

To snipe or not to snipe, that's the question!

Transitions in sniping behaviour among competing high-frequency traders

Somayeh Kokabisaghi · André B. Dorsman ·
Eric J. Pauwels

Received: date / Accepted: date

Abstract In this paper we extend the investigation into the transition from sure to probabilistic sniping as introduced in Menkveld and Zoican [16]. In that paper, the authors introduce a stylized version of a competitive game in which high frequency traders (HFTs) interact with each other and liquidity traders. The authors then show that risk aversion plays an important role in the transition from sure to mixed (or probabilistic) sniping. In this paper, we re-interpret and extend these conclusions in the context of repeated games and highlight some differences in results. In particular, we identify situations in which probabilistic sniping is genuinely profitable that are qualitatively different from the ones obtained in [16]. It turns out that beyond a specific risk aversion threshold the game resembles the well-known prisoner's dilemma, in that probabilistic sniping becomes a way to cooperate among the HFTs that leaves all the participants better off. In order to turn this into a viable strategy for the repeated game, we show how compliance can be monitored through the use of sequential statistical testing.

Keywords algorithmic trading · bandits · , high-frequency exchange · Nash equilibrium · repeated games · sniping · subgame-perfect equilibrium · Sequential probability ratio · transition

Somayeh Kokabisaghi
Centrum Wiskunde & Informatica, Amsterdam, Netherlands
E-mail: karen.kokabisaghi@cwi.nl, ORCID: 0000-0002-4589-7638

André B. Dorsman
Vrije University, Amsterdam, Netherlands.
E-mail: a.b.dorsman@vu.nl

Eric J. Pauwels
Centrum Wiskunde & Informatica, Amsterdam, Netherlands.
E-mail: Eric.Pauwels@cwi.nl, ORCID: 0000-0001-7518-6856

1 Introduction

1.1 Motivation

The burgeoning of algorithmic and high-frequency trading on modern financial exchanges has give rise to theoretical investigations scrutinizing the mechanisms that underpin these markets. To analyse this complexity of continuous-time trading, a number of authors (e.g. [2, 16]) have proposed stylized models of the stock market. Although these stylized models can not capture all the intricacies of the markets in detail, they are nonetheless very useful as they often hint at interesting behaviour. As a case in point, Menkveld and Zoican [16] designed a simple model (referred to as the MZ game from here on) to study the impact of exchange latency on liquidity. They show that faster exchanges reduce payoff risk and spread, but also give rise to speculative behaviours among high frequency traders (HFT).

More specifically, they envisage a situation in which a group of competing but comparably equipped HFTs is operating alongside a large crowd of traditional liquidity traders (LT). A market is created by selecting a single market maker (MM) from among the HFTs. This market maker then publishes an ask-bid quote for a single financial asset on the exchange, and stands to earn the spread when this quote is hit by a liquidity trader. However, what makes this game interesting is that fresh news items, published at random times, push the intrinsic value of the asset either up or down. The ensuing jump in value causes all the HFTs to race to the exchange, albeit it with different intentions. The market maker attempts to cancel his own now stale quotes¹, while the other HFTs (dubbed *bandits* hereafter) race to take advantage of the arbitrage opportunity created by the outdated quotes. This type of opportunistic behaviour on the part of the bandits, is called **sniping**. The authors in [16] conclude (among other things) that increasing risk aversion in the market induces a qualitative transition in the sniping behaviour of HFTs.

While these results are intriguing and highly non-trivial, they hinge on some subtle but important features of the game. In particular, the authors study this problem in the context of a single-shot game, and focus on Nash equilibria as their main solution concept. In this paper we espouse an alternative view by interpreting the market activities of the HFTs as a *repeated game with an infinite horizon*. In this alternative setting a player can adopt more far-sighted strategies that attempt to maximise long-term gains. This change of viewpoint has a number of important consequences:

- It identifies situations in which probabilistic sniping is genuinely profitable and therefore extends the conclusions in [16] by introducing qualitatively different solutions. More specifically, it shows that forms of informal collaboration in which bandits voluntary decrease their sniping frequency, results in better utility for everyone. This explains why bandits are willing to engage in probabilistic sniping, for which there seems to be no compelling reason in the original MZ model.

¹ A quote for an asset is **stale** when it no longer reflects the most recent information.

- It allows us to identify two thresholds (for risk aversion) that govern the transitions between the three different sniping regimes: from *sure sniping* to *probabilistic sniping* to *no sniping*.
- Finally, we identify a way for an individual HFT to monitor the behaviour of the other HFTs in order to detect non-compliance. This detection mechanism is necessary to deter devious HFTs from exploiting trustworthy colleagues by sniping more than is optimal.

1.2 Overview and contributions of this paper

The remainder of this paper is organised as follows.

- In **section 2** we introduce the **MZ-game**, and show how one can compute the utility of each possible outcome. Combining these outcomes with the probabilities of the corresponding events yields the expected utility for both the market maker and any one of the bandits. It turns out that both these utilities are linear functions of the spread s . The intersection of these two utility lines determines a point of indifference (s^*, u^*) (see Fig. 2) that predicts the behaviour of the agents. More precisely, since at the start of the MZ game, HFTs are still unsure about their role in the game (market maker or bandit) they will advertise a spread s that makes them indifferent (in terms of expected utility) to the outcome of this selection. The point of indifference (s^*, u^*) is therefore a focal point and constitutes a Nash equilibrium of the game whenever $u^* > 0$. However, this is no longer the case when $u^* \leq 0$ and this is where probabilistic sniping becomes relevant. .
- **Section 3** takes a detailed look at the concept of **probabilistic sniping**: rather than attempting to snipe at every possible opportunity (*sure sniping*), bandits voluntarily restrain themselves and only enter the sniping race with a probability $p < 1$. Less sniping is obviously beneficial for the market maker but, under certain circumstances, this also protects bandits from adverse effects. As a consequence, depending on the choice of the sniping probability p , the point of indifference shifts to a (potentially) more favourable position $(s^*(p), u^*(p))$ (see Fig. 3). In this section we characterize in detail how utilities (and derived quantities such as the point of indifference) depend on the sniping probability.
- Whether or not probabilistic sniping results in better utility (for all HFTs involved) depends ultimately on the game parameters (see Fig. 4): changing these parameters might induce a *transition in optimal behaviour*, e.g. from *sure sniping* to *probabilistic sniping* to *no sniping* (see sections 4.1 and 4.2). We then highlight the specific role of **risk aversion** γ in these sniping transitions and identify the relevant **thresholds** $\bar{\gamma}_K$ and $\bar{\gamma}_L$ (section 4.3). Finally, we introduce the optimal sniping probability p_K^* (section 4.4) that results in the best possible utility u_K^* .

- Up to this point in the paper, the equilibria obtained under various conditions were optimal, but not necessarily Nash equilibria for the single-shot MZ game. In **section 5** we explain what the implications are of shifting our viewpoint from one-shot to **repeated games**. In particular, we argue that probabilistic sniping will give rise to a new set of playable subgame-perfect Nash equilibria. To turn this into a workable solution, we need to be able to reliably detect *non-compliance* on the part of the other agents, in order to be able to take retaliatory measures. We therefore quantify the effect that devious agents (agents that pretend to snipe probabilistically, but in fact snipe for sure) have on the utility of trustworthy agents (see e.g. Figs. 12 or 13). We conclude this section by discussing how Wald's sequential probability ratio test can be used to reliably detect non-compliance (see Fig. 14).
- Finally, in **section 6**, we discuss how our work compares to related research and finish by offering some conclusions and suggestions for further research in **section 7**.
- The main part of this paper focuses on the conceptual framework that underpins our results. All the computational details are therefore relegated to sections A through J in the **appendix**. Since the mathematical characterisation of the results often requires tedious but straightforward algebra, we also provide Python notebooks to support all the computations. These can be found:

KPD_supplementary_material link to GitHub

Contributions The main contributions in this paper can be found in sections 3 through 5, and can be summarized as follows:

1. We show how the MZ model of the stock exchange can be re-interpreted as a *repeated (sequential) game against nature*. In addition, we think of the choice of sniping probability (for the bandits) as a choice of a pure action from a continuous action space, rather than a probabilistic mixing of two pure strategies (i.e. *sniping* and *non-sniping*). This change of viewpoint allows us to identify a wider set of conditions under which probabilistic sniping makes sense (see section 3.4). Put succinctly, although these new equilibria are not Nash equilibria for a one-shot stage game, they are genuine subgame-perfect equilibria in the corresponding repeated game with infinite horizon.
2. We identify two threshold values (viz. $\bar{\gamma}_K$ and $\bar{\gamma}_L$) for the risk aversion factor γ that govern the transition to probabilistic sniping. More specifically, when γ increases beyond $\bar{\gamma}_K$, probabilistic sniping becomes advantageous for all HFTs involved. Conversely, when γ exceeds $\bar{\gamma}_L$, even probabilistic sniping is no longer profitable.
3. We show that under conditions for which probabilistic sniping is advantageous, there is an optimal sniping probability p_K^* that maximises expected utilities for both the market maker and the bandits. Both the sniping probability and the corresponding optimal utility vary continuously as functions of the game parameters (see Fig. 9).

4. When the game conditions call for probabilistic sniping in order to optimise utility, the agents are tempted to snipe for sure as this maximises their own utility to the detriment of their compliant colleagues whose utility is negatively affected. As a consequence, agents participating in such a probabilistic sniping regime need a strategy to detect the non-compliance of deceptive agents. We show how this can be done using Wald's sequential probability ratio test (section 5.3.2).
5. **Reproducible research:** We provide Python notebooks that will allow the diligent reader to check the straightforward but tedious calculations underpinning our arguments, as well as simulations that illustrate the fluctuations about the mean values that result from the theoretical approach.

2 The Menkveld-Zoican (MZ) game revisited

2.1 A bird's eye view of the MZ game

As this paper builds on the MZ paper [16], we briefly highlight some of the main arguments and conclusions from that paper. In [16], the authors define a stylized version of the behaviour and strategies of high frequency traders (HFT) interacting with an (high-frequency) exchange. We defer some details to section 2.2 but, roughly speaking, they envisage a game in which, during an initial pre-game stage, the HFTs have to pick the value of the (half) spread s in order to post a single bid-ask quote ($v \pm s$) for a financial asset of (commonly perceived) value v . One of the HFTs is then selected as market maker, whereas the others are relegated to the role of bandit. Only the quotes of the market maker will enter the order book, so he is the only one who can benefit from trading with liquidity traders who visit the exchange at random times (with rate μ). The bandits will try to gain some payoff by attempting to snipe stale quotes (see section 2.2 for more details).

The game is initiated by a trigger event which can be either the publication of news item (changing the intrinsic value v of the asset), or the interaction of a liquidity trader with the market maker (resulting in income for the latter). The latter event does not evoke any reaction from the HFTs, but the former alters the intrinsic value of the asset and causes all the HFTs to race towards the exchange, albeit for different reasons. The market maker attempts to cancel his now stale quotes, whereas the bandits hope to obtain financial gain by sniping the market maker. The outcome of this race is further complicated by the fact that, during the race, another event (news or liquidity trade) might occur. This is less likely when the exchange is fast (low exchange latency δ).

In [16] the authors investigate in detail how the interplay between various game parameters affect the behaviour and strategies of the HFTs. Although this model is highly stylized and therefore somewhat unrealistic, it is the contention of the MZ authors that even under these simplifying assumptions, interesting and non-trivial conclusions can be drawn. This suggests that a more realistic model will give rise to even more intriguing insights.

2.2 Primitives for the MZ game

Because they are important to understand the rest of the paper, we briefly recapitulate the main ingredients and assumptions governing the MZ game. For more details we refer to the original paper [16].

- **Agents or Players** The players in this game is the group of $H > 2$ high-frequency traders (HFTs) who operate in an environment populated by an infinite number of (uninformed) liquidity traders (LT). All HFTs have simultaneous and instantaneous access to all public information affecting the market. At the start of the game, one HFT is assigned to the role of market maker. The remaining $(H - 1)$ HFTs take up the role of high-frequency bandits (HFB), intent on financial gain by sniping the stale quotes whenever an opportunity presents itself. A final important characteristic of all HFTs is that they are *risk-averse*, i.e. the utility of *negative* pay-offs is inflated by a factor $\gamma > 1$.
- **Exchange latency δ** measures the time delay δ between the trigger event (that sets off the game) and the arrival of one HFT at the matching engine where each order is actually processed. This delay determines the time span over which the race game is played out and, as such, has a direct impact on the expected number of additional events that might occur during the race.
- **Exogenous events** The behaviour of the HFTs is affected by two types of exogenous events that happen independent of the HFTs actions:
 1. **Publication of news item causing a common value shock:** the publications of news (either good or bad) is observed by everyone, and changes the intrinsic value of an asset. This process is modelled as a Poisson process with rate α (number of news events per time unit). The size of the shock is fixed at $\pm\sigma$. Throughout this paper we are assuming that the (half) spread (s) is (strictly) less than σ ($0 < s < \sigma$).
 2. **Private value shock:** These are only known to the LTs and cause one of them them to join the matcher queue to hit the outstanding quote (resulting in income for the market maker). These events are modelled as an independent Poisson process with rate μ (number of LT arrivals per time unit).
- **Detailed chronology of the MZ game** The description below is based on the MZ paper, but attempts to clarify the process in more detail (cf. Fig 1).




Fig. 1 Chronology of the MZ game

- At $t = -1$, Strategy selection: all HFTs are informed about the game's parameters (i.e. $\alpha, \mu, \delta, \sigma$ and H). Based on this information they can

then compute the expected utility as a function of the spread s (and the sniping behaviour) for both the market maker and bandits (more details in section 2.3), but are still unsure about which one applies to them. Taking this uncertainty into account, each HFT decides on a spread s , and posts a single bid and ask quote in an empty order book (i.e. initially each HFT has a zero position²).

- At $t = -1/2$, **Type assignation:** one of the HFTs is chosen to be the market maker, whereas all the others will become "bandits" (eager to snipe). The selection process of the market maker proceeds in two steps:
 1. First, the HFT(s) that posted the smallest spread s are selected;
 2. If there are multiple HFTs that posted the same minimal spread, then one of them is selected at random (uniformly).

The fact that the outcome of this selection process is uncertain, is an important aspect of the game that the HFTs need to take into account when selecting a strategy at $t = -1$ as it is impossible for them to change their position at a later time.

- At $t = 0$, **Initial trigger starts game:** one of two possible *trigger events* occurs:
 1. *News event* (public) that changes the value of the asset (i.e. value shock of $\pm\sigma$). This immediately triggers a race among the HFTs: the market maker will try to update his stale quotes, while the bandits will attempt to snipe. The winner of this race is randomly (uniform) chosen among all contestants. The winner is known at time $t = \delta$.
 2. *LT arrives at matcher (queue)*. This will result in a transaction, which is invariably profitable for the market maker as he will cash in the spread (s). This event does not elicit any reaction from the HFTs.
- Interval $0 < t < \delta$ During this interval one (and only one) of three things can happen:
 1. An additional news event becomes (publicly) known (rate α).
 2. Another LT arrives at the matcher, intent on interacting with the market maker (rate μ);
 3. Nothing happens.
- At $t = \delta$, **Conclusion and pay-off:** The game concludes and the positions are used to compute the pay-off and corresponding utility for both the bandits and market maker (see section 2.3 for more details). Importantly, the utility of a negative pay-off is further inflated by the **risk-aversion factor** γ :

$$\text{utility} = \begin{cases} \text{pay-off} & \text{if pay-off} \geq 0 \\ \gamma \cdot \text{pay-off} & \text{if pay-off} < 0 \end{cases} \quad (1)$$

where $\gamma \geq 1$.

In the next sections we will explore some of the concepts in more detail, starting with the detailed computation of the pay-off.

² A zero position is when traders don't own any amount of financial assets or commodity

2.3 Payoff and utility computation

2.3.1 General principles

To compute the payoff for specific game-outcomes for each HFT we observe that there are three sources of payoff:

- *Changes in position*: obtaining more (or less) of a financial asset results in a corresponding change in wealth;
- *Change of the intrinsic value of an asset*: if the intrinsic value of an asset changes (due to good or bad news), the owner benefits (or suffers) accordingly;
- *Income* from a successful transaction (either buy or sell). The income can be both positive or negative depending on whether the HFT gets paid or needs to pay to complete the transaction.

These observations can be summarized in the following equations:

$$\text{payoff} = \underbrace{\text{position}_{(t=\delta)} \times \text{value}_{(t=\delta)}}_{\text{final}} + \underbrace{\text{income}}_{\text{during}} - \underbrace{\text{position}_{(t=0)} \times \text{value}_{(t=0)}}_{\text{initial}}$$

but since we are assuming that all HFTs start with a zero position, the last term vanishes and we therefore get the simplified formula:

$$\text{payoff} = \text{position}_{(t=\delta)} \times \text{value}_{(t=\delta)} + \text{income}_{(t=\delta)} \quad (2)$$

The corresponding utilities U_M (for the market maker) and U_B (for a bandit) can then be inflated negative payoffs with the risk aversion factor γ (cf. eq. 1).

In addition to the primary MZ-parameters introduced above (viz. H, α, μ, δ and γ) we will introduce some useful shorthand notation for derived quantities that recur throughout this paper:

- $\bar{\alpha} := \alpha\delta/2$ and $\bar{\mu} := \mu\delta/2$ are (half) the expected number of news and LT arrivals, respectively, in the interval between the start of game ($t = 0$) and its conclusion (at $t = \delta$). Since we are assuming that δ is sufficiently small so that we expect less than one event to happen in each interval, we have the following strict inequality:

$$(\alpha + \mu)\delta < 1 \quad \text{or again} \quad \bar{\alpha} + \bar{\mu} < \frac{1}{2}. \quad (3)$$

- $\beta := \alpha/(\alpha + \mu)$ is the probability that the trigger event will be arrival of news rather than an LT arrival. Obviously, the latter therefore has probability $1 - \beta$.
- Some additional notation to streamline the equations below:

$$q := \gamma - 1 \quad \text{and} \quad m = 1 - \bar{\mu} \quad \text{and} \quad \bar{\theta} = 2 \frac{\bar{\alpha}\bar{\mu}}{\bar{\alpha} + \bar{\mu}}. \quad (4)$$

Notice that $\bar{\theta}$ is the harmonic mean of $\bar{\alpha}$ and $\bar{\mu}$ therefore tends to be closer to the smaller of the two.

