

**On the Trail of Lost Pennies:
Player-Funded Tug-Of-War on the Integers**

Alan Hammond

A. HAMMOND, DEPARTMENT OF MATHEMATICS AND STATISTICS, U.C. BERKELEY, 899 EVANS
HALL, BERKELEY, CA, 94720-3840, U.S.A.

Email address: `alanmh@berkeley.edu`

1991 *Mathematics Subject Classification.* 82C22, 82B23 and 60H15.

Key words and phrases. Bidding games on graphs, Colonel Blotto games, infinite games, multi-turn games, non-zero-sum two-player games, player-funded stake-governed games, random-turn games, strategic move evaluation, tug-of-war, Tullock contests.

The author is supported by NSF grant DMS-2153359.

ABSTRACT. We study a class of two-player, random-turn games in which a valued resource must be allocated over the lifetime of the game so that a favourable outcome is likely secured without wasteful expense. The Trail of Lost Pennies is the game played on \mathbb{Z} . It is specified by a parameter $x \in (0, \infty)$. A counter begins at $0 \in \mathbb{Z}$ and lies at some integer at the start of any given \mathbb{N} -indexed turn. At this turn, Maxine and Mina offer non-negative real-valued stakes. If Maxine stakes a and Mina b , the counter moves one unit to the right with probability $\frac{a}{a+b}$; otherwise, it moves to the left by one unit. If the counter location tends to minus infinity over the course of the game, Mina receives a terminal payment of one, and Maxine zero; if the counter tends to plus infinity, then these respective receipts are zero and x . Thus the net receipt to a given player is $-A + B$, where A is the sum of the stakes she offers during the game, and B is her terminal receipt. The game was inspired by random-turn tug-of-war analysed in the mathematical treatment [49] from 2009 but in fact closely resembles the original version of tug-of-war, introduced [26] in the economics literature in 1987. We show that the game has surprising features. For the class of strategies that satisfy a natural time-invariance, Nash equilibria exist precisely when x lies in the interval $[\lambda, \lambda^{-1}]$, for a certain $\lambda \in (0, 1)$. We indicate that the value of λ is remarkably close to one, proving that $\lambda \leq 0.999904$ and presenting clear numerical evidence that $\lambda \geq 1 - 10^{-4}$. For each $x \in [\lambda, \lambda^{-1}]$, we exhibit countably many Nash equilibria for the game, showing that each is roughly characterized by an integral *battlefield* index: when the counter is around this value, both players stake intensely, with rapid but asymmetric decay in stakes as the counter moves away from the value. Our results validate and quantify hypotheses, for strategic resource management and the effect of discouragement in determining the relation between incentive and outcome, which are similar to premises advanced in economics [26, 37], and which plausibly hold more generally among player-funded stake-governed random-turn games. Alongside a companion treatment [24] of games with allocated budgets, we thus offer a detailed mathematical treatment of an illustrative class of tug-of-war games governed by limited resources. We also review the separate developments of tug-of-war in economics and mathematics in the hope that mathematicians may direct further attention to tug-of-war games, in their original resource-allocation guise.

Contents

Chapter 1. Introduction	5
1.1. Rough guide: motivations, model, results	5
1.1.1. Allocating resources strategically in a multi-turn game	5
1.1.2. Tug-of-war, constant-bias and stake-governed	6
1.1.3. An informal summary of the main conclusions	8
1.1.4. Overview: allocated-budget stake games and the Present-Future Compromise	9
1.2. The two strands of tug-of-war	10
1.2.1. Tug-of-war in economics	10
1.2.2. Tug-of-war in mathematics	14
1.2.3. Acknowledgments	16
Chapter 2. Main results	17
2.1. Game setup, strategies and Nash equilibria	17
2.2. Time-invariant Nash equilibria and the ABMN equations	19
2.3. Explicit ABMN solutions	24
2.4. The Mina margin map	26
2.5. The Trail of Lost Pennies on a finite interval	29
2.5.1. Gameplay, strategies and Nash equilibria for the finite trail	29
2.5.2. The ABMN equations	30
2.5.3. A result and a conjecture for the finite trail	30
2.6. Some further formulas	32
2.6.1. Alternative formulas for standard solutions and their Mina margins	32
2.6.2. The game with symmetric boundary data	34
2.7. The article's structure	35
Chapter 3. Some basic tools	36
3.1. Role-reversal symmetry	36
3.2. Penny Forfeit	38
3.3. The game with a delayed start	39
3.4. Lack of escape entails infinite costs	39
3.5. Relating the finite and infinite trail games	40
Chapter 4. The structure of time-invariant Nash equilibria	43
4.1. Time-invariant Nash equilibria result in positive ABMN solutions	43
4.2. The reverse implication	49
Chapter 5. Explicit ABMN solutions and their consequences	56
5.1. Explicit ABMN solutions	56
5.2. Asymptotic decay of solutions	59
5.3. Consequences of asymptotic decay	61

Chapter 6. The Mina margin map	64
6.1. Approximating the Mina margin by its finite trail counterpart	64
6.1.1. An explicit form for the finite-trail Mina margin map	64
6.1.2. Estimates for the finite trail Mina margin map	66
6.1.3. Proofs via the finite trail Mina margin map	69
6.1.4. Some further consequences	69
6.2. The Mina margin map after domain coordinate change	71
6.3. The Mina margin map is not identically equal to one	73
Chapter 7. Trail prospects	78
7.1. Properties of the Mina margin map and prospective routes to conjectures	78
7.2. The possible existence of further Nash equilibria	78
7.2.1. Time-invariant random Nash equilibria	78
7.2.2. Nash equilibria that are not time-invariant	78
7.3. Gameboards beyond \mathbb{Z}	80
7.4. The trail game and “Incentive Inch, Outcome Mile”	81
7.5. The game of chicken in the Trail of Lost Pennies	82
7.6. Play between people and algorithms	84
7.6.1. Cooperative behaviour	84
7.6.2. Tit-for-tat	84
7.6.3. The loadsamoney bully	84
Appendix A. Proofs of the further formulas	85
A.1. Deriving alternative formulas for standard solutions and Mina margins	85
A.2. The symmetric game and the A equations	85
A.2.1. The AM equations	85
A.2.2. Deriving Theorem 2.32	87
Appendix. Bibliography	89

CHAPTER 1

Introduction

In tug-of-war random-turn games governed by stakes, two players must each decide how to allocate a precious resource over the lifetime of the game in order to gain often enough temporary control of the movement of a counter, and how to move the counter when their resource expenditure and good fortune permit them to do so. These games have been studied by economists for forty years and by mathematicians for twenty. The former study games that emphasise the aspect of resource allocation in player decisions; the latter, those where these decisions concern strategic movement. The two research strands have been vigorous but remarkably disjoint. The principal aim of this monograph is to present a detailed inquiry into the Trail of Lost Pennies, an infinite-turn stake-governed tug-of-war game, in whose infinite set of Nash equilibria we will find delicate and surprising features. We further aim to direct mathematicians' attention to stake-governed tug-of-war games and to advocate these games as objects for further inquiry.

The first part of the introduction offers a rough conceptual overview of the motivations and form of the player-funded tug-of-war game that we study and of our main results; and we compare these conclusions to those for a class of allocated-budget tug-of-war games that have been treated in the companion article [24]. In the second part, we trace the course of the two long, unwoven strands of tug-of-war research, in economics and mathematics.

1.1. Rough guide: motivations, model, results

The first of four subsections presents three broad premises for how resources may be allocated strategically in a random multi-turn game. The second conceptually presents tug-of-war games, first in the simpler setting of constant bias, and then with bias governed by stakes drawn from budgets. These budgets may be allocated by the bank or funded by the players themselves. Our principal game of study, the Trail of Lost Pennies, is of the latter type. In the third subsection, we explain several of our main conclusions about this game, finding it to exemplify two of the three broad premises. The remaining premise is instead illustrated by a class of allocated-budget games analysed in [24], as the fourth subsection explains.

1.1.1. Allocating resources strategically in a multi-turn game. When political opponents buy advertising before an election or rival firms rent stalls in a market place, they allocate precious resources in an effort to win a favourable outcome, if need be at the expense of the opponent. Such competition may be modelled by multi-turn two-player games at each of whose turns each player allocates a resource that she values in an effort to improve her strategic position and raise the chance of a favourable outcome when the game ends. In seeking to play such a game effectively, a player must address the basic question: how much of the resource should I allocate for play at the impending turn? To find the answer, the player is led to evaluate the strategic importance of the next turn relative to such a measure for the accumulation of all subsequent turns during the likely lifetime of the game. So broadly specified, this question may elicit a wide range of conceptual responses. We distinguish three:

- (1) **Compromise between the short and long terms.** In facing a choice of level of resource allocation for the upcoming turn, a player is tempted to offer a large amount, in the hope of securing an improved position at the turn. However, such expenditure depletes the resource, making it more difficult to stake competitively later in the game. In this view, there is a tension between the short and the long term, with the two time-scales offering upward and downward pressure on resource allocation at a given turn. Presumably then under optimal play the present level of expenditure will maximize a unimodal function.
- (2) **An inch in incentive may be a mile for outcome.** Suppose that the two players will receive different rewards for winning the game. The player with the greater incentive may be led to commit greater resources at any given stage of the game. If the difference in incentive is small, then the resulting increase in stage-win probability for the more committed player may be only slight. But if there are many stages in the game, then these advantages will accumulate to become pronounced, just as a random walk with a slight bias to the right will proceed linearly in that direction in the long term. Moreover, the thus disadvantaged player, in becoming aware of her poor prospects, will be discouraged from committing resources, leading her to stake very little, leaving the other player, even if only slightly incentivised, able to win the game without any real opposition.
- (3) **The battlefield. Cut your losses/Foot on gas.** In competing for dominance in some form—say, over a network or geographical space— in a game taking many turns, both players may begin to recognise as gameplay advances that one has secured a certain advantage. Perhaps one player has won a key battle over a significant part of space, or is simply blessed with a favourable terrain. In such a circumstance, the player in the seemingly weaker position may hesitate to commit much of the valued resource at an upcoming turn. If he makes big stakes consistently, and performs better at several forthcoming turns, he may for example become able to refight a lost battle, but he knows that his determined opponent may well commit again to win that battle; so his route to victory will likely entail high expenses in the resource, and he may prefer to *cut his losses*, accepting a bad terminal game result but consoled by minimal resource loss in the later game. To the opposing player we might send a different message, namely *don't take your foot off the gas*: if the weakly positioned player will stake very little, his opponent hardly needs to stake intensely, but she should make sure to outspend her opponent significantly, because she wishes to convert a promising position into a terminal victory without setbacks.

The three premises will be called the *Present-Future Compromise*; “*Incentive Inch, Outcome Mile*”; and the *Battlefield Cyl Fog*.

The topic of strategic importance of positions in multi-turn games is of course natural and important: [32, 11, 30] have explored such themes in the framework of iterative network-bargaining problems. To interpret and investigate the broad and perhaps vague premises that we have just outlined, a suitable class of games is clearly needed. Such a class is offered by stake-governed tug-of-war games, in which players win the right to move (a counter along the edges of a graph) according to the flip of a coin with a bias determined by their commitment of resources at the turn.

1.1.2. Tug-of-war, constant-bias and stake-governed. We present tug-of-war first in the simple case of constant bias; and then versions governed by stakes from budgets that are either allocated or funded by players. The Trail Of Lost Pennies, a player-funded game on the integers, is then introduced.

1.1.2.1. *Constant-bias tug-of-war.* Let $G = (V, E)$ be a finite graph equipped with a set $B \subset V$ of boundary vertices on which a function $f : B \rightarrow \mathbb{R}$ is defined. Further suppose given a bias parameter $\lambda \in (0, \infty)$. The initial counter location X_0 of λ -biased tug-of-war is taken to be a given $v \in V$. At the i^{th} turn (where $i \geq 0$), a coin with head's probability $\frac{\lambda}{1+\lambda}$ is flipped. If it lands heads, Maxine moves the counter to a neighbour of its present location of her choosing, specifying $X_{i+1} \in V$ with $X_{i+1} \sim X_i$. If it lands tails, then Mina does so. The game ends at time $\tau = \min\{i \in \mathbb{N} : X_i \in B\}$, when the counter first reaches the boundary, with a payment of $f(X_\tau)$ made from Mina to Maxine. The unbiased case $\lambda = 1$ is considered in [49], and the biased $\lambda \in (0, \infty) \setminus \{1\}$ in [47]; either way, the value of tug-of-war—in essence, the mean terminal payment under optimal play—exists and equals $h(\lambda, v)$, where $h(\lambda, \cdot) : V \rightarrow \mathbb{R}$ is the λ -biased infinity harmonic extension of the boundary data $f : B \rightarrow \mathbb{R}$; namely, the unique function satisfying the system of equations

$$h(\lambda, v) = \frac{\lambda}{1+\lambda} \max_{u \sim v} h(\lambda, u) + \frac{1}{1+\lambda} \min_{u \sim v} h(\lambda, u) \quad (1)$$

for $v \in V \setminus B$ and $h(\lambda, \cdot)|_B = f$.

1.1.2.2. *Allocated-budget stake-governed tug-of-war.* Let the data $G = (V, E)$, $B \subset V$, $f : B \rightarrow \mathbb{R}$ and $\lambda \in (0, \infty)$ again be given. The two players are now allocated a budget at the outset of the game: λ for Maxine; and one for Mina. Take $X_0 = v \in V$ as above. At the outset of each turn of the new game, each player retains some part of her initial budget. Each stakes some part of her remaining budget, as she sees fit. These staked funds are withdrawn from the players, so that the new budget of a player, for use at the next turn, is given by subtracting her stake from the budget's present value. If Maxine stakes a and Mina b , a biased coin is then flipped, whose probability of landing heads equals $\frac{a}{a+b}$. The turn victor moves the counter according to the rules of constant-bias tug-of-war, and the game ends just as it did originally, with a terminal payment from Mina to Maxine equal to the evaluation of the function f at the terminal counter location in the boundary B . Any remaining budget held by the players is taken from them.

1.1.2.3. *Player-funded stake-governed tug-of-war.* In the allocated-budget stake game, the initially awarded budgets are an irredeemable resource whose sole role is to afford capacity to the respective players to win moves throughout the lifetime of the game. Maxine and Mina's initial budgets are given finite quantities whose values are part of the game design. The finiteness of these values is what makes the resource precious. In the variant we now specify, a different means is used to ensure that the players value this resource: each player can spend without constraint, but she must spend *her own money*. Mina and Maxine are supposed to be wealthy, so each may stake as freely as she sees fit, but the sequence of stakes made by a player constitutes a running cost that must be offset against the potential benefit that higher expenditure will bring the counter to a more favourable terminal location in the boundary set.

The game data now takes the form $G = (V, E)$, $B \subset V$ and $f, g : B \rightarrow \mathbb{R}$. Thus a second boundary function $g : B \rightarrow \mathbb{R}$ is introduced (and the parameter λ is absent, reflecting the players' wealth). At the start of the i^{th} turn, Maxine stakes $a_i \in [0, \infty)$ and Mina $b_i \in [0, \infty)$. A coin is flipped whose head's probability is $\frac{a_i}{a_i+b_i}$. As above, heads means Maxine wins the right to move the counter to an adjacent vertex; tails, and Mina does. The game ends, as above, when the counter reaches B at time τ . The terminal payments are $f(X_\tau)$ to Maxine, and $g(X_\tau)$ to Mina. Thus the net receipt in the game equals

$$f(X_\tau) - \sum_{i=0}^{\tau-1} a_{X_i} \text{ for Maxine, and } g(X_\tau) - \sum_{i=0}^{\tau-1} b_{X_i} \text{ for Mina.} \quad (2)$$

In the special case that g equals $-f$, the terminal payments coincide with those in the games already discussed. Unlike these games, the player-funded stake game is not zero-sum, even when $g = -f$; since zero-sum tools will not be used to analyse the game, it is natural to generalize it with separate terminal payment functions f and g .

1.1.2.4. *The Trail of Lost Pennies.* The constant-bias, and allocated-budget stake-governed, tug-of-war games make little sense on certain infinite graphs. Consider constant-bias tug-of-war on the integers \mathbb{Z} with nearest-neighbour edges, and a payment of one unit from Mina to Maxine if the counter tends to ∞ , and of minus one unit if instead it tends to $-\infty$. The players have no choices for strategy and the counter evolves as simple random walk, so no terminal payment is made (or perhaps a default rule stipulates the payment). And likewise the game on \mathbb{Z} makes no real sense for the allocated-budget stake game: roughly put, since the game will require infinitely many turns, any positive expenditure of the globally finite budget of a given player is unjustified at any given turn; but if both players consistently stake nothing, then (at least if a symmetric rule is adopted for this circumstance) the counter will again evolve as simple random walk.

It may seem that the self-funded stake game would be equally trivial when played on \mathbb{Z} . Since this gameboard is an infinite transitive graph, there seems to be nothing to distinguish one move from any other in the infinite game, and this seems to make any positive stake by either player at any turn unjustified. This consideration in fact permits one in essence to disprove the uniqueness of a Nash equilibrium in the game, but it is quite false insofar as it suggests that the game is trivial. Indeed, self-funded stake-governed tug-of-war has an intricate theory on \mathbb{Z} . In this article, we investigate the new game when the underlying graph is either \mathbb{Z} or a finite interval therein. We call the game on these graphs the Trail of Lost Pennies—either on \mathbb{Z} , or on a given finite integer interval.

The attentive reader will have noticed imprecisions in our specification of stake games. What happens if both players stake nothing at a given turn? Or if the counter fails to reach the boundary in a suitable sense? We defer resolving these points from the present conceptual overview to Chapter 2, where the Trail of Lost Pennies will be precisely specified.

1.1.3. An informal summary of the main conclusions. Our principal focus will be on the Trail of Lost Pennies on \mathbb{Z} , a game that is necessarily of infinite duration. Our main theorems show that this simple game has a remarkably rich and surprising structure of Nash equilibria, which witnesses two of the broad premises, “Incentive Inch, Outcome Mile” and the Battlefield Cyl Fog.

