1)-y(z)=1 \text{ for all } z\geq 1. \end{cases}$$

Substituting $\tilde{u}(z)=y(M+1-z)$, we obtain the following for $\tilde{u}(\cdot)$:

$$\begin{cases} \tilde{u}(M+1)=0,\, \tilde{u}(M)=\tilde{\beta}, \\ \left[\frac{1}{2}-\rho(z)\right] \tilde{u}(z-1)+\left[\frac{1}{2}+\rho(z)\right] \tilde{u}(z+1)-\tilde{u}(z)=1 \text{ for all } \bar{Z}-1<z\leq M. \end{cases}$$

Thus, for all z satisfying $\bar{Z}-1<z\leq M$, $\left[\frac{1}{2}+\rho(z)\right]\tilde{u}(z+1)+\left[\frac{1}{2}-\rho(z)\right]\tilde{u}(z-1)-\tilde{u}(z)=1$.

4. When $z=\bar{Z}-1$, we show that $\tilde{u}(\bar{Z})-\tilde{u}(\bar{Z}-1)=1$ (this case is valid only if $\bar{Z}>-\infty$). Note that

$$\begin{aligned} \tilde{u}(\bar{Z}-1)-\tilde{u}(\bar{Z}) &= \tilde{\beta}\prod_{m=\bar{Z}}^{M}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}+\sum_{k=\bar{Z}}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=\bar{Z}}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} \\ &\overset{(L.3)}{=} -1-\sum_{k=\bar{Z}}^{M}\frac{1}{\frac{1}{2}-\rho(k)}\prod_{m=\bar{Z}}^{k-1}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}+\sum_{k=\bar{Z}}^{M}\frac{1}{\frac{1}{2}+\rho(k)}\prod_{m=\bar{Z}}^{k}\frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}=-1. \end{aligned}$$

Combining our findings in the four cases above, we conclude that $\tilde{\mathbb{E}}_1^z\tilde{u}(Z_2)-\tilde{u}(z)=\mathbb{I}\{z\leq M\}$ for all $z\in\mathbb{Z}$ satisfying $z>\bar{Z}-2$. ■

L.5. Proofs of Auxiliary Lemmas for Regret Analysis

Proof of Lemma L.3. We first derive a recursive relation for $\{a_n\}$ as follows:

$$a_{n+1}=\sum_{k=1}^{n+1}\frac{x_k^2}{\frac{1}{2}-x_k}\prod_{m=k}^{n+1}\frac{\frac{1}{2}-x_m}{\frac{1}{2}+x_m}=\frac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}\underbrace{\sum_{k=1}^{n}\frac{x_k^2}{\frac{1}{2}-x_k}\prod_{m=k}^{n}\frac{\frac{1}{2}-x_m}{\frac{1}{2}+x_m}}_{=a_n}+\frac{x_{n+1}^2}{\frac{1}{2}+x_{n+1}}=\frac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}a_n+\frac{x_{n+1}^2}{\frac{1}{2}+x_{n+1}}. \tag{L.10}$$

We break the rest of the proof into two steps.

Step 1. We claim that $\sum_n a_n$ diverges. To prove this claim, we deduce from (L.10) the following for sufficiently large n:

$$\begin{aligned} \left(\frac{a_n}{a_{n+1}}-1\right)n &= \left(\frac{\frac{1}{2}+x_{n+1}}{\frac{1}{2}-x_{n+1}}-\frac{x_{n+1}^2}{\frac{1}{2}-x_{n+1}}\frac{1}{a_{n+1}}-1\right)n \\ &= \left(\frac{2x_{n+1}}{\frac{1}{2}-x_{n+1}}-\frac{x_{n+1}^2}{\frac{1}{2}-x_{n+1}}\frac{1}{a_{n+1}}\right)n \\ &= \left[-2+\frac{1}{a_{n+1}}-\left(\frac{\frac{1}{2}-x_{n+1}}{a_{n+1}}+\frac{\frac{1}{4a_{n+1}}-1}{\frac{1}{2}-x_{n+1}}\right)\right]n && \text{[rearranging terms]} \\ &\leq \left[-2+\frac{1}{a_{n+1}}-2\sqrt{\frac{1}{a_{n+1}}\left(\frac{1}{4a_{n+1}}-1\right)}\right]n && [a+b\geq 2\sqrt{ab}\ \forall a,b>0\ ] \\ &= \frac{4n}{\frac{1}{a_{n+1}}-2+2\sqrt{\frac{1}{4a_{n+1}^2}-\frac{1}{a_{n+1}}}}\leq\frac{4n}{\frac{1}{a_{n+1}}-2}=\frac{4a_{n+1}n}{1-2a_{n+1}}. \end{aligned}$$