2.3.2 Payoffs for specific event-sequences

Since the chronology of the MZ game is strictly defined, there are only 20 possible event sequences that can happen between the initial trigger (at $t = 0$) and the final conclusion of the game at $t = \delta$. All the possible combinations and their corresponding payoffs (and utilities) are given in [16] and summarized for the reader's convenience in Table 1. It is important to realise that bandits can only benefit from the game if their sniping attempt is successful. Hence, if there is no race (i.e. trigger event is the arrival of a liquidity trader), or if a bandit loses the race, then their payoff equals zero. Furthermore, if a race ensues which the market maker loses, *only the successful sniper* will gain utility (2nd utility column in table).

For a detailed explanation of each event's payoff and probability we refer to the appendix C, but we will illustrate the general ideas using NG-LA (first event is good news and the second event is arrival of LT on ask) as a specific example:

- **Probability of event:** Since, in general, the first (trigger) event is either the publication of a news item (rate α) or the arrival of a liquidity trader (LT, rate μ), the probability of the former equals $\beta := \alpha/(\alpha + \mu)$. Since for half of these events the news is good, the probability of the trigger event being NG equals $\beta/2$. The probability of a second news event arriving in the interval $[0, \delta]$ equals $\alpha\delta$ and if we insist on this news item also being good, the probability halves to $\alpha\delta/2 = \bar{\alpha}$. Since the two events are independent, we can simply multiply the two probabilities.
- **Payoff of event:** In the case of NG-LA, the first (i.e. trigger) event is good news and hence increases the intrinsic value of the asset (by an amount σ). As a consequence, the market maker (MM) will race to the exchange to cancel his (now stale) order in order to prevent losses, while all other HFTs (bandits) race to snap up the higher value asset at the outdated low price (sniping). The second event is the arrival of LT on *ask* (value +s) and means that an uninformed liquidity trader (LT) happens to join the queue at the matching engine ahead of all the HFTs (who are still racing) and manages to buy the asset at the outdated (low) ask price. This is of course bad news for the market maker who ends up selling an asset below its current market value. The payoff for all the bandits is zero as they failed to snipe. By referring to equation (2), we can compute payoff for market maker:

$$\begin{aligned} \text{payoff}_M &= (\text{position})_{(t=\delta)} \times (\text{value})_{(t=\delta)} + \text{income} \\ &= (-1) \times (v + \sigma) + (v + s) \\ &= -(\sigma - s) \end{aligned} \tag{5}$$

Since the payoff is negative, the corresponding utility (u_M) is inflated by the risk aversion factor γ , hence:

$$u_M = -\gamma(\sigma - s) \quad \text{and} \quad u_B = 0. \tag{6}$$

For more details, see:

KPD_supplementary_material_1.ipynb (section.1)

Utilities						
			MM loses race (Prob = h)	MM wins race (Prob = $1 - h$)		
Event code	1st and 2nd event prob		MM	B (sniper!)	MM	B (all!)
NG-NG	$\frac{1}{2}\beta$	$\bar{\alpha}$	$-\gamma(2\sigma - s)$	$2\sigma - s$	0	0
NG-NB	$\frac{1}{2}\beta$	$\bar{\alpha}$	s	$-\gamma s$	0	0
NG-LA	$\frac{1}{2}\beta$	$\bar{\mu}$	$-\gamma(\sigma - s)$	0	$-\gamma(\sigma - s)$	0
NG-LB	$\frac{1}{2}\beta$	$\bar{\mu}$	$2s$	$\sigma - s$	$\sigma + s$	0
NG-no	$\frac{1}{2}\beta$	$1 - 2(\bar{\alpha} + \bar{\mu})$	$-\gamma(\sigma - s)$	$\sigma - s$	0	0
NB-NG	$\frac{1}{2}\beta$	$\bar{\alpha}$	s	$-\gamma s$	0	0
NB-NB	$\frac{1}{2}\beta$	$\bar{\alpha}$	$-\gamma(2\sigma - s)$	$2\sigma - s$	0	0
NB-LA	$\frac{1}{2}\beta$	$\bar{\mu}$	$2s$	$\sigma - s$	$\sigma + s$	0
NB-LB	$\frac{1}{2}\beta$	$\bar{\mu}$	$-\gamma(\sigma - s)$	0	$-\gamma(\sigma - s)$	0
NB-no	$\frac{1}{2}\beta$	$1 - 2(\bar{\alpha} + \bar{\mu})$	$-\gamma(\sigma - s)$	$\sigma - s$	0	0
No race!						
			MM	B (all!)	-	-
LA-NG	$\frac{1}{2}(1 - \beta)$	$\bar{\alpha}$	$-\gamma(\sigma - s)$	0	-	-
LA-NB	$\frac{1}{2}(1 - \beta)$	$\bar{\alpha}$	$\sigma + s$	0	-	-
LA-LA	$\frac{1}{2}(1 - \beta)$	$\bar{\mu}$	s	0	-	-
LA-LB	$\frac{1}{2}(1 - \beta)$	$\bar{\mu}$	$2s$	0	-	-
LA-no	$\frac{1}{2}(1 - \beta)$	$1 - 2(\bar{\alpha} + \bar{\mu})$	s	0	-	-
LB-NG	$\frac{1}{2}(1 - \beta)$	$\bar{\alpha}$	$\sigma + s$	0	-	-
LB-NB	$\frac{1}{2}(1 - \beta)$	$\bar{\alpha}$	$-\gamma(\sigma - s)$	0	-	-
LB-LA	$\frac{1}{2}(1 - \beta)$	$\bar{\mu}$	$2s$	0	-	-
LB-LB	$\frac{1}{2}(1 - \beta)$	$\bar{\mu}$	s	0	-	-
LB-no	$\frac{1}{2}(1 - \beta)$	$1 - 2(\bar{\alpha} + \bar{\mu})$	s	0	-	-
Legend	NG/NB: Good/Bad News, LA/LB = liquidity trader on ask/bid, no = no event					

Table 1 Payoff table for all possible events in MZ-game. For instance, in NG-LA, the first event is good news and the second event is arrival of LT on ask. The negative payoffs are made explicit using a minus sign and are inflated with the risk aversion factor $\gamma \geq 1$ in order to obtain the actual utility. The trigger event for the first ten cases is the publication of a news item which sets off a race. A bandit can only receive payoff if he succeeds in winning that race, in all other cases there is no payoff. For that reason we only list the utility of the bandit that was successful as a sniper. For the explanation of the probabilities that the market maker will lose or win the race, we refer the reader to section 3.2). Notice that for the last ten events for which the trigger event is the arrival of liquidity trader, there is no race, and so no need to distinguish between the possible fates of the market maker. For detailed computation of the table, see appendix C.

2.3.3 Expected utilities

Table 1 lists the utilities of specific events for both the market maker and any of the bandits once we specify the spread s . However, in the actual game these events are generated according to a stochastic process and the utilities are therefore stochastic variables (denoted by capital letters U_M and U_B). To compute the *expected* utilities ($\mathbb{E}U_M$ and $\mathbb{E}U_B$) we need to sum the utilities of all relevant events weighted by the corresponding probabilities. The event probabilities are listed explicitly in the table. But we also have to take into account the probabilities with which any of the HFTs will win or lose the race as this changes their pay-off. Below we explore this in more detail for the market maker and bandit separately.

Market maker A glance at Table 1 shows that the market maker receives non-trivial utility for all the events. To streamline the exposition we will introduce a wildcard and denote all event codes that correspond to a news trigger event as N^{***} . Each of these events occurs with probability $\beta/2$ and provokes a race. Similarly, events triggered by an LT arrival will be denoted as L^{***} and occur with probability $(1 - \beta)/2$ but don't spark a race. Furthermore, when there is an actual race, the market maker's utility depends on whether or not he beats the other HFTs to the finish. To take this into account, we introduce h to denote the probability that the market maker will *lose* the race. Since all the HFTs enter the race and have the same probability of winning it, we conclude:

$$h := P\{\text{MM loses race} \mid \text{there is a race}\} = \frac{H - 1}{H}. \quad (7)$$

Gathering all this information we can now express the expected utility of the market maker as a weighted sum over the possible *second* event e_2 where $p(e_2)$ and $u_M(e_2)$ are the corresponding *event* probability and utility, respectively:

$$\begin{aligned} \mathbb{E}U_M(s) = & \frac{\beta}{2} \left[\sum_{e_2 \in N^{***}} p(e_2) \{h u_M(e_2 \mid \text{MM loses}) + (1 - h) u_M(e_2 \mid \text{MM wins})\} \right] \\ & + \left(\frac{1 - \beta}{2} \right) \left[\sum_{e_2 \in L^{***}} p(e_2) u_M(e_2) \right]. \end{aligned} \quad (8)$$

In section 3.2 we will see how we need to re-interpret this expression in the case of probabilistic sniping.

Individual bandit The expected utility for an *individual* bandit can be computed along similar lines. In fact, the calculation is easier as a bandit only receives payoff if his sniping attempt is successful. If there is a race (event code N^{***}), the probability that a specific, individual bandit will win equals

$$g := P\{\text{bandit wins race} \mid \text{there is a race}\} = \frac{1}{H} \quad (9)$$

as there are H agents in the race. Notice that this is identical to the probability $1 - h$ that the market maker will win the race – for the obvious reason. It might therefore seem an unnecessary complication of the notation, but once we introduce the concept of probabilistic sniping, these two quantities will diverge in meaning.

Since all L^{***} events result in zero utility for a bandit (no race), we only need to sum over the N^{***} events, resulting in the following expression:

$$\mathbb{E}U_B(s) = \frac{\beta}{2}g \sum_{e_2 \in N^{***}} p(e_2) u_B(e_2 \mid \text{MM loses race}) \quad \text{where} \quad g = \frac{1}{H}. \quad (10)$$

See `KPD_supplementary_material_1.ipynb` (section.2.1)

Some simplifications Expanding eqs. 8 and 10 is straightforward but tedious and therefore deferred to appendix D. However, even without plowing through the algebra, a cursory inspection of table 1 shows that all event utilities are proportional to either the spread s or the difference ($w\sigma \pm s$, where $w = 1, 2$) of spread and jumpsize σ . From this, we can introduce the following notational simplifications:

- **Rescaling** We can factor out σ by dividing both the spread s and utility U by σ . This amounts to rescaling both spread and utility by using the jumpsize σ as a natural yardstick:

$$s \longrightarrow \tilde{s} := \frac{s}{\sigma} \quad \text{and} \quad U \longrightarrow \tilde{U} := \frac{U}{\sigma}$$

This rescaling allows us to fix $\sigma \equiv 1$ which slightly simplifies the equations in the remainder of the paper without having any impact on the qualitative nature of the conclusions. Notice that we will drop the tildes for notational convenience, and simply use s and U to denote spread and utility measured using σ as unit.

- **Utilities are linear in s** The expected utilities for both the market maker ($\mathbb{E}U_M$) and bandit ($\mathbb{E}U_B$) are linear combinations of the utilities in the table and therefore result in linear (or more precisely, affine) functions in the spread s :

$$\mathbb{E}U_B(s) = A(1-s) + Bs, \quad \text{where} \quad A := \mathbb{E}U_B(s=0) \quad \text{and} \quad B := \mathbb{E}U_B(s=1), \quad (11)$$

and similarly:

$$\mathbb{E}U_M(s) = C(1-s) + Ds, \quad \text{where} \quad C := \mathbb{E}U_M(s=0) \quad \text{and} \quad D := \mathbb{E}U_M(s=1). \quad (12)$$

Hence, the values A, B, C and D are the endpoints of the linear line-segments that represent the expected utilities for the market maker and a bandit as a function of spread $0 \leq s \leq 1$ (also see Fig. 2). Expressing the utilities in terms of these endpoints is useful as it simplifies some the formulae in the remainder of this paper.

2.4 Computing the point of indifference

Recall from the description of the MZ-game in section 2.2 that at the inception of the game (i.e. at $t = -1$ when they need to specify the spread s for their quotes), the HFTs are agnostic about their eventual role in the game as this is only assigned at $t = -\frac{1}{2}$. The intersection point between the two utilities will therefore play a pivotal role in the HFT's decision making as it represents the spread (hereafter denoted by s^*) that makes them indifferent as to which role they will play. In terms of the quantities A, B, C and D as defined above, it is straightforward to determine s^* and the corresponding utility $u^* := \mathbb{E}U_B(s^*) \equiv \mathbb{E}U_M(s^*)$ by computing the intersection of the linear utility functions (also see Fig. 2):

$$s^* = \frac{(A - C)}{(A - C) + (D - B)} \quad \text{and} \quad u^* = \frac{(AD - BC)}{(A - C) + (D - B)} \quad (13)$$

It is shown in [16] that, if $u^* > 0$ (intersection lies above x-axis), the MZ game has a unique **pure Nash equilibrium** in which all HFTs will pick the optimal spread s^* and expect to earn the utility u^* .

For more details, see:

KPD_supplementary_material_numeric.ipynb (section.2)

However, if $u^* \leq 0$ (intersection lies on or below x-axis), things are less straightforward as pursuing the above strategy will now result in losses for both the market maker and the bandits. One of the contribution in [16] is the realisation that probabilistic sniping might suggest an additional equilibrium. This will be explored in detail in the next section.

3 Probabilistic versus Sure Sniping

3.1 The rationale for probabilistic sniping

Up till now we have assumed that all bandits will race to the exchange whenever an opportunity arises (i.e. whenever the trigger event is the publication of a news item). This seems reasonable as a successful snipe is the only way a bandit can gain positive utility. However, a careful inspection of Table 1 reveals that sniping has a potential down-side. For instance, if a positive news fact is followed by a negative one (NG-NB), or vice versa (NB-NG), the sniper acquires an asset that has no added value and, as a consequence, incurs negative utility. Hence, under these circumstances it might be more profitable to refrain from sniping, or to snipe only occasionally. This is translated into the idea of **probabilistic sniping** in which an individual bandit enters the race with **sniping probability** p .

Obviously, less sniping is always advantageous for the market maker as sniping reduces his payoff. So it seems that there are situations (e.g. high news rates α , see caption of Fig. 7) for which probabilistic sniping could actually be advantageous for all parties. In the next sections we will quantify in more detail how probabilistic sniping affects the various utilities.

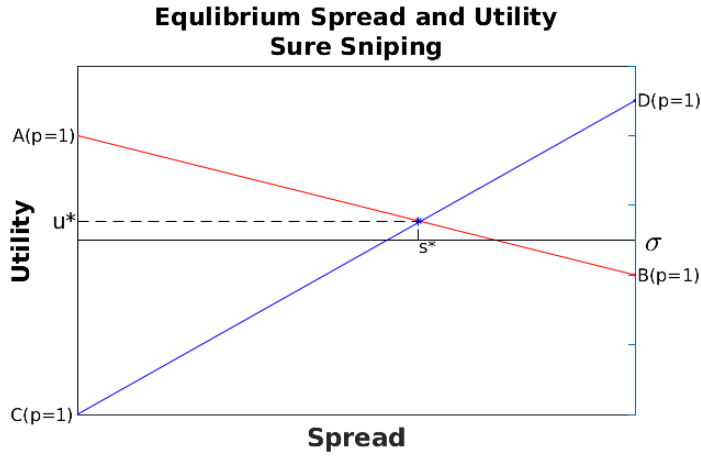


Fig. 2 Graph representing the expected utilities $\mathbb{E}U_B(s)$ (line AB) and $\mathbb{E}U_M(s)$ (line CD) for a bandit and the market maker, respectively. The HFTs will select spread s^* (resulting in expected utility u^*) by determining the intersection of the two utilities (point of indifference). For the time being, we only consider sure sniping which corresponds to $p = 1$. For more details on probabilistic sniping (for which $p < 1$), see section 3.

3.2 The impact of probabilistic sniping

We start by defining more formally the sniping probability p to be the probability that an individual bandit will participate in the race, given that there is a race, i.e. given that the initial trigger event is "news":

$$p := P\{\text{Individual bandit will race} \mid \text{Trigger} = \text{News}\} \quad (14)$$

Obviously, $p = 1$ in the case of sure sniping. Notice also that this probability is only relevant for bandits, as the market maker will always enter the race! The corresponding change in sniping behaviour implies that we have to re-evaluate the probabilities h (probability that the market maker will lose the race) and g (probability that an individual bandit who entered the race, will actually win it) as both quantities will now depend on p .

1. **Probability $h(p)$ that market maker will lose the race** In the case of sure sniping, all HFTs enter the race and since they all have equal probability of winning, the probability that the market maker will *lose* equals $h = (H-1)/H$. However, when individual bandits snipe (independently!) with a probability $p < 1$, the probability that the market maker will get sniped depends on the number of bandits that do attempt to out-race him, and therefore depends on p . We will therefore denote it as $h(p)$, defined formally as:

$$h(p) := P\{\text{MM will lose MZ race} \mid \text{individual bandits enter race with prob } p\} \quad (15)$$

A straightforward calculation (see Appendix E) shows that

$$h(p) = \frac{(1-p)^H - (1-Hp)}{Hp} \quad (16)$$

which indeed reduces to $h \equiv h(1) = (H-1)/H$ in the case of sure sniping ($p = 1$). Furthermore, using the expansion

$$(1-p)^H = 1 - Hp + \frac{H(H-1)}{2}p^2 - \frac{H(H-1)(H-2)}{3!}p^3 + \dots + (-p)^H \quad (17)$$

shows that we have the following asymptotics for small values of $p \downarrow 0$:

$$h(p) = \left(\frac{H-1}{2}\right)p + o(p^2), \quad (18)$$

whence $\lim_{p \downarrow 0} h(p) = 0$. The probability $h(p)$ plays a crucial role in the utility computation for both the market maker and the bandits.