In Chapter 2, we will carefully state our principal conclusions about the game. In summary now of the game rules and our results about it, we may say that our game data is \mathbb{Z} with nearest-neighbour edges; two boundary points, at $-\infty$ and ∞ ; and a parameter $x \in (0, \infty)$. If the counter tends to $-\infty$ over the course of the game, then a terminal payment of one is made to Mina, and of zero to Maxine. If the counter tends to ∞ , Maxine will receive x , and Mina zero. We will prove that Nash equilibria exist in the game precisely when $x \in [\lambda, \lambda^{-1}]$ for a certain value $\lambda \in (0, 1)$. The quantity λ is canonically associated to \mathbb{Z} via a simple enough game, making its sheer closeness to the value one remarkable: we will prove that $\lambda \leq 0.999904$ and present clear numerical evidence that $\lambda \geq 1 - 10^{-4}$. (On long finite interval intervals, the rise of x roughly from λ to λ^{-1} is a transition from a phase of Mina’s dominance to one of Maxine’s via a brief phase with many equilibria. An incentive gain of 10^{-4} leads to a unit-order gain in outcome: the order of magnitude of the ratio of mile to inch is respected.) When $x \in [\lambda, \lambda^{-1}]$ is given, we will exhibit countably many Nash equilibria in the game. Each of these equilibria will be computed explicitly and will be found to witness the Battlefield Cyl Fog: each has an integral battlefield index around which both players stake at unit order; as the counter moves to the right from this battlefield, Mina cuts her losses and

stakes minutely, while Maxine also rapidly reduces spending but at a much slower rate (so that her foot is, in a relative sense, on the gas); the roles are reversed to the left of the battlefield. Note that the countably many Nash equilibria break the translation symmetry of \mathbb{Z} : the consideration mooted above, that this symmetry would render the game trivial, is thus invalidated, because it is the collection of equilibria that is translation invariant, even as each equilibrium is a non-trivial object, centred on an integral battlefield, and lacking translation invariance.

By finding the two premises to be exemplified, our study offers a detailed answer to a question concerning the strategic importance of positions and the optimal allocation of valued resources in multi-turn games. (The question is of course very general, and many answers have been offered. As we will review later in the introduction, economists' investigations of tug-of-war have often addressed the premise of a site of intense competition with rapid and asymmetric decay of effort rates away from the site.) The essence of the answer we offer may well be valid more generally: as we will discuss further in Section 7.3, player-funded stake-governed games may on many graphs have a space of multiple Nash equilibria that exist only in a narrow range of values of the relative terminal reward and whose elements are roughly characterized by a small 'battlefield' zone that roughly divides the gameboard into two regions in which the subsequent victory of one or other player is close to assured.

1.1.4. Overview: allocated-budget stake games and the Present-Future Compromise. The guiding principle governing allocated-budget stake games stands in contrast to those of the player-funded games that are the focus of this article. The difference is stark: a staccato burst of stakes that punctuates near peace in player-funded games; a much smoother flow of expenditure when fixed budgets are allocated. In order to put the contrast into relief, we here give a brief overview of the sense in which [24] shows that the allocated-budget games verify the Present-Future Compromise. Suppose that, in such a game, Maxine adopts a strategy that involves placing deterministic stakes, and that Mina learns this strategy. She may mimic Maxine's strategy, by staking the same proportion of her remaining budget at any given turn as Maxine does of hers. The relative fortune of the players is thus maintained at its initial value, and this permits Mina in effect to reduce the game to the λ -biased tug-of-war. And Maxine can do likewise if the roles are reversed. Under the premise that players stake deterministic amounts, we find that, under optimal play, the two players stake at any given turn the same proportion of their remaining budget. Call this proportion the 'stake function' $S(\lambda, v)$ when the counter is at $v \in V$, Mina's budget is one, and Maxine's is λ . We may seek to determine the value of $S = S(\lambda, v)$ by a perturbative argument: under optimal play, Mina will stake S and Maxine λS at the upcoming turn. If Maxine slightly increases her stake, she will gain from an increased probability of winning the next turn, with a resulting improvement in position; and she will lose, because she will have a smaller budget from the start of the next turn. These gain and loss terms cancel to first order under infinitesimal perturbation, since λS is Maxine's optimal stake. The formula that results is

$$S(\lambda, v) = \frac{\max_{u \sim v} h(\lambda, u) - \min_{u \sim v} h(\lambda, u)}{(\lambda + 1)^2 \frac{\partial}{\partial \lambda} h(\lambda, v)}, \quad (3)$$

where $h(\lambda, v)$ is the λ -biased infinity harmonic function in (1). This formula witnesses the Present-Future Compromise, with the right-hand numerator representing the short-term demand to stake big and win the next turn, and the denominator reflecting the long-term via the need to retain budget for later turns. But the just sketched derivation is non-rigorous. Two notable problems are that the proposed formula depends on the differentiability in λ of $h(\lambda, v)$; and that the perturbative argument identifies merely a local saddle point, where a global one is needed. Both of these problems are real. The function $h(\lambda, v)$ does fail to be differentiable at certain values of λ for some simple

graphs. And the local saddle point often fails to be global, essentially because, when the counter is at one step from the boundary B , one or other player may be tempted to *go for broke*, staking all of her budget in an effort to close out the game at the present turn. In [24], these problems are respectively addressed by restricting the class of game data and by modifying the game rules. The class of root-reward trees, in which $G = (V, E)$ is a tree with B its set of leaves, and $f = \mathbf{1}_r$ for a distinguished leaf r , is considered, because $\lambda \rightarrow h(\lambda, v)$ is then differentiable for each $v \in V$. And the game is modified to a leisurely version, in which, after stakes are offered at any given turn, it is decided that a move will take place only with a small positive probability. The leisurely game disables the efficacy of the go-for-broke deviation and renders the local saddle point global. The principal result of [24] operates under these assumptions and proves in essence that the stake function indeed takes the form (3). It is in this sense that [24] shows that a class of allocated-budget stake-governed tug-of-war games satisfies the Present-Future Compromise.

1.2. The two strands of tug-of-war

In 1987, Harris and Vickers [26] proposed a model of two firms that compete to secure a patent. They called their model tug-of-war, and the notion and this name appear to have been introduced by them. The model is quite close to player-funded stake-governed tug-of-war played on a finite integer interval, and will explain it shortly. The model was intended to describe how firms may compete over time when the effect of a technological advantage is significant but uncertain, and when the firms' decisions are influenced by their relative progress. Their work has inspired a wealth of research, mainly in economics but also among computer scientists, on tug-of-war and related bidding games: [26] has a total of almost six hundred citations as of early 2023 according to google scholar.

In 2009, Peres, Schramm, Sheffield and Wilson's [49] considered constant-bias tug-of-war with a fair coin. The relation they found between this game in a Euclidean setting and infinity harmonic functions sparked a wave of attention among analysts and probabilists, with almost five hundred works citing this article.

These two contributions study variants of the same basic concept—the two sets of authors even appear to have arrived independently at the name ‘tug of war’—and each has influenced many researchers, but they seem to have done so with a remarkable disjointness: none of the 1100 references that cite one or other article cites both. The author of this monograph was until recently quite unaware of the origin and development of tug-of-war in economics. Mathematicians may benefit from learning of this literature, and there is surely value in communication between the concerned communities of researchers. Next then we review developments for tug-of-war. First economics, and then mathematics.

1.2.1. Tug-of-war in economics. We begin by reviewing [26], whose treatment is accompanied by the proof-supplying mimeograph [25] and Vickers' 1985 DPhil thesis [58].

1.2.1.1. *Harris and Vickers' 'Racing with uncertainty'*. In 1980, Lee and Wilde [41] proposed a single-stage model for the relationship between research effort and outcome. Mina and Maxine (as we call them) select respective effort rates $x, y \in (0, \infty)$. Maxine's discovery time is M_+ , an exponentially distributed random variable of mean x^{-1} ; Mina's is M_- , an independent exponential of mean y^{-1} . The player with the lower discovery time $M = \min\{M_-, M_+\}$ wins the race, and receives a reward. Maxine and Mina have respective cost functions $c_-(x)$ and $c_+(y)$, and the payments they make are $Mc_-(x)$ and $Mc_+(y)$. In other words, c_- and c_+ describe the rate of payment incurred by the given player as a function of her effort rate, so the payments reflect that

the efforts are sustained for time M . Since M has mean $(x + y)^{-1}$, Maxine and Mina's mean costs with joint effort rates (x, y) are $\frac{c_-(x)}{x+y}$ and $\frac{c_+(y)}{x+y}$.

Lee and Wilde's model is a single-turn player-funded tug-of-war. It fits in the framework of Subsection 1.1.2.3 if we take G equal to $\{-1, 0, 1\}$ with the adjacency relation, $B = \{-1, 1\}$, and provided that in the player payments (2) we replace a_{X_0} by $\frac{c_-(a_0)}{a_0+b_0}$ and b_{X_0} by $\frac{c_+(b_0)}{a_0+b_0}$, (where note that $X_0 = 1$ and $\tau = 1$).

The original tug-of-war model of Harris and Vickers has gameboard¹ $\llbracket 0, N + 1 \rrbracket$ and boundary set $B = \{0, N + 1\}$. It is the player-funded tug-of-war whose single-step rule is Lee and Wilde's: the players offer effort rates at each turn, and the above rule decides the outcome according to given turn-independent cost functions c_- and c_+ . Formally then we may specify the model by taking G to be an integer interval and replacing $a_{X_i} \rightarrow \frac{c_-(a_{X_i})}{a_0+b_{X_i}}$ and $b_{X_0} \rightarrow \frac{c_+(b_{X_i})}{a_{X_i}+b_{X_i}}$ in (2).

Also formally, player-funded stake-governed tug-of-war is the special case of Harris and Vickers' model if we take $c_-(x) = x(x + y)$ and $c_+(y) = x(x + y)$, though the dependence of these cost functions on the opponent's strategy make them inadmissible for this framework.

Harris and Vickers use a fixed point theorem prove that, for a fairly general class of cost functions c_- and c_+ , there exists a Nash equilibrium that is *stationary* in the sense that players' effort rates are determined merely by present counter location. If Maxine and Mina's effort rates are labelled x_i and y_i when the counter is at $i \in \llbracket 1, N \rrbracket$, then the equilibrium may be called symmetric if $x_i = y_{N-i}$. And the game itself is called symmetric if $c_- = c_+$ and the terminal payments to the two players are equal after interchange of endpoints. Harris and Vickers prove that the symmetric game has a symmetric equilibrium, and that there is only one symmetric equilibrium when $c_-(z) = c_+(z) = z^\eta$ for some $\eta > 1$ or when the gameboard is short and under weaker hypotheses on the cost functions. The symmetric equilibrium is computed explicitly in the example where N is even and $\eta = 2$. The gameboard has a central edge $\llbracket N/2, N/2 + 1 \rrbracket$. It is found that Maxine's effort x_i peaks at $x_i = N/2 + 1$, the right end of this edge; that the ratio x_n/x_{n-1} equals $2^{1/2+(N/2+1)-n}$; and that $y_n/y_{n-1} = 2^{-1}x_n/x_{n-1}$. Thus Maxine's effort drops sharply to the right of $N/2 + 1$, with a ratio of efforts at consecutive vertices having exponential decay in the distance to $N/2$; Mina's effort drops even more rapidly, with the ratio of her effort to Maxine's decaying exponentially in the same sense. This computation illustrates the main message of Harris and Vickers: as they write in the abstract, 'in our principal model, which is of a one-dimensional race, it is shown that the leader in the race makes greater efforts than the follower, and efforts increase as the gap between competitors decreases'. In summary, we may say that tug-of-war was first introduced principally to illustrate this effect. In the language of Section 1.1.1, the Battlefield Cyl Fog is exemplified, with the battlefield at the vertices $N/2$ and $N/2 + 1$.

Harris and Vickers briefly comment on the possible existence of asymmetric equilibria, mentioning that, for certain cost functions, computer simulations appear to indicate a total of $2N - 1$ equilibria: one symmetric, and remaining asymmetric in N pairs.

1.2.1.2. *Further treatments of player-funded tug-of-war in economics.* In 1980, Tullock [57] considered two-player single-stage contests. Player A stakes $x \in [0, \infty)$ and player B , $y \in [0, \infty)$, and the contest is won by A with probability $\frac{x^\gamma}{x^\gamma + y^\gamma}$, where $\gamma \in (0, \infty)$ is now known as the Tullock exponent. The case $\gamma = 1$ is called a lottery; this is the contest used in each step of the stake-governed tug-of-war games specified in Section 1.1.2. Taking $\gamma = 0$ yields constant-bias tug-of-war with a fair

¹For $i, j \in \mathbb{Z}$, $i \leq j$, we write $\llbracket i, j \rrbracket = \{k \in \mathbb{Z} : i \leq k \leq j\}$.

coin. And the limit $\gamma \rightarrow \infty$ leads to all-pay auctions, in which the higher staking player wins (with a tie-break rule needed for equal stakes), and both stakes are forfeit.

Konrad and Kovenock [36], Agastya and McAfee [1] and Konrad [37], considered player-funded tug-of-war on finite integer intervals with the all-pay auction rule used to decide turn victor. Konrad and Kovenock found the battlefield effect to be manifest in a strong form with this rule: players offer stakes in essence at just two vertices. Agastya and McAfee reach rather different conclusions, because of use of a different tie-break rule and of discounting of terminal payments. There are some equilibria where no stakes are ever placed. In others, a battlefield exists, with momentum for the stronger player away from this location. But, contrary to the Battlefield Cyl Fog, stakes rise near the boundary: discounting future payments makes the imminence of victory or defeat significant. Konrad [37] presents a discussion involving player-funded games with Tullock lottery contests, emphasising the discouraging effect of early losses on a player. (Here, we may compare how the premise ‘‘Incentive Inch, Outcome Mile’’ evokes discouragement that effects the less motivated player.)

Häfner [22] considered a ‘tug-of-war team contest’: player-funded tug-of-war on an integer interval with all-pay auction rules at each turn, but in which each player is in fact a team of countably many people, with each individual on a team responsible only for the payment at some given turn (but receiving nonetheless the full terminal payment for his team). It is shown that efforts diminish as the stronger team approaches victory: this conclusion, shared by Harris and Vickers (and by the Battlefield Cyl Fog), arises in this case in view of free-rider and effort-discouragement effects for team players staking when victory or defeat is close to assured. In [23], Häfner and Konrad consider a modified setup in which, for each team, a member is assigned to each site in open play, and is responsible for all stake payments made for play at that site. It is found that members of the leading team outspend their opponents (as in the Battlefield Cyl Fog) when there is little discounting of the terminal payment, and that the reverse holds when there is much discounting.

1.2.1.3. *The player-funded game with the majoritarian objective.* In tug-of-war on an integer interval gameboard symmetric about the origin, one player wins when she secures a number of turn victories that exceeds her opponent’s by a certain value. Harris and Vickers further considered the game where the game winner is the first to secure a certain number of victories. This *majoritarian objective* model fits in the framework of Subsection 1.1.2.3 if we consider what may be called *directed rectangle graphs*: that is, if we take $G = (V, E)$ to be a directed graph with $V = \llbracket 0, m \rrbracket \times \llbracket 0, n \rrbracket$, and E the set of north and east pointing nearest-neighbour edges; $f(x, y) = \lambda \mathbf{1}_{y=m}$ and $g(x, y) = \mathbf{1}_{x=n}$; and initially $X_0 = (0, 0)$. (Harris and Vickers’ stake-rules entail further changes as described in Subsection 1.2.1.1.) In this model, backward induction readily shows the equilibrium is unique under reasonable hypotheses on the Harris-Vickers cost functions c_- and c_+ . Numerical work in [26] with z^η -cost functions indicate some expected effects, such as the rapidly emerging dominance of the stronger player (the presence then of a battlefield); but also of surprises, such as stark oscillations in effort rates in the case $m = n$ as a function of $n \in \llbracket 1, 10 \rrbracket$.

The Democratic and Republican primary contests to determine the party’s nominee in a U.S. presidential election are run on a calendar of several months. In the ‘New Hampshire’ effect, candidates typically spend a significant fraction of their resources in the earliest contests. Klumpp and Polborn [35] considered as a model of this phenomenon majoritarian objective tug-of-war with a Tullock contest rule for deciding turns. The early burst of campaign expenditure and momentum effects discussed are similar to the Battlefield Cyl Fog, and contrast with a more uniform resource allocation in simultaneous contests, so that a sequential rather than simultaneous primary season is typically cheaper for the winning candidate.

Kovenock and Konrad [38] considered the majoritarian objective game with all-pay auction rules at the turn, with the introduction of intermediate prizes, or payoffs to players at non-terminal times. Gelder [21] shows that discounting dissipates the accrual of momentum to the stronger player in the majoritarian game with all-pay auction rules. Best-of-three contests have been examined by Malueg and Yates [45], and by Sela [55]. Fu, Lu and Pan [20] considered a team contest version of the majoritarian game.

Ewerhart and Teichgräber [19] introduce a player-funded tug-of-war framework that is more general than integer interval or first-to- N contests. Operating under assumptions including that gameplay before termination has a certain exchangeability of move order, they show that gameplay takes place in an initial phase, in which no state may be revisited, followed by a later tug-of-war phase. They consider symmetric equilibria, showing that there exists a unique such that is Markov perfect (or stationary, in the above usage), for a broad range of contest rules; and they discuss the problem of designing contests that are optimal for the tournament organizer by maximizing combined player expenditure. Existence and uniqueness of the symmetric stationary equilibrium in player-funded tug-of-war with a Tullock contest of exponent one-half at each turn were proved and this equilibrium analysed in [31].

1.2.1.4. *The allocated-budget games.* In 1921, Emile Borel considered a game [15] that pits a protagonist now commonly called Colonel Blotto against an opponent. Borel gives an example: players A and B each choose an ordered triple of non-negative reals that sum to one. Player A wins the game if at least two of his values exceed the counterparts for B . In general, the colonel and his opponent each allocate a given resource across a fixed number of battlegrounds, with a given battle won by the better committed player, and the battle-total determining the winner of the war. The label ‘Blotto’ is sometimes attached to games where players have fixed resources, whether they be committed simultaneously as in Borel’s game, or for example sequentially, as in budget-allocated tug-of-war from Subsection 1.1.2.2.

In Section 1.1.4, we advanced a premise for how to play the Blotto game with the lottery contest function $\frac{x}{x+y}$: stake a proportion (3) of your remaining fortune, whether you be Mina or Maxine. But we also indicated that even tug-of-war on $\llbracket 0, 3 \rrbracket$ is a counterexample to this premise. Rather than altering the game (as [24] did), we may ask whether a class of graphs can be found for which the premise is correct. The majoritarian objective games—tug-of-war on directed rectangle graphs—are a good starting point. In this game on $\llbracket 0, 1 \rrbracket \times \llbracket 0, n \rrbracket$, there are n contests, with one player needing to win every one, and the other requiring only one victory. A symmetry consideration points to a stake proportion of $1/n$ being optimal when $X_0 = (0, 0)$. This choice in fact validates (3), as [24, Section 2.6] explains. Klumpp, Konrad and Solomon [34] generalise this example to show that the game on any directed rectangle graph, and for a broad range of turn contest functions, has an equilibrium with an even split of resources across turns.

Allocated-budget tug-of-war on a finite integer interval has been examined by Klumpp [33]. In note 1 on page 27, he notes an example, similarly as does [24], in which a formula equivalent to (3) is invalidated when the weaker player is one step from winning. Where [24] turned to a leisurely version of the game to regularise this example, Klumpp considers lowering the Tullock exponent from the lottery value $\gamma = 1$. He surmises from a numerical investigation that the formula in question is likely to offer an equilibrium precisely when γ at most one-half.