Suppose towards a contradiction that $\sum_n a_n$ converges. This implies that $\frac{4a_{n+1}n}{1-2a_{n+1}}\to 0$ as $n\to\infty$ because $\sum_n\frac{1}{n}$ is a divergent series. However, due to the derivations above, this means $\left(\frac{a_n}{a_{n+1}}-1\right)n\to 0$ as $n\to\infty$. Applying Raabe's test of convergence of series with positive terms (Bromwich 1908, p. 33) to $\{a_n\}$, we note that $\sum_n a_n$ diverges, leading to a contradiction as desired.

Step 2. We claim that if $\lim_{n\to\infty}\{x_n\} > 0$, then $a_n = O(x_n)$. Observe that when $\lim_{n\to\infty}\{x_n\} > 0$, there exists $\varepsilon_0 \in (0, \frac{1}{2})$ and $N_0 \in \mathbb{N}$ such that $x_n \geq \varepsilon_0 > 0$ for all $n \geq N_0$. Thus, to prove our claim in this step, it suffices to show that $\{a_n\}$ is bounded from above. For $n \geq N_0$,

$$\begin{aligned} a_{n+1} = \tfrac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}} a_n + \tfrac{x_{n+1}^2}{\frac{1}{2}+x_{n+1}} &\leq \tfrac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0} a_n + \tfrac{1}{2} \\ &\leq \tfrac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0}\left(\tfrac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0} a_{n-1} + \tfrac{1}{2}\right) + \tfrac{1}{2} \\ &\vdots \\ &\leq \limsup_{k\to\infty}\left\{\left(\tfrac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0}\right)^k a_{N_0} + \tfrac{1}{2}\left(1 + \cdots + \left(\tfrac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0}\right)^{k-1}\right)\right\}, \end{aligned}$$

which is a finite constant independent of n.

Step 3. We claim that if (i) $\liminf_{n\to\infty}\{nx_n\} > \frac{1}{4}$, (ii) $\lim_{n\to\infty}\left\{\left(\frac{x_n}{x_{n+1}} - 1\right)n\right\}$ exists, and (iii) $\lim_{n\to\infty}\{x_n\} = 0$, then $a_n = O(x_n)$. Let us define $c_n := \frac{a_n}{x_n}$. To show that $a_n = O(x_n)$, it suffices to prove that $\{c_n\}$ is bounded from above. Invoking the recursive relation between a_{n+1} and a_n in (L.10), we obtain a recursive relation between c_{n+1} and c_n. That is, $c_{n+1} = \frac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}\frac{x_n}{x_{n+1}}c_n + \frac{x_{n+1}}{\frac{1}{2}+x_{n+1}}$. We seek finite constants $M_1, M_0, N_1 \in \mathbb{N}$ such that for all $n \geq N_1$, (i) $c_n \geq M_0$ implies that $c_{n+1} \leq c_n$, and (ii) $c_n \leq M_0$ implies that $c_{n+1} \leq M_1$. If such a tuple (M_0, M_1, N_1) exists, the sequence $\{c_n\}$ is bounded by the value $\max\{c_1, \ldots, c_{N_1+1}, M_0, M_1\}$. To find (M_0, M_1, N_1), let us evaluate the following quantity for every $M > 0$:

$$\begin{aligned} (*) &= \tfrac{1}{x_n}\left(\tfrac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}\tfrac{x_n}{x_{n+1}}M + \tfrac{x_{n+1}}{\frac{1}{2}+x_{n+1}} - M\right) \\ &= n\left(\tfrac{x_n}{x_{n+1}} - 1\right)\tfrac{M}{nx_n} - \tfrac{4}{1+2x_{n+1}}M + \tfrac{x_{n+1}}{x_n}\tfrac{1}{\frac{1}{2}+x_{n+1}} \\ &< n\left(\tfrac{x_n}{x_{n+1}} - 1\right)\tfrac{M}{nx_n} - \tfrac{4}{1+2x_{n+1}}M + 2\left(\tfrac{x_{n+1}}{x_n}\right). \end{aligned}$$

We claim that $(*)$ is negative for sufficiently large n and M. To show this claim, recall that $\sum_n x_n$ diverges. Because $\lim_{n\to\infty}\left\{\left(\frac{x_n}{x_{n+1}} - 1\right)n\right\}$ exists, we apply Raabe's test to the sequence $\{x_n\}$ and conclude that $\lim_{n\to\infty}\left\{\left(\frac{x_n}{x_{n+1}} - 1\right)n\right\} \leq 1$. That is, there exists a constant $A \leq 1$ such that $\frac{x_n}{x_{n+1}} = 1 + \frac{A}{n} + o\left(\frac{1}{n}\right)$. This implies that $\lim_{n\to\infty}\left\{\frac{x_n}{x_{n+1}}\right\} = \lim_{n\to\infty}\left\{\frac{x_{n+1}}{x_n}\right\} = 1$. Moreover, $\lim_{n\to\infty}\{x_n\} = 0$. Combining all the pieces, we know that there exist $\varepsilon_1 > 0$ and $N_1 \in \mathbb{N}$ such that for all $n \geq N_1$, (i) $n\left(\frac{x_n}{x_{n+1}} - 1\right) \leq 1 + \frac{2\varepsilon_1}{1+\varepsilon_1}$, (ii) $nx_n \geq \frac{1}{4} + \varepsilon_1$, (iii) $x_{n+1} \leq \frac{\varepsilon_1}{2(1+4\varepsilon_1)}$, and (iv) $\frac{1}{2} \leq \frac{x_{n+1}}{x_n} \leq 2$. Thus, for all $n \geq N_1$,

$$\begin{aligned} (*) &< \left(1 + \tfrac{2\varepsilon_1}{1+\varepsilon_1}\right)\tfrac{M}{\frac{1}{4}+\varepsilon_1} - \tfrac{4}{1+\frac{\varepsilon_1}{1+4\varepsilon_1}}M + 4 \\ &= 4M\left[\left(1 + \tfrac{2\varepsilon_1}{1+\varepsilon_1}\right)\tfrac{1}{1+4\varepsilon_1} - 1 + \tfrac{\frac{\varepsilon_1}{1+4\varepsilon_1}}{1+\frac{\varepsilon_1}{1+4\varepsilon_1}}\right] + 4 \\ &< 4M\left[1 - \tfrac{2\varepsilon_1}{1+4\varepsilon_1} - 1 + \tfrac{\varepsilon_1}{1+4\varepsilon_1}\right] + 4 = -\tfrac{4\varepsilon_1}{1+\varepsilon_1}M + 4. \end{aligned} \tag{L.11}$$

Choose M_0 such that $-\frac{4\varepsilon_1}{1+\varepsilon_1}M + 4 < 0$ for all $M \geq M_0$. This is possible because $\frac{4\varepsilon_1}{1+\varepsilon_1} > 0$. By construction, for all n satisfying $n \geq N_1$ and $c_n \geq M_0$,

$$c_{n+1} - c_n = \tfrac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}\tfrac{x_n}{x_{n+1}}c_n + \tfrac{x_{n+1}}{\frac{1}{2}+x_{n+1}} - c_n$$

$$= x_n \tfrac{1}{x_n}\left(\tfrac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}\tfrac{x_n}{x_{n+1}}c_n + \tfrac{x_{n+1}}{\frac{1}{2}+x_{n+1}} - c_n\right)$$
$$< x_n\left(-\tfrac{4\varepsilon_1}{1+\varepsilon_1}c_n + 4\right) < 0. \qquad \text{[choosing } M = c_n \geq M_0 \text{ in (L.11)]}$$