2. **Conditional probability $g(p)$ that an individual bandit wins race, given he enters race** In the case of probabilistic sniping, the probability that an individual bandit will win the race needs to be conditioned on him actually entering the race, as this is no longer a sure thing. Therefore, we introduce the following formal definition of this conditional probability:

$$g(p) := P\{\text{individual bandit will win race} \mid \text{he participates in the race}\} \quad (19)$$

The actual computation $g(p)$ is deferred to Appendix E but yields:

$$g(p) = \frac{(1-p)^H - (1-pH)}{H(H-1)p^2} = \frac{h(p)}{(H-1)p}. \quad (20)$$

Intuitively, this result makes sense: $h(p)$ is the probability that a bandit rather than the market maker will win the race, which is then divided by $(H-1)p$, i.e. the expected number of bandits in the race.

Another way to see the connection between $h(p)$ and $g(p)$ is to observe that $pg(p)$ is the **un**-conditional probability that an individual bandit will win the race. For this to happen, the market maker needs to lose the race (probability $h(p)$) and each of the $(H-1)$ bandits has the same odds. Hence:

$$pg(p) = \frac{h(p)}{H-1} \quad (21)$$

Notice also that if every bandit races for sure ($p = 1$), then the probability that an individual bandit will win, simplifies to $g \equiv g(1) = 1/H$ as expected. In the same vein, using the asymptotics in eq. (18), it follows that

$$\lim_{p \downarrow 0} g(p) = 1/2. \quad (22)$$

This also makes sense: if $p \approx 0$ then any bandit that is in the race will probably be the only bandit in that race, and therefore will have a probability of 1/2 to outrun the racing market maker.

For more details, see:

(KPD_supplementary_material_2.ipynb)

3.3 Expected utilities under probabilistic sniping

Now that we have a precise definition of the probabilities $h(p)$ and $g(p)$ we are in a position to determine the expected utilities for both the market maker and individual bandits under probabilistic sniping.

3.3.1 Expected utility for bandit

As pointed out above, the bandit will only race if the trigger event (Tr) is the arrival of a news item (N). Recall that this happens with probability $\beta = \alpha/(\alpha+\mu)$. Hence, we can compute the expected utility by conditioning on the trigger event being news. Furthermore, we need to take into account that the bandit will only participate in the race with probability p , and needs to win the race in order to gain utility. We therefore find:

$$\begin{aligned}\mathbb{E}U_B &= \mathbb{E}(U_B \mid Tr = N) P(Tr = N) \\ &= \mathbb{E}(U_B \mid Tr = N) \beta \\ &= \mathbb{E}(U_B \mid Tr = N \ \& \ \text{bandit enters race}) P\{\text{bandit enters race}\} \beta \\ &= \mathbb{E}(U_B \mid Tr = N \ \& \ \text{bandit enters and wins race}) P\{\text{bandit wins race} \mid \text{he races}\} p \beta \\ &= \mathbb{E}(U_B \mid Tr = N \ \& \ \text{bandit enters and wins race}) g(p) p \beta\end{aligned}$$

From this we conclude:

$$\mathbb{E}U_B(s, p) = \beta g(p) p \mathbb{E}(U_B \mid Tr = N \ \& \ \text{bandit enters and wins race}) \quad (23)$$

To further elaborate the expression for $\mathbb{E}U_B$ we need to condition the RHS on the second event. This amounts to summing the utilities (weighted with the probabilities of the second event) in the sniper column of Table, which yields:

and a straightforward but tedious calculation shows:

$$\mathbb{E}(U_B \mid Tr = N \ \& \ \text{bandit enters and wins race}) = (1 - \bar{\mu})(1 - s) - \bar{\alpha}qs.$$

Plugging this into eq. 23 shows how the expected utility of each bandit depends on spread s , sniping probability p and risk aversion $\gamma \geq 1$:

$$\mathbb{E}U_B(s, p, \gamma) = \beta p g(p) \{(1 - \bar{\mu})(1 - s) - \bar{\alpha}(\gamma - 1)s\} = \beta p g(p) \{m(1 - s) - \bar{\alpha}qs\}. \quad (24)$$

In particular, this means that the end points of the line segment, as functions of p and γ , are given by:

$$\begin{aligned}A \equiv A(p) &:= \mathbb{E}U_B(s = 0) = m\beta p g(p) \\ B \equiv B(p, \gamma) &:= \mathbb{E}U_B(s = 1) = -\bar{\alpha}(\gamma - 1)\beta p g(p)\end{aligned} \quad (25)$$

We conclude this paragraph with a straightforward but useful observation. The value s^{**} for which the bandit's utility becomes zero (i.e. the intersection point

with the x-axis, see Fig. 3) does not depend on the sniping probability. Indeed, a straightforward computation shows:

$$s^{**} = \frac{A}{A+B} = \frac{m}{m - \bar{\alpha}q} \quad (26)$$

This means that geometrically speaking, probabilistic sniping makes the bandit's utility function pivot about the intersection point s^{**} towards the horizontal axis (see Fig. 3).

For more details, see:

KPD_supplementary_material_1.ipynb(section.2.1)

3.3.2 Expected utility for market maker

The computation of the expected utility $\mathbb{E}U_M$ proceeds along similar lines but is somewhat more involved. Again, we start by conditioning on the trigger event (Tr) which can either be the arrival of news (N , prob = β) or a liquidity trader (LT , prob = $1 - \beta$):

$$\mathbb{E}U_M = \mathbb{E}(U_M \mid Tr = N)\beta + \mathbb{E}(U_M \mid Tr = LT)(1 - \beta)$$

The second term (trigger event is the *arrival of a liquidity trader*) is relatively simple as no sniping is involved. A straightforward computation yields:

$$\mathbb{E}(U_M \mid Tr = LT) = (1 + \bar{\mu})s - \bar{\alpha}q(1 - s) \quad (27)$$

In case the trigger event is *news*, the market maker's payoff depends on whether or not he loses the ensuing race and therefore gets sniped. Using the notation introduced above we conclude:

$$\begin{aligned} \mathbb{E}(U_M \mid Tr = N) &= \mathbb{E}(U_M \mid Tr = N \ \& \ \text{MM loses race})h(p) \\ &\quad + \mathbb{E}(U_M \mid Tr = N \ \& \ \text{MM wins race})(1 - h(p)) \end{aligned}$$

Again, a straightforward but tedious calculation shows that the market maker's utility in case he gets sniped equals:

$$\mathbb{E}(U_M \mid Tr = N \ \& \ \text{MM loses race}) = (2\bar{\mu} - \bar{\alpha}q)s - \gamma(1 - \bar{\mu})(1 - s) \quad (28)$$

Similarly:

$$\mathbb{E}(U_M \mid Tr = N \ \& \ \text{MM wins race}) = ((\gamma + 1)s - \gamma + 1)\bar{\mu} \quad (29)$$

From the expressions above we can constitute the linear function $\mathbb{E}U_M(s)$ but in the remainder of the paper we only need explicit expressions for the endpoints:

$$\begin{aligned} C \equiv C(p, \gamma) &:= \mathbb{E}U_M(s = 0) = -\{q\bar{\theta} + \beta(m\gamma - \bar{\mu}q)h(p)\} \\ D \equiv D(p, \gamma) &:= \mathbb{E}U_M(s = 1) = (1 + \bar{\mu}) - \beta(m + \bar{\alpha}qh(p)) \end{aligned} \quad (30)$$

See also: KPD_supplementary_material_1.ipynb (section.2.2)

3.3.3 Indifference spread and utility for probabilistic sniping

From the discussion above we know that for probabilistic sniping the end-points of the linear utility functions are replaced by their probabilistic counterparts:

$$A \rightarrow A(p), \quad B \rightarrow B(p), \quad C \rightarrow C(p), \quad \text{and} \quad D \rightarrow D(p).$$

Using the same logic as before we conclude that the point of indifference for probabilistic sniping is given by (cf. eqs. 13):

$$s^*(p) = \frac{A(p) - C(p)}{(A(p) - C(p)) + (D(p) - B(p))} \quad (31)$$

$$u^*(p) = \frac{A(p)D(p) - B(p)C(p)}{(A(p) - C(p)) + (D(p) - B(p))} \quad (32)$$

For ease of reference, we will denote the common denominator by

$$Q(p) := A(p) - C(p) + D(p) - B(p), \quad (33)$$

and introduce

$$N(p) := A(p)D(p) - B(p)C(p). \quad (34)$$

Hence, we arrive at the useful shorthand

$$u^*(p) = \frac{N(p)}{Q(p)}. \quad (35)$$

See also: `KPD_supplementary_material_1.ipynb` (section.4)

3.3.4 Impact of risk aversion and probabilistic sniping on utilities

Since the point of indifference is completely determined by the utility endpoints (A, B, C and D), it is helpful to make their dependence on the parameters γ (risk aversion) and p (individual sniping probability) explicit. To this end we hark back to eqs. (25) and (30), from which we can compute:

$$\begin{aligned} \frac{dA}{d\gamma} &= 0, & \frac{dB}{d\gamma} &= -\bar{\alpha}\beta p g(p) < 0, \\ \frac{dC}{d\gamma} &= -(\bar{\theta} + \beta(1 - 2\bar{\mu})h(p)) < 0, & \frac{dD}{d\gamma} &= -\bar{\alpha}\beta h(p) < 0. \end{aligned} \quad (36)$$

It therefore transpires that increasing risk aversion pushes the utility endpoints (except for A) down, dragging the utility u^* along. We will return to this observation in section 4 where we will show how increasing risk-aversion will induce transitions in sniping behaviour.

The dependence of the utilities on the sniping probability p can be classified in a similar fashion:

$$\begin{aligned} \frac{dA}{dp} &= m\beta(g(p) + pg'(p)) > 0, & \frac{dB}{dp} &= -\bar{\alpha}q\beta(g(p) + pg'(p)) < 0 \\ \frac{dC}{dp} &= -m\beta h'(p) < 0, & \frac{dD}{dp} &= -\bar{\alpha}\beta qh'(p) < 0. \end{aligned} \quad (37)$$

Notice that since $g(p) + pg'(p) = (pg(p))' = h'(p)/(H - 1) > 0$, it follows that (see Fig 3) probabilistic sniping (i.e. reducing p) will pivot the bandit's utility towards smaller values about the fixed intersection point s^{**} . Similarly, since both $dC/dp < 0, dD/dp < 0$ the utility for the market maker (unsurprisingly) shifts upwards when p is lowered and the bandits snipe only occasionally.

See also: KPD_supplementary_material_1.ipynb (section.3)

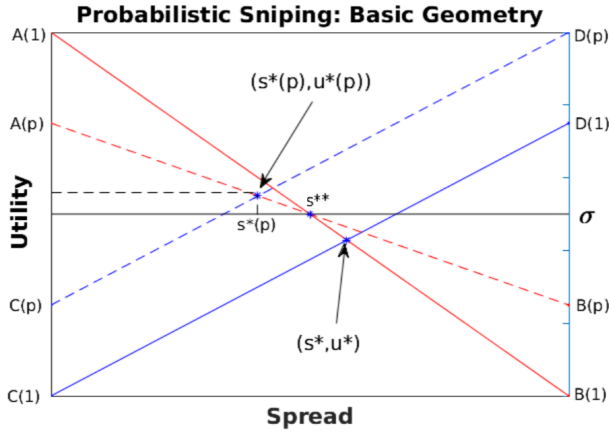


Fig. 3 Basic geometry of probabilistic sniping: The solid lines (AB for the bandit, and CD for the market maker) represent the expected utilities (as function of the spread s) in the case of sure sniping. The resulting point of indifference (s^*, u^*) is non-playable since $u^* < 0$. Reducing the probability of sniping to $p < 1$ pushes the MM utility function CD upwards (as the MM benefits from less sniping), while the bandit's utility pivots about its zero crossing. As a consequence, the point of indifference moves to the new intersection point $(s^*(p), u^*(p))$ which is now playable as $u^*(p) > 0$.

3.4 When is probabilistic sniping advantageous?

3.4.1 Geometric interpretation: Probabilistic sniping shifts point of indifference

Is there an advantage in switching from sure sniping to probabilistic sniping? Obviously, less sniping is advantageous for the market maker as his pay-off is reduced by sniping. In section 3.1 we argued that less sniping might also protect bandits

from acquiring assets whose values are unstable (e.g. due to subsequent emergence of contradictory news). Some further intuition is gleaned from Fig. 4, which illustrates the difference between sure and probabilistic sniping. More specifically, let us assume that sure sniping results in an equilibrium E such that $u^* := u(s^*) = 0$. Sure sniping ($p = 1$) is represented by the two blue lines (AB and CD, expected utility for bandit and market maker, respectively). Changing to probabilistic sniping ($p < 1$) results in two different lines ($A'B'$ and $C'D'$, respectively). Notice that $A'B'$ pivots about the equilibrium point E , while $C'D'$ shifts upwards (as less sniping means less risk for the market maker). As a consequence, the new intersection point E' will yield a (strictly) positive utility.

Another way to put this is that the derivative du^*/dp (evaluated at $p = 1$) is negative: smaller values for p result in higher utility (for a concrete, numerical example, see Fig 5 below). In section 4 below, we will recast this geometric insight into an algebraic expression, but first we will show that this new equilibrium is not a Nash equilibrium for the one-shot MZ¹.




Fig. 4 LEFT: Geometric representation that illustrates why probabilistic sniping is advantageous when $u^* = 0$. The blue lines AB (bandit) and CD (market maker) represent the expected utilities under sure sniping yielding a point of indifference E corresponding to zero utility $u^*(p = 1) = 0$. Probabilistic sniping (with sniping probability $p < 1$) shift these utilities to $A'B'$ and $C'D'$ respectively, resulting in a new point of indifference E' that corresponds to strictly better utility $u^*(p < 1) > 0$. RIGHT: Configuration in which probabilistic sniping results in lower expected utility and sure sniping is to be preferred. See main text for more information.

3.4.2 Does $(s^*(p), u^*(p))$ constitute a Nash equilibrium of the MZ¹ game?

One of the reasons why for $p < 1$ the new intersection point $(s^*(p), u^*(p))$ is not considered in the MZ paper (even if $u^*(p) > 0$), is that it does not constitute a Nash equilibrium for the (single-shot) MZ game. To see this, it suffices to realise that when $p < 1$, a bandit will always benefit from unilaterally deviating to sure sniping (assuming $u^*(p) \equiv U_B(s^*(p)) > 0$). As a consequence, these intersection points are uninteresting in terms of the single-shot version of the MZ game.

However, the reason that probabilistic sniping (and the corresponding role of $s^*(p)$ and $u^*(p)$) is of interest nonetheless, is that things change when we will consider (in section 5) the infinite horizon repeated game version (MZ[∞]). The

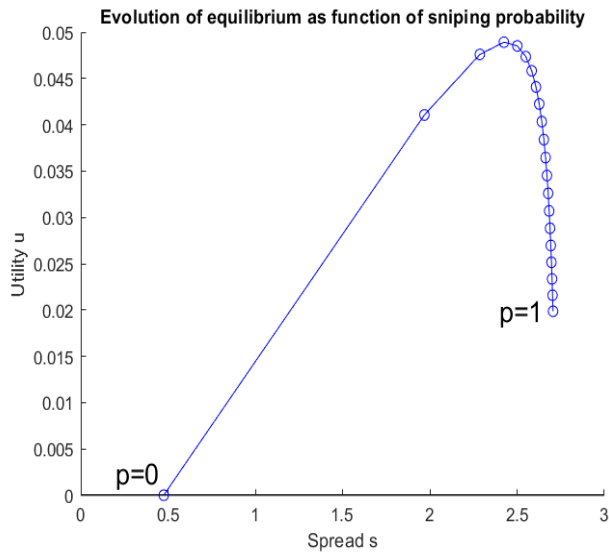


Fig. 5 An example where probabilistic sniping is advantageous even though sure sniping results in a strictly positive utility. In this example, we depict the evolution of the point of indifference as a function of the sniping probability p , i.e. $(s^*(p), u^*(p))$, where p starts at $p = 1$ and is decremented with steps of 0.05 to $p = 0$. Notice how the expected utility rises from 0.02 for sure sniping ($p = 1$) to a maximum of 0.05 for $p \approx 0.15$, before dropping to zero when there is no sniping ($p = 0$).

promise of higher pay-offs means that even purely selfish agents are incentivised to collaborate if they expect this cooperation to be profitable. This new equilibrium is stabilized by the implicit threat that, if anyone breaks the tacit plan, opponents will do the same and opportunities will instantly evaporate. Working out whether cooperation is the rational choice involves estimating what is most favourable: the exceptional but one-off payoff one stands to gain from defecting, or the smaller but accumulating payoffs that result from continuing collaboration that over time, will eclipse the former. For more details, we refer to section 5.

4 Characterising the transitions in sniping behaviour

In this section we use a geometric argument to deduce the general conditions under which, first, the transition from sure to probabilistic sniping occurs, and second, probabilistic sniping ceases to be profitable (i.e. transition from probabilistic to no sniping). We will then show how these general conditions can be related to risk aversion γ and how they give rise to two thresholds ($\bar{\gamma}_K$ and $\bar{\gamma}_L$) that govern these transitions.

4.1 Transition from pure to probabilistic sniping

Geometric interpretation As argued above, the equilibrium (s^*, u^*) can be seen as a function of the sniping probability p , i.e. $(s^*(p), u^*(p))$. To decide whether probabilistic sniping is advantageous, we need to determine when the slope of the tangent to the curve $u^*(p)$ at $p = 1$ changes sign (see Fig. 6):

$$\left. \frac{du^*(p)}{dp} \right|_{p=1} > 0 \text{ (sure sniping)} \quad \longrightarrow \quad \left. \frac{du^*(p)}{dp} \right|_{p=1} < 0 \text{ (probabilistic sniping)} \quad (38)$$

Indeed, if the slope of this tangent is positive, it means that reducing p from 1 (sure sniping) to a lower value $p < 1$ decreases the value of $u^*(p)$. Hence, sure sniping is better than probabilistic sniping. Conversely, if the slope of the tangent is negative, moving from sure to probabilistic sniping (i.e. reducing p) does improve $u^*(p)$ and therefore probabilistic sniping is to be preferred.