1.2.1.5. *An experimental study.* In 2002, Zizzo [60] reported outcomes of a study of human players in a majoritarian objective, first-to-10 game. The experimental setting was designed to replicate as closely as possible the Harris-Vickers first-to-10 game with cost functions $c_-(x) =$

$c_+(x) = x^2$. Leaders clearly outspent their opponents only when a significant gap had opened, and the Harris-Vickers predictions appeared not to be well supported by the tests. Two differences between theory and setup should be noted, however.

First, although players retained unspent budgets and thus in effect lost spent monies as in a player-funded game, they also operated under a budget constraint, which equalled one-half of the terminal payment to the victor. Zizzo acknowledges this difference and writes “to minimize the possibility of distortions that might occur ... as a result of a binding budget constraint ... instructions stressed the need for subjects to retain enough money to fund the later rounds of the competition in an unconstrained way.” Such an admonition was likely intended to prevent player bankruptcy, the most blatant manifestation of budget constraint. But it also illustrates the role of the future in the Present-Future Compromise seen in allocated-budget games. The setup was a hybrid of player-funded and allocated-budget rules, and it is conceivable that the outcomes felt the effects of both regimes.

Second, recall that Harris and Vickers’ cost functions $c_-(z)$ and $c_+(z)$ are rates of charge to players as a function of their effort rates z . Zizzo’s implementation sets costs equal to z^2 but treats them simply as the payments made by players at the turn. When Mina and Maxine stake x and y at a given turn, their mean payments at the turn are $\frac{x^2}{x+y}$ and $\frac{y^2}{x+y}$ for quadratic Harris-Vickers; for Zizzo, the payments are deterministic quantities, x^2 and y^2 . Zizzo’s framework at a given turn is readily seen to be equivalent to a Tullock contest from Subsection 1.2.1.2, with exponent γ equal to one-half. The player-funded majoritarian objective game with this framework may well have qualitatively similar theoretical features to those that Harris and Vickers identified and predicted, but this premise deserves scrutiny in interpreting Zizzo’s work as an experimental test of Harris and Vickers’ theory.

1.2.2. Tug-of-war in mathematics. Peres, Schramm, Sheffield and Wilson [49] considered constant-bias tug-of-war with a fair coin. Under optimal play, Maxine moves to an h -maximizing neighbour of the counter’s present location, and Mina to an h -minimizing one, where h is the discrete infinity harmonic extension of boundary data $f : B \rightarrow \mathbb{R}$; namely, $h(\cdot)$ equals $h(\lambda, \cdot)$ from (1) with $\lambda = 1$. In this way, h is shown to be the game’s value. This is perhaps the simplest connection between discrete infinity harmonic functions and tug-of-war games, but it was not the first. Richman games were introduced in [40] and further studied in [39]. A first-price Richman game is an allocated-budget tug-of-war game with an auction to determine the turn victor (so that the higher staking player wins), and with the turn victor paying the opponent his stake, the opponent making no payment at the turn. The threshold ratio is a value t such that, if the initial relative ratio of Maxine’s budget and the total budget for the two players exceeds t , then Maxine will almost surely win the game if she plays correctly, while Mina will win in this sense if this ratio is less than t . In [40], it is proved that the threshold ratio exists as a function of initial counter location and equals the discrete infinity harmonic extension of $f : B \rightarrow \mathbb{R}$. Several variants of these games—with all-pay rules, or infinite duration, or in the poormen variant, where the turn victor pays the bank, not the opponent—have been studied [5, 7, 8, 9] by theoretical computer scientists, who also address approximate algorithms and computation complexity: see [6] for a survey that treats these directions.

We have recalled constant-bias tug-of-war in the discrete context of given graphs in part because the framework of our results is discrete, but the principal aim of [49] was to relate tug-of-war in a low mesh limit to infinity harmonic functions mapping domains in \mathbb{R}^d to \mathbb{R} , these being continuum counterparts to (1). In ‘ ϵ tug-of-war’ for given $\epsilon > 0$, the gameboard is a domain D in \mathbb{R}^d , with

boundary data $f : B \rightarrow \mathbb{R}$, $B = \partial D$. Mina and Maxine (as we name them) move a counter—a point in D —a distance at most ϵ at each turn, according to the outcome of a fair coin flip. Mina pays Maxine the evaluation of f at the counter location when the game ends as the counter reaches B . In [49], the value of ϵ tug-of-war is shown in a low- ϵ limit to converge to the infinity harmonic function that extends f to D . This is the viscosity solution $h : D \rightarrow \mathbb{R}$ of the infinity Laplace equation $\sum_{i,j} \partial_{x_i} \partial_{x_i x_j} \partial_{x_j} h = 0$ subject to $h|_B = f$. The infinity Laplace operator is a degenerate second-order differential operator that was introduced and studied by Aronsson [3, 4] in 1967 and that has a subtle uniqueness [29] and regularity [54, 18] theory. Its viscosity solutions [16] are absolutely minimizing Lipschitz functions extending the given boundary data [29].

As [46] surveys, the game theory connection identified in [49] sparked much attention from the PDE community. In [2], boundary rules in ϵ tug-of-war were altered from [49] to obtain more regular game value functions. In [13], random move sizes with a heavy tail (judged on scale ϵ) were introduced, thus connecting to the infinity fractional Laplacian. Suppose that ϵ tug-of-war is played with a random displacement of magnitude $c\epsilon$ of the counter at the end of each move. For $p \in (2, \infty)$, the p -Laplacian [44] is an interpolation of the usual Laplacian operator (for $p = 2$) and the infinity version. In [50], it is shown that the value of this game converges to a p -harmonic function for p suitably chosen as a function of c : see [42] for a survey centred around this perspective. A variant of this game has been used to study p -Laplacian obstacle problems in [43]. The book [14] reviews the abundant connections between tug-of-war and PDE.

It is the fair case $\lambda = 1$ of constant-bias tug-of-war that [49] considers. The biased case $\lambda \in (0, \infty)$ was addressed in [47]: ϵ tug-of-war was modified so that a suitable drift is maintained in the low- ϵ limit, and a continuum biased infinity harmonic function is then obtained as limiting game value. In [51], a rapid algorithm for computing biased discrete infinity harmonic functions (1) was found in terms of a path decomposition of the graph (V, E) . As the bias parameter λ varies, this decomposition changes at certain values. It is these changes that cause the lack of differentiability in $(0, \infty) \rightarrow \mathbb{R} : \lambda \rightarrow h(\lambda, v)$ to which we alluded in Subsection 1.1.4.

Tug-of-war played on directed graphs has also garnered attention among mathematicians. Many combinatorial games, such as Go and Hex, are zero-sum two-player games in which players alternate in make directed moves by placing counters that cannot be moved later. In [48], random-turn versions of these ‘selection games’ were considered. Hex is a game discovered by Piet Hein in 1942 in which two players alternately place red or blue hexagons on a hexagonal grid, each seeking to forge a connection between a pair of diagonally opposed boundary segments in his own colour. It is a complex game [27] which has no explicit solution on 11 by 11 boards on which it is played competitively. In [48], the solution to random-turn Hex (among other selection games) was found explicitly in terms of the maximum pivotal probability for critical percolation in the unplayed region. This work offers a perspective on the strategic importance of positions in multi-turn games.

To conclude, we first mention two broad antecedents to stake-governed random-turn games. Rufus Isaacs’ differential games [28], include examples where a pursuer seeks to capture an evader in Euclidean space, and each instantaneously selects certain control variables such as velocity, analogously to Mina and Maxine’s ongoing selection of stakes. In Shapley’s stochastic games [56], transition probabilities governed step-by-step by decisions of two players encompass the stake game framework in a broad sense.² These are fairly loose connections, however. Regarding tug-of-war itself, we have seen how mathematicians have tended to focus on constant-bias models. They have

²Both differential and stochastic games have attracted ongoing attention. Several chapters of a recent handbook [12] treat dynamic game theory, and [53] collects articles about stochastic games.

explored a much wider geometric setting than the finite integer intervals and (what we have called) the directed rectangle graphs that are common elements in the economics literature, exploring these games' implications in probability and PDE. Economists originated tug-of-war, considering from the outset bidding rules beyond the trivial constant-bias, and have sustained attention on these games for almost four decades, putting the theme of resource allocation at the heart of their investigations throughout. This article and [24] offer detailed mathematical treatments of resource-allocation tug-of-war games: with allocated budgets in [24], and with player funding in the present article. Raising the prospect of uncovering further beautiful mathematical structure, and informing theoretical analysis of economic behaviour, tug-of-war in its original resource-allocation guise would surely repay the further attention of mathematicians.

1.2.3. Acknowledgments. The author thanks Gábor Pete for many discussions about stake-governed games. He thanks Judit Zádor for help with Mathematica and in preparing the article's figures. He is supported by the National Science Foundation under DMS grants 1855550 and 2153359 and by the Simons Foundation as a 2021 Simons Fellow.

CHAPTER 2

Main results

In this chapter, we specify precisely the Trail of Lost Pennies and state our main conclusions. The structure of the remainder of the article is explained at the end of the chapter.

2.1. Game setup, strategies and Nash equilibria

The Trail of Lost Pennies on \mathbb{Z} will be denoted by $\text{Trail}(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$. The data that specifies the game takes the form of

$$\text{a quadruple } (m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4 \text{ that satisfies } m_{-\infty} < m_{\infty} \text{ and } n_{\infty} < n_{-\infty}. \quad (4)$$

For any given $k \in \mathbb{Z}$, we will specify the gameplay of $\text{Trail}(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ where the initial location of the counter is equal to k . The counter's location X will evolve from its initial location $X_0 = k$ in discrete time-steps, the result being a stochastic process $X : \mathbb{N} \rightarrow \mathbb{Z}$ whose law is determined by k (we take \mathbb{N} to include zero). This random process is specified by the pair of strategies adopted by Mina (who plays to the left) and Maxine (who plays to the right).

For either player, a strategy is a map $S : \mathbb{Z} \times \mathbb{N}_+ \rightarrow [0, \infty)$. A player who follows the strategy S stakes $S(X_{i-1}, i)$ at the turn with index $i \geq 1$. Let \mathcal{S} denote the set of strategies. An element of \mathcal{S}^2 is called a strategy pair. A generic element of \mathcal{S}^2 will be written (S_-, S_+) , where the respective components are the strategies of Mina and Maxine.

We wish then to specify the gameplay process $X : \mathbb{N} \rightarrow \mathbb{Z}$ as a function of a given element $(S_-, S_+) \in \mathcal{S}^2$ and the initial location $X_0 = k \in \mathbb{Z}$. We will write \mathbb{P}_{S_-, S_+}^k for a probability measure that specifies this gameplay process and accompanying aspects of the game; the associated expectation operator will be written $\mathbb{E}_{S_-, S_+}^k[\cdot]$. At the turn with index $i \in \mathbb{N}_+$, Mina stakes $S_-(X_{i-1}, i)$ and Maxine stakes $S_+(X_{i-1}, i)$. By sampling of independent randomness, the turn victor is declared to be Maxine with probability $\frac{S_+(X_{i-1}, i)}{S_-(X_{i-1}, i) + S_+(X_{i-1}, i)}$; in the other event, it is declared to be Mina. Maxine will elect to move the counter one place to the right if she is the turn victor; Mina, one place to the left. (It is intuitive given the rules of the game that we are specifying that the two players will always elect to move the counter in the said directions, and we will not furnish the straightforward details to the effect that permitting other options changes nothing essential about the game.) Should neither player make a stake at the given turn—that is, if $S_-(X_{i-1}, i) = S_+(X_{i-1}, i) = 0$ —then a further rule is needed to permit play to continue. We will declare that, in this event, each player wins the right to move with equal probability (with Maxine moving right, and Mina left, as usual).

Formally, then, our counter evolution satisfies the condition that, for $(k, i, \ell) \in \mathbb{Z} \times \mathbb{N}_+ \times \mathbb{Z}$,

$$\mathbb{P}_{S_-, S_+}^k \left(X_i - X_{i-1} = \ell \mid X_j, j \in \llbracket 0, i-1 \rrbracket \right) = \frac{S_-(X_{i-1}, i)}{S_-(X_{i-1}, i) + S_+(X_{i-1}, i)} \mathbf{1}_{\ell=-1} + \frac{S_+(X_{i-1}, i)}{S_-(X_{i-1}, i) + S_+(X_{i-1}, i)} \mathbf{1}_{\ell=1},$$

where recall that we use the integer-interval notation $\llbracket i, j \rrbracket = \{\ell \in \mathbb{Z} : i \leq \ell \leq j\}$, $i, j \in \mathbb{Z}$. Note that, in reading the ratios on the right-hand side in the display, we adopt the convention that $0/0 = 1/2$.

We further wish to specify the other pertinent features of the game when the strategy pair (S_-, S_+) is played. These features are the resulting payoffs to Mina and Maxine. Mina's payoff P_- is the sum of a negative term given by the total costs incurred to Mina during gameplay, and a further term that is the terminal payment that is made to her. Indeed, we may write

$$P_- = - \sum_{t=1}^{\infty} C_-(t) + T_-, \quad (5)$$

where $C_-(t)$ denotes the cost incurred to Mina at the turn with index $t \in \mathbb{N}_+$, and T_- equals the terminal payment to Mina. We have then that the cost $C_-(t)$ is equal to Mina's stake $S_-(X_{t-1}, t)$.

The terminal payment T_- is in essence given by $n_{-\infty}$ if Mina wins the game by eventually bringing the counter infinitely far to the left; and to n_{∞} in the opposing event. However, a precise formulation is needed to make sense of this. We define the *escape* event E according to

$$E = \left\{ \lim_n |X_n| = \infty \right\}. \quad (6)$$

The *left* and *right* escape events are given by

$$E_- = \left\{ \limsup_n X_n = -\infty \right\} \quad \text{and} \quad E_+ = \left\{ \liminf_n X_n = \infty \right\}. \quad (7)$$

Note that E_- and E_+ are disjoint events whose union equals E . We regard them as victory events for Mina and Maxine, and accordingly set the terminal payment to Mina as follows:

$$T_- = \begin{cases} n_{-\infty} & \text{when } E_- \text{ occurs,} \\ n_{\infty} & \text{when } E_+ \text{ occurs,} \\ n_* & \text{when } E^c \text{ occurs.} \end{cases} \quad (8)$$

Here n_* is a given real value that is at most n_{∞} . By assigning this terminal payment to Mina in the event of non-escape, we ensure that this payment no more generous than that made in the event E_+ of her defeat.

We may specify Maxine's payoff

$$P_+ = - \sum_{t=1}^{\infty} C_+(t) + T_+ \quad (9)$$

with counterpart interpretations for the right-hand terms: the cost $C_+(t)$ incurred to Maxine at the turn with index $t \in \mathbb{N}_+$ equals Maxine's stake $S_+(X_{t-1}, t)$, while the terminal payment T_+ that she receives is given by

$$T_+ = \begin{cases} m_{-\infty} & \text{when } E_- \text{ occurs,} \\ m_{\infty} & \text{when } E_+ \text{ occurs,} \\ m_* & \text{when } E^c \text{ occurs,} \end{cases} \quad (10)$$

where m_* is a given real value¹ that is at most $m_{-\infty}$.

The quantities labelled P , C and T for Mina and Maxine are determined by the gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$. The gameplay and these other random variables are thus coupled together under the law \mathbb{P}_{S_-, S_+}^k . Of course, starting location $X_0 = k$ and strategy pair (S_-, S_+) are fundamental for determining game

¹Note that $(m_*, n_*) \in \mathbb{R}^2$ and $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$ are the parameters that specify the Trail of Lost Pennies on \mathbb{Z} . We thus speak imprecisely when we refer to the latter quadruple as the game's data. Given the upper bounds that we impose on them, the values of m_* and n_* will be immaterial for our analysis.

outcome including the above described quantities. In our notation, this dependence is communicated by the labels of the law \mathbb{P}_{S_-, S_+}^k , rather than in the annotations P_-, T_- , and so on.

A strategy is *time-invariant*² if $S(i, j)$ is independent of $j \in \mathbb{N}_+$ for every $i \in \mathbb{Z}$. The set of time-invariant strategies will be denoted by \mathcal{S}_0 . A time-invariant strategy pair $(S_-, S_+) \in \mathcal{S}^2$ may be identified with a pair of sequences $\{a_i : i \in \mathbb{Z}\}$ and $\{b_i : i \in \mathbb{Z}\}$, where $a_i = S_+(i, j)$ and $b_i = S_-(i, j)$ for $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$.

A strategy pair $(S_-, S_+) \in \mathcal{S}^2$ is a Nash equilibrium if

$$\mathbb{E}_{S_-, S_+}^k[P_+] \geq \mathbb{E}_{S_-, S}^k[P_+] \quad \text{and} \quad \mathbb{E}_{S_-, S_+}^k[P_-] \geq \mathbb{E}_{S, S_+}^k[P_-]$$

for all $S \in \mathcal{S}$ and $k \in \mathbb{Z}$. Note that this condition takes a strong form³, in that it stipulates the displayed bound for any initial condition $X_0 = k \in \mathbb{Z}$ for the counter location.

Let $\mathcal{N}(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty) \subset \mathcal{S}^2$ denote the set of Nash equilibria. Consider a time-invariant Nash equilibrium, namely an element (S_-, S_+) of \mathcal{S}_0^2 that satisfies the above condition: when such a strategy pair is played, neither player would gain by altering strategy, even if the proposed alternative strategy is not time-invariant.

In an abuse of notation, generic elements of \mathcal{S}_0 , for respective use by Maxine and Mina, will be called $(a_i : i \in \mathbb{Z})$ and $(b_i : i \in \mathbb{Z})$. In a further abuse, the accompanying element of \mathcal{S}_0^2 will be denoted⁴ $\{(b_i, a_i) : i \in \mathbb{Z}\}$.

2.2. Time-invariant Nash equilibria and the ABMN equations

We now begin to present our main results. We first introduce the ABMN equations, which will be fundamental to this study. Theorems 2.3 and 2.4 present basic properties of the equations' solution set, and Theorem 2.6 is the result that bridges between the equations and the trail game (as we will sometimes informally call the Trail of Lost Pennies). In Theorem 2.8, these theorems are leveraged to characterize when the trail game has time-invariant Nash equilibria in terms a condition on boundary data involving an important basic quantity, the *Mina margin*, which is introduced here. The section ends with Theorem 2.14, which offers precise asymptotic decay estimates for Nash equilibria as the index varies away from the *battlefield index*, at which the players are most likely to decide the ultimate outcome of a given game; and with its consequence Theorem 2.15, which describes gameplay at any Nash equilibrium.