Lastly, it is useful to observe that for all n such that $n \geq N_1$ and $c_n \leq M_0$,

$$c_{n+1} = \underbrace{\tfrac{\frac{1}{2}-x_{n+1}}{\frac{1}{2}+x_{n+1}}}_{\leq 1}\underbrace{\tfrac{x_n}{x_{n+1}}}_{\leq 2}c_n + \underbrace{\tfrac{x_{n+1}}{\frac{1}{2}+x_{n+1}}}_{\leq 1} \leq 2M_0 + 1.$$

Thus, it suffices to choose $M_1 = 2M_0 + 1$. We have found (N_1, M_0, M_1) as desired. ■

Proof of Lemma L.4. With the initial point $z \in \{\bar{Z}-1, \bar{Z}, \ldots\}$, the Markov chain Z_t is contained in the region $\{\bar{Z}-1, \bar{Z}, \ldots\}$. In particular, if $\bar{Z}$ is finite, the state $\bar{Z}-1$ is a reflecting boundary. Invoking Lemma 4 with $S = \{\bar{Z}-1, \bar{Z}, \ldots\}$, it suffices to solve the difference equation $\tilde{\mathbb{E}}_1^z v(Z_2) - v(z) = \frac{\tilde{l}(z)}{4-2c}$ for all $z \in S$. The remainder of the proof verifies that (L.7) provides such a solution with the desired monotonicity and growth properties. We complete the rest of the proof in three steps.

Step 1. We claim that when $v(\cdot)$ is as in (L.7), $\tilde{\mathbb{E}}_1^z v(Z_2) - v(z) = \frac{\tilde{l}(z)}{4-2c}$ for all $z \in \{\bar{Z}-1, \bar{Z}, \ldots\}$. Observe that $v(\cdot)$ solves the difference equation (F.1) with $\hat{z} = \bar{Z}$, $\delta = \frac{\rho(\bar{Z}-1)}{2}$, $x(z) = \rho^2(z)$, and $p(z) = \frac{1}{2} + \rho(z)$. That is,

$$\begin{cases} v(\bar{Z}-1) = 0,\ v(\bar{Z}) = \frac{\rho(\bar{Z}-1)}{2}, \\ \left[\frac{1}{2}+\rho(z)\right]v(z+1) + \left[\frac{1}{2}-\rho(z)\right]v(z-1) - v(z) = \rho^2(z) \ \text{ for all } z \geq \bar{Z}. \end{cases}$$

As a result, for all $z \in \{\bar{Z}-1, \bar{Z}, \ldots\}$,

$$\begin{aligned} \tilde{\mathbb{E}}_1^z[v(Z_2)] - v(z) &= \begin{cases} v(z+1) - v(z) & \text{if } z = \bar{Z}-1, \\ \left[\frac{1}{2}+\rho(z)\right]v(z+1) + \left[\frac{1}{2}-\rho(z)\right]v(z-1) - v(z) & \text{if } z \geq \bar{Z}, \end{cases} \\ &= \begin{cases} \frac{\rho(\bar{Z}-1)}{2} & \text{if } z = \bar{Z}-1, \\ \rho^2(z) & \text{if } z \geq \bar{Z}, \end{cases} \\ &\overset{(L.4)}{=} \frac{\tilde{l}(z)}{4-2c}. \end{aligned}$$

Step 2. We claim that with $v(z)$ increases in z. This follows from the fact that $v(z)$ defined in (L.7) is a partial sum of nonnegative terms.