Using the notation introduced in eq. (35) we know that $u^*(p) = N(p)/Q(p)$, and hence

$$\left. \frac{du^*(p)}{dp} \right|_{p=1} = \frac{N'(1)Q(1) - N(1)Q'(1)}{Q^2(1)} \quad (39)$$

Hence the threshold is determined by the equation:

$$\left. \frac{du^*(p)}{dp} \right|_{p=1} = 0 \quad \iff \quad N'(1)Q(1) - N(1)Q'(1) = 0 \quad (40)$$

See : KPD_supplementary_material_1.ipynb (section.5)

4.2 Transition from probabilistic to cessation of sniping (non-sniping)

If the expected utility associated with the sure sniping equilibrium is sufficiently negative (i.e. $u^*(s^*, p = 1) \ll 0$), even probabilistic sniping will not be able to turn this into profits. Geometrically, this corresponds to a situation in which the path traced by $(s^*(p), u^*(p))$ will never enter (strictly) positive territory before it hits the x-axis at $p = 0$ (see Fig. 6 bottom-right panel).

This transition can therefore be characterised by a condition which is analogous to the characterisation of the transition *from pure to probabilistic sniping* in eq.(38), this time however focusing on the tangent at $p = 0$:

$$\left. \frac{du^*(p)}{dp} \right|_{p=0} > 0 \text{ (probabilistic sniping)} \quad \longrightarrow \quad \left. \frac{du^*(p)}{dp} \right|_{p=0} < 0 \text{ (no sniping)} \quad (41)$$

This transition therefore occurs when

$$\left. \frac{du^*(p)}{dp} \right|_{p=0} = 0 \quad (42)$$

Expanding

$$\left. \frac{du^*(p)}{dp} \right|_{p=0} = \frac{N'(0)Q(0) - N(0)Q'(0)}{Q^2(0)}$$

and using that $N(0) = 0$ while $Q(0) = D - C > 0$, we see that eq. (42) simplifies to:

$$N'(0) = 0. \quad (43)$$

Expanding equation (43):

$$N'(0) = A'(0)D(0) + A(0)D'(0) - B'(0)C(0) - B(0)C'(0)$$

and using $A(0) = B(0) = 0$, condition (43) can be further simplified to:

$$N'(0) \equiv A(0)D'(0) - B'(0)C(0) = 0. \quad (44)$$

See : KPD_supplementary_material_1.ipynb (section.6)

4.3 Increasing risk aversion induces transitions

In section 3.3.4 we showed how the endpoints of the utilities (except for A) are pushed down by increasing risk aversion γ , dragging the point of indifference (s^*, u^*) down in their wake. Since u^* decreases, there will be a value for γ beyond which probabilistic sniping becomes more favourable than sure sniping. Increasing γ even further will eventually result in a utility that is so low that even probabilistic sniping is no longer profitable. Hence we see that increasing risk aversion γ is one way to induce the two transitions in sniping behaviour discussed above. In this section we will translate the general transition conditions (40) and (44) into precise transition thresholds for γ .

4.3.1 Threshold $\bar{\gamma}_K$: Transition from sure to probabilistic sniping

Both N and Q depend on risk aversion γ and expanding the transition condition eq. (40) as a function of γ yields a cubic polynomial (denoted by $K(\gamma)$), for which we need to find a zero-crossing:

$$N'(1)Q(1) - N(1)Q'(1) \equiv \underbrace{K_3\gamma^3 + K_2\gamma^2 + K_1\gamma + K_0}_{:=K(\gamma)} = 0 \quad (45)$$

The computation of the polynomial coefficients is cumbersome and uninteresting and is therefore relegated to the supplementary material.

See: KPD_Supplementary_material_1.ipynb (section.5.1)

In terms of this notation, the threshold condition in eq. (40) amounts to finding the appropriate zero-crossing $K(\gamma) = 0$. Recall that a cubic equation has at least one solution (zero-crossing). However, we need to make sure that this solution satisfies $\gamma \geq 1$ and can therefore be interpreted as a risk aversion factor. This is proven in the following theorem in which we also define the transition threshold:

Proposition 1 (Risk aversion threshold for transition from pure to probabilistic sniping)

The cubic polynomial $K(\gamma)$ specified in eq. (45) has a unique zero-crossing greater than 1. We will denote this unique zero-crossing by $\bar{\gamma}_K \geq 1$ as it is the threshold that governs the transition from sure to probabilistic sniping.

The proof of this proposition can be found in appendix F.




Fig. 6 Examples of transition thresholds in sniping behaviour. In the above example we fix the parameters $\alpha = .45$, $\mu = .5$, $\delta = .5$, $H = 5$. From these values we can compute the threshold values $\bar{\gamma}_K = 2.515$ and $\bar{\gamma}_L = 7.831$ (for a definition of $\bar{\gamma}_K$ and $\bar{\gamma}_L$ see section 4). Next we set γ to four different values: $\gamma = 1.7575$ (slightly risk-averse, sure sniping), $\gamma = \bar{\gamma}_K = 2.5150$ (risk averse, knife edge), $\gamma = 3.5$ (probabilistic sniping), and finally, $\gamma = \bar{\gamma}_L = 7.8313$ (strongly risk-averse, no more sniping). Notice that in the bottom left picture probabilistic sniping yields a better utility, even though the utility for sure sniping is positive!

See also: `KPD_supplementary_material_numeric.ipynb` (section.3.4)

4.3.2 Threshold $\bar{\gamma}_L$: Transition from probabilistic sniping to non-sniping

A similar approach can be taken to quantify how a further increase in risk aversion γ gives rise to the threshold $\bar{\gamma}_L$ beyond which even probabilistic sniping is no longer profitable, i.e. both sure and probabilistic sniping results in negative utility $u^*(s^*, p) < 0$, $(\forall 0 \leq p \leq 1)$.

In order to determine this threshold value, we expand transition condition eq.(44) as a function of γ and obtain a quadratic polynomial (denoted by $L(\gamma)$) for which a zero-crossing needs to be found:

$$N'(0) = 0 \quad \iff \quad L(\gamma) := L_2(\gamma - 1)^2 + L_0 = 0, \quad (46)$$

which can easily be solved explicitly. We can summarize all of these observations in the following proposition.

Proposition 2 (Risk aversion threshold for transition from probabilistic to non-sniping)

There is a unique threshold $\bar{\gamma}_L$ for risk aversion γ that governs the transition from probabilistic to no sniping. Specifically:

$$\bar{\gamma}_L = 1 + \sqrt{\frac{(1 - \bar{\mu})Z}{\bar{\alpha}\bar{\theta}}} \quad \text{where} \quad Z = 1 + \bar{\mu} - \beta(1 - \bar{\mu}). \quad (47)$$

The proof of this proposition is straightforward and can be found in Appendix section G.

See: KPD_supplementary_material_1.ipynb (section.6)

4.4 Optimal sniping probability p_K^*

In the preceding sections we have shown that increasing risk aversion ($\gamma > 1$) creates multiple sniping regimes for the bandits:

- $1 \leq \gamma < \bar{\gamma}_K$: Sure sniping is most profitable; This situation corresponds to the top left panel in Fig. 6.
- $\bar{\gamma}_K \leq \gamma < \bar{\gamma}_L$: Probabilistic sniping results in higher utilities; in fact, there is an optimal sniping probability p_K^* (see below) which results in a maximal (positive) utility. This situation corresponds to the bottom left panel in Fig. 6.
- $\gamma \geq \bar{\gamma}_L$: No sniping (resulting in zero utility) is the best option. This situation corresponds to the bottom right panel in Fig. 6.

The situation in the second case $\bar{\gamma}_K < \gamma < \bar{\gamma}_L$ is schematically illustrated in Fig. 8. It follows that there is an optimal sniping probability that yields the largest utility:

$$p_K^* := \arg \max_p u^*(p) \quad (48)$$

We introduce the following notation for the corresponding utility u_K^* and spread s_K^* . Since this is a unique optimal point it provides a natural **focal point** for the game [19]. In addition, Fig. 9 compares the expected utilities under both sure and (optimal) probabilistic sniping as a function of risk aversion γ , showing clearly that profitability can be sustained for higher values of γ .

See: u_depend_on_gamma_p_v2.m




Fig. 7 These figures show the evolution of transition thresholds $\bar{\gamma}_K$ (blue) and $\bar{\gamma}_L$ (red) as functions of the rates for the trigger event: arrival rate of news (α , TOP) and arrival rate of LT (μ , bottom). Notice (top) how increasing the news rate α results a faster transition to probabilistic sniping (as explained in section 3.1).

5 Optimal probabilistic sniping as SPE in repeated MZ^∞ game

In the previous section we have shown that if the risk aversion factor γ satisfies $\bar{\gamma}_K < \gamma < \bar{\gamma}_L$, all bandits are better off if they implement probabilistic sniping, i.e. engage in racing with probability p_K^* only. However, this is not a Nash equilibrium for the stage game, as bandits will be tempted to snipe more frequently. This is a situation reminiscent of the classical prisoner's dilemma (PD) where the Nash




Fig. 8 Definition of p_K^* : When probabilistic sniping is advantageous (i.e. $\bar{\gamma}_K < \gamma < \bar{\gamma}_L$) reducing the sniping probability p shifts the equilibrium position $(s^*(p), u^*(p))$ along a concave arc, the top of which corresponds to the optimal sniping probability p_K^* .

equilibrium forces the players into a strategy that results in the worst social welfare possible [12].

However, it is well-known that if we repeat the prisoner’s dilemma an indefinite number of times, it becomes possible to devise strategies that do yield the optimal utility. This reversal of fortune is due to the fact that agents can build up reputations that encourage collaboration that is guaranteed to be profitable for all parties. Exactly the same rationale will also play out in the (infinite horizon) repeated MZ-game (MZ^∞).

5.1 The MZ stage game as a sequential game

Game Tree Before we explore the strategies that are available in the (indefinitely) repeated MZ game (a.k.a. MZ^∞), we need to pause for a minute and clarify the nature of the strategies that are available to the players in the single-shot stage game (a.k.a. MZ^1). To this end, it is helpful to realise that the stage game involves different interdependent steps and can therefore be seen as a sequential game involving a (possibly) random move by “nature”. Below, we give an explicit representation of the decision tree for the game and clarify a subtle but important point in the interpretation of probabilistic sniping.

In section 2.2, we already outlined the detailed chronology of the various defining events in the game. For the interpretation of the MZ game as a sequential game it is also helpful to highlight the various steps for an individual HFT in the corresponding decision tree (also see Fig. 10):

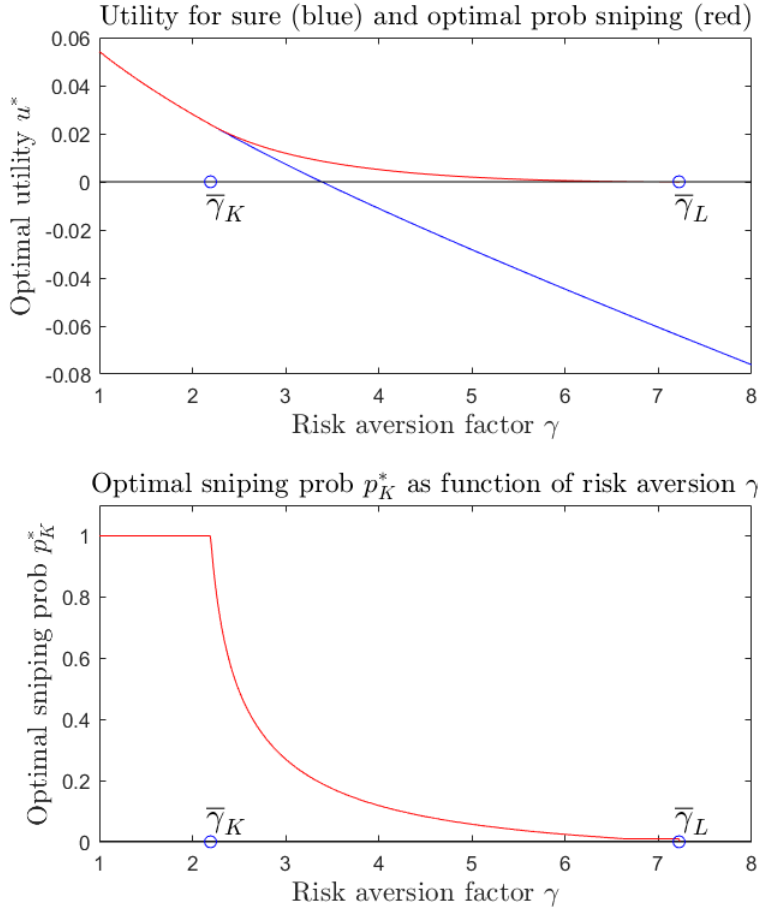


Fig. 9 TOP: The figure shows the equilibrium utility spread as a function of risk aversion in both sure sniping ($u_K^*(p = 1)$ blue line) and optimal probabilistic sniping ($u_K^* = u(p = p_K^*)$, red line). BOTTOM: corresponding optimal sniping probability p_K^* , which drops from $p_K^* = 1$ (sure sniping) when $\gamma \leq \bar{\gamma}_K$ to 0 (no sniping) when $\gamma \geq \bar{\gamma}_L$.

1. HFT (at time $t = -1$) has to choose a spread s from the continuous action space $0 \leq s \leq \sigma$. Alternatively, he can refrain from playing altogether (NULL action resulting in zero utility).
2. Nature (at time $t = -1/2$) selects one of the HFTs as MM (using the rule specified earlier in section 2.2), relegating the others to the role of bandit;
3. At $t = 0$, each HFT knows its type (market maker or bandit) and can therefore make his reaction to the trigger event (news or LT arrival) contingent on his type. This results in a new bifurcation of the decision tree:

- if he's bandit: he has to choose between quitting the game (NULL-action with utility 0), or some form of sniping: *sure* $p = 1$, *no* ($p = 0$) or *probabilistic* sniping ($0 < p < 1$) resulting in a utility $u_B(s, p)$.
- if he's MM: he either proceeds and the game ends with utility $u_M(s, p)$ (which actually also depends on p), or he withdraws from the game (NULL with utility 0):

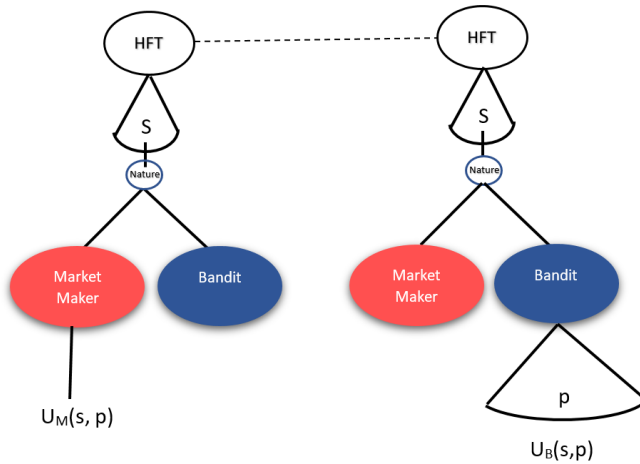


Fig. 10 The MZ game as a sequential game against nature. Notice that in the last decision point, the MM is uncertain about his utility $u_M(s, p)$ as it depends on the bandit's sniping probability which at the moment of decision is unknown to him!

Actions and Strategies Disregarding the trivial NULL action, what actions and strategies are available to the players? At the first decision point, it is clear that each HFT can choose a spread s , which corresponds to a **pure but continuous action**.

At the second HFT decision point, the HFTs know their type. The strategy for the market maker is straightforward: if the trigger event is the publication of a *news item*, the intrinsic value of the asset jumps and he will race to cancel his stale quotes. The choice of strategies for the other HFTs (bandits) is more subtle. Indeed, in the MZ paper [16], the authors describe mixed sniping as randomly mixing between the **pure and discrete strategies** of *sniping* and *not sniping*. As a consequence, mixed sniping will only feature as a Nash equilibrium when the bandits are indifferent between sniping and not sniping! In our interpretation, we think of each bandit selecting a sniping probability p as him choosing a *pure strategy* from a **continuous action space**. This allows us to introduce probabilistic sniping in situations where the expected utilities at the extremes (i.e. $p = 0$ or 1) are not necessarily the same.

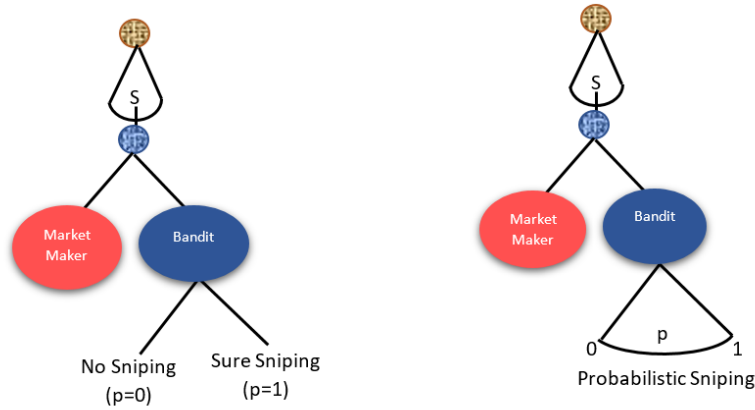


Fig. 11 Interpretation of the sniping probability. **Left:** In the MZ paper, the pure actions available to the bandit are *sniping* and *not sniping*. Mixed or (probabilistic) sniping is therefore only possible when the bandit is indifferent between the pure actions. **Right:** In this paper, we interpret the sniping probability as a *pure but continuous* action that the bandit can take. In this interpretation the bandit need not be indifferent between the extremes in order to adopt a probabilistic strategy.

5.2 Subgame-Perfect Equilibria (SPE) in repeated games

5.2.1 The emergence of collaboration in repeated games

In the sections above we identified the risk aversion threshold $\bar{\gamma}_K$ above which probabilistic sniping has the potential of yielding better utilities for both the market maker and the bandits. When $\gamma > \bar{\gamma}_K$ there is a corresponding optimal sniping probability p_K^* that yields the most favourable outcome for all parties. The problem with this strategy is that it does not constitute a Nash equilibrium for the single-shot MZ¹ game as bandits will be tempted to snipe more often than is allowed (see section 3.4). However, things are different when we consider the infinite horizon repeated game version (MZ[∞]). Indeed in this case, we can invoke the so-called *folk theorems* that show that any equilibrium that is strictly better than a Nash equilibrium in the single-shot (stage) game gives rise to a new Nash equilibrium (NE) in the repeated game (at least, under the mild assumption that all the players are sufficiently patient) [1].