DEFINITION 2.1. Let $\{(b_i, a_i) : i \in \mathbb{Z}\}$ denote a time-invariant strategy pair: namely,

$$\{(b_i, a_i) : i \in \mathbb{Z}\} \in \mathcal{S}_0^2.$$

Let $S_-, S_+ \in \mathcal{S}$ be strategies such that $S_-(i, j) = b_i$ and $S_+(i, j) = a_i$ whenever $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$.

Set $m_i = \mathbb{E}_{S_-, S_+}^i[P_+]$ and $n_i = \mathbb{E}_{S_-, S_+}^i[P_-]$ for $i \in \mathbb{Z}$. By this means, we have associated to any element $\{(b_i, a_i) : i \in \mathbb{Z}\} \in \mathcal{S}_0^2$ a \mathbb{Z} -indexed quadruple $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ of elements taking values in $[0, \infty)^2 \times (\mathbb{R} \cup \{-\infty\})^2$.

²This property of strategies corresponds to that called stationary by Harris and Vickers: see Subsection 1.2.1.1. And it is sometimes called Markov, since decisions depend on history only through the present state of play.

³The strengthened form is closely related to the notions of subgame perfect, and Markov perfect, equilibrium.

⁴In the strategy-pair notation $(S_-, S_+) \in \mathcal{S}^2$, governed by $- < +$, Mina precedes Maxine. Thus the notation (b, a) for strategy pairs will be standard. We will shortly introduce an (a, b, m, n) -quadruple notation for stakes and mean payoffs, in which Maxine precedes Mina (in the sense of 'a before b' and 'm before n'). As a result, usages of the form '(a, b, m, n) is the quadruple associated to the Nash equilibrium (b, a)' will be made.

DEFINITION 2.2. The ABMN system on \mathbb{Z} is the set of equations in the four real variables a_i , b_i , m_i and n_i , indexed by $i \in \mathbb{Z}$,

$$(a_i + b_i)(m_i + a_i) = a_i m_{i+1} + b_i m_{i-1} \quad \text{ABMN(1)}$$

$$(a_i + b_i)(n_i + b_i) = a_i n_{i+1} + b_i n_{i-1} \quad \text{ABMN(2)}$$

$$(a_i + b_i)^2 = b_i(m_{i+1} - m_{i-1}) \quad \text{ABMN(3)}$$

$$(a_i + b_i)^2 = a_i(n_{i-1} - n_{i+1}) \quad \text{ABMN(4)},$$

where i ranges over \mathbb{Z} . We will refer to the above equations throughout in the form ABMN(i), for $i \in \{1, 2, 3, 4\}$, rather than by a conventional numerical labelling. It is always supposed that a_i and b_i are non-negative for $i \in \mathbb{Z}$. A solution is said to have boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ when

$$\lim_{k \rightarrow \infty} m_{-k} = m_{-\infty}, \quad \lim_{k \rightarrow \infty} m_k = m_{\infty}, \quad \lim_{k \rightarrow \infty} n_{-k} = n_{-\infty} \quad \text{and} \quad \lim_{k \rightarrow \infty} n_k = n_{\infty}. \quad (11)$$

For such a solution, the *Mina margin* is set equal to $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}}$. A solution is called *positive* if $a_i > 0$ and $b_i > 0$ for all $i \in \mathbb{Z}$. It is called *strict* if $m_{i+1} > m_i$ and $n_i > n_{i+1}$ for such i .

THEOREM 2.3. *Let $\{(a_i, b_i, m_i, n_i) \in (0, \infty)^2 \times \mathbb{R}^2 : i \in \mathbb{Z}\}$ be a positive ABMN solution.*

- (1) *The solution is strict.*
- (2) *The solution has boundary conditions $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ that satisfy $m_{\infty} > m_{-\infty}$ and $n_{-\infty} > n_{\infty}$.*
- (3) *The values $m_{-\infty}$, m_{∞} , $n_{-\infty}$ and n_{∞} are real numbers. As such, the Mina margin $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}}$ exists and is a positive and finite real number.*

The Mina margin has a fundamental role to play in determining whether the ABMN system can be solved, as we now see.

THEOREM 2.4. *Invoking Theorem 2.3(3), we may set $I \subset (0, \infty)$ equal to the set of values of the Mina margin $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}}$, where $\{(a_i, b_i, m_i, n_i) \in (0, \infty)^2 \times \mathbb{R}^2 : i \in \mathbb{Z}\}$ ranges over the set of positive ABMN solutions.*

- (1) *There exists a value $\lambda \in (0, 1]$ such that the set I is equal to the interval $[\lambda, \lambda^{-1}]$.*
- (2) *Moreover, a positive ABMN solution exists with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$ if and only if $m_{-\infty} < m_{\infty}$, $n_{\infty} < n_{-\infty}$ and $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} \in [\lambda, \lambda^{-1}]$.*
- (3) *The value of λ is at most 0.999904.*

CONJECTURE 2.5. The value of λ is at least 0.999902.

Evidence for this conjecture will be presented in Section 6.3. It is perhaps surprising that a value so close to one may emerge from the ABMN system. The next result makes it all the more remarkable.

THEOREM 2.6. *Let $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$ satisfy $m_{-\infty} < m_{\infty}$ and $n_{\infty} < n_{-\infty}$.*

- (1) *Suppose that an element $\{(b_i, a_i) : i \in \mathbb{Z}\}$ of \mathcal{S}_0^2 lies in $\mathcal{N}(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$. Then the quadruple $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ associated to the element by Definition 2.1 is a positive ABMN solution with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$.*

- (2) *Conversely, suppose that $\{(a_i, b_i, m_i, n_i) \in (0, \infty)^2 \times \mathbb{R}^2 : i \in \mathbb{Z}\}$ is a positive ABMN solution with boundary data $(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$. Then $\{(b_i, a_i) : i \in \mathbb{Z}\} \in \mathcal{S}_0^2$ lies in $\mathcal{N}(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$.*

In view of this result, Theorem 2.4(3) eliminates the possibility that $\lambda = 1$, that Nash equilibria exist only when the players have symmetric roles. Conjecture 2.5 implies that the bound on λ proved in Theorem 2.4(3) is close to sharp: this quantity, which is canonically associated to the Trail of Lost Pennies on \mathbb{Z} —a natural and simple enough game—would then be remarkably close to one, differing from it by less than 10^{-4} .

DEFINITION 2.7. In its *standard* form, the trail game has boundary data that satisfies $m_{-\infty} = 0$, $n_\infty = 0$ and $m_\infty = 1$. For a game in this form, the game’s data is thus specified by one parameter, $n_{-\infty} \in (0, \infty)$. This parameter equals the Mina margin $\frac{n_{-\infty} - n_\infty}{m_\infty - m_{-\infty}}$.

Let $x \in (0, \infty)$. By $\text{Standard}(x)$, we denote the Trail of Lost Pennies on \mathbb{Z} in its standard form, with the Mina margin equal to x . That is, $\text{Standard}(x)$ equals $\text{Trail}(0, 1, x, 0)$, as this game has been specified in Section 2.1.

We briefly recapitulate in the present case the rough argument that motivated the premise “Incentive Inch, Outcome Mile” in Subsection 1.1.1. Suppose that x exceeds one. In playing $\text{Standard}(x)$, Maxine has more to play for than does Mina. Maxine may be tempted to outstake Mina, perhaps staking a certain constant multiple $f(x) > 1$ of the stake that Mina offers at any given turn. The resulting gameplay is a walk with a constant bias to the right, making Mina’s defeat inevitable—she may as well (or better) have staked nothing. If instead it is x^{-1} that exceeds one, then it is of course Mina who may be tempted by such an approach. Perhaps an argument can be fashioned along these lines to the effect that the game is competitive precisely when x lies in an interval of the form $[\mu, \mu^{-1}]$ for some $\mu \in (0, 1]$. This heuristic hardly lacks shortcomings, and it is quite unclear what the value of μ should be. However, the next result, which anyway follows from Theorems 2.4 and 2.6, validates its conclusion in a certain sense, with the value of μ equal to λ .

THEOREM 2.8. *Recall the quantity $\lambda \in (0, 1)$, which is specified and described by Theorem 2.4. For $x \in (0, \infty)$, the game $\text{Standard}(x)$ has a time-invariant Nash equilibrium precisely when x lies in $[\lambda, \lambda^{-1}]$.*

Admitting Conjecture 2.5, it is tempting to claim that this theorem validates the premise “Incentive Inch, Outcome Mile” in a precise and quantitative way. In fact, we can surmise that the premise holds for long finite intervals more clearly than we can for \mathbb{Z} , as we will discuss further in Section 7.4.

The shift operator on \mathbb{Z} has a basic role to play as we analyse the Trail of Lost Pennies on this set.

DEFINITION 2.9. Consider two time-invariant strategy pairs $\{(b_i, a_i) : i \in \mathbb{Z}\}$ and $\{(b'_i, a'_i) : i \in \mathbb{Z}\}$. These pairs are called *shift equivalent* if there exists $k \in \mathbb{Z}$ for which $(b_i, a_i) = (b'_{i+k}, a'_{i+k})$ for all $i \in \mathbb{Z}$. It is straightforward to see that an element $\{(b_i, a_i) : i \in \mathbb{Z}\}$ of \mathcal{S}_0^2 lies in $\mathcal{N}(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$ if and only if every shift equivalent element does so.

Let $Q : (0, \infty) \rightarrow \mathbb{N}$ be such that $Q(x)$ is the maximum cardinality of a set of mutually shift inequivalent time-invariant Nash equilibria for the game $\text{Standard}(x)$ for $x \in (0, \infty)$. The preceding result implies that the set of $x \in (0, \infty)$ for which $Q(x) > 0$ is equal to the interval $[\lambda, \lambda^{-1}]$ —which

interval is non-degenerate in view of Theorem 2.4(3). In the next result, we assert that a pair of shift inequivalent solutions exist when the Mina margin lies in the interval's interior.

THEOREM 2.10. *For $x \in (\lambda, \lambda^{-1})$, $Q(x) \geq 2$.*

We conjecture that no further time-invariant Nash equilibria exist.

CONJECTURE 2.11. We have that $Q(x) = 2$ when $x \in (\lambda, \lambda^{-1})$ and $Q(x) = 1$ when $x \in \{\lambda, \lambda^{-1}\}$.

This conjecture will be discussed in Section 7.1.

The next result describes precise asymptotic estimates on four sequences associated to any positive ABMN solution. In light of Theorem 2.6, it also describes decay rates for the pair of sequences given by any time-invariant Nash equilibrium.

DEFINITION 2.12. Let (a, b, m, n) be an ABMN solution. For $i \in \mathbb{Z}$, set $\phi_i = \frac{n_{i-1} - n_i}{m_i - m_{i-1}}$.

DEFINITION 2.13. For an ABMN solution (a, b, m, n) , the *battlefield index* is the unique value $k \in \mathbb{Z}$ such that $\phi_k \in (1/3, 3]$.

In Lemma 5.10, we will prove the existence and uniqueness claims implicit in the last definition, thus showing that the battlefield index is well-defined.

THEOREM 2.14. *Let $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ be a positive ABMN solution, and let $k \in \mathbb{Z}$ denote its battlefield index.*

(1) *There exist positive constants A and F such that, for $j \geq k$,*

$$\begin{aligned} a_j &= (m_k - m_{k-1}) \cdot 2F \cdot 2^{2(j-k)} \exp\{-2 \cdot 2^{j-k} A\} (1 + e^{-O(1)2^{j-k}}); \\ b_j &= (m_k - m_{k-1}) \cdot 4F \cdot 2^{2(j-k)} \exp\{-3 \cdot 2^{j-k} A\} (1 + e^{-O(1)2^{j-k}}); \\ m_j - m_{j-1} &= (m_k - m_{k-1}) \cdot F \cdot 2^{2(j-k)} \exp\{-2^{j-k} A\} (1 + e^{-O(1)2^{j-k}}); \text{ and} \\ n_{j-1} - n_j &= (m_k - m_{k-1}) \cdot 2F \cdot 2^{2(j-k)} \exp\{-2^{j-k+1} A\} (1 + e^{-O(1)2^{j-k}}). \end{aligned}$$

The constants A and F may be chosen to lie in a compact interval of $(0, \infty)$ that does not depend on the choice of the solution $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$. The positive constant that is implicit in the O -notation in the four displayed expressions may be chosen independently of this solution.

(2) *There exist positive constants B and G such that, for $j \leq k - 1$,*

$$\begin{aligned} a_j &= (n_{k-1} - n_k) \cdot 4G \cdot 2^{2(k-j)} \exp\{-3 \cdot 2^{k-j} B\} (1 + e^{-O(1)2^{k-j}}); \\ b_j &= (n_{k-1} - n_k) \cdot 2G \cdot 2^{2(k-j)} \exp\{-2 \cdot 2^{k-j} B\} (1 + e^{-O(1)2^{k-j}}); \\ m_j - m_{j-1} &= (n_{k-1} - n_k) \cdot 2G \cdot 2^{2(j-k)} \exp\{-2^{k-j+1} B\} (1 + e^{-O(1)2^{k-j}}); \text{ and} \\ n_{j-1} - n_j &= (n_{k-1} - n_k) \cdot G \cdot 2^{2(j-k)} \exp\{-2^{k-j} B\} (1 + e^{-O(1)2^{k-j}}). \end{aligned}$$

The conditions on B and G satisfy those set out for A and F in the preceding part; the constant implicit in the O -notation satisfies the condition recorded in this part.

When $X = k$ —when the counter is at the battlefield index—both players spend big to try to win the next move. For example, when $m_\infty - m_{-\infty} = n_{-\infty} - n_\infty = 1$, so that the difference in terminal receipt between victory and defeat is one unit for each player, then values of Maxine's stake a_k lie on the interval $[0.12, 0.20]$ and values of Mina's stake b_k lie in $[0.025, 0.18]$. (We will shortly present explicit solutions to the ABMN equations, which validate this assertion: we may use Theorem 2.31(1), for example. Maxine's expense interval is displaced to the right from Mina's, but the situation is reversed if the counter reaches $k - 1$, one place to the left.) These are big expenditures in a single turn of a game with infinitely many. The expenditures drop rapidly as the counter moves away from the battlefield, however. Indeed, if we write $g_i \ll h_i$ to denote that $g_i \leq \exp\{-e^{ci}\}h_i$ for $i \in \mathbb{N}_+$ (where c is some given positive constant), then Theorem 2.14 implies that,

$$\text{for } i \in \mathbb{N}_+, \quad 0 < b_{k+i} \ll a_{k+i} \ll 1 \text{ and } 0 < a_{k-i} \ll b_{k-i} \ll 1 \quad :$$

to the right of the battlefield, both expenditures drop suddenly; but Maxine, eyeing victory, keeps her foot on the gas, making sure to vastly outspend the loss-cutting Mina; while to the left of the battlefield, the roles are reversed. The Battlefield Cyl Fog of Section 1.1.1 is thus exemplified. We also have that $0 < n_{k+i} - n_{k+i+1} \ll m_{k+i+1} - m_{k+i} \ll 1$ and $0 < m_{k-i} - m_{k-i-1} \ll n_{k-i-1} - n_{k-i} \ll 1$; and, by extension,

$$0 < n_{k+i} - n_\infty \ll m_\infty - m_{k+i} \ll 1 \text{ and } 0 < m_{k-i} - m_{-\infty} \ll n_{-\infty} - n_{k-i} \ll 1.$$

Indeed, in the left part of the last display, which is to the right of the battlefield, Mina has essentially (but not absolutely!) thrown in the towel, and her expected payoff n_{k+i} is minutely above her defeat terminal receipt of n_∞ . Maxine's average payoff m_{k+i} is just slightly below her victory receipt of m_∞ , but her need to keep moving the counter rightwards provides some lower bound on the difference. In the right part of the display, roles are naturally reversed.

The players may dread the return of the counter to the battlefield index because this is an expensive occasion for both of them. The next result, a consequence of Theorem 2.14, shows that they are typically saved from witnessing this event repeatedly when a Nash equilibrium is played.

Let $(S_-, S_+) \in \mathcal{S}^2$ and $i \in \mathbb{Z}$. Under \mathbb{P}_{S_-, S_+}^i , the *unanimity* event U occurs when all but finitely many of the differences $X_{j+1} - X_j \in \{-1, 1\}$, $j \in \mathbb{N}$, of the gameplay process $X : \mathbb{N} \rightarrow \mathbb{Z}$, $X_0 = i$, adopt a given value. Writing U_- and U_+ for the respective events specified when the given value is -1 and 1 , the occurrence of these events correspond to victories for Mina and Maxine, and U is the disjoint union of U_- and U_+ .

THEOREM 2.15. *Let (a, b, m, n) denote a positive ABMN solution on \mathbb{Z} with given boundary data of the form (4). Suppose that the solution has battlefield index $k \in \mathbb{Z}$, and let $i \in \mathbb{Z}$. Let $(S_-, S_+) \in \mathcal{S}_0^2$ be given by (b, a) , with the usual abuse of notation.*

$$(1) \text{ We have that } \mathbb{P}_{S_-, S_+}^i(U) = 1.$$

There exist positive constants C and c that may be chosen independently of the element (S_-, S_+) and the index $i \in \mathbb{Z}$ for which the following hold.

$$(2) \text{ If } i \geq k \text{ then } \mathbb{P}_{S_-, S_+}^i(U_-) \leq C \exp\{-c2^{k-i}\}.$$

$$(3) \text{ If } i \leq k - 1 \text{ then } \mathbb{P}_{S_-, S_+}^i(U_+) \leq C \exp\{-c2^{i-k}\}.$$

Consider the Trail of Lost Pennies on \mathbb{Z} with given boundary data (4). Redefine (S_-, S_+) to be an element of $\mathcal{S}_0^2 \cap \mathcal{N}(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$. Writing (b, a) for (S_-, S_+) , the data (a, b, m, n) specified by Definition 2.1 determines the battlefield index. Suppose that this index is $k \in \mathbb{Z}$, and let $i \in \mathbb{Z}$.

(4) *The preceding three parts remain valid in this framework.*

We see then that the player who wins a local victory at (or around) the battlefield index typically comes to entirely dominate the later moves of the game. By playing at a time-invariant Nash equilibrium, players thereby forge an implicit consensus to avoid the mutually destructive circumstance of many returns to the battlefield.

We recapitulated the conclusions of [24] regarding Nash equilibria in allocated-budget stake-governed games in Section 1.1.4. The formula (3) indicates that, for suitable graphs and at Nash equilibrium, each budget must be spent via staking in a more-or-less regular flow, so that the concerned player is competitive throughout the lifetime of the game. Game behaviour under the Battlefield Cyl Fog is very different from under the Present-Future Compromise. The empirical stake process of a player at any Nash equilibrium is punctuated by a few brief intense periods as the counter passes through the battlefield. In the large, the only concern for outcome is the answer to the question: on which side of the battlefield does the counter lie?