Step 3. We claim that $v(z) \to \infty$ as $z \to \infty$ with a rate such that $\limsup_{z\to\infty}\left\{\frac{v(z)}{\sum_{k=1}^{z}\rho(k)}\right\} < \infty$. To prove this claim, we analyze $v(z)$ by introducing two auxiliary sequences. First, let $a_k := \sum_{n=k}^{\infty}\prod_{m=k}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}$ for $k \in \{\bar{Z}-1, \bar{Z}, \ldots\}$. Because ρ is slowly vanishing (see Definition 2), we invoke Lemma L.1 with $x_m = \rho(m)$, and deduce that $a_1 < \infty$. Moreover, since a_k satisfies the inductive relation $a_k = \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}(1 + a_{k+1})$, we also deduce that $a_k < \infty$ for all $k \in \mathbb{Z}$. We introduce our second auxiliary sequence as follows: for $z \geq \bar{Z} \vee 2$, let

$$b_z := \sum_{n=\bar{Z}}^{z-1}\sum_{k=\bar{Z}}^{n}\tfrac{\rho^2(k)}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{n}\tfrac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} \overset{(a)}{=} \sum_{\check{n}=1}^{z-\bar{Z}}\sum_{\check{k}=1}^{\check{n}}\tfrac{\rho^2(\check{k}+\bar{Z}-1)}{\frac{1}{2}-\rho(\check{k}+\bar{Z}-1)}\prod_{\check{m}=\check{k}}^{\check{n}}\tfrac{\frac{1}{2}-\rho(\check{m}+\bar{Z}-1)}{\frac{1}{2}+\rho(\check{m}+\bar{Z}-1)},$$

where (a) holds by the following change of variables: $\check{k} := k - \bar{Z} + 1$, $\check{n} := n - \bar{Z} + 1$ and $m' := m - \bar{Z} + 1$. We claim that $\{b_z\}$ is a sequence that diverges at a rate satisfying $b_z = O(\rho(\bar{Z}) + \ldots + \rho(z-1))$. Invoking the first part of Lemma L.3 with $x_m = \rho(m+\bar{Z}-1)$, we know that $b_z \uparrow \infty$ as $z \uparrow \infty$. Moreover, since ρ is regular and slowly vanishing, the following three conditions hold:

1. $\liminf_{m\to\infty}\{m\rho(m+\bar{Z}-1)\} \geq \liminf_{m\to\infty}\{\frac{m}{m+\bar{Z}-1}\}\liminf_{m\to\infty}\{(m+\bar{Z}-1)\rho(m+\bar{Z}-1)\} > \frac{1}{4}$, because ρ is slowly vanishing.
2. $\lim_{m\to\infty}\left\{\left(\frac{\rho(m+\bar{Z}-1)}{\rho(m+\bar{Z})}-1\right)m\right\} = \lim_{m\to\infty}\left\{\left(\frac{\rho(m+\bar{Z}-1)}{\rho(m+\bar{Z})}-1\right)(m-\bar{Z}+1)\right\}\lim_{m\to\infty}\left\{\frac{m}{m-\bar{Z}+1}\right\}$ exists, because ρ is regular.
3. $\lim_{m\to\infty}\{\rho(m+\bar{Z}-1)\} = \lim_{m\to\infty}\{\rho(m)\} = 0$, because ρ is slowly vanishing.

Therefore, we deduce from the second part of Lemma L.3 that $b_z = O(\rho(\bar{Z}) + \ldots + \rho(z-1))$. With the introduction of $\{a_k\}$ and $\{b_z\}$, let us evaluate $v(z)$ below for every $z \geq \bar{Z} \vee 2$:

$$\begin{aligned} v(z) =& \left(1+\sum_{n=\bar{Z}}^{z-1}\prod_{m=\bar{Z}}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\frac{\rho(\bar{Z}-1)}{2} + \sum_{n=\bar{Z}}^{z-1}\sum_{k=\bar{Z}}^{n}\frac{\rho^2(k)}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} \\ \leq& \left(1+\sum_{n=\bar{Z}}^{\infty}\prod_{m=\bar{Z}}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)\frac{\rho(\bar{Z}-1)}{2} + \sum_{n=\bar{Z}}^{z-1}\sum_{k=\bar{Z}}^{n}\frac{\rho^2(k)}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} \\ =& (1+a_{\bar{Z}})\frac{\rho(\bar{Z}-1)}{2} + b_z. \end{aligned}$$