However, for this to become a viable strategy we need to make the threat of rescinding the cooperation implicit in probabilistic sniping, credible by devising a method to detect non-compliance among the other agents. This is non-trivial since even deceptive agents would still advertise spread s_K^* , as spreads are publicly observable, and deviating from this would immediately flag their non-compliance and trigger a collapse of the (implicit) agreement. Deceptive agents would therefore stick to spread s_K^* but snipe for sure, rather than with probability p_K^* . However, the sniping behaviour of agents is hidden to the other agents as it is not publicly known who entered the race and who won it. An agent only has access to its own utility (aggregated over many repetitions of the stage game). So each agent will

have to base his decision (on whether or not to continue with the game) on the impact that sniping by the deceptive agents has on his own utility.

5.2.2 Impact of deceptive sniping

When $\bar{\gamma}_K < \gamma < \bar{\gamma}_L$ and all agents snipe with optimal probability p_K^* they stand to earn $u_K^* > 0$ (on average). If, however, there are rogue agents that engage in sniping at every possible occasion, then the utility for the trustworthy agents will go down for two reasons:

1. as market makers they will suffer from more intense sniping;
2. as bandits they will have to race against an (on average) larger group of competing bandits.

The combination of these two effects will affect their outcomes in various but related ways: It will reduce the number of times they succeed in winning the race, and this in turn will push down their aggregated utility.

This is illustrated in Fig. 12 where the blue error-bars show the average utility (averaged over 10000 simulations of the stage game) when *all* agents snipe probabilistically (with probability p_K^*). These results are clearly consistent with the theoretically predicted value $u_K^* = 0.015$ (blue dotted line). However, things change significantly when there is even a single devious agent (HFT 5 in the figure) that snipes for sure (red error-bars). The utility for this deceptive agent more than doubles while the expected utility for the trustworthy agents (1 through 4) dips below zero ($u = -0.0016$, red dotted line).

See: `KPD_supplementary_material_game_simulations.ipynb` (section 7.2)

Observing this unexpected downturn in utility an agent can suspect (some of) the other agents are not complying, and consequently decide to renege on his earlier commitment to the probabilistic sniping strategy. To turn this intuition into a reliable operational strategy, we need to quantify the effect that deceptive agents have on the expected utility of the trustworthy agents. This is taken up in the next section.

5.3 Expected utility in the presence of deceptive agents

5.3.1 Generalising probabilities $h(p)$ and $g(p)$ when there are deceptive agents

To evaluate the expected utility of a specific trustworthy agent amidst deceptive agents, we consider the case where the total of H agents is divided into H_t trustworthy agents (who stick to probabilistic sniping) and H_d deceptive agents who pretend to restrict their sniping but in reality enter the sniping race at every possible occasion (obviously, $H_t + H_d = H$). We can now build on the earlier exposition in section 3.3 where, in eqs. (25) and (30), we expressed the expected utilities in terms of the probabilities $h(p)$ and $g(p)$. Closer inspection of these equations shows that only these probability functions need to be modified in the presence of

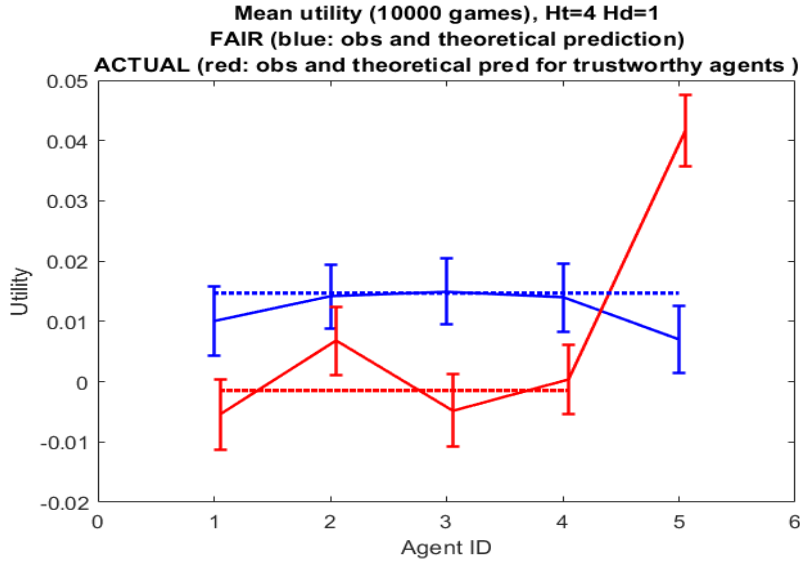


Fig. 12 Mean utility for agent group of $H_t = 4$ trustworthy agents (agents 1 through 4) and $H_d = 1$ deceptive agent (agent 5). The mean and error-bars (± 1 standard error) show the results for a simulation of a large number of stage games, while the dotted lines indicate the theoretically predicted value. When everyone adheres to probabilistic sniping, the mean is the same for all agents. When the deceptive agent snipes for sure, the theoretical mean (dotted red line) corresponds to the expected utility for the trustworthy agents (1 through 4). The deceptive agent (agent 5) is clearly able to increase his gains at the expense of the trustworthy agents.

deceptive agents as the other factors only depend on the game parameters. We therefore introduce the following notation which extends the probabilities $h(p)$ and $g(p)$:

- $h_{tm}(p, H_t, H_d)$ is the probability that a specific individual **trustworthy** agent, who is acting as **market maker**, will **lose** the MZ-race given that the H_d deceptive agents snipe for sure, while the remaining $H_t - 1$ trustworthy agents enter the sniping race with probability p (denoted as: wp. p):

$$h_{tm}(p, H_t, H_d) := P\{\text{MM will lose race} \mid H_d \text{ agents snipe for sure, } \dots \\ \dots \text{ while } (H_t - 1) \text{ remaining trustworthy agents snipe w.p. } p\} \quad (49)$$

In addition to the market maker, the number of agents in the race is H_d (all the deceptive agents snipe for sure) plus a sample (of stochastic size $N_t \sim \text{Bin}(H_t - 1, p)$) of the remaining $H_t - 1$ trustworthy agents who snipe with probability p . Hence:

$$h_{tm}(p, H_t, H_d) = \mathbb{E} \left(\frac{H_d + N_t}{1 + H_d + N_t} \right) \quad \text{where} \quad N_t \sim \text{Bin}(H_t - 1, p) \quad (50)$$

- $g_{tb}(p, H_t, H_d)$ is the conditional probability that a specific **trustworthy agent**, acting as **bandit**, **wins** the sniping race, given that *he actually enters the race*, and H_d deceptive agents snipe for sure, while the remaining $H_t - 1$ trustworthy agents enter the race with probability p :

$$g_{tb}(p, H_t, H_d) := P \{ \text{agent will win race} \mid \begin{array}{l} \text{he enters the race as bandit, } \dots \\ \dots \text{ while the remaining } H_t - 1 \text{ trustworthy agents snipe w.p. } p, \dots \\ \dots \text{ and the } H_d \text{ deceptive agents snipe for sure} \end{array} \} \quad (51)$$

A similar reasoning implies that $g_{tb}(p)$ is a weighted average of the following two terms:

- if the market maker happens to be selected from the group of trustworthy agents (probability $(H_t - 1)/(H - 1)$) then the list of agents in the race is composed of: the agent under scrutiny, the market maker, all the deceptive agents (H_d) and a random selection of the remaining $H_t - 2$ trustworthy agents;
- if the market maker happens to be selected from the group of deceptive agents (probability $H_d/(H - 1)$), the list of agents in the race comprises, in addition to the trustworthy agent under scrutiny, the market maker, all the remaining deceptive agents ($H_d - 1$), as well as a random selection from the remaining $H_t - 1$ trustworthy agents.

Recasting this logic in mathematical parlance, we arrive at the following expression:

$$\begin{aligned} g_{tb}(p, H_t, H_d) &= \left(\frac{H_t - 1}{H - 1} \right) \mathbb{E} \left(\frac{1}{2 + H_d + N_t^{(m)}} \right) + \left(\frac{H_d}{H - 1} \right) \mathbb{E} \left(\frac{1}{2 + (H_d - 1) + N_t^{(b)}} \right) \\ &= \left(\frac{H_t - 1}{H - 1} \right) \mathbb{E} \left(\frac{1}{2 + H_d + N_t^{(m)}} \right) + \left(\frac{H_d}{H - 1} \right) \mathbb{E} \left(\frac{1}{1 + H_d + N_t^{(b)}} \right) \\ &\quad \text{where } N_t^{(m)} \sim \text{Bin}(H_t - 2; p) \quad \text{and} \quad N_t^{(b)} \sim \text{Bin}(H_t - 1; p) \end{aligned} \quad (52)$$

(For the sake of completeness we add that the binomial random variables are independent).

An alternative characterisation of g_{tb} which, from a computational point of view, is slightly more convenient is presented in appendix H.

See: `KPD_Supplementary_material_2.ipynb` (section.3)

5.3.2 Non-compliance detection based on sequential utility testing

Rather than basing the non-compliance test on the aggregated utility (which is a rather crude measure), we monitor the rate of occurrence of each individual utility outcome. To this end, we utilise the fact that since the game parameters

are common knowledge, each agent can use Table 1 to compute the values of the $J = 9$ different possible utilities. In addition, combining the event probabilities in that table with the generalisation h_{tm} and g_{tb} specified in eqs. (50) and (52), he will be able to compute the corresponding probability of each utility outcome, both under the assumption of compliance ($H_d = 0$) and non-compliance ($H_d > 0$). These probabilities are listed in Table 2.

utility	probability
0	$\frac{1}{H} \{\beta(1 - 2\bar{\mu})(1 - h_{tm}(p))\} + \frac{H-1}{H} \{1 - \beta p g_{tb}(p)(1 - \bar{\mu})\}$
$-\gamma(2\sigma - s)$	$\bar{\alpha}\beta \cdot \left(\frac{1}{H}\right) \cdot h_{tm}$
s	$\left(\frac{1}{H}\right) \{\bar{\alpha}\beta h_{tm}(p) + (1 - \beta)(1 - 2\bar{\alpha} - \bar{\mu})\}$
$-\gamma(\sigma - s)$	$\left(\frac{1}{H}\right) \{\bar{\alpha}(1 - \beta) + \beta h_{tm}(p)(1 - 2(\bar{\alpha} + \bar{\mu})) + \beta\bar{\mu}\}$
$s + \sigma$	$\left(\frac{1}{H}\right) \{\beta\bar{\mu}(1 - h_{tm}(p) + \bar{\alpha}(1 - \beta))\}$
$2s$	$\bar{\mu} \cdot \left(\frac{1}{H}\right) \{\beta h_{tm}(p) + (1 - \beta)\}$
$2\sigma - s$	$\frac{H-1}{H} \{\bar{\alpha}\beta p g_{tb}(p)\}$
$\sigma - s$	$\left(\frac{H-1}{H}\right) \cdot \beta p g_{tb}(p) \{1 - 2\bar{\alpha} - \bar{\mu}\}$
$-\gamma s$	$\left(\frac{H-1}{H}\right) \cdot \frac{\beta}{2} \cdot \bar{\alpha} p g_{tb}(p)$

Table 2 Overview of unique utilities and corresponding probabilities (for a trustworthy agent).

So, in essence, we have a situation in which each stage game produces a utility outcome u_j ($j = 1, 2, \dots, J$) with a known probability distribution $\pi^0 = (\pi_1^0, \dots, \pi_J^0)$ under the assumption of overall compliance (null-hypothesis of fair behaviour with everyone sniping probabilistically at the optimal rate p_K^*) and an alternative known distribution $\pi^1 = (\pi_1^1, \dots, \pi_J^1)$ that applies to a trustworthy agent, assuming that there are a specific number $H_d > 0$ of deceptive agents that snipe for sure (see Appendix J). This is illustrated in Fig 13 where we have plotted the probability of occurrence of the nine possible outcomes under two different regimes (with and without deceptive agents).

When, after stage game t , an individual trustworthy agent receives utility u_t , he can compute the likelihood of this outcome under the two alternative hypotheses i.e. without (null hypothesis \mathcal{H}_0) or with (alternative hypothesis \mathcal{H}_1) deceptive snipers, and determine which one is more likely:

$$P(u_t \mid \mathcal{H}_0) \leq ?? \geq P(u_t \mid \mathcal{H}_1)$$

For illustrative purposes we look at a concrete example, See, `KPD_supplementary_material_game_simulations.ipynb` (section 9)

Wald's sequential probability ratio test (SPRT) [21] provides a principled way of aggregating the evidence collected over a sequence u_1, u_2, \dots, u_t of utility

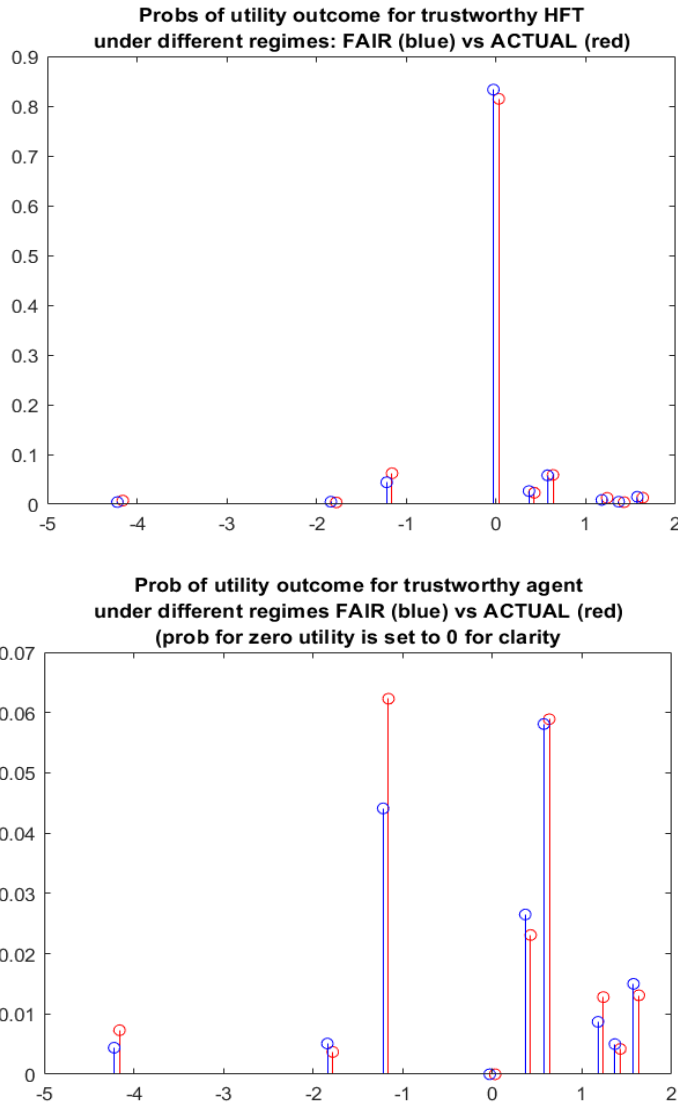


Fig. 13 Probability distribution (taken from Table 2) for the finite set of utility outcomes under two different regimes (BLUE: Fair ($H_d = 0$), RED: $H_d = 1$). Bottom: Same figure as above, but probability of zero utility has been suppressed for reasons of clarity.

observations. More specifically, at every time (stage game) t we construct the following test statistic:

$$S(t) = \sum_{\tau=1}^t \log \frac{P(u_\tau | \mathcal{H}_1)}{P(u_\tau | \mathcal{H}_0)}. \quad (53)$$

To implement the sequential test procedure we first need to compute the following two threshold values:

$$a := -\log\left(\frac{1-\alpha}{\beta}\right) \quad b := \log\left(\frac{1-\beta}{\alpha}\right) \quad (54)$$

where α is type I error (related to significance) and β is the type II error (related to power). We then get the following stopping rule:

- if $a < S_t < b$: continue testing;
- if $S_t < a$: accept \mathcal{H}_0 (all HFTs are trustworthy)
- if $S_t > b$: accept \mathcal{H}_1 (i.e. reject \mathcal{H}_0 , there is at least one devious HFT)

So, any trustworthy agent that is monitoring his accumulated utility will compose the Wald test-statistic as follows: at time (stage game) τ he will receive utility $u_{j(\tau)}$. Once this result is known, he will compute the probability ratio $\pi_{j(\tau)}^1/\pi_{j(\tau)}^0$ of this event under both hypotheses (fair \mathcal{H}_0 and deceptive \mathcal{H}_1) and then sum the logs of these ratios. He will then continue testing until hitting one of the two thresholds in eq.(54), whereupon he can draw a statistically significant conclusion regarding the compliance of the other agents. This is illustrated in Fig. 14 where the top row shows the fate of a trustworthy agent in a group of three additional agents one of which is cheating. The bottom row shows the fate of a trustworthy agent when each of his three colleagues is honouring the implicit contract of optimal probabilistic sniping.

See, `KPD_supplementary_material_game_simulations.ipynb` (section.10)

6 Related work

High-frequency trading Over the last decades, financial markets have changed significantly and become fragmented. Hence, traders can search across many markets, but this requires costly infrastructure and technology to make trading profitable [8]. In addition, high speed trading is now considered to be a crucial part of the trading technology. In 2010, the speed of round-trip between NASDAQ and the Chicago Mercantile Exchange decreased from over 14.5 milliseconds to under 8.1 milliseconds in 2014 [5]. Also NASDAQ supports high frequency traders by offering faster access to their infrastructure and trade data transmission [13]. In response to these drastic changes in financial markets, *algorithmic trading* harnesses advanced technologies to connect to different markets directly and trade ever faster at lower costs. Therefore, high-frequency traders (HFTs) can trade continuously and benefit from serial order processing. This continuous time trading has given rise to a speed race and sniping has emerged as an opportunity for HFTs to benefit from a stale quote [11]. In [2] it is shown that HFTs will tend to post thin but strictly positive bid-ask spreads to compensate the costs of getting sniped even without asymmetric information about fundamentals.