2.3. Explicit ABMN solutions

Here we present an explicit form for all positive ABMN solutions. It is useful to begin by classifying the solutions into classes, where members of a given class differ in simple ways. If one or other player receives, or must pay, some given amount before a game begins, play will be unaffected—or at least the Nash equilibria will not be. If the unit currency is revalued before play, the outcome will be a mere scaling of all quantities. We identify ABMN solutions that differ according to translations $\chi_{x,0}$ or $\chi_{0,y}$ or dilations τ_u (where $x, y \in \mathbb{R}$ and $u \in (0, \infty)$) that correspond to such operations. If we can describe one element in each equivalence class, we will be able to describe all solutions. Equivalence classes are naturally parametrized by the positive real quantity $\phi_0 = \frac{n-1-n_0}{m_0-m-1}$, which we call the *central ratio*, specified by Definition 2.12. So there is a one-parameter family of essentially different positive ABMN solutions. In each equivalence class, we will distinguish two special solutions—the default solution, which has a simpler explicit formula; and the standard solution, which corresponds to a convenient choice of boundary data for the trail game. We will set up this structure and then state the explicit form of the default solution in each equivalence class.

We consider \mathbb{Z} -indexed sequences $g = \{g_i : i \in \mathbb{Z}\}$. A sequence is *monotone* if it is non-decreasing or non-increasing. A bounded monotone sequence g has left and right limits

$$g_{-\infty} = \lim_{k \rightarrow \infty} g_{-k} \quad \text{and} \quad g_{\infty} = \lim_{k \rightarrow \infty} g_k$$

that are elements of \mathbb{R} . We will specify certain bounded monotone sequences g by giving one of the limiting values, $g_{-\infty}$ or g_{∞} , alongside the difference sequence $\{g_{i+1} - g_i : i \in \mathbb{Z}\}$.

Let $u \in (0, \infty)$ and $v \in \mathbb{R}$. For any sequence $g : \mathbb{Z} \rightarrow \mathbb{R}$, we write $u \cdot g : \mathbb{Z} \rightarrow \mathbb{R}$ for the sequence given by $(u \cdot g)_i = u \cdot g_i$.

Let Θ denote the space of quadruples of sequences; thus, when $(a, b, m, n) \in \Theta$, each component $*$ $\in \{a, b, m, n\}$ has the form $* : \mathbb{Z} \rightarrow \mathbb{R}$. For $u \in (0, \infty)$ and $v_1, v_2 \in \mathbb{R}$, define $\tau_u, \chi_{v_1, v_2} : \Theta \rightarrow \Theta$ so that $\tau_u(a, b, m, n) = (u \cdot a, u \cdot b, u \cdot m, u \cdot n)$ and $\chi_{v_1, v_2}(a, b, m, n) = (a, b, m + v_1, n + v_2)$.

Two solutions of the ABMN equations on \mathbb{Z} are called *equivalent* if one is the image of the other under a composition of the form $\tau_u \circ \chi_{v_1, v_2}$ for such u, v_1 and v_2 as above. The relation of two such solutions will be denoted by \sim ; Proposition 2.17 asserts that \sim is indeed an equivalence relation.

Let (a, b, m, n) be an ABMN solution on \mathbb{Z} . Recall from Definition 2.2 that the solution's *Mina margin* is defined to be $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}}$. The solution's *central ratio* CenRatio is set equal to $\frac{n_{-1} - n_0}{m_0 - m_{-1}}$. The solution is called *standard* if $m_{-\infty} = 0$, $n_{\infty} = 0$ and $m_{\infty} = 1$. It is called *default* if $m_{-\infty} = 0$, $n_{\infty} = 0$ and $m_0 - m_{-1} = 1$. Compatibly with the usage of Definition 2.7, the Mina margin of a standard solution equals $n_{-\infty}$; note further that the central ratio of a default solution equals $n_{-1} - n_0$.

PROPOSITION 2.16. *For any default ABMN solution, the value of the central ratio CenRatio lies in $(0, \infty)$. For any $x \in (0, \infty)$, there is exactly one default solution for which CenRatio equals x .*

PROPOSITION 2.17.

- (1) *The space of ABMN solutions is partitioned into equivalence classes by the relation \sim .*
- (2) *Each equivalence class contains a unique standard solution and a unique default solution.*

Propositions 2.16 and 2.17 provide a natural labelling of ABMN solution equivalence classes: any given class is labelled by the value of the central ratio of the unique default solution in the class. The labelling parametrizes the equivalence classes by a copy of $(0, \infty)$.

According to the latter assertion of Proposition 2.16, there is a unique default solution to the ABMN equations whose central ratio equals a given value $x \in (0, \infty)$. The next definitions will enable us to record the form of this solution in Theorem 2.21.

DEFINITION 2.18. Set $\omega : (0, \infty) \rightarrow (1, \infty)$, $\omega(x) = \sqrt{8x + 1}$. Writing $\omega = \omega(x)$, we further set

$$c(x) = \frac{(\omega + 3)^2}{16}, \quad d(x) = \frac{(\omega + 3)^2}{8(\omega + 1)} \quad \text{and} \quad s(x) = \frac{(\omega - 1)^2}{4(\omega + 7)} \quad \text{for } x \in (0, \infty).$$

DEFINITION 2.19. Let $s_{-1} : (0, \infty) \rightarrow (0, \infty)$ be given by $s_{-1}(x) = 1/s(1/x)$. We now define a collection of functions $s_i : (0, \infty) \rightarrow (0, \infty)$ indexed by $i \in \mathbb{Z}$. We begin by setting $s_0(x) = x$ for $x \in (0, \infty)$. We then iteratively specify that $s_i(x) = s(s_{i-1}(x))$ and $s_{-i}(x) = s_{-1}(s_{-(i-1)}(x))$ for $i \in \mathbb{N}_+$ and $x \in (0, \infty)$. Note that s_1 equals s and that the two specifications of s_{-1} coincide.

Set $c_j, d_j : (0, \infty) \rightarrow (0, \infty)$, $j \in \mathbb{Z}$, by means of $c_j(x) = c(s_j(x))$ and $d_j(x) = d(s_j(x))$.

To get a sense of the maps s_i , $i \in \mathbb{Z}$, a few points are worth noting. First, as we will see in Proposition 3.4, s_{-1} is the inverse of s . Second, as Lemma 5.3(5) attests, $s(x) < x$ for $x \in (0, \infty)$; the orbit $s_i(x)$ thus decreases or increases from x as i grows to the right or the left. And third, note that $s(3) = 1/3$. In view of the second point, we see that $(0, \infty) = \cup_{k \in \mathbb{Z}} s_k[1/3, 3)$ is a partition whose interval elements are arranged in decreasing order in the index k .

DEFINITION 2.20. For a sequence h , we may naturally write $\prod_{i=0}^k h_i = h_0 \cdots h_k$ for $k \in \mathbb{N}$. A convenient device extends this notation to cases where $k \in \mathbb{Z}$ is negative: we set

$$\prod_{i=0}^k h_i = \begin{cases} 1 & \text{for } k = -1 \\ h_{k+1}^{-1} \cdots h_{-1}^{-1} & \text{for } k \leq -2. \end{cases}$$

Let $x \in (0, \infty)$. This parameter will index four real-valued sequences

$$a^{\text{def}}(x), b^{\text{def}}(x), m^{\text{def}}(x), n^{\text{def}}(x) : \mathbb{Z} \rightarrow (0, \infty)$$

which we denote in the form $\{*_i^{\text{def}}(x) : i \in \mathbb{Z}\}$ for $* \in \{a, b, m, n\}$.

We begin by specifying $m^{\text{def}}(x) : \mathbb{Z} \rightarrow \mathbb{R}$. This is the increasing sequence such that

$$m_{-\infty}^{\text{def}}(x) = 0, \quad \text{and} \quad m_{k+1}^{\text{def}}(x) - m_k^{\text{def}}(x) = \prod_{i=0}^k (c_i(x) - 1) \quad \text{for } k \in \mathbb{Z}.$$

Note that $m_0^{\text{def}}(x) - m_{-1}^{\text{def}}(x) = 1$ in view of the notation for products.

Next we set $n^{\text{def}}(x) : \mathbb{Z} \rightarrow \mathbb{R}$. This is the decreasing sequence with

$$n_{\infty}^{\text{def}}(x) = 0, \quad \text{and} \quad n_k^{\text{def}}(x) - n_{k+1}^{\text{def}}(x) = x \prod_{i=0}^k (d_i(x) - 1) \quad \text{for } k \in \mathbb{Z}.$$

Note that $n_{-1}(x) - n_0(x) = x$.

To specify $a^{\text{def}}(x), b^{\text{def}}(x) : \mathbb{Z} \rightarrow (0, \infty)$, we set

$$M_i(x) = m_{i+1}^{\text{def}}(x) - m_{i-1}^{\text{def}}(x) \quad \text{and} \quad N_i(x) = n_{i-1}^{\text{def}}(x) - n_{i+1}^{\text{def}}(x)$$

for $i \in \mathbb{Z}$. For such i , we take

$$a_i^{\text{def}}(x) = \frac{M_i(x)^2 N_i(x)}{(M_i(x) + N_i(x))^2} \quad \text{and} \quad b_i^{\text{def}}(x) = \frac{M_i(x) N_i(x)^2}{(M_i(x) + N_i(x))^2}.$$

THEOREM 2.21. *Let $x \in (0, \infty)$. The unique default ABMN solution with $\text{CenRatio} = x$ is the quadruple $(a_i^{\text{def}}(x), b_i^{\text{def}}(x), m_i^{\text{def}}(x), n_i^{\text{def}}(x) : i \in \mathbb{Z})$ specified in Definition 2.20.*

For $x \in (0, \infty)$, we write $\mathcal{C}(x)$ for the equivalence class of ABMN solutions that contains the element $(a_i^{\text{def}}(x), b_i^{\text{def}}(x), m_i^{\text{def}}(x), n_i^{\text{def}}(x) : i \in \mathbb{Z})$. Let $(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z})$ denote the unique standard solution in $\mathcal{C}(x)$.

Remark. Let $x \in (0, \infty)$. Set $Z(x) = m_{\infty}^{\text{def}}(x)$; which is to say, $Z(x) = \sum_{k \in \mathbb{Z}} \prod_{i=0}^k (c_i(x) - 1)$. It is straightforward that

$$\left(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z} \right) \quad \text{equals} \quad Z(x)^{-1} \cdot \left(a_i^{\text{def}}(x), b_i^{\text{def}}(x), m_i^{\text{def}}(x), n_i^{\text{def}}(x) : i \in \mathbb{Z} \right). \quad (12)$$

2.4. The Mina margin map

According to Theorem 2.21, the central ratio ϕ_0 is a convenient parameter for indexing ABMN solution equivalence classes. And Theorem 2.8 tells us that the Mina margin is a fundamental parameter for locating Nash equilibria in the trail game. The map $(0, \infty) \rightarrow (0, \infty)$ from equivalence class index to the Mina margin of any member solution is a natural object that we will use to organize and prove results. We call this function the *Mina margin map*. Here, we define it, and state basic properties in Theorem 2.23. Theorem 2.24 shows how to solve the trail game with given boundary data by finding time-invariant Nash equilibria indexed by the map's level sets. Theorem 2.28 states that a reparametrization of the Mina margin map's domain leads to a periodic form for the map that commutes with the shift operator on \mathbb{Z} .

DEFINITION 2.22. Let the Mina margin map $\mathcal{M} : (0, \infty) \rightarrow (0, \infty)$ be given by $\mathcal{M}(x) = n_{-\infty}^{\text{st}}(x)$ for $x \in (0, \infty)$. Namely, $\mathcal{M}(x)$ is the Mina margin of $(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z})$.

THEOREM 2.23.

- (1) The function $\mathcal{M} : (0, \infty) \rightarrow (0, \infty)$ satisfies $\mathcal{M}(s(x)) = \mathcal{M}(x)$ for $x \in (0, \infty)$.
- (2) The function \mathcal{M} is continuous on $(0, \infty)$ and satisfies

$$\mathcal{M}(x) = \left(\sum_{k \in \mathbb{Z}} \prod_{i=0}^k (c_i(x) - 1) \right)^{-1} \cdot x \sum_{k \in \mathbb{Z}} \prod_{i=0}^k (d_i(x) - 1).$$

- (3) The range $\mathcal{M}(0, \infty)$ takes the form $[\lambda, \lambda^{-1}]$, where $\lambda \in (0, 0.999904]$ is specified in Theorem 2.4.

THEOREM 2.24. Let $x \in [\lambda, \lambda^{-1}]$. Set $X = \{z \in (0, \infty) : \mathcal{M}(z) = x\}$, and let $Y = X \cap (1/3, 3]$, so that, as noted after Definition 2.19, $X = \cup_{k \in \mathbb{Z}} s_k(Y)$.

- (1) The collection of time-invariant Nash equilibria in the game $\text{Standard}(x)$ is given by the set of maps

$$\mathbb{Z} \rightarrow (0, \infty)^2 : i \rightarrow (b_i^{\text{st}}(z), a_i^{\text{st}}(z))$$

indexed by z in X .

- (2) Alternatively, this collection is the set of maps

$$\mathbb{Z} \rightarrow (0, \infty)^2 : i \rightarrow (b_{i+j}^{\text{st}}(x), a_{i+j}^{\text{st}}(x)),$$

where now the index (x, j) ranges over $Y \times \mathbb{Z}$.

We now develop the notation for the symbolic shift map that was mooted in Definition 2.9.

DEFINITION 2.25. We let \mathcal{S}_1 denote the left shift by one place: this is the map that sends the space of quadruples $(\mathbb{R}^4)^{\mathbb{Z}}$ to itself by the action

$$\mathcal{S}_1 \{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\} = \{(a_{i+1}, b_{i+1}, m_{i+1}, n_{i+1}) : i \in \mathbb{Z}\}.$$

By iterating this map, we specify the left shift \mathcal{S}_k by k places, for $k \geq 2$; and by specifying $\mathcal{S}_{-1} = \mathcal{S}_1^{-1}$ and iterating, we specify the right shift \mathcal{S}_{-k} by k places, for $k \geq 1$.

What is the effect of applying the shift \mathcal{S}_k to a standard solution? It takes the form of a replacement $x \rightarrow s_k(x)$ in the CenRatio-variable, as we will see in the short proof of the next result.

PROPOSITION 2.26. For $x \in (0, \infty)$ and $k \in \mathbb{Z}$,

$$\mathcal{S}_k(a^{\text{st}}(x), b^{\text{st}}(x), m^{\text{st}}(x), n^{\text{st}}(x)) = \left(a^{\text{st}}(s_k(x)), b^{\text{st}}(s_k(x)), m^{\text{st}}(s_k(x)), n^{\text{st}}(s_k(x)) \right).$$

Proof. The symbolic shift map leaves invariant the boundary quadruple of any ABMN solution. Thus, the displayed left-hand quadruple is a standard solution of the ABMN system. To identify it as the right-hand quadruple, it is thus enough to show that its CenRatio-value equals $s_k(x)$. But this amounts to $\frac{n_{k-1} - n_k}{m_k - m_{k-1}} = s_k(x)$, because $x = \frac{n_{-1} - n_0}{m_0 - m_{-1}}$. \square

Theorem 2.23(1) leads directly to $\mathcal{M}(s_k(x)) = \mathcal{M}(x)$ for $x \in (0, \infty)$ and $k \in \mathbb{Z}$. To understand the map \mathcal{M} , we see then that the asymptotics in highly positive and negative k of the orbits $s_k(x)$ are important. As we will see in Lemma 5.3(2,3), $s(x) \sim x^2/2$ for $0 < x \ll 1$ and $s(x) \sim 2^{-1/2}x^{1/2}$ for

$x \gg 1$. Thus the forward orbit $s_k(x)$, $k \rightarrow \infty$, converges rapidly to zero, while the backward orbit, $k \rightarrow -\infty$, grows quickly towards infinity.

We now undertake a change of coordinates of the Mina margin map $\mathcal{M} : (0, \infty) \rightarrow (0, \infty)$. The domain $(0, \infty)$ will be identified with \mathbb{R} by an increasing bijection θ^{-1} . The goal of the coordinate change is to ensure that the original action $(0, \infty) \rightarrow (0, \infty) : x \rightarrow s_1(x)$ becomes the map $\mathbb{R} \rightarrow \mathbb{R} : x \rightarrow x - 1$. The action of the symbolic sequence shift \mathcal{S}_1 on the x -variable, as stated in Proposition 2.26, comes to correspond to a left shift by a unit in the new real variable. This leads to an attractive representation of the Mina margin map in the guise $\mathbb{R} \rightarrow (0, \infty) : x \rightarrow \mathcal{M}(\theta^{-1}(x))$.

DEFINITION 2.27. Let $q : [1/3, 3) \rightarrow [0, 1)$ be an increasing surjection; for definiteness, we may take $q(x) = 3(x - 1/3)/8$. We specify $\theta : (0, \infty) \rightarrow \mathbb{R}$ so that, for $x \in (0, \infty)$, $\theta(x) = k + q(s_k(x))$, where $k \in \mathbb{Z}$ is the unique integer such that $s_k(x) \in [1/3, 3)$. Since $\theta : (0, \infty) \rightarrow \mathbb{R}$ is an increasing surjection, the inverse $\theta^{-1} : \mathbb{R} \rightarrow (0, \infty)$ is well defined. We may thus represent the Mina margin map after domain coordinate change by the function ψ , where

$$\psi : \mathbb{R} \rightarrow (0, \infty) \quad , \quad \psi(x) = \mathcal{M}(\theta^{-1}(x)) .$$

We define the *standard solution* map $\text{StSol} : \mathbb{R} \rightarrow (\mathbb{R}^4)^{\mathbb{Z}}$,

$$\text{StSol}(x) = \left(a^{\text{st}}(\theta^{-1}(x)), b^{\text{st}}(\theta^{-1}(x)), m^{\text{st}}(\theta^{-1}(x)), n^{\text{st}}(\theta^{-1}(x)) \right) \quad \text{for } x \in \mathbb{R} .$$

For $u \in (0, \infty)$ and $j \in \mathbb{Z}$, $\theta(s_{-j}(u)) - \theta(u) = j$. For $z = \Theta(1)$, the value of $\theta^{-1}(z + k)$ thus tracks that of $s_{-k}(z)$ as k rises, either by growing to infinity (if k is positive); or by decaying to zero (if k is negative). To understand the transformation θ^{-1} , it is thus useful to introduce a simple explicit function $\Theta : \mathbb{R} \rightarrow (0, \infty)$ which is designed so that $s_{-k}(z)$ grows or decays roughly as does $\Theta(k)$ for $|k|$ large; here $z \in [1/3, 3)$, say.