Now, note that $\bar{Z}$ and $a_{\bar{Z}}$ are finite constants independent of z, and $\{b_z\}$ is a sequence that diverges at a rate satisfying $b_z = O(\rho(\bar{Z}) + \ldots + \rho(z-1))$. This completes the proof. ∎

Proof of Lemma L.5. Because ρ is slowly vanishing, it is also strictly upper bounded. Thus, by Lemma E.3, there exists $\varepsilon_0 > 0$ such that $\rho(z) \geq \varepsilon_0$ for all $z \in \mathbb{N}_-$. Let

$$B := \frac{\rho^2(0)}{\frac{1}{2}+\rho(0)} + \frac{\frac{1}{2}-\rho(0)}{\frac{1}{2}+\rho(0)}\frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0} \quad \text{and} \quad \tilde{B} := -\frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0}. \tag{L.12}$$

To complete the proof, we invoke Lemma 4 with $S = \mathbb{Z}$, and deduce that it suffices solve the difference equation $\tilde{\mathbb{E}}_1^z\tilde{v}(Z_2) - \tilde{v}(z) = \frac{\tilde{l}(z)}{4-2c}$ for $z \in \mathbb{Z}$. Due to the dynamics of $\{Z_t\}$ and the definition of $\tilde{l}(\cdot)$ in (L.4), this simplifies to $\left[\frac{1}{2}-\rho(z)\right]\tilde{v}(z-1) + \left[\frac{1}{2}+\rho(z)\right]\tilde{v}(z+1) - \tilde{v}(z) = \rho^2(z)$ for $z \in \mathbb{Z}$. The rest of the proof confirms that (L.9) presents such a solution with the desired monotonicity and growth properties. We complete the remainder of the proof in three steps.

Step 1. We claim that when $\tilde{v}(\cdot)$ is as in (L.9), $\left[\frac{1}{2}-\rho(z)\right]\tilde{v}(z-1) + \left[\frac{1}{2}+\rho(z)\right]\tilde{v}(z+1) - \tilde{v}(z) = \rho^2(z)$ for all $z \in \mathbb{Z}$. We can view this function as a concatenation of two one-sided functions: one defined on $\mathbb{N}$ and the other one defined on $\mathbb{Z}_- \cup \{0\}$. For the first piece (defined on $\mathbb{N}$), we invoke Lemma F.1 and observe that $\tilde{v}(\cdot)$ restricted to $\mathbb{N}$ satisfies the difference equation (F.1) with $\hat{z} = 1$, $\delta = B$, $x(z) = \rho^2(z)$, and $p(z) = \frac{1}{2} + \rho(z)$. That is,

$$\begin{cases} \tilde{v}(0) = 0,\ \tilde{v}(1) = B, \\ \left[\frac{1}{2}+\rho(z)\right]v(z+1) + \left[\frac{1}{2}-\rho(z)\right]v(z-1) - v(z) = \rho^2(z) \ \text{ for all } z \geq 1. \end{cases}$$

For the second piece (defined on $\mathbb{Z}_- \cup \{0\}$), let us introduce an auxiliary sequence $\{y(z), z \in \mathbb{N}\}$ as follows:

$$y(z) := \begin{cases} 0 & \text{if } z = 0, \\ \left(1+\sum_{n=1}^{z-1}\prod_{m=1}^{n}\frac{\frac{1}{2}+\rho(-m)}{\frac{1}{2}-\rho(-m)}\right)\tilde{B} + \sum_{n=1}^{z-1}\sum_{k=1}^{n}\frac{\rho^2(-k)}{\frac{1}{2}+\rho(-k)}\prod_{m=n}^{k}\frac{\frac{1}{2}+\rho(-m)}{\frac{1}{2}-\rho(-m)} & \text{if } z \geq 1. \end{cases}$$