The emergence of high-speed trading has spurred on a lot of research regarding the effects on important market parameters such as liquidity. The seminal paper by

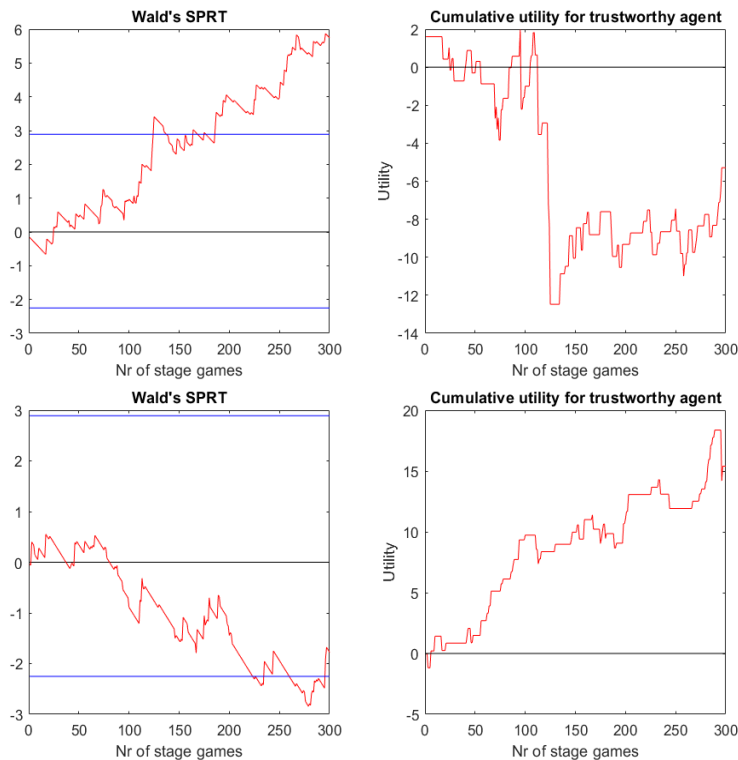


Fig. 14 TOP: Simulation results for a group of 4 agents, three of which are trustworthy and snipe with optimal probability p_K^* while one of them is deceptive and snipes for sure. The evolution of the utility of one of the trustworthy agents during the repeated game is plotted on the right, and the corresponding Wald SPRT statistic on the left. Around stage game 115 the Wald statistic hits the upper limit (b in eq. 54) indicating that the null hypothesis of FAIR game can be rejected. This would therefore be the moment the trustworthy agent would choose to quit the game. BOTTOM: Same statistic and data but this time for a group of four agents that are all abiding by implicit probabilistic sniping contract. In this case the Wald statistic hits the lower limit (a in eq. 54) around stage game 175, allowing the agent to accept the null hypothesis of FAIR game.

Menkveld and Zoican [16] has been the starting point for this research and has been discussed extensively in this paper. The focus in both this and the MZ paper has been on a stylized version of the interactions occurring at high frequency electronic exchanges. However, real market are competitive and continuous, and algorithmic traders need to choose their strategies from scratch and for a longer period in order to benefit from the dynamic of the market and speed (arms race). Moreover, traders should have a set of predefined strategies [6] and traders' behaviour may be influenced by observing their competitors [18]. Therefore, we explain agent's strategies and the reason why repeated game instead of one shot Nash equilibrium is a better choice to explain behaviours in the electric exchange. The case for repeated games is further strengthened by research by Bruce (2007) that shows

that a one-shot stage game brings significant surplus loss to the traders in a dynamic setting [4]. More recently, some studies showed that repetition in the game allows decision makers (agents, players in the market) to construct their strategies explicitly and benefit in the long-run [7,18].

Breitmoser(2015) analyzed individual strategies that are cooperate after mutual cooperation, defect after mutual defection, otherwise randomize. He finds a semi-grim equilibrium as a threshold that agents start cooperating in the first round by knowing treatment parameters and switch to semi-Grim with equal probability in the repeated game. At the end, sustain mutual cooperation led to the long-run welfare [3] and [10]. Furthermore, agents gain experience from the history of the game and trust is built [14] and an indefinitely repeated games with high continuation probability helps agents to mix strategies with minimal restrictions on the type and lengths of pure strategies [18]. While repeated game helps sustain cooperation, it is important to monitor the game in case of any defections. In 1947, Wald [21] develops a statistical test called Sequential probability ratio test (SPRT) based on likelihood ratio. This test is applied repeatedly during the sampling process and terminate whenever there is sufficient evidence in the data for one of the hypotheses of interest. In this paper, we use SPRT to detect non-compliance agent in different scenarios.

7 Conclusions and further research

In this paper we re-analysed the MZ game, a stylized version of a market interacting with a high frequency exchange as originally described in [16], and argued that a recasting of the problem in terms of repeated games offers a natural interpretation in which the solution space of equilibria accommodates genuine probabilistic sniping. Probabilistic sniping was introduced in the original paper to account for changes in the behavior of HFTs when risk aversion increases, but its practical significance in the single-shot game is limited. However, **in the context of repeated games** probabilistic sniping provides genuinely new and interesting insights. In particular:

- We outlined a simple geometric argument that allows one to deduce the general conditions that govern the transitions in sniping behaviour. By translating these conditions in terms of the risk aversion parameter γ we obtained two new thresholds ($\bar{\gamma}_K$ and $\bar{\gamma}_L$).
- Contrary to the situation in the single-shot game, we predict that probabilistic sniping will start playing a role even before the equilibrium utility for sure sniping is reduced to zero (by increasing risk aversion).
- We defined unfavourable scenario in which there are deceptive agents in the game. In comparison with favorable (fair) situation, the presence of non-compliant agents reduces trustworthy agent's expected utility. We showed how one can use Wald's SPRT to monitor the sequential game and flag when the utility for trustworthy agents deviates from the optimum point.

Some notes on extensions for further research In both this paper and the original MZ paper it was shown how simple and highly stylized models can already provide unexpected insights and interesting conclusions. We therefore expect that more realistic models will yield even more intriguing insights. Some of the challenges we intend to tackle in an upcoming paper are:

- We showed how important quantities that determine the qualitative transitions in the results (such as s_K^* or $\bar{\gamma}_K$) can be computed based on the knowledge of the game parameters, i.e. α, μ, δ and H . However, in most realistic situations, these parameters are unknown and need to be estimated from observations. As a result there is an amount of uncertainty and noise associated with these results. How stable are the conclusions with respect to noisy game parameters?
- In the repeated game scenario expounded in this paper, each stage game starts from zero position. It can be modified to a more realistic strategy, as the traders want to regress back to zero-positions.
- Risk aversion varies depending on position. One reason for this is that traders dread reporting bad news. Hence, when they find themselves in a negative position, they are willing to take more risk since, if they are lucky, they can redeem themselves, and if they are unlucky, the "size" of their bad news does not make all that much difference.

Acknowledgements The authors are very grateful to Profs. Menkveld and Zoican for their detailed feedback on earlier drafts of this paper.

References

1. Abreu, D, Dutta, P, K, & Smith, L, The folk theorem for repeated games: a NEU condition. *Econometrica*, 62, 939–939 (1994).
2. Budish, E, Cramton, P, & Shim, J, The high-frequency trading arms race: Frequent batch auctions as a market design response, *The Quarterly Journal of Economics*, 130, 1547-1547 (2015).
3. Breitmöser, Y, Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma, *The American Economic Review*, 105, 2882-2910 (2015).
4. Bruce Ian Carlin, Miguel Sousa Lobo, & Viswanathan, S, Episodic Liquidity Crises: Cooperative and Predatory Trading, *The Journal of Finance*, 62, 2235-2274 (2007).
5. Baldauf, M, & Mollner, J, High-Frequency Trading and Market Performance, conference proceeding (2014).
6. Cason, T, & Mui, V, Individual versus group choices of repeated game strategies: A strategy method approach, *Games and Economic Behavior*, 114, 128-145 (2019).
7. Dal Bo, P, & Frechette, G, R, Strategy Choice in the Infinitely Repeated Prisoners' Dilemma, WZB Berlin Social Science Center (2015).
8. Delaney, L, Investment in high-frequency trading technology: A real options approach, *European Journal of Operational Research*, 270, 375-385 (2018).
9. D, Fudenberg, & J, Tirole: *Game Theory*, 148, MIT Press (1991-2009).
10. Fudenberg, D, Rand, D, & Dreber, A, Slow to anger and fast to forgive: Cooperation in an uncertain world, *The American Economic Review*, 102, 720-749 (2012).
11. Hagströmer, B, & Nordé, L, The diversity of high-frequency traders, *Journal of Financial Markets*, 16, 741-770 (2013).
12. Holt, C A, Roth, A E, & Smith, V L, The nash equilibrium: a perspective. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 3999–4002 (2004).

13. Kauffman, R, Hu, Y, & Ma, D, Will high-frequency trading practices transform the financial markets in the asia pacific region?, *Financial Innovation*, 1, 1-27 (2015).
14. Kagel, J, & McGee, P, Team versus individual play in finitely repeated prisoner dilemma games, *American Economic Journal. Microeconomics*, 8, 253-276 (2016).
15. Laclau, M, A folk theorem for repeated games played on a network, *Games and Economic Behavior* 76, 711–37 (2012).
16. Menkveld, A, & Zoican, M, Need for speed? exchange latency and liquidity, *The Review of Financial Studies*, 30, 1188-1228 (2017).
17. Pinto, B, Repeated games and the reciprocal dumping model of trade, *Journal of International Economics*, 20, 357-366 (1986).
18. Romero, J, & Rosokha, Y, Constructing strategies in the indefinitely repeated prisoner's dilemma game *European Economic Review*, 104, 185-219 (2018).
19. Schelling, T, C, *The strategy of conflict*, Cambridge: Harvard University Press (1960).
20. Tomala, T, Fault reporting in partially known networks and folk theorems, *Operations Research*, 59, 754-763 (2011).
21. Wald, A, Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16, 117–186 (1945).
22. Link to KPD_2021 GitHub repository, <https://github.com/Karen-20/Karen.git> (2020)

Appendix: Supplementary Material

A List of notation and abbreviations

MZ or MZ ¹	Menkveld-Zoican (single-shot) stage game
MZ [∞]	infinite horizon repeated MZ game
HFT	High Frequency Trader, also called <i>agent</i>
LT	Liquidity trader (operates in background by hitting available quotes)
MM or M	market maker (selected from among the HFTs)
B or HFB	(high frequency) bandit (all HFTs that were not selected as market maker)
H	total number of HFT agents ($H = H_t + H_d$)
H_t	total number of <i>trustworthy</i> agents (snipe probabilistically)
H_d	total number of <i>deceptive</i> agents (snipe for sure)
v	intrinsic value of a (financial) asset
s	(half) spread ($0 \leq s \leq \sigma$), simply referred to as <i>spread</i>
p	sniping probability
α	arrival rate for news events (Poisson process)
σ	jump size upon news
μ	arrival rate for LT arrivals (Poisson process)
δ	exchange latency (duration of MZ race)
γ	risk aversion factor ($\gamma > 1$)
β	probability that trigger (i.e. first) event is news: $\beta = \alpha/(\alpha + \mu)$
$1 - \beta$	probability that trigger (i.e. first) event is LT arrival: $1 - \beta = \mu/(\alpha + \mu)$
$\bar{\alpha}$	(half) the expected number of news arrivals in interval $[0, \delta)$
$\bar{\mu}$	(half) the expected number of LT arrivals in interval $[0, \delta)$
m	$1 - \bar{\mu}$
$\bar{\theta}$	harmonic mean of $\bar{\alpha}$ and $\bar{\mu}$
s^*, u^*	spread and utility at point of indifference
$s^*(p), u^*(p)$	same as above under probabilistic sniping
p_K^*	optimal sniping probability (maximises expected utility)
s_K^*	corresponding optimal spread $s_K^* = s^*(p_K^*)$
$\bar{\gamma}_K$	γ -threshold for transition from <i>sure</i> to <i>probabilistic</i> sniping
$\bar{\gamma}_L$	γ -threshold for transition from <i>probabilistic</i> to <i>no</i> sniping
$h(p)$	probability that market maker will lose race (when HFTs snipe wp. p)
$g(p)$	probability that bandit will win the race, given he enters race and all HFTs snipe wp. p
$h_{tm}(p, H_t, H_d)$	$h(p)$ for trustworthy market maker when there are H_t trustworthy and H_d devious HFTs
$g_{tb}(p, H_t, H_d)$	$g(p)$ for trustworthy bandit when there are H_t trustworthy and H_d devious HFTs

B Computer code

1. Simulation experiments:
`KPD_supplementary_material_game_simulations.ipynb`
2. Check the existence of the solutions:
`KPD_supplementary_material_1.ipynb`
3. Check the sniping probability properties and numerical test:
`KPD_supplementary_material_2.ipynb`
4. Check the numerical analysis:
`KPD_supplementary_material_numeric.ipynb`

C Details of payoff computation

The payoff for each agent is based on eq. (2) which we here restate:

$$\text{payoff} = \text{position}_{(t=\delta)} \times \text{value}_{(t=\delta)} + \text{income}_{(t=\delta)}$$

This allows us to compute the payoff for each event-code. See table (3) and (4) for detailed payoff and utility computation.

Event codes when there is a race	Values at $t = \delta$	MM loses race		MM wins race	
		MM	B	MM	B
NG-NG	Position =	-1	+1	0	0
	Value =	2σ	2σ	2σ	2σ
	Income =	s	$-s$	0	0
	payoff =	$-2\sigma + s$	$2\sigma - s$	0	0
	Utility =	$-\gamma(2\sigma - s)$	$2\sigma - s$	0	0
NG-NB	Position =	-1	+1	0	0
	Value =	0	0	0	0
	Income =	s	$-s$	0	0
	payoff =	$0 + s$	$0 - s$	0	0
	Utility =	s	$-\gamma s$	0	0
NG-LA	Position =	-1	0	-1	0
	Value =	σ	σ	σ	σ
	Income =	s	0	s	0
	payoff =	$-(\sigma - s)$	0	$-(\sigma - s)$	0
	Utility =	$-\gamma(\sigma - s)$	0	$-\gamma(\sigma - s)$	0
NG-LB	Position =	0	+1	+1	0
	Value =	σ	σ	σ	σ
	Income =	$2s$	$-s$	$-(-s)$	0
	payoff =	$(0)(\sigma) + 2s$	$\sigma - s$	$\sigma + s$	0
	Utility =	$2s$	$\sigma - s$	$\sigma + s$	0
NG-no	Position =	-1	+1	0	0
	Value =	σ	σ	σ	σ
	Income =	s	$-s$	0	0
	payoff =	$-(\sigma - s)$	$\sigma - s$	0	0
	Utility =	$-\gamma(\sigma - s)$	$\sigma - s$	0	0
NB-NG	Position =	+1	-1	0	0
	Value =	0	0	0	0
	Income =	s	$-s$	0	0
	payoff =	$0 + s$	$0 - s$	0	0
	Utility =	s	$-\gamma s$	0	0
NB-NB	Position =	+1	-1	0	0
	Value =	2σ	2σ	2σ	2σ
	Income =	s	$-2s$	0	0
	payoff =	$-2\sigma + s$	$2\sigma - s$	0	0
	Utility =	$-\gamma(2\sigma - s)$	$2\sigma - s$	0	0
NB-LA	Position =	0	-1	-1	0
	Value =	$-\sigma$	$-\sigma$	$-\sigma$	$-\sigma$
	Income =	$2s$	$-s$	s	0
	payoff =	$(0)(-\sigma) + 2s$	0	$\sigma + s$	0
	Utility =	$2s$	$\sigma - s$	$\sigma + s$	0
NB-LB	Position =	+1	0	+1	0
	Value =	$-\sigma$	$-\sigma$	$-\sigma$	$-\sigma$
	Income =	s	0	s	0
	payoff =	$-\sigma + s$	0	$-\sigma + s$	0
	Utility =	$-\gamma(\sigma - s)$	0	$-\gamma(\sigma - s)$	0
NB-no	Position =	+1	-1	0	0
	Value =	$-\sigma$	$-\sigma$	$-\sigma$	$-\sigma$
	Income =	$-s$	s	0	0
	payoff =	$-(\sigma - s)$	$\sigma - s$	0	0
	Utility =	$-\gamma(\sigma - s)$	$\sigma - s$	0	0

Table 3 This table represents a list of possible events, when there is a race, from $t = 0$ to $t = \delta$ with detailed computation of payoff and utility based on eq. (2)

Event code		No race		Event code		No race	
		MM	B			MM	B
LA-NG	Position =	-1	0	LB-NG	Position =	+1	0
	Value =	σ	σ		Value =	σ	σ
	Income =	s	0		Income =	s	0
	payoff =	$-\sigma + s$	0		payoff =	$\sigma + s$	0
	Utility =	$-\gamma(\sigma - s)$	0		Utility =	$\sigma + s$	0
LA-NB	Position =	-1	0	LB-NB	Position =	+1	0
	Value =	$-\sigma$	$-\sigma$		Value =	$-\sigma$	$-\sigma$
	Income =	s	0		Income =	s	0
	payoff =	$\sigma + s$	0		payoff =	$-\sigma + s$	0
	Utility =	$\sigma + s$	0		Utility =	$-\gamma(\sigma - s)$	0
LA-LA	Position =	-1	0	LB-LA	Position =	+1	0
	Value =	0	0		Value =	0	0
	Income =	s	0		Income =	$2s$	0
	payoff =	s	0		payoff =	$2s$	0
	Utility =	s	0		Utility =	$2s$	0
LA-LB	Position =	0	0	LB-LB	Position =	+1	0
	Value =	0	0		Value =	0	0
	Income =	$2s$	0		Income =	s	0
	payoff =	$2s$	0		payoff =	s	0
	Utility =	$2s$	0		Utility =	s	0
LA-no	Position =	-1	0	LB-no	Position =	+1	0
	Value =	0	0		Value =	0	0
	Income =	s	0		Income =	s	0
	payoff =	s	0		payoff =	s	0
	Utility =	s	0		Utility =	s	0