Let $\text{Sign} : \mathbb{R} \rightarrow \{-1, 1\}$ equal $\text{Sign}(x) = \mathbf{1}_{x \geq 0} - \mathbf{1}_{x < 0}$. Then set

$$\Theta : \mathbb{R} \rightarrow (0, \infty) \quad , \quad \Theta(x) = 2^{\text{Sign}(x)(2^{|x|} - 1)} .$$

We now present our result concerning the Mina margin map after domain coordinate change. The transformed function ψ is periodic, of unit period; symbolic shift by one place corresponds to unit translation of the domain; and coordinate change asymptotics are, crudely at least, described by Θ .

THEOREM 2.28.

- (1) For $x \in \mathbb{R}$, $\psi(x + 1) = \psi(x)$.
- (2) For $x \in (0, \infty)$ and $k \in \mathbb{Z}$,

$$\text{StSol}(x + k) = \mathcal{S}_{-k} \circ \text{StSol}(x) .$$

- (3) There exists a positive constant C such that, for $z \geq 0$,

$$2^{2z-C} \leq \theta^{-1}(z) \leq 2^{2z+C} ;$$

and, for $z < 0$,

$$2^{-2|z|+C} \leq \theta^{-1}(z) \leq 2^{-2|z|-C} .$$

The map Θ is a simple and explicit surrogate for θ^{-1} , and the transformed Mina margin map $\mathbb{R} \rightarrow (0, \infty) : x \rightarrow \mathcal{M}(\Theta(x))$ shares the periodicity property of ψ in Theorem 2.28(1) up to a domain perturbation that decays rapidly away from zero. And this surrogate has a more practical version, in which the Mina margin map \mathcal{M} is replaced by a counterpart for a trail that is a finite interval, rather than all of \mathbb{Z} . These counterpart Mina margin maps $\mathcal{M}_{j+1, k+1}$ will be presented in the next section. Plots of several of these maps, indexed by different finite trails, appear in Figure 2.3.

2.5. The Trail of Lost Pennies on a finite interval

Even if much of our focus lies with the trail game in the infinite setting, with gameboard \mathbb{Z} , it is instructive to introduce and discuss the game whose trail is a finite interval. This is the setting used in economists' treatments of tug-of-war that we surveyed in Section 1.2.1. And it is a more practical setting if two people are to play the game, taking decisions turn-by-turn because, at least for short intervals, the game will end (by the token reaching one end of the interval or the other) in a limited number of moves. The theoretical aspects of the game—time-invariant Nash equilibria; ABMN solutions and their standard solutions; the Mina margin map—share many basic aspects between the infinite and finite settings. The finite setting permits important objects, such as the Mina margin map, to be plotted in Mathematica, and such investigation has informed several of our main results (in the infinite setting). Our goal then in this section is to communicate the principal aspects of the finite setting so that the reader can interpret pertinent Mathematica plots and understand how these suggest some of our principal results and conjectures. We will also present a conjecture concerning the number of time-invariant Nash equilibria in a symmetric version of the finite game; we will seek to explain why we believe it during the section. The section contains one result, Proposition 2.29, which we will use and whose proof appears in Section 3.1. Our basic aim is heuristic, however, and at times our presentation will be informal.

2.5.1. Gameplay, strategies and Nash equilibria for the finite trail. Let $j, k \in \mathbb{N}$. The Trail of Lost Pennies with trail (or gameboard) $\llbracket -j - 1, k + 1 \rrbracket$ is specified by

$$(m_{-j-1}, m_{k+1}, n_{-j-1}, n_{k+1}) \in \mathbb{R}^4,$$

boundary data on which the conditions $m_{-j-1} < m_{k+1}$ and $n_{k+1} < m_{-j-1}$ are imposed. Begun from ℓ , an element in the field of open play $\llbracket -j, k \rrbracket$, gameplay is a stochastic process $X : \llbracket 0, K \rrbracket \rightarrow \llbracket -j - 1, k + 1 \rrbracket$, $X_0 = \ell$, where

$$K = \inf \left\{ i \in \mathbb{N}_+ : X_i \in \{-j - 1, k + 1\} \right\}.$$

Indeed, with Mina and Maxine playing to the left and right, the game will end with victory to these respective players when the counter arrives at $-j - 1$ or at $k + 1$.

The gameplay is specified by a strategy pair, where a strategy is a map $S : \llbracket -j, k \rrbracket \times \mathbb{N}_+ \rightarrow [0, \infty)$. The construction of X from a given location $X_0 = \ell \in \llbracket -j, k \rrbracket$ coincides with that explained in Section 2.1, where instances of the trail \mathbb{Z} are replaced by $\llbracket -j, k \rrbracket$, it being understood that the construction stops when X arrives in $\{-j - 1, k + 1\}$.

A strategy S for which $S(\ell, i)$ is independent of $i \in \mathbb{N}_+$ for all $\ell \in \llbracket -j, k \rrbracket$ is said to be *time-invariant*. Let $\mathcal{S}[j, k]$ denote the space of strategies. For a strategy pair $(S_-, S_+) \in \mathcal{S}[j, k]$, we may reuse notation from the \mathbb{Z} -indexed trail game, and speak of the law \mathbb{P}_{S_-, S_+}^i of gameplay $X : \mathbb{N} \rightarrow \llbracket -j - 1, k + 1 \rrbracket$, $X_0 = i$, governed by the pair (S_-, S_+) , and stopped on arrival in $\{-j - 1, k + 1\}$.

Counterpart to (5) and (9) are the \mathbb{P}_{S_-, S_+}^i -almost sure payoff identities

$$P_-^{j,k} = - \sum_{i=1}^{\infty} C_-^{j,k}(i) + T_-^{j,k} \quad \text{and} \quad P_+^{j,k} = - \sum_{i=1}^{\infty} C_+^{j,k}(i) + T_+^{j,k}, \quad (13)$$

where the cost $C_{\pm}^{j,k}(i)$ incurred to each player at the i^{th} turn, $i \in \mathbb{N}_+$, equals $S_{\pm}(X_{i-1}, i)$, as in the original case. To specify the terminal payments $T_{\pm}^{j,k}$, we permit E_- to denote the event that X arrives at the vertex $-j-1$ at some positive time, and E_+ to denote the event that this process instead reaches $k+1$ at some such time. We then adopt (8) and (10) for $T_{\pm}^{j,k}$, where m_* and n_* denote given real values that satisfy $m_* \leq m_{-j-1}$ and $n_* \leq n_{k+1}$.

Definitions concerning Nash equilibria continue to be specified as they are at the end of Section 2.1. A collection of quadruples $\{(a_i, b_i, m_i, n_i) : i \in \llbracket -j, k \rrbracket\}$ is associated to any element $\{(b_i, a_i) : i \in \llbracket -j, k \rrbracket\}$ by Definition 2.1 after evident changes in notation have been made.

2.5.2. The ABMN equations. Recall Definition 2.2. Let $j, k \in \mathbb{N}$.

The ABMN system on $\llbracket -j, k \rrbracket$ is the set of equations ABMN(1, 2, 3, 4) in the real variables a_i, b_i, m_i and n_i , where the index i varies over $\llbracket -j, k \rrbracket$. These equations refer to the components of the quadruple $(m_{-j-1}, m_{k+1}, n_{-j-1}, n_{k+1}) \in \mathbb{R}^4$ which acts as boundary data and for which we suppose a fixed value that satisfies $m_{-j-1} < m_{k+1}$ and $n_{-j-1} > n_{k+1}$. Similarly to Definition 2.2, a solution is *positive* if a_i and b_i exceed zero for $i \in \llbracket -j, k \rrbracket$.

2.5.3. A result and a conjecture for the finite trail. The basic relation between time-invariant Nash equilibria $\{(b_i, a_i) : i \in \llbracket -j, k \rrbracket\}$ and positive ABMN solutions $\{(a_i, b_i, m_i, n_i) : i \in \llbracket -j, k \rrbracket\}$ embodied in Theorem 2.6 is maintained.

The trail game on $\llbracket -j-1, k+1 \rrbracket$ is in its *standard form* when its boundary data satisfies $m_{-j-1} = n_{k+1} = 0$ and $m_{k+1} = 1$. This class of games is thus parametrized by the Mina margin $n_{-j-1} \in (0, \infty)$. If further $n_{-j-1} = 1$, then we speak of the *symmetric standard game*. Likewise a solution of the ABMN equations on $\llbracket -j, k \rrbracket$ is *standard* when $m_{-j-1} = n_{k+1} = 0$ and $m_{k+1} = 1$. The space of standard solutions may be parametrized by the central ratio $\text{CenRatio} = \frac{n_{-1} - n_0}{m_0 - m_{-1}} \in (0, \infty)$. The Mina margin map $\mathcal{M}_{j+1, k+1} : (0, \infty) \rightarrow (\infty)$ associates to $x \in (0, \infty)$ the value of the Mina margin n_{-j-1} of the unique standard ABMN solution on $\llbracket -j, k \rrbracket$ for which $\text{CenRatio} = x$.

Standard solutions may be computed explicitly, similarly as was (12) in the infinite setting. To obtain the standard solution on $\llbracket -j, k \rrbracket$ with $\text{CenRatio} = x \in (0, \infty)$, we start with the restriction of the default solution from Theorem 2.21 to $\llbracket -j-1, k+1 \rrbracket$. By adding a suitable constant to each m -term, and another such to each n -term, and then multiplying the result by a suitable scaling factor, we obtain a standard solution whose CenRatio remains equal to x because the additions and the scaling leave this value unchanged. We thus see that, for $x \in (0, \infty)$,

$$\mathcal{M}_{j+1, k+1}(x) = \frac{n_{-j-1} - n_{k+1}}{m_{k+1} - m_{-j-1}}, \quad (14)$$

where $\{(a_i, b_i, m_i, n_i) : i \in \llbracket -j, k \rrbracket\}$ is any ABMN solution on $\llbracket -j, k \rrbracket$ such that $\frac{n_{-1} - n_0}{m_0 - m_{-1}} = x$.

The trail game on trails $\llbracket -k, k \rrbracket$ of even length differs from that on trails $\llbracket -k-1, k \rrbracket$ of odd length, because the trails in the two classes are reflection symmetric about different objects (the vertex 0 or the edge $\llbracket -1, 0 \rrbracket$). The next result records outcomes of these symmetries for the finite trail Mina margin map.

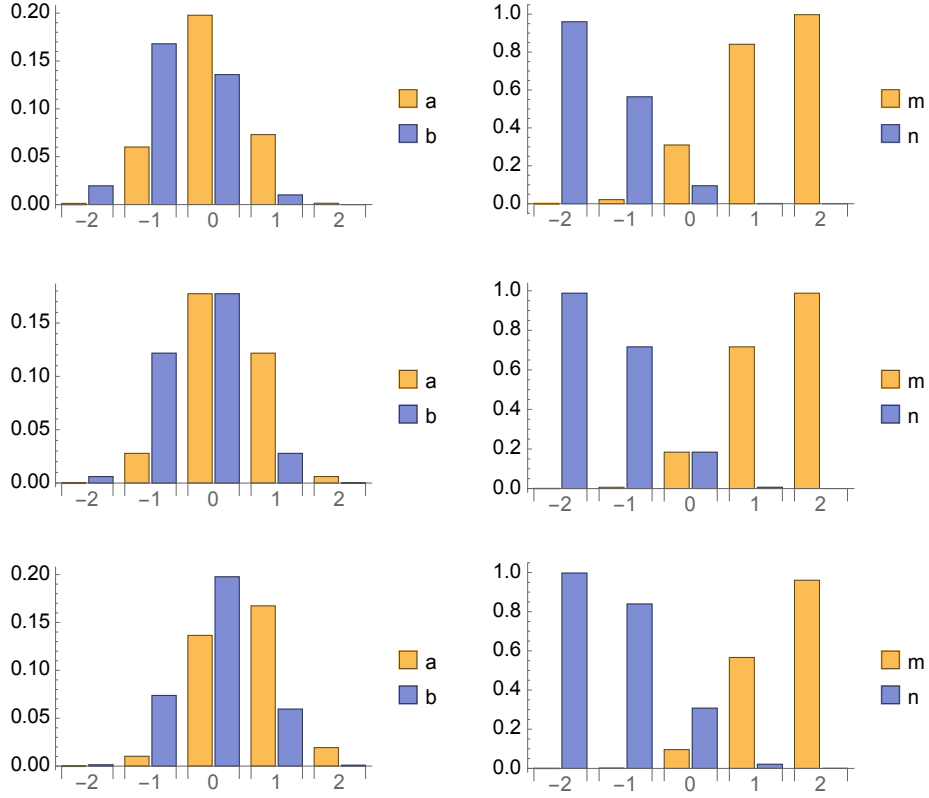


FIGURE 2.1. The shortest trail with non-unique Nash equilibria for at least some boundary conditions has length six, with five sites in open play. The values $x_1 = 1.63$, $x_2 = 3$ and $x_3 = 5.64$ approximate the three solutions of $\mathcal{M}_{3,3}(x) = 1$. (There is no error for x_2 .) The (a, b) and (m, n) data on $[-2, 2]$ for the standard solution on $[-3, 3]$ corresponding to x_1 appears in the top row; to x_2 in the middle; and to x_3 in the lower row. The left column thus depicts the three Nash equilibria in standard symmetric game on the shortest trail for which this game may be expected to have several equilibria. Note that the x_3 -solution is formed from the x_1 -solution by role-reversal: that is, by interchanging the roles of a and b , and of m and n , and by reflecting in the origin.

PROPOSITION 2.29. Let $k \in \mathbb{N}_+$ and $x \in (0, \infty)$.

- (1) We have that $\mathcal{M}_{k,k}(x) \cdot \mathcal{M}_{k,k}(1/s(x)) = 1$.
- (2) And that $\mathcal{M}_{k+1,k}(x) \cdot \mathcal{M}_{k+1,k}(x^{-1}) = 1$.

Here is our conjecture concerning the symmetric form of the finite trail game.

CONJECTURE 2.30. Consider the Trail of Lost Pennies on $[-j-1, k+1]$ in its symmetric standard form. The number of time-invariant Nash equilibria equals $\max\{2(j+k) - 5, 1\}$.

Roughly similar is Harris and Vickers' numerically based claim (recalled in Subsection 1.2.1.1) of order $2n$ equilibria in their game. We mention also that the number of Nash equilibria is odd in almost all finite games [59].

We offer an explanation of why we believe Conjecture 2.30. By a counterpart to Theorem 2.6 (which we have roughly indicated), any time-invariant Nash equilibrium of the symmetric trail game on $\llbracket -j - 1, k + 1 \rrbracket$ corresponds to a positive ABMN solution on $\llbracket -j, k \rrbracket$. This solution must have $m_{-j-1} = n_{k+1} = 0$ and $m_{k+1} = 1$, as well as $n_{-j-1} = 1$. That is, the solution must be standard, and it must satisfy $\mathcal{M}_{-j-1, k+1}(x) = 1$, where $x \in (0, \infty)$ is the solution's value of CenRatio. We may thus obtain the set of time-invariant Nash equilibria by recording, for each solution $x \in (0, \infty)$ of the equation $\mathcal{M}_{-j-1, k+1}(x) = 1$, the reverse-ordered (a, b) -component pair of the unique standard ABMN solution on $\llbracket -j, k \rrbracket$ whose CenRatio-value equals x . The case for Conjecture 2.30 thus rests on advancing an argument for the equality

$$\#\{x \in (0, \infty) : \mathcal{M}_{j+1, k+1}(x) = 1\} = \max\{2(j+k) - 5, 1\}. \quad (15)$$

Plots of several finite-trail Mina margin maps $(0, \infty) \rightarrow (0, \infty) : x \rightarrow \mathcal{M}_{j+1, k+1}(x)$ led to the conjecture. The pattern begins to emerge in the four plots displayed in Figure 2.2, for which $j+k \in \llbracket 2, 5 \rrbracket$. To see the pattern continue, we need higher values of $j+k$. For these, a suitable device is the finite-trail Θ -transformed Mina margin map $(0, \infty) \rightarrow (0, \infty) : x \rightarrow (\mathcal{M}_{j+1, k+1} \circ \Theta)(x)$ mentioned at the end of Section 2.4: see Figure 2.3 for four depictions, where $j+k \in \llbracket 6, 10 \rrbracket$.

2.6. Some further formulas

In this article, we study a new game, presenting conjectures as well as results. We have derived some formulas of which we do not make use, and we choose to present them as our final results in the introduction because they appear interesting and could be of value in further study of the Trail of Lost Pennies. First we state Theorem 2.31, an alternative explicit form for standard ABMN solutions. Then we present the A system, which is a closed \mathbb{Z} -indexed set of equations that we find in Theorem 2.32 to describe the a - (or b -)variables in any time-invariant Nash equilibrium in the special case of the game with a symmetric form of boundary data.

2.6.1. Alternative formulas for standard solutions and their Mina margins. Recall the function $Z : (0, \infty) \rightarrow (0, \infty)$, $Z(x) = m_\infty^{\text{def}}(x) = \sum_{k \in \mathbb{Z}} \prod_{i=0}^k (c_i(x) - 1)$, from the remark that concludes Section 2.3.

THEOREM 2.31. *Let f, g and h mapping $(0, \infty)$ to itself be specified by*

$$f(x) = \frac{xc(x)d(x)}{(c(x) + xd(x))^2}; \quad g(x) = Z(x)^{-1}c(x)f(x); \quad \text{and} \quad h(x) = Z(x)^{-1}xd(x)f(x).$$

Let $x \in (0, \infty)$.

- (1) For $k \in \mathbb{Z}$, $a_k^{\text{st}}(x) = g(s_k(x))$ and $b_k^{\text{st}}(x) = h(s_k(x))$.
- (2) For $j, k \in \mathbb{Z}$ such that $j < k$,

$$m_k^{\text{st}}(x) - m_j^{\text{st}}(x) = \sum_{i=j+1}^k \frac{1}{Z(s_i(x))} \quad \text{and} \quad n_j^{\text{st}}(x) - n_k^{\text{st}}(x) = \sum_{i=j+1}^k \frac{s_i(x)}{Z(s_i(x))}.$$

In particular, $m_k^{\text{st}}(x) - m_{k-1}^{\text{st}}(x) = Z(s_k(x))^{-1}$ and $n_k^{\text{st}}(x) - n_{k-1}^{\text{st}}(x) = s_k(x)Z(s_k(x))^{-1}$.