First, we verify that $y(\cdot)$ is a "reflected" version of $\tilde{v}(\cdot)$ restricted to $\mathbb{Z}_- \cup \{0\}$. That is, for all $z \in \mathbb{Z}_-$,

$$\begin{aligned}
\tilde{v}(z) &= \left(1 + \sum_{n=z+1}^{-1} \prod_{m=n}^{-1} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right)\tilde{B} + \sum_{n=z+1}^{-1} \sum_{k=n}^{-1} \frac{\rho^2(k)}{\frac{1}{2}+\rho(k)} \prod_{m=n}^{k} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)} \\
&\overset{(a)}{=} \left(1 + \sum_{n=1}^{\check{z}-1} \prod_{\check{m}=1}^{\check{n}} \frac{\frac{1}{2}+\rho(-\check{m})}{\frac{1}{2}-\rho(-\check{m})}\right)\tilde{B} + \sum_{\check{n}=1}^{\check{z}-1} \sum_{\check{k}=1}^{\check{n}} \frac{\rho^2(-\check{k})}{\frac{1}{2}+\rho(-\check{k})} \prod_{\check{m}=\check{n}}^{\check{k}} \frac{\frac{1}{2}+\rho(-\check{m})}{\frac{1}{2}-\rho(-\check{m})} \\
&= y(-z)
\end{aligned}$$

where (a) holds by the following change of variables: $\check{m} := -m, \check{n} := -n, \check{z} := -z$. Second, note that $y(\cdot)$ solves the difference equation (F.1) with $\hat{z} = 1$, $\delta = \tilde{B}$, $x(z) = \rho^2(-z)$, and $p(z) = \frac{1}{2} - \rho(-z)$; i.e.,

$$\begin{cases} y(0) = 0,\ y(1) = \tilde{B}, \\ \left[\frac{1}{2} - \rho(-z)\right] y(z+1) + \left[\frac{1}{2} + \rho(-z)\right] y(z-1) - y(z) = \rho^2(-z) \ \text{ for all } z \geq 1. \end{cases}$$

Because $\tilde{v}(z) = y(-z)$ for all $z \in \mathbb{Z}_-$, we have

$$\begin{cases} \tilde{v}(0) = 0,\ \tilde{v}(-1) = \tilde{B}, \\ \left[\frac{1}{2} - \rho(z)\right] \tilde{v}(z-1) + \left[\frac{1}{2} + \rho(z)\right] \tilde{v}(z+1) - \tilde{v}(z) = \rho^2(z) \ \text{ for all } z \leq -1. \end{cases}$$

Lastly, it suffices to verify that B and $\tilde{B}$ are such that $\left[\frac{1}{2} - \rho(0)\right] \tilde{v}(-1) + \left[\frac{1}{2} + \rho(0)\right] \tilde{v}(1) - \tilde{v}(0) = \rho^2(0)$. In fact,

$$\begin{aligned}
&\left[\tfrac{1}{2} - \rho(0)\right]\tilde{v}(-1) + \left[\tfrac{1}{2} + \rho(0)\right] \tilde{v}(1) - \tilde{v}(0) \\
&\quad= \left[\tfrac{1}{2} - \rho(0)\right] \tilde{B} + \left[\tfrac{1}{2} + \rho(0)\right] B \\
&\quad= -\left[\tfrac{1}{2} - \rho(0)\right] \frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0} + \left[\tfrac{1}{2} + \rho(0)\right] \left[\frac{\rho^2(0)}{\frac{1}{2}+\rho(0)} + \frac{\frac{1}{2}-\rho(0)}{\frac{1}{2}+\rho(0)} \frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0}\right] \quad \text{[by (L.12)]} \\
&\quad= \rho^2(0).
\end{aligned}$$

This completes the proof of our claim in Step 1.