Table 4 This table represents a list of possible events, when there is no race, from $t = 0$ to $t = \delta$ with detailed computation of payoff and utility based on eq. (2)

D Expected utility for market maker and bandit

Expected utility for market maker To compute the expected value $\mathbb{E}U_M$ we first condition on the nature of the trigger event: news (prob = β) or LT arrival (prob = $1 - \beta$). Next, in case the trigger event is news, a race ensues, and we additionally condition on whether (prob: $1 - h$) or not (prob. h) the market maker wins. To streamline the notation we split an eventcode e into the first and second event $e = (e_1, e_2)$ (e.g. for $e = \text{LA-NG}$ we have, $e_1 = \text{LA}$ and $e_2 = \text{NG}$). Referring to the probability tree in Fig. 15) we see that the probability of each eventcode e can be decomposed as:

$$p(e) = p(e_1)p(e_2) = \begin{cases} \left(\beta \cdot \frac{1}{2}\right) \cdot p(e_2) & \text{if } e \in \mathbf{N^{***}} \\ \left((1 - \beta) \cdot \frac{1}{2}\right) \cdot p(e_2) & \text{if } e \in \mathbf{L^{***}} \end{cases}$$

Hence, we get the following expansion in terms of possible event-codes e :

$$\begin{aligned} \mathbb{E}U_M(s) &= \beta \underbrace{\left[\sum_{e \in \mathbf{N^{***}}} p(e) u_M(s | e) \right]}_{\text{Trigger = News (race!)}} + (1 - \beta) \underbrace{\left[\sum_{e \in \mathbf{L^{***}}} p(e) u_M(s | e) \right]}_{\text{Trigger = LT (no race!)}} \\ &= \beta \left[\frac{1}{2} \sum_{e_1 \in \mathbf{N^*}} \sum_{e_2} p(e_2) u_M(s | e) \right] + (1 - \beta) \left[\frac{1}{2} \sum_{e_1 \in \mathbf{L^*}} \sum_{e_2} p(e_2) u_M(s | e) \right] \\ &= \frac{\beta}{2} \sum_{e_1 \in \mathbf{N^*}} \sum_{e_2} \left[p(e_2) \left\{ h \cdot u_M(s | e \& \text{MM loses}) + (1 - h) \cdot u_M(s | e \& \text{MM wins}) \right\} \right] \\ &\quad + \frac{(1 - \beta)}{2} \sum_{e_1 \in \mathbf{L^*}} \sum_{e_2} \{ p(e_2) u_M(s | e) \} \end{aligned}$$

Based on the equation above, the following three sums need to be evaluated (using Table 1).

1. **There is a race which MM loses:**

$$\begin{aligned} S_1 &:= \frac{1}{2} \sum_{e_1 \in \mathbf{N^*}} \sum_{e_2} p(e_2) u_M(s | e \& \text{MM loses}) \\ &= \frac{1}{2} [2\bar{\alpha} s(1 - \gamma) - 2\bar{\mu}(\gamma(s - 1) - 2s) + 2\gamma(s - 1)] \end{aligned} \quad (55)$$

2. **There is a race which MM wins:**

$$S_2 := \frac{1}{2} \sum_{e_1 \in \mathbf{N^*}} \sum_{e_2} p(e_2) u_M(s | e_2 \& \text{MM wins}) = \frac{1}{2} [2\bar{\mu}((\gamma + 1)s - \gamma + 1)] \quad (56)$$

3. **No race:**

$$S_3 := \frac{1}{2} \sum_{e_1 \in \mathbf{L^*}} \sum_{e_2} p(e_2) u_M(s | e) = \frac{1}{2} [2\bar{\alpha}(s(\gamma - 1) + 1 - \gamma) + 2s(\bar{\mu} + 1)] \quad (57)$$

Recombining we get:

$$\mathbb{E}U_M(s) = \beta (hS_1 + (1 - h)S_2) + (1 - \beta)S_3$$

In addition, it is useful to compute the values at the endpoints ($s = 0$ and $s = 1$):

$$\begin{aligned}
C &= \mathbb{E}U_M(s=0) = \beta h(-\gamma m) + (1-h)(-\bar{\mu}(\gamma-1)) + (1-\beta)(\bar{\alpha}(1-\gamma)) \\
&= \beta h(\bar{\mu}q - m\gamma) - q\bar{\theta} \\
&= \beta h(2\bar{\mu}q + \bar{\mu} - q - 1) - q\bar{\theta} \\
D &= \mathbb{E}U_M(s=1) = -\beta(\bar{\alpha}hq + m) + (\bar{\mu} + 1)
\end{aligned} \tag{58}$$

See: KPD_Supplementary_material_1.ipynb (section.2.3)

Expected utility for bandit Since the bandit only gets utility when he wins the race, we have to condition three events:

1. there is a race: prob = β
2. the bandit enters the race: prob = p
3. he wins the race (given that he entered): prob = $g(p)$

Using this information we can make the following factorisation:

$$\mathbb{E}U_B(s) = \beta pg(p) \mathbb{E}U_B(s \mid \text{bandit enters and wins race}) \tag{59}$$

$$\begin{aligned}
&= \frac{1}{2} \beta pg(p) \sum_{e_1 \in N^*} \sum_{e_2} p(e_2) u_B(s \mid e_2 \ \& \ \text{bandit wins race}) \\
&= \beta pg(p) S
\end{aligned} \tag{60}$$

where we can use Table 1 to compute

$$S := \frac{1}{2} \sum_{e_1 \in N^*} \sum_{e_2} p(e_2) u_B(s \mid e \ \& \ \text{bandit wins}) = \bar{\alpha} s(1-\gamma) + (\bar{\mu}(s-1) - s + 1) \tag{61}$$

From this we get:

$$\begin{aligned}
A &:= \mathbb{E}U_B(s=0) = (1-\bar{\mu}) \beta pg(p) = m \beta pg(p) \\
B &:= \mathbb{E}U_B(s=1) = -\bar{\alpha}(\gamma-1) \beta pg(p) = -\bar{\alpha}q \beta pg(p)
\end{aligned} \tag{62}$$

See: KPD_Supplementary_material_1.ipynb (section.2.1)

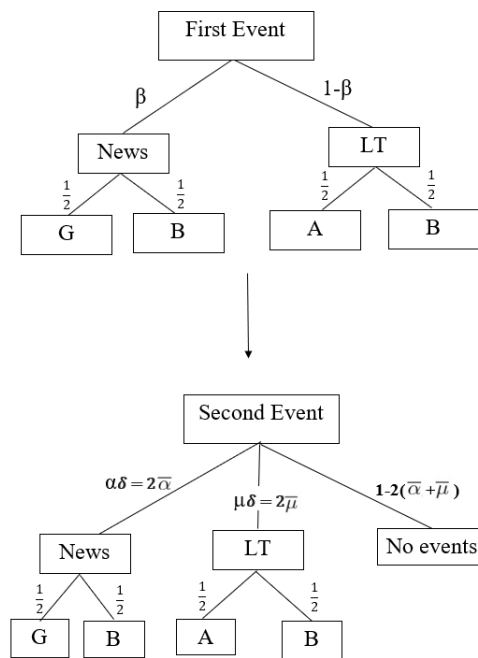


Fig. 15 Schematic representation of the probabilities of the first and second event. Since these two events are independent, the total probability of any event code is the product of two of these probabilities.

E Probability of sniping computation

E.1 Probability $h(p)$

Recall from eq.(15) that $h(p)$ is defined as the probability that the market maker will **lose** the race given that there is a race and all the bandits will attempt to snipe with probability p :

$$h(p) := P \{ \text{MM will lose MZ race} \mid \text{individual bandits enter race with prob } p \}$$

It is conceptually easier to focus on the complementary probability $\bar{h}(p) := 1 - h(p)$ that the market maker will *win* the MZ-race and therefore thwart the sniping attempts of the bandits. Specifically, since the $H - 1$ bandits decide independently and with equal probability p , whether or not they will snipe, the number of bandits that will enter the race is a random variable (N_B say) that is distributed according to a binomial distribution:

$$N_B \sim \text{Bin}(H - 1, p).$$

This is equivalent to saying that the probability of having $b = 0, 1, \dots, H - 1$ bandits participating in a race, equals:

$$P(N_B = b) = \binom{H-1}{b} p^b (1-p)^{H-1-b} \quad \text{for } b = 0, 1, \dots, H-1 \quad (63)$$

When there are b bandits racing, the probability that the market maker will win the race equals $1/(1+b)$ since each participant has the same probability of winning. As a consequence, the probability $\bar{h}(p)$ that market maker will win the race (and therefore escape sniping) equals

$$\bar{h}(p) = \sum_{b=0}^{H-1} \frac{1}{b+1} \binom{H-1}{b} p^b (1-p)^{H-1-b} \quad (64)$$

which can be recast as:

$$\bar{h}(p) = \mathbb{E} \left(\frac{1}{N_B + 1} \right) \quad \text{where } N_B \sim \text{Bin}(H-1, p). \quad (65)$$

Computation of $\bar{h}(p)$ and $h(p)$ To prove that eqs. (64) and (65) can be expanded to give rise to

$$\bar{h}(p) = \frac{1 - (1-p)^H}{pH}, \quad (66)$$

we proceed as follows (we introduce $n = H - 1$ for notational convenience).

Recall that it is a straightforward consequence of the definition of binomial coefficients that:

$$\binom{n+1}{b+1} = \frac{n+1}{b+1} \binom{n}{b} \quad \text{whence} \quad \frac{1}{b+1} \binom{n}{b} = \frac{1}{n+1} \binom{n+1}{b+1}$$

Substituting this in eq. (64) yields:

$$\begin{aligned}
\bar{h}(p) &= \sum_{b=0}^n \frac{1}{(b+1)} \binom{n}{b} p^b (1-p)^{n-b} \\
&= \frac{1}{p(n+1)} \sum_{b=0}^n \binom{n+1}{b+1} p^{b+1} (1-p)^{(n+1)-(b+1)} \\
&= \frac{1}{p(n+1)} \sum_{k=1}^{n+1} \binom{n+1}{k} p^k (1-p)^{(n+1)-k} \quad (\text{subs: } k = b+1) \\
&= \frac{1}{p(n+1)} \left(\underbrace{\sum_{k=0}^{n+1} \binom{n+1}{k} p^k (1-p)^{(n+1)-k}}_{=1} - p^0 (1-p)^{n+1} \right) \\
&= \frac{1 - (1-p)^{n+1}}{p(n+1)} = \frac{1 - (1-p)^H}{pH} \quad (n = H-1).
\end{aligned}$$

The computation above also implies that

$$h(p) = 1 - \bar{h}(p) = \mathbb{E} \left(\frac{N_B}{1 + N_B} \right) = \frac{(1-p)^H - (1-Hp)}{Hp} \quad (67)$$

where $N_B \sim \text{Bin}(H-1, p)$.

Some special cases Notice that for sure sniping (i.e. $p = 1$) we get the familiar factor $h := h(1) = (H-1)/H$. Similarly, we can differentiate eq.(67) to obtain:

$$h'(p) = \frac{1 - (1-p)^{H-1}(Hp + 1 - p)}{Hp^2} \implies h'(1) = \frac{1}{H} \quad (68)$$

To investigate the behaviour for small values $p \downarrow 0$ we use the standard binomial expansion:

$$(1-p)^H = 1 - Hp + \frac{H(H-1)}{2} p^2 - \frac{H(H-1)(H-2)}{6} p^3 + \dots$$

from which we conclude:

$$h(p) = \frac{H-1}{2} p - \frac{(H-1)(H-2)}{6} p^2 + o(p^3) \quad \text{as } p \rightarrow 0. \quad (69)$$

and therefore:

$$h(0) = 0 \quad \text{and} \quad h'(0) = \frac{H-1}{2}. \quad (70)$$

E.2 Probability $g(p)$

Recall eq.(19) and (20)

$$g(p) := P\{\text{individual bandit will win race} \mid \text{he participates in the race}\}$$

Since the market maker always races, and we condition on this bandit also being in the race, the number of agents in the race equals $2 + N_B$ where $N_B \sim \text{Bin}(H-2, p)$. Hence:

$$g(p) = \mathbb{E} \left(\frac{1}{2 + N_B} \right)$$

Using the binomial identity:

$$\binom{n+2}{b+2} = \frac{(n+2)(n+1)}{(b+2)(b+1)} \binom{n}{b}$$

a computation completely analogous to the one for $h(p)$ shows that:

$$g(p) = \sum_{b=0}^{H-2} \frac{1}{b+2} \binom{H-2}{b} p^b (1-p)^{H-2-b} = \frac{(1-p)^H - (1-p)H}{H(H-1)p^2} = \frac{h(p)}{(H-1)p} \quad (71)$$

Special cases In the paper we need the following special cases:

– **Limit for sure sniping** $p \rightarrow 1$

$$g(1) = \frac{h(1)}{H-1} = \frac{1}{H} \quad (72)$$

and also:

$$g'(p) = \frac{h'(p)p - h(p)}{(H-1)p^2} \implies g'(1) = \frac{h'(1) - h(1)}{H-1} = -\frac{H-2}{H(H-1)} \quad (73)$$

– **Limit for no sniping** $p \rightarrow 0$ Using the expansion eq. 69 for $h(p)$, we get a similar expansion for $g(p)$:

$$g(p) = \frac{1}{(H-1)p} \left\{ \frac{H-1}{2}p - \frac{(H-1)(H-2)}{6}p^2 + o(p^3) \right\} = \frac{1}{2} - \frac{H-2}{6}p + o(p^2) \quad (74)$$

From this we can conclude:

$$g(0) = \frac{1}{2} \quad \text{and} \quad g'(0) = -\frac{H-2}{6}. \quad (75)$$

See: `KPD_Supplementary_material_2.ipynb`

F Proof of Proposition 1: Existence and uniqueness of $\bar{\gamma}_K$

To prove the existence and uniqueness of the transition threshold $\bar{\gamma}_K$ we need to show that the cubic polynomial (cf. eq. 45):

$$K(\gamma) = K_3\gamma^3 + K_2\gamma^2 + K_1\gamma + K_0$$

has a unique zero-crossing greater than 1. To this end we prove the following three steps (also see Fig 16):

1. $K_3 < 0$ indicating that $\lim_{\gamma \rightarrow \infty} K(\gamma) = -\infty$, i.e. $K(\gamma)$ will be negative for sufficiently large values of γ ;
2. $K(1) = K_3 + K_2 + K_1 + K_0 > 0$, which in combination with $K_3 < 0$, shows that there is at least one zero-crossing greater than 1;
3. Finally, we show that $K''(1) = 6K_3 + 2K_2 < 0$ which proves that the cubic polynomial is concave to the right of $\gamma = 1$. This implies that the zero-crossing $\bar{\gamma}_K$ is in fact unique.

The detailed calculations that prove the three steps above can be found in the Python notebook:

See: `KPD_Supplementary_material_1.ipynb` (section.5)




Fig. 16 The proof of the existence and uniqueness of the threshold $\bar{\gamma}_K$ is based on the following three properties of the $K(\gamma)$ cubic polynomial: (i) $K_3 < 0$, $K(1) > 0$, and K is concave to the right of 1.

G Proof of Proposition 2: Existence and uniqueness of $\bar{\gamma}_L$

To prove that the transition threshold $\bar{\gamma}_L$ exists and is unique, we point out that (cf. eq. (42))

$$\left. \frac{du^*(p)}{dp} \right|_{p=0} = 0 \iff N'(0) = 0.$$

The right hand side can be expanded as a function of γ , or even more conveniently, as a function of $q = \gamma - 1$. It turns out :

$$N'(0) = -\frac{1}{2}\bar{\alpha}\bar{\beta}\bar{\theta}q^2 + \frac{1}{2}\beta(1-\bar{\mu})(\beta\bar{\mu} - \beta + \bar{\mu} + 1).$$

Computing the zero-crossings for the quadratic equation on the right hand side, we conclude that there is a unique solution greater than 1, viz.:

$$\bar{\gamma}_L = 1 + \sqrt{\frac{(1-\bar{\mu})Z}{\bar{\alpha}\bar{\theta}}} \quad \text{where} \quad Z = 1 + \bar{\mu} - \beta(1-\bar{\mu}).$$

This concludes the proof.

The detailed calculations that prove the three steps above can be found in the Python notebook:

KPD_Supplementary_material_1.ipynb (section.6)

H Alternative computation of g_{tb}

Let A be the stochastic variable representing the number of agents in the race. Denote by $M \in \mathcal{T}(\mathcal{D})$ the event that the market maker is one of the trustworthy (\mathcal{T}), or deceptive (\mathcal{D}) agents, respectively. Then

$$A = (A \mid M \in \mathcal{T})P(M \in \mathcal{T}) + (A \mid M \in \mathcal{D})P(M \in \mathcal{D})$$

We can make the following decomposition:

- Since we condition on a specific trustworthy agent in the race, there are $H - 1$ agents left, $H_t - 1$ of them are trustworthy.
- We can visualise the selection process as follows. Line up the agents. The first H_t are trustworthy, the next H_d are the devious ones. Indicate who will race or not. Agent 1 will race (we condition on this event), from the remaining $H_t - 1$ trustworthy agent a stochastic number $N_t \sim B(H_t - 1, p)$ will race, as well as all the devious ones. Then select a market maker uniformly from agent 2 through H (we have conditioned on the first agent to race as *bandit*).
 - If the choice of market maker is among the devious ones, nothing changes in the list of racing participants.
 - If the choice of market maker is one of the $H_t - 1$ remaining trustworthy agents there are two possibilities. If the agent was already singled out for racing (prob = $N_t / ((H_t - 1))$), nothing changes. If however, the agent was not planning on racing (prob = $(H_t - 1 - N_t) / (H_t - 1)$), he will now do so and the number of racers is incremented by one to $1 + N_t + H_d + 1 = 2 + N_t + H_d$. Combining these observations we get:

$$\left(\frac{N_t}{H_t - 1}\right) \left(\frac{1}{1 + N_t + H_d}\right) + \left(\frac{H_t - 1 - N_t}{H_t - 1}\right) \left(\frac{1}{2 + N_t + H_d}\right)$$

or again:

$$\frac{1}{(H_t - 1)} \left\{ \frac{N_t}{1 + N_t + H_d} + \frac{H_t - 1 - N_t}{2 + N_t + H_d} \right\}$$

- $P(M \in \mathcal{T}) = (H_t - 1) / (H - 1)$ the market maker is chosen among the remaining trustworthy
- $P(M \in \mathcal{D}) = H_d / (H - 1)$

Combining all these together we get the following alternative for eq.(52):

$$\begin{aligned}
g_{tb}(p) &= \left(\frac{1}{H-1}\right) \mathbb{E} \left\{ \frac{N_t}{1+N_t+H_d} + \frac{H_t-1-N_t}{2+N_t+H_d} \right\} \\
&\quad + \left(\frac{H_d}{H-1}\right) \mathbb{E} \left\{ \frac{1}{1+N_t+H_d} \right\} \\
&= \left(\frac{1}{H-1}\right) \mathbb{E} \left\{ \frac{N_t}{1+N_t+H_d} + \frac{H_t-1-N_t}{2+N_t+H_d} + \frac{H_d}{1+N_t+H_d} \right\} \quad (76)
\end{aligned}$$

where $N_t \sim \text{Bin}(H_t - 1, p)$.