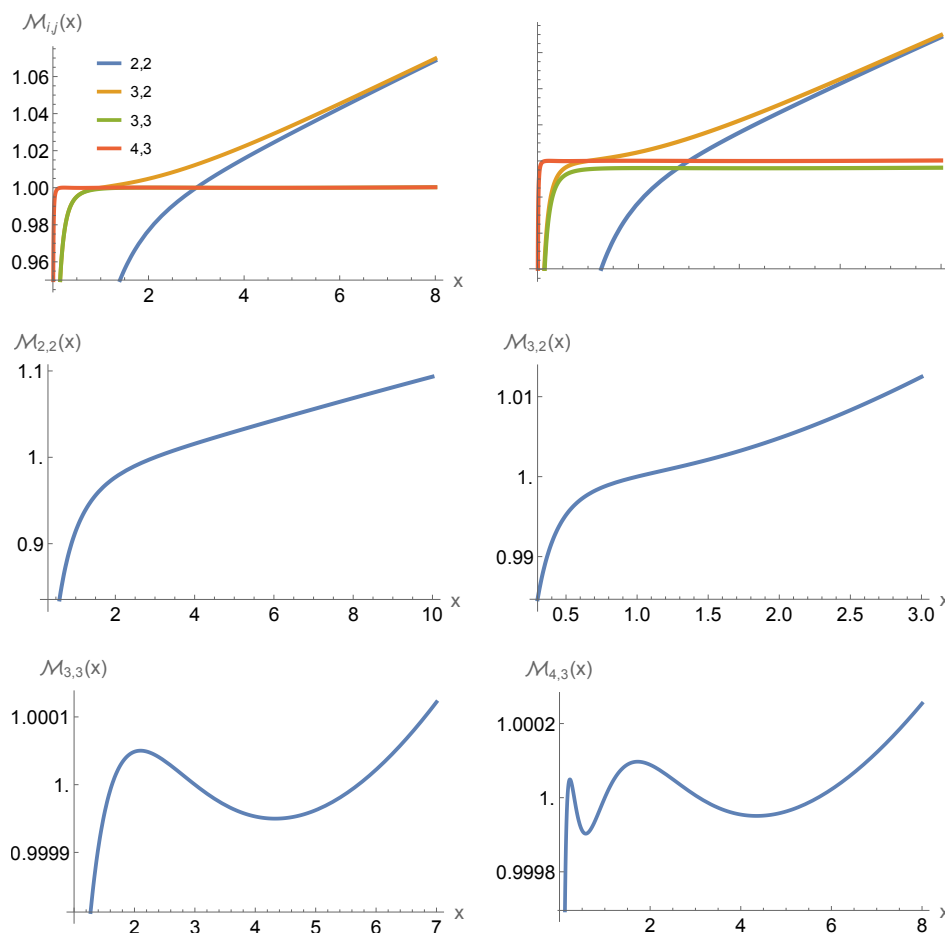


FIGURE 2.2. Four finite-trail Mina margin maps $\mathcal{M}_{j+1,k+1} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$ are depicted, for values of $j+k$ in $\llbracket 2, 5 \rrbracket$. *Top left.* The four functions are plotted together. *Top right.* This is a ‘Tube Map’ of the left-hand graph (a distorted but practical depiction), in which the green curve has been artificially displaced to separate it, so that the viewer may watch the different lines as they run.

The green and red curves seem to suggest that the curves converge to the constant function one as $j+k$ rises, but this impression is false. Indeed, the middle and lower graphs plot the four functions in turn, each on a scale that shows the finer journey of the map as it rises through height one. The maps lose injectivity in the $(j+1, k+1)$ -index change $(3, 2) \rightarrow (3, 3)$.

- (3) For $j, k \in \mathbb{N}$, the finite trail Mina margin map $\mathcal{M}_{j+1,k+1} : (0, \infty) \rightarrow (0, \infty)$ satisfies the equation

$$\mathcal{M}_{j+1,k+1}(x) = \left(\sum_{i=-j}^{k+1} \frac{1}{Z(s_i(x))} \right)^{-1} \cdot \sum_{i=-j}^{k+1} \frac{s_i(x)}{Z(s_i(x))}.$$

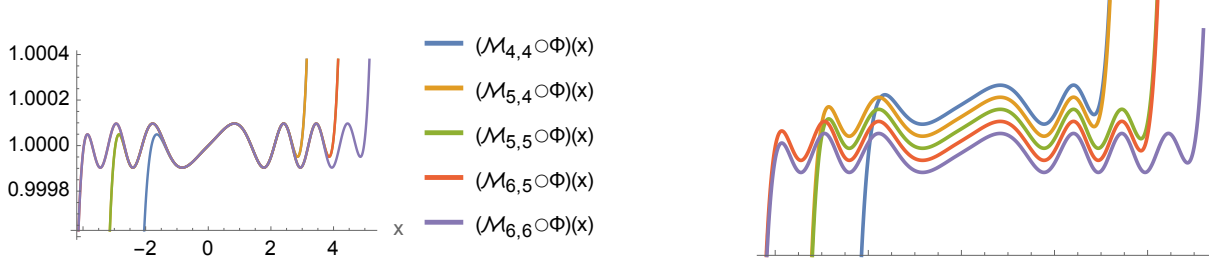


FIGURE 2.3. *Left:* Five Θ -transformed finite-trail Mina margin maps $\mathcal{M}_{j+1,k+1} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$ are depicted, for increasing values of $j+k$ in $\llbracket 6, 10 \rrbracket$. The graphs join and leave a shared highway, which is (up to visually negligible discrepancies) the graph of the limiting transformed map $\mathcal{M} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$. *Right:* As in Figure 2.2(top,right), curves have been artificially displaced so that their routes can be clearly seen.

(4) *The Mina margin map $\mathcal{M} : (0, \infty) \rightarrow (0, \infty)$ satisfies*

$$\mathcal{M}(x) = \sum_{i \in \mathbb{Z}} \frac{s_i(x)}{Z(s_i(x))} .$$

2.6.2. The game with symmetric boundary data. As Section 1.2.1 reviewed, economists have often considered player-funded tug-of-war games in a symmetric form (with players receiving equal terminal rewards), finding and analysing symmetric equilibria (in which the stake of the leading player is determined by the position of the counter relative to the gameboard midpoint rather than by the identity of the player). In our case, these symmetries significantly simplify analysis, while leading to equilibria that exemplify generic features. Here we present our conclusions about them. The A system on \mathbb{Z} is the set of equations⁵ in the real variables A_i , $i \in \mathbb{Z}$:

$$A_{-i-1}(2A_i + A_{-i}) = A_{i+1}^2, \quad (16)$$

where the index ranges over \mathbb{Z} . We will also speak of the A system on $\mathbb{Z} + 1/2$. In this case, the real variables A_i are indexed by i in the one-half-offset lattice $\mathbb{Z} + 1/2$; the set of equations is given by (16) with the index ranging over $\mathbb{Z} + 1/2$.

By the Trail of Lost Pennies on \mathbb{Z} in its symmetric form is meant the game $\text{Trail}(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$, where the boundary parameters are supposed to satisfy $m_{-\infty} = n_{\infty} = 0$ and $m_{\infty} = n_{-\infty}$. There is thus a one-parameter family of such games, indexed by $m_{\infty} \in (0, \infty)$.

THEOREM 2.32.

- (1) *For $\lambda \in (0, \infty)$, there is a unique solution $\{a_i(\lambda) : i \in \mathbb{Z}\}$ of the A system on \mathbb{Z} such that $a_0(\lambda) = \lambda$. The solutions satisfy $a_i(\lambda) = \lambda a_i(1)$ for $\lambda \in (0, \infty)$ and $i \in \mathbb{Z}$.*

⁵In Section III.3 of Vickers' thesis [58], symmetric equilibria are found in the Harris-Vickers tug-of-war game reviewed in Subsection 1.2.1.1, with cost functions $c_{-}(x) = c_{+}(x) = x^2/2$. Effort rates x_n satisfy $x_{-n}x_{n-1} = x_n^2/2$ for positive and negative n . This cousin of the A-system is solved $x_n = x_0 2^{-n(n-1)/2}$. As Subsection 1.2.1.1 reviewed, the elements of the Battlefield Cyl Fog are found, but decay rates are $e^{-\Theta(1)n^2}$, rather than doubly exponential in n .

- (2) For $\lambda \in (0, \infty)$, there is a unique solution $\{A_i(\lambda) : i \in \mathbb{Z} + 1/2\}$ of the A system on $\mathbb{Z} + 1/2$ such that $A_{1/2}(\lambda) = \lambda$. The solutions satisfy $A_i(\lambda) = \lambda A_i(1)$ for $\lambda \in (0, \infty)$ and $i \in \mathbb{Z} + 1/2$.
- (3) In the notation of the first part, let S_1 denote the set of strategy pairs $(a_{-i+k}(\lambda), a_{i+k}(\lambda) : i \in \mathbb{Z})$ indexed by $k \in \mathbb{Z}$ and $\lambda \in (0, \infty)$. In the notation of the second part, let S_2 denote the set of the strategy pairs $(A_{-i-1/2+k}(\lambda), A_{i+1/2+k}(\lambda) : i \in \mathbb{Z})$ with the same index set. The elements of $S_1 \cup S_2$ are pairwise distinct time-invariant Nash equilibria for the Trail of Lost Pennies on \mathbb{Z} in its symmetric form.
- (4) Admit Conjecture 2.11 in the special case that $x = 1$: namely, suppose that $Q(1) = 2$. Then there are no other time-invariant Nash equilibria for the Trail of Lost Pennies on \mathbb{Z} in its symmetric form than those identified in the preceding part.

2.7. The article's structure

There are five further chapters and an appendix. Two basic aspects of later use are treated in Chapter 3: a role-reversal symmetry satisfied by the ABMN system; and the solution of the simplest of the finite trail games, with just one site in open play. The fundamental relationship Theorem 2.6 between Nash equilibria and the ABMN equations is proved in Chapter 4. In Chapter 5, the explicit solution Theorem 2.21 and asymptotic decay estimates Theorem 2.14 are derived, along with the eventual gameplay unanimity Theorem 2.15. The Mina margin map \mathcal{M} is addressed in Chapter 6: its approximation by finite-trail counterparts, Theorem 2.23 and several consequences among our main results; the map's Θ -transformed version, and Theorem 2.28; and an explicitly recorded computation that \mathcal{M} evaluated at 0.58 is bounded away from one above, and thus Theorem 2.4(3). In Chapter 7, we discuss several aspects of our results and proofs and some prospects for further study. The appendix contains the proofs of the further formulas from Section 2.6.

CHAPTER 3

Some basic tools

Role-reversal symmetry is treated in Section 3.1 and the trail game on $\llbracket -1, 1 \rrbracket$ in Section 3.2. Later sections introduce some further basic notation and properties.

3.1. Role-reversal symmetry

DEFINITION 3.1. The *role-reversal map* \mathcal{R} sends the space of quadruples $\mathbb{Z} \rightarrow \mathbb{R}^4$ to itself by mapping $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ to $\{(b_{-i}, a_{-i}, n_{-i}, m_{-i}) : i \in \mathbb{Z}\}$.

PROPOSITION 3.2. *Suppose that $(a, b, m, n) = \{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ is an ABMN solution. Then so is $\mathcal{R}(a, b, m, n)$.*

Proof. The result may be verified by inspecting the ABMN equations. We instead indicate in rough terms a more conceptual, game-theoretic, argument which is available for positive ABMN solutions if we admit their connection to the trail game via Theorem 2.6. Suppose that a time-invariant Nash equilibrium $(b, a) : \mathbb{Z} \rightarrow (0, \infty)$ is played in the first instance. If Mina and Maxine swap roles, so that the strategy pair (a, b) is played, each acts in diametric opposition to her interests. But if the gameboard is then reflected through the origin, these interests are reversed, and each plays optimally once more. It is the strategy pair $\mathbb{Z} \rightarrow (0, \infty)^2 : i \rightarrow (a_{-i}, b_{-i})$ that is now being played. This pair is a Nash equilibrium (for the game whose boundary data in \mathbb{R}^4 is specified by this pair), and the associated quadruple is an ABMN solution. This quadruple is $\mathcal{R}(a, b, m, n)$. \square

We will obtain Proposition 2.29 by using the role-reversal map \mathcal{R} on quadruples whose index set is finite; to do so, we extend our notation to handle this circumstance. Let $j, k \in \mathbb{N}_+$ and let $\{(a_i, b_i, m_i, n_i) \in \mathbb{R}^4 : i \in \llbracket -j, k \rrbracket\}$. We set $\mathcal{R}(a, b, m, n)$ equal to $\{(b_{-i}, a_{-i}, n_{-i}, m_{-i}) \in \mathbb{R}^4 : i \in \llbracket -k, j \rrbracket\}$. Proposition 3.2 has a counterpart in the finite case which asserts that

$$\begin{aligned} (a, b, m, n) : \llbracket -j, k \rrbracket &\rightarrow \mathbb{R}^4 \text{ is an ABMN solution} \\ \implies \mathcal{R}(a, b, m, n) : \llbracket -k, j \rrbracket &\rightarrow \mathbb{R}^4 \text{ is an ABMN solution.} \end{aligned} \tag{17}$$

The ABMN equations can again be inspected to verify this statement.

We will further consider the left shift \mathcal{S}_1 , which sends any quadruple $(a, b, m, n) : \llbracket -j, k \rrbracket \rightarrow \mathbb{R}^4$ to the quadruple $\llbracket -j-1, k-1 \rrbracket \rightarrow \mathbb{R}^4 : i \rightarrow (a_{i+1}, b_{i+1}, m_{i+1}, n_{i+1})$.

Proof of Proposition 2.29(1). For $x \in (0, \infty)$, let $\{(a_i, b_i, m_i, n_i) : i \in \llbracket -k, k \rrbracket\}$ be an ABMN solution on $\llbracket -k, k \rrbracket$ such that $\frac{n_{-1}-n_0}{m_0-m_{-1}}$ equals x ; that such a solution may be found has been explained in Subsection 2.5.3. We *claim* that

$$\mathcal{M}_{k,k} \left(\frac{m_1-m_0}{n_0-n_1} \right) = \mathcal{M}_{k,k} \left(\frac{n_{-1}-n_0}{m_0-m_{-1}} \right)^{-1}. \tag{18}$$

Admitting this claim, we see that $s(x) = \frac{n_0-n_1}{m_1-m_0}$ by (17); thus, $1/s(x) = \frac{m_1-m_0}{n_0-n_1}$. Using the claim, we confirm Proposition 2.29(1).

To confirm (18), we let $\hat{\phi}_i$ denote the ϕ_i -value of $\hat{R}(a, b, m, n)$. The claim follows from

$$\mathcal{M}_{k,k}\left(\frac{m_1-m_0}{n_0-n_1}\right) = \mathcal{M}_{k,k}(\hat{\phi}_0) = \frac{\hat{n}_{-k} - \hat{n}_k}{\hat{m}_k - \hat{m}_{-k}} = \frac{m_k - m_{-k}}{n_{-k} - n_k} = \mathcal{M}_{k,k}\left(\frac{n_{-1}-n_0}{m_0-m_{-1}}\right)^{-1},$$

where the second and fourth equalities are due to (14).

(2). We now let $\{(a_i, b_i, m_i, n_i) : i \in \llbracket -k-1, k \rrbracket\}$ be an ABMN solution on $\llbracket -k-1, k \rrbracket$ such that $\frac{n_{-1}-n_0}{m_0-m_{-1}} = x$. We consider the operator $\mathcal{A} = \mathcal{S}_1 \circ \mathcal{R}$; note that, directly from (17), $\mathcal{A}(a, b, m, n)$ is an ABMN solution, also on the index set $\llbracket -k-1, k \rrbracket$. We denote $\mathcal{A}(a, b, m, n) = \{(\tilde{a}_i, \tilde{b}_i, \tilde{m}_i, \tilde{n}_i) : i \in \llbracket -k-1, k \rrbracket\}$; and we let $\tilde{\phi}_i$ denote the ϕ_i -value of $\mathcal{A}(a, b, m, n)$ for $i \in \llbracket -k, k-1 \rrbracket$.

By (14),

$$\mathcal{M}_{k+1,k}(\tilde{\phi}_0) = \frac{\tilde{n}_{-k-1} - \tilde{n}_k}{\tilde{m}_k - \tilde{m}_{-k-1}} = \frac{m_k - m_{-k-1}}{n_{-k-1} - n_k}.$$

And again by (14), $\mathcal{M}_{k+1,k}\left(\frac{n_{-1}-n_0}{m_0-m_{-1}}\right) = \frac{n_{-k-1}-n_k}{m_k-m_{-k-1}}$. Hence, we obtain

$$\mathcal{M}_{k+1,k}(\tilde{\phi}_0) = \mathcal{M}_{k+1,k}\left(\frac{n_{-1}-n_0}{m_0-m_{-1}}\right)^{-1}. \quad (19)$$

Note further that

$$\tilde{\phi}_0 = \frac{\tilde{n}_{-1} - \tilde{n}_0}{\tilde{m}_0 - \tilde{m}_{-1}} = \frac{m_0 - m_{-1}}{n_{-1} - n_0}.$$

Since $x = \frac{n_{-1}-n_0}{m_0-m_{-1}}$, we have that $\tilde{\phi}_0 = 1/x$. From (19), we thus obtain Proposition 2.29(2). \square

COROLLARY 3.3. For $k \in \mathbb{N}_+$, $\mathcal{M}_{k,k}(3) = \mathcal{M}_{k+1,k}(1) = 1$.

Proof. By Proposition 2.29(1) and $s(3) = 1/3$, we have that $\mathcal{M}_{k,k}(3)^2 = 1$. Since $\mathcal{M}_{k,k} > 0$, we obtain $\mathcal{M}_{k,k}(3) = 1$. By Proposition 2.29(2), $\mathcal{M}_{k+1,k}(1)^2 = 1$. Since $\mathcal{M}_{k+1,k} > 0$, we confirm that $\mathcal{M}_{k+1,k}(1) = 1$. \square

The form of the inverse of the map s may be obtained by use of role-reversal symmetry.

PROPOSITION 3.4. The function $s : (0, \infty) \rightarrow (0, \infty)$ from Definition 2.18 is invertible, and its inverse is given by

$$s^{-1}(x) = \frac{1}{s(1/x)}, \quad \text{for } x \in (0, \infty).$$

Proof. It is enough to show that $h : (0, \infty) \rightarrow (0, \infty)$ given by $h(x) = 1/s(1/x)$ satisfies

$$(s \circ h)(x) = (h \circ s)(x) = x. \quad (20)$$

Set

$$(a_i, b_i, m_i, n_i) = (a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x)), \quad \text{for } i \in \mathbb{Z}.$$

We have that $\phi_0 = \frac{n_{-1}-n_0}{m_0-m_{-1}} = x$. First note that, by Proposition 2.26, $s(\phi_{-1}) = \phi_0$; or, in other words,

$$s\left(\frac{n_{-2} - n_{-1}}{m_{-1} - m_{-2}}\right) = \frac{n_{-1} - n_0}{m_0 - m_{-1}} = x. \quad (21)$$

Let $\hat{\phi} : \mathbb{Z} \rightarrow (0, \infty)$ be such that, for $i \in \mathbb{Z}$, $\hat{\phi}_i$ is the value of ϕ_i for the quadruple $\mathcal{R}(a, b, m, n)$. Note then that

$$\hat{\phi}_1 = \frac{\hat{n}_0 - \hat{n}_1}{\hat{m}_1 - \hat{m}_0} = \frac{m_0 - m_{-1}}{n_{-1} - n_0} = 1/x. \quad (22)$$

Thus, note that

$$s(1/x) = s(\hat{\phi}_1) = \hat{\phi}_2 = \frac{\hat{n}_1 - \hat{n}_2}{\hat{m}_2 - \hat{m}_1} = \frac{m_{-1} - m_{-2}}{n_{-2} - n_{-1}},$$

where the second equality is justified by Propositions 3.2 and 2.26. Applying s , we find from (21) that $s(1/s(1/x)) = x$. We have confirmed that $s(h(x)) = x$ for $x \in (0, \infty)$.