Step 2. We claim that $\tilde{v}(z)$ increases in z. To prove this claim, we directly verify that $\tilde{v}(\cdot)$ increases on $\mathbb{N}$, as $\tilde{v}(z)$ is a partial sum of nonnegative terms for $z \in \mathbb{N}$ according to (L.9). For all $z \in \mathbb{Z}_-$, note that

$$\begin{aligned}
\frac{\tilde{v}(z) - \tilde{v}(z+1)}{\prod_{m=z+1}^{-1} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}} &= \frac{1}{\prod_{m=z+1}^{-1} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}} \left[\left(\prod_{m=z+1}^{-1} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right)\tilde{B} + \sum_{k=z+1}^{-1} \frac{\rho^2(k)}{\frac{1}{2}+\rho(k)} \prod_{m=z+1}^{k} \frac{\frac{1}{2}+\rho(m)}{\frac{1}{2}-\rho(m)}\right] \\
&= \tilde{B} + \sum_{k=z+1}^{-1} \frac{\rho^2(k)}{\frac{1}{2}+\rho(k)} \prod_{m=k+1}^{-1} \frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)} \\
&\overset{(b)}{\leq} \tilde{B} + \sum_{k=-\infty}^{-1} \frac{\rho^2(-\infty)}{\frac{1}{2}+\rho(-\infty)} \prod_{m=k+1}^{-1} \frac{\frac{1}{2}-\varepsilon_0}{\frac{1}{2}+\varepsilon_0} \\
&= \tilde{B} + \frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0} \overset{(c)}{=} 0,
\end{aligned}$$

where: (b) follows because $\rho(z) \geq \varepsilon_0$ for all $z \in \mathbb{N}_-$ and the function $x \mapsto \frac{x^2}{\frac{1}{2}+x}$ increases in x when $x \geq 0$, and (c) follows because $\tilde{B} = -\frac{\frac{1}{2}+\varepsilon_0}{8\varepsilon_0}$ by definition; see (L.12).

Step 3. We claim that $\tilde{v}(z) \to \infty$ as $z \to \infty$ with a rate such that $\limsup_{z\to\infty} \left\{\frac{\tilde{v}(z)}{\sum_{k=1}^{z} \rho(k)}\right\} < \infty$. To prove this claim, we first note that because ρ is regular and slowly vanishing, the following three conditions hold:

1. $\liminf_{m\to\infty}\{m\rho(m)\} > \frac{1}{4}$, because ρ is slowly vanishing.

2. $\lim_{m\to\infty}\left\{\left(\frac{\rho(m)}{\rho(m+1)}-1\right)m\right\}\in[0,\infty]$ exists, because ρ is regular.
3. $\lim_{m\to\infty}\{\rho(m)\}=0$, because ρ is slowly vanishing.

Hence, invoking Lemma L.1 and Lemma L.3 with $x_n=\rho(n)$, we have the following for $z\geq\mathbb{Z}_+$,

$$\begin{aligned}\sum_{n=1}^{z-1}\sum_{k=1}^{n}\frac{\rho^2(k)}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\leq\tilde{v}(z)&=\left(1+\sum_{n=1}^{z-1}\prod_{m=1}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)B+\sum_{n=1}^{z-1}\sum_{k=1}^{n}\frac{\rho^2(k)}{\frac{1}{2}-\rho(k)}\prod_{m=k}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\\&\leq\left(1+\sum_{n=1}^{\infty}\prod_{m=1}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)B+M\sum_{n=1}^{z-1}\rho(n)\\&\leq\left(1+\sum_{n=1}^{\infty}\prod_{m=1}^{n}\frac{\frac{1}{2}-\rho(m)}{\frac{1}{2}+\rho(m)}\right)B+M\sum_{n=1}^{z}\rho(n)\\&\leq\check{C}+M\big(\rho(1)+\cdots+\rho(z)\big)\end{aligned}$$

where $\check{C},M$ are finite constants independent of z. This completes the proof. ■