This characterisation of g_{tb} has the advantage that it involves only a single random variable, which makes sense whenever $H_t \geq 1$. The original characterisation in eq. (52) — although conceptually simpler — involves two binomial random variables and requires $H_t \geq 2$ to be well-defined..

See: `KPD_supplementary_material_2.ipynb`

I Change in mean utility due to deceptive agents

We are now in a position to combine all these results and write down the equivalents for eqs. (25) and (30). More specifically, if a trustworthy agent acts as bandit, his utility is a linear function determined by the endpoints:

$$\begin{aligned}
A_{tb}(p, H_t, H_d) &= m\beta p g_{tb}(p, H_t, H_d) \\
B_{tb}(p, H_t, H_d) &= -\bar{\alpha}(\gamma - 1)\beta p g_{tb}(p, H_t, H_d) \\
C_{tm}(p, H_t, H_d) &= -\left\{ q\bar{\theta} + \beta(m\gamma - \bar{\mu}q) h_{tm}(p, H_t, H_d) \right\} \\
D_{tm}(p, H_t, H_d) &= (1 + \bar{\mu}) - \beta \{m + \bar{\alpha}q h_{tm}(p, H_t, H_d)\}
\end{aligned} \quad (77)$$

To evaluate for a given spread the expected utility $u_t(s)$ of a trustworthy HFT one simply needs to evaluate the expected utilities in both roles (bandit and market maker):

$$u_{tb}(s) = A_{tb}(1-s) + B_{tb}s \quad \text{and} \quad u_{tm}(s) = C_{tm}(1-s) + D_{tm}s$$

which can then be combined (assuming that all agents advertise the same spread s and therefore have equal probability of being selected as market maker):

$$u_t(s) = \frac{1}{H} u_{tm}(s) + \frac{H-1}{H} u_{tb}(s) \quad (78)$$

See: `KPD_supplementary_material_numeric.ipynb` (section.6)

J Probability distribution of unique utilities

In general, there are only 9 unique utilities (see table 1). This is still the case in the presence of deceptive agents since even deceptive agents will publish the optimal spread. We can compute the probability with which these occur explicitly. Notice that there are three different cases:

1. Utilities that occur only for MM $(-\gamma(2\sigma - s), s, -\gamma(\sigma - s), 2s, \sigma + s)$
2. Utilities that occur only for bandit $(2\sigma - s, -\gamma s, \sigma - s)$
3. Zero utility that occurs for both market maker and bandit

Let's do so for some concrete cases:

- **Utility $2\sigma - s$ for successful sniper** (example of utility that occurs only for bandit): this utility occurs twice (NG-NG and NB-NB + MM loses race to successful sniper). To compute the probability we need to multiply the probability of each row ($(1/2)\bar{\alpha}\beta$) with that of the column and add it all together. Since the column is the same for both events (let's denote it as p_{col} we get:

$$p\{2\sigma - s\} = \bar{\alpha}\beta \cdot p_{col}$$

where

$$\begin{aligned} p_{col} &= P(\text{agent get this utility} \mid \text{he's bandit})P(\text{he's bandit}) \\ &\quad + P(\text{agent get this utility} \mid \text{he's MM})P(\text{he's MM}) \\ &= P(\text{agent get this utility} \mid \text{he's bandit})((H-1)/H) \\ &\quad + 0 \cdot (1/H) \\ &= \frac{(H-1)}{H} P(\text{agent get this utility} \mid \text{he's bandit}) \end{aligned}$$

The last probability can be expanded further:

$$\begin{aligned} &P(\text{agent get this utility} \mid \text{he's bandit}) \\ &= P(\text{agent enters race and wins} \mid \text{he's bandit}) \\ &= P(\text{agent enters race} \mid \text{he's bandit}) \cdot P(\text{agent wins race} \mid \text{he's bandit and enters race}) \\ &= p g_{tb}(p) \end{aligned}$$

where we need to use g_{tb} since we are only interested in the fate of trustworthy agents. Combining both results we see that

$$P\{2\sigma - s\} = \bar{\alpha}\beta \cdot \left(\frac{H-1}{H}\right) p g_{tb}(p)$$

- **Utility $\sigma - s$ for successful sniper** this utility occurs 4 times when there is NG-NO, NB-NO and MM loses the race, and NG-LB and NB-LA and bandit benefits from the news arrival (MM loses the race in the news side).

In all cases, the probability that successful sniper (bandit) enters the race $(H-1)/H$ and win the race is $p g_{tb}(p)$ (As we explained in detail in the utility $P\{2\sigma - s\}$).

In the case of NG-NO and NB-NO, the row probability is the same for both events. Combining with column probability (bandit is selected $\left(\frac{H-1}{H}\right)$ and is the race winner $p g_{tb}(p)$):

$$\text{prob}_{NG-NO, NB-NO} = \beta \cdot (1 - 2(\bar{\alpha} + \bar{\mu})) \cdot \left(\frac{H-1}{H}\right) p g_{tb}(p)$$

and the same for NG-LB, NB-LA:

$$\text{prob}_{NG-LB, NB-LA} = \beta \cdot \bar{\mu} \cdot \left(\frac{H-1}{H}\right) p g_{tb}(p)$$

The total probability of getting the utility $(\sigma - s)$:

$$P\{\sigma - s\} = \beta \cdot \left(\frac{H-1}{H}\right) p g_{tb}(p) \cdot (1 - 2\bar{\alpha} - \bar{\mu})$$

- **Utility $-\gamma s$ for successful sniper** The only two events that produces this utility are NG-NB and NB-NG for a successful sniper (event probability in each case $\bar{\alpha}\beta/2$). In this case, the trustworthy agent acts as a bandit (prob: $(H-1)/H$), he enters the race (prob: p) and wins (prob: $g_{tb}(p)$). However, he doesn't benefit from it (make loss) as additional bad news decreases the intrinsic value. Combining all this yields:

$$P\{-\gamma s\} = \bar{\alpha}\beta \left(\frac{H-1}{H}\right) p g_{tb}(p)$$

- **Utility s for market maker.** This utility occurs on six occasions: NG-NB, NB-NG, LA-LA, LA-no, LB-LB and LB-no. For each of these cells in the table we can compute the probability by multiplying the probability of the row and the column. For the last four, the column probability is that of being market maker ($1/H$). Adding the probabilities of the rows yields:

$$(1/2)(1 - \beta) \cdot 2 \cdot (\bar{\mu} + 1 - 2(\bar{\alpha} + \bar{\mu})) = (1 - \beta)(1 - 2\bar{\alpha} - \bar{\mu})$$

Hence the total probability of the last four cells equals:

$$\text{prob}_{LA-LA, LA-No, LB-LB, LB-No} = \frac{1}{H}(1 - \beta)(1 - 2\bar{\alpha} - \bar{\mu})$$

To complete the computation we need to compute the probability of the first two cells. Only the probability of the column needs attention. It requires being selected as MM (prob = $1/H$) but losing the race $h_{tm}(p)$. Hence the combined total probability of these two cells equals:

$$\text{prob}_{NG-NB, NB-NG} = 2 \cdot (1/2)\bar{\alpha}\beta \cdot \left(\frac{1}{H}\right) h_{tm}(p) = \left(\frac{1}{H}\right) \bar{\alpha}\beta h_{tm}(p).$$

Combining the two results we get the following total probability for the utility s :

$$P\{s\} = \left(\frac{1}{H}\right) [\bar{\alpha}\beta h_{tm}(p) + (1 - \beta)(1 - 2\bar{\alpha} - \bar{\mu})]$$

- **Utility $-\gamma(\sigma - s)$ for market maker.**

This utility occurs when there is: NG-LA, NB-LB, LA-NG, LB-NB, NG-No, and NB-No. The column probability for all cases is that of being selected as a market maker ($1/H$). Adding the probabilities of events in the rows, we see :

Probability for NG-LA and NB-LB:

In these events, MM wins and loses the race, so this utility happens four times and h_{tm} cancels out!.

$$\text{prob}_{NG-LA} = \left(\frac{1}{H}\right) 2 \cdot \frac{\beta}{2} \cdot \bar{\mu} \cdot h_{tm} + \left(\frac{1}{H}\right) 2 \cdot \frac{\beta}{2} \cdot \bar{\mu} \cdot (1 - h_{tm})$$

Therefore :

$$\text{prob}_{NG-LA} = \left(\frac{1}{H}\right) \cdot \frac{\beta}{2} \cdot \bar{\mu}$$

It is the same for NB-LB and combining two events:

$$\text{prob}_{NG-LA, NB-LB} = \left(\frac{1}{H}\right) \cdot \beta \cdot \bar{\mu}$$

Probability for NG-No and NB-No + MM loses the race with probability $h_{tm}(p)$: In both events, MM loses the race as bandit hits the stale quote:

$$\text{prob}_{NG-No, NB-No} = \left(\frac{1}{H}\right) 2 \cdot \frac{\beta}{2} \cdot (1 - 2(\bar{\alpha} + \bar{\mu})) \cdot h_{tm}(p)$$

Probability for LA-NG and LB-NB (The utility in these events happens four times):

$$\text{prob}_{LA-NG, LB-NB} = \left(\frac{1}{H}\right) (1 - \beta) \cdot \bar{\alpha}$$

Combining probabilities of six events:

$$P\{-\gamma(\sigma - s)\} = \left(\frac{1}{H}\right) \{\bar{\alpha}(1 - \beta) + \beta h_{tm}(1 - 2\bar{\alpha} - 2\bar{\mu}) + \beta\bar{\mu}\}$$

- **Utility $s + \sigma$ for market maker.** There are four corresponding events for this utility: NG-LB, NB-LA, LA-NB and LB-NG.

To compute the probability:

To compute row probabilities for the first two events (NG-LB and NB-LA), we need to multiply the probability of each row:

$$\begin{aligned}\text{prob}_{NG-LB, NB-LA} &= \bar{\mu}\beta \\ \text{prob}_{LA-NB, LB-NG} &= \bar{\alpha}(1 - \beta)\end{aligned}$$

Then, to compute the column probability, we need the probability of $(\frac{1}{H})$ MM being selected and $(1 - h_{tm}(p))$ for winning the race (for NG-LB and NB-LA) hence the total probability for this utility equals :

$$P\{s + \sigma\} = \left(\frac{1}{H}\right) \{\beta\bar{\mu}(1 - h_{tm}(p)) + \bar{\alpha}(1 - \beta)\}$$

- **Utility $2s$ for market maker** This utility befalls the market maker in four events : NG-LB, NB-LA, LA-LB and LB-LA.

Row probability for NG-LB and NB-LA + MM loses the race:

$$\bar{\mu}\beta$$

Adding $1/H$ the probability of being selected as MM and losing the race h_{tm} yields:

$$\text{prob}_{NG-LB, NB-LA} = \bar{\mu}\beta \cdot \left(\frac{1}{H}\right) \cdot h_{tm}$$

Row probability for LA-LB and LB-LA:

$$\text{prob}_{LA-LB, LB-LA} = \bar{\mu}(1 - \beta) \cdot \left(\frac{1}{H}\right)$$

Adding the probabilities of events in the rows, we see :

$$P\{2s\} = \bar{\mu} \left(\frac{1}{H}\right) (\beta \cdot h_{tm}(p) + (1 - \beta))$$

- **Utility $-\gamma(2\sigma - s)$ for market maker.** this utility occurs in events: NG-NG and NB-NB.

Adding row probability with column probability that is MM being selected ($1/H$) and loses the race h_{tm} yields:

$$P\{-\gamma(2\sigma - s)\} = \bar{\alpha}\beta \cdot \left(\frac{1}{H}\right) \cdot h_{tm}$$

- **Zero utility $P\{0\}$ for bandit**

This utility occurs in all events when the first trigger event is the arrival of LT. In these events, bandit has no chance of winning the race. Also, this utility occurs whenever bandit loses the race, either MM wins the race or the arrival of LT.

LA-LA, LA-LB, LB-LA, LB-LB:

$$\bar{\mu}(1 - \beta)(2) \left(\frac{H - 1}{H}\right)$$

LA-NG, LA-NB, LB-NG, LB-NB:

$$\bar{\alpha}(1 - \beta)(2) \left(\frac{H - 1}{H}\right)$$

LA-No and LB-No :

$$(1 - 2(\bar{\alpha} + \bar{\mu}))(1 - \beta) \left(\frac{H - 1}{H}\right)$$

NG-LA, NB-LB (bandit lose and win the race) :

bandit loses the race :

$$P_{NG-LA} = \bar{\mu}\beta \left(\frac{H-1}{H} \right) \cdot p(1 - g_{tb})$$

bandit wins the race :

$$P_{NG-LA} = \bar{\mu}\beta \left(\frac{H-1}{H} \right) \cdot p(g_{tb})$$

Therefore:

$$P_{NG-LA} = \bar{\mu}\beta/2 \left(\frac{H-1}{H} \right)$$

the same for NB-LB:

$$P_{NB-LB} = \bar{\mu}\beta/2 \left(\frac{H-1}{H} \right)$$

combining four events:

$$P_{NB-LB} = \bar{\mu}\beta \left(\frac{H-1}{H} \right)$$

When bandit loses the race in all other events if the first trigger event is news.
(events = NG-LB, NB-LA, NG-NG, NG-No, NB-No, NB-NG, NB-NB)

$$P_{events} = \left(\frac{H-1}{H} \right) \beta(1 - pg_{tb})(1 - \bar{\mu})$$

Total probability for $P\{0\}$ is:

$$P\{0\}_{bandit} = \left(\frac{H-1}{H} \right) (1 - \beta + \beta\bar{\mu} + \beta(1 - pg_{tb})(1 - \bar{\mu}))$$

This can further be simplified to:

$$P\{0\}_{bandit} = \left(\frac{H-1}{H} \right) (1 - \beta pg_{tb}(1 - \bar{\mu}))$$

- **Zero utility $P\{0\}$ for MM**

e1 = NG-NG, NG-NB, NB-NG, NB-NB:

$$P_{e1} = \left(\frac{1}{H} \right) 2\beta\bar{\alpha}(1 - h_{tm})$$

e2 = NG-No, NB-No:

$$P_{e2} = \left(\frac{1}{H} \right) \beta(1 - 2(\bar{\alpha} + \bar{\mu}))(1 - h_{tm})$$

Therefore :

$$P\{0\}_{MM} = \left(\frac{1}{H} \right) \beta(1 - h_{tm})(1 - 2\bar{\mu})$$

Alternative derivation

$P_B\{u = 0\} := \text{Prob that trustworthy bandit receives zero utility } (u = 0)$

$$\begin{aligned} P_B\{u = 0\} &= P_B(u = 0 \mid \text{race})P(\text{race}) + P_B(u = 0 \mid \text{no race})P(\text{no race}) \\ &= P_B(u = 0 \mid \text{race}) \cdot \beta + 1 \cdot (1 - \beta) \end{aligned} \quad (79)$$

$$P_B(u = 0 \mid \text{race}) = P_B(u = 0 \mid \text{race and B enters race})P(\text{B enters race} \mid \text{race}) \quad (80)$$

utility	probability
0	$\frac{1}{H} \{\beta(1 - 2\bar{\mu})(1 - h_{tm}(p))\} + \frac{H-1}{H} \{1 - \beta pg_{tb}(p)(1 - \bar{\mu})\}$
$-\gamma(2\sigma - s)$	$\bar{\alpha}\beta \cdot \left(\frac{1}{H}\right) \cdot h_{tm}$
s	$\left(\frac{1}{H}\right) \{\bar{\alpha}\beta h_{tm}(p) + (1 - \beta)(1 - 2\bar{\alpha} - \bar{\mu})\}$
$-\gamma(\sigma - s)$	$\left(\frac{1}{H}\right) \{\bar{\alpha}(1 - \beta) + \beta h_{tm}(p)(1 - 2(\bar{\alpha} + \bar{\mu})) + \beta\bar{\mu}\}$
$s + \sigma$	$\left(\frac{1}{H}\right) \{\beta\bar{\mu}(1 - h_{tm}(p)) + \bar{\alpha}(1 - \beta)\}$
2s	$\bar{\mu} \cdot \left(\frac{1}{H}\right) \{\beta h_{tm}(p) + (1 - \beta)\}$
$2\sigma - s$	$\frac{H-1}{H} \{\bar{\alpha}\beta pg_{tb}(p)\}$
$\sigma - s$	$\left(\frac{H-1}{H}\right) \cdot \beta pg_{tb}(p) \{1 - 2\bar{\alpha} - \bar{\mu}\}$
$-\gamma s$	$\left(\frac{H-1}{H}\right) \cdot \frac{\beta}{2} \cdot \bar{\alpha} pg_{tb}(p)$

Table 5 Overview of unique utilities and corresponding probabilities.