Next we note that $s(x) = \phi_1 = \frac{n_0 - n_1}{m_1 - m_0}$, so that $1/s(x) = \frac{m_1 - m_0}{n_0 - n_1} = \hat{\phi}_0$. But $s(\hat{\phi}_0) = \hat{\phi}_1 = 1/x$, by Propositions 2.26 and 3.2, and (22). Which is to say, $1/s(1/s(x)) = x$, or $h(s(x)) = x$ for $x \in (0, \infty)$. This completes the derivation of (20) and thus the proof of Proposition 3.4. \square

3.2. Penny Forfeit

The simplest case of the finite trail game from Section 2.5 has $j = k = 0$, when the first move is the last. The straightforward solution of this case is already instructive, and we provide it now, calling this game Penny Forfeit.

Here is an explicit description of this one-turn game. Maxine and Mina are asked to stake non-negative quantities a and b . After these stakes have been submitted, the game victor is declared: this will be Maxine, with probability $\frac{a}{a+b}$; otherwise, it will be Mina. If Maxine wins, she receives m_1 , and Mina n_1 ; if Mina wins, Maxine receives m_{-1} and Mina n_{-1} . These four values act as boundary data. They are supposed to be real values that satisfy $m_{-1} < m_1$ and $n_1 < n_{-1}$. Maxine and Mina's mean winnings in the game are

$$\frac{a}{a+b}m_1 + \frac{b}{a+b}m_{-1} - a \quad \text{and} \quad \frac{b}{a+b}n_{-1} + \frac{a}{a+b}n_1 - b, \quad (23)$$

where in each expression the respective terms are mean terminal receipt in the event of turn victory; such receipt in the event of turn defeat; and the negative contribution from the forfeited stake. The pair (b, a) is a Nash equilibrium—a notion that is specified by suitably adapting the definition in Section 2.1—when these last two expressions are both global maxima as the variables a and b are respectively varied over $[0, \infty)$.

LEMMA 3.5. *There is a unique solution in $(a, b) \in [0, \infty)^2$ in which the pair of expressions in (23) are both global maxima as the variables a and b are respectively varied over $[0, \infty)$. It is given by*

$$(a, b) = \left(\frac{M^2 N}{(M + N)^2}, \frac{M N^2}{(M + N)^2} \right), \quad \text{with } M = m_1 - m_{-1} \quad \text{and} \quad N = n_{-1} - n_1, \quad (24)$$

Note that a and b are strictly positive.

Proof. A critical point (a, b) is given by setting the respective partial derivatives in a and b of the two expressions equal to zero: the conditions are

$$\frac{b}{(a+b)^2}(m_1 - m_{-1}) - 1 = \frac{a}{(a+b)^2}(n_{-1} - n_1) - 1 = 0.$$

At least one component in the desired pair (a, b) must be non-zero. Indeed, if for example a equals zero, then an infinitesimal increase of b from zero will increase Mina's expected payoff from $\frac{n_{-1} + n_1}{2}$ to n_{-1} . Restricting then, as we may, to solutions with at least one positive component, we see that there exists a unique solution in $(a, b) \in [0, \infty)^2$ of the last displayed equations, and that this solution is given by (24). This is indeed a global maximum for the pair of expressions in (23) under respective variation of a and b in $[0, \infty)$. \square

Remark. We see then that Penny Forfeit has a unique Nash equilibrium (b, a) , with (a, b) as just specified. It is straightforward to see that this Nash equilibrium is unique even if we permit the players to offer random stakes.

3.3. The game with a delayed start

We will wish to consider the finite and infinite trail games begun at a turn whose index $\ell \in \mathbb{N}_+$ is general. For $(i, \ell) \in \mathbb{Z} \times \mathbb{N}_+$ and $(S_-, S_+) \in \mathcal{S}$ we will write $\mathbb{P}_{S_-, S_+}^{i, \ell}$ and $\mathbb{E}_{S_-, S_+}^{i, \ell}[\cdot]$ for the law and expectation operator of gameplay $X : \llbracket \ell, \infty \rrbracket = \mathbb{Z} \cap [\ell, \infty) \rightarrow \mathbb{Z}$, $X_\ell = i$, in the Trail of Lost Pennies on \mathbb{Z} , begun at the $(\ell + 1)^{\text{st}}$ turn at ℓ . Payoffs, costs and terminal receipts P_\pm , $C_\pm(u)$ (for $u \in \llbracket \ell, \infty \rrbracket$) and T_\pm remain as specified by Section 2.1. Mina and Maxine's payoff identities (5) and (9) now take the $\mathbb{P}_{S_-, S_+}^{i, \ell}$ -almost sure form

$$P_\pm = - \sum_{j=\ell+1}^{\infty} C_\pm(j) + T_\pm. \quad (25)$$

Note that \mathbb{P}_{S_-, S_+}^i equals $\mathbb{P}_{S_-, S_+}^{i, 0}$.

3.4. Lack of escape entails infinite costs

LEMMA 3.6. *Let $(S_1, S_2) \in \mathcal{S} \times \mathcal{S}_0$ be a strategy pair whose second component is time-invariant. Writing $a_i = S_2(i, j)$ for $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$, suppose that $a_i > 0$ for $i \in \mathbb{Z}$. For $(i, \ell) \in \mathbb{Z} \times \mathbb{N}_+$, $\mathbb{P}_{S_1, S_2}^{i, \ell}(E^c) > 0$ implies that $\mathbb{E}_{S_1, S_2}^{i, \ell}[P_-] = -\infty$.*

Proof. For $j \in \mathbb{Z}$, let $W_-(j)$ denote the event that Mina wins infinitely many turns at which the counter is at j . We claim that, up to a $\mathbb{P}_{S_1, S_2}^{i, \ell}$ -null set,

$$E^c \subseteq \bigcup_{j \in \mathbb{Z}} W_-(j). \quad (26)$$

To see this, set V_j denote the event that the counter visits $j \in \mathbb{Z}$ infinitely often. The occurrence of E^c entails that of $\bigcup_{j \in \mathbb{Z}} V_j$. If V_j occurs and Mina wins infinitely many of the turns at which X visits j , then $W_-(j)$ occurs. If V_j occurs but Mina does not thus succeed, there are infinitely many occasions on which X leaves j to the right, only to return to j at some later time. Consider the set of turns that occur just before each of these returns. At each, X is at $j + 1$ and Mina wins the turn, so that X passes to j . Thus, $W_-(j + 1)$ occurs. We have derived (26).

For $j \in \mathbb{Z}$, let $\text{TotalCost}_-(j) = \sum_{t=\ell}^{\infty} \mathbf{1}_{X_t=j} C_-(t + 1)$ denote Mina's running cost expended at j under $\mathbb{P}_{S_1, S_2}^{i, \ell}$. Let $N_-(j, j - 1) = \sum_{t=\ell}^{\infty} \mathbf{1}_{X_t=j, X_{t+1}=j-1}$ denote the number of turns with index at least $\ell + 1$ that are won by Mina and at whose start X visits j . Since $C_-(t + 1) = S_1(X_t, t + 1)$, we have that

$$\begin{aligned} \mathbb{E}_{S_1, S_2}^{i, \ell}[N_-(j, j - 1)] &= \sum_{t=\ell}^{\infty} \mathbb{P}_{S_1, S_2}^{i, \ell}(X_t = j) \cdot \frac{S_1(j, t+1)}{S_1(j, t+1) + a_j} \leq a_j^{-1} \sum_{t=\ell}^{\infty} \mathbb{P}_{S_1, S_2}^{i, \ell}(X_t = j) S_1(j, t + 1) \\ &= a_j^{-1} \mathbb{E}_{S_1, S_2}^{i, \ell}[\text{TotalCost}_-(j)] \leq a_j^{-1} \mathbb{E}_{S_1, S_2}^{i, \ell} \sum_{t=\ell}^{\infty} C_-(t). \end{aligned}$$

By (26) and $a_j > 0$ for $j \in \mathbb{Z}$, we see then that, if $\mathbb{P}_{S_1, S_2}^{i, \ell}(E^c) > 0$, then $\mathbb{E}_{S_1, S_2}^{i, \ell}[N_-(j, j - 1)]$ is infinite for some $j \in \mathbb{Z}$, and thus so is Mina's mean total running cost $\mathbb{E}_{S_1, S_2}^{i, \ell} \sum_{t=\ell}^{\infty} C_-(t)$. Applying $\mathbb{E}_{S_1, S_2}^{i, \ell}$

to (25) with $\pm = -1$ and noting that terminal receipts T_- are almost surely bounded, we find that Mina's mean payoff $\mathbb{E}_{S_1, S_2}^{i, \ell}[P_-]$ equals minus infinity. This completes the proof of Lemma 3.6. \square

3.5. Relating the finite and infinite trail games

Let $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$ satisfy $m_{-\infty} < m_{\infty}$ and $n_{\infty} < n_{-\infty}$. It is useful to specify a coupling of the Trail of Lost Pennies Trail $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ and its finite trail counterparts.

DEFINITION 3.7. Let $i \in \mathbb{Z}$ and $(S_-, S_+) \in \mathcal{S}^2$. Recall that the gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$, $X_0 = i$, of the infinite trail game governed by (S_-, S_+) is specified under the law \mathbb{P}_{S_-, S_+}^i . For $j, k \in \mathbb{N}_+$, strategy pairs in $\mathcal{S}[j, k]^2$ for the game with trail $\llbracket -j - 1, k + 1 \rrbracket$ result by restricting the domain of S_- and S_+ to $\llbracket -j, k \rrbracket$. Copies of the gameplay $X^{j, k} : \mathbb{N} \rightarrow \llbracket -j - 1, k + 1 \rrbracket$, $X_0^{j, k} = i$, that result from use of these restricted pairs may be coupled under \mathbb{P}_{S_-, S_+}^i whenever $j, k \in \mathbb{N}_+$ are such that $i \in \llbracket -j, k \rrbracket$. To specify these copies, set

$$\tau^{j, k} = \inf \{i \in \mathbb{N}_+ : X_i \in \{-j - 1, k + 1\}\}. \quad (27)$$

Writing \wedge for minimum, we then take $X^{j, k}(u) = X(u \wedge \tau^{j, k})$ for $u \in \mathbb{N}$.

The finite and infinite trail payoffs, costs and terminal receipts $*_{\pm}^{j, k}$ and $*_{\pm}$, $* \in \{P, C, T\}$, are coupled under \mathbb{P}_{S_-, S_+}^i by this definition. We note some basic relationships that result.

LEMMA 3.8. Let $(S_-, S_+) \in \mathcal{S}^2$. Suppose that $i \in \mathbb{Z}$ and $j, k \in \mathbb{N}_+$ satisfy $i \in \llbracket -j, k \rrbracket$.

- (1) We have that $P_- - P_-^{j, k} = T_- - T_-^{j, k} - \sum_{t=\tau^{j, k}}^{\infty} C_-(t)$.
- (2) And that $P_- - P_-^{j, k} \leq T_- - T_-^{j, k}$.
- (3) For $\ell \in \mathbb{N}$, it is $\mathbb{P}_{S_-, S_+}^{i, \ell}$ -almost certain that $P_- - P_-^{j, k} \leq n_{-\infty} - n_{\infty}$.

Proof: (1). This follows from (5), (13) and $C_-^{j, k}(t) = C_-(t)$ for $u \in \llbracket 0, \tau^{j, k} - 1 \rrbracket$.

(2). Due to the preceding part and the non-negativity of costs $C_-(t)$.

(3). The receipt $T_-^{j, k}$ is a weighted average of n_{-j-1} and n_{k+1} . Since $\{n_i : i \in \mathbb{Z}\}$ is decreasing (this due to Theorem 2.3(1), because this sequence is the n -component of a positive ABMN solution), we find that $T_-^{j, k} \geq n_{\infty}$. Also note that $P_- \leq n_{-\infty}$. The preceding part of the lemma thus implies the stated result. \square

LEMMA 3.9. Let $(S_-, S_+) \in \mathcal{S}^2$, $k \in \mathbb{Z}$ and $\ell \in \mathbb{N}$. Let $Q \in \mathbb{N} \cup \{\infty\}$ be a stopping time with respect to gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$ under the law $\mathbb{P}_{S_-, S_+}^{k, \ell}$ specified in Section 3.3. Then

$$\mathbb{E}_{S_-, S_+}^{k, \ell}[P_-] = -\mathbb{E}_{S_-, S_+}^{k, \ell} \sum_{t=\ell+1}^{Q-1} C_-(t) + \mathbb{E}_{S_-, S_+}^{k, \ell} [\mathbb{E}_{S_-, S_+}^{X(Q)}[P_-]]. \quad (28)$$

In reading this display in the event that $Q = \infty$, we adopt the conventions that $\mathbb{E}_{S_-, S_+}^{\infty, \ell}[P_-] = n_{\infty}$ and $\mathbb{E}_{S_-, S_+}^{-\infty, \ell}[P_-] = n_{-\infty}$, as well as $Q - 1 = \infty$. We also have the counterpart identity for Maxine, given by $P_- \rightarrow P_+$ and $C_- \rightarrow C_+$.

Proof. The right-hand side of (25) with $\pm = -1$ may be written $A_1 + A_2$, where A_1 is the sum of costs $C_-(t)$ with $\ell + 1 \leq t < Q$; and A_2 is the sum of the higher indexed costs (in the case that Q is finite) and the terminal receipt T_- . Since T_- equals $n_{-\infty}$ or n_{∞} when the events E_- or E_+ occur, we find that, when the mean $\mathbb{E}_{S_-, S_+}^{k, \ell}$ of (25) thus represented is taken, the two right-hand terms in the lemma result. \square

We have used Theorem 2.3(1), and we will use it again in a moment. We now give the simple proofs of Theorem 2.3(1,2).

Proof of Theorem 2.3(1). Since $a_i + b_i > 0$, ABMN(3) implies that $m_{i+1} > m_{i-1}$. We may rearrange ABMN(1) in the form $m_i = \frac{a_i}{a_i + b_i} m_{i+1} + \frac{b_i}{a_i + b_i} m_{i-1} - a_i$. Using $m_{i-1} < m_{i+1}$ and $b_i > 0$, we find that $m_i < m_{i+1} - a_i$. Since $a_i > 0$, $m_i < m_{i+1}$. That $n_{i+1} < n_i$ follows similarly. We have shown that the ABMN solution (a, b, m, n) is strict.

(2). The sequences $\{m_i : i \in \mathbb{Z}\}$ and $\{n_i : i \in \mathbb{Z}\}$ are increasing and decreasing, by the preceding part. Thus, the limiting values (11) exist, at least as elements of $\mathbb{R} \cup \{\infty\} \cup \{-\infty\}$; they satisfy $m_{\infty} > m_{-\infty}$ and $n_{-\infty} > n_{\infty}$. \square

Note that Theorem 2.3(1,2) do not exclude the possibilities that m_{∞} or $n_{-\infty}$ equals ∞ or that n_{∞} or $m_{-\infty}$ equals $-\infty$. We will deduce this when we prove Theorem 2.3(3). This result will be derived in Section 5.3 as a consequence of the asymptotic decay estimate Theorem 2.14.

The next result interprets the m - and n -components of a ABMN solution as mean payoffs. It is couched in the notation of delayed-start games from Section 3.3.

LEMMA 3.10. *Let $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ denote a positive solution of the ABMN equations with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$. Let $S_-, S_+ \in \mathcal{S}$ satisfy $S_-(i, j) = b_i$ and $S_+(i, j) = a_i$ for $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$.*

(1) *Let $i \in \mathbb{Z}$ and $\ell \in \mathbb{N}$. Then $\mathbb{P}_{S_-, S_+}^{i, \ell}(E) = 1$.*

(2) *For $i \in \mathbb{Z}$ and $\ell \in \mathbb{N}$,*

$$m_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_+] \quad \text{and} \quad n_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_-].$$

(3) *Let $j, k \in \mathbb{N}$ and $\ell \in \mathbb{N}$. For $i \in \llbracket -j, k \rrbracket$,*

$$m_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_+^{j, k}] \quad \text{and} \quad n_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_-^{j, k}].$$

Proof: (1). By Theorem 2.3(1), $n_i > n_{\infty}$. But $n_{\infty} > -\infty$ by assumption. Thus Lemma 3.6 implies the sought statement.

(2). Since $a_i + b_i > 0$, ABMN(1) may be written in the form $m_i = \frac{a_i}{a_i + b_i} m_{i+1} + \frac{b_i}{a_i + b_i} m_{i-1} - a_i$ or equivalently $m_i = \mathbb{E}_{S_-, S_+}^i[m(X_1)] - a_i$. Iterating, we find that

$$m_i = \mathbb{E}_{S_-, S_+}^{i, \ell}[m(X_{u+1})] - \mathbb{E}_{S_-, S_+}^{i, \ell} \sum_{i=\ell}^u b_{X(i)} \quad (29)$$

for any $u \in \mathbb{N}$, $u \geq \ell$. The value of $\lim_{u \rightarrow \infty} m(X_u)$ exists on the event E , equalling m_{∞} or $m_{-\infty}$ according to whether E_+ or E_- occurs. By Lemma 3.10(1), we see that $\lim_{u \rightarrow \infty} \mathbb{E}_{S_-, S_+}^{i, \ell}[m(X_{u+1})]$ equals $m_{\infty} \cdot \mathbb{P}_{S_-, S_+}^{i, \ell}(E_+) + m_{-\infty} \cdot \mathbb{P}_{S_-, S_+}^i(E_-)$. In the notation of Lemma 3.9, we find by taking the high- u limit of the preceding display that m_i equals the right-hand side of (28) with $k = i$

and Q identically equal to infinity. Thus, Lemma 3.9 implies that $m_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_+]$. That $n_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_-]$ is similarly proved.

(3). We may obtain (29) with X replaced by its stopped version $X^{j, k}$. By taking the high- u limit, we find that m_i equals the right-hand side of (28) with $k = i$ and $Q = \tau^{j, k}$. From Lemma 3.9 we thus find that $m_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_+^{j, k}]$. That $n_i = \mathbb{E}_{(S_-, S_+)}^{i, \ell}[P_-^{j, k}]$ follows similarly. This completes the proof of Lemma 3.10(3). \square

Some simple relationships between escape in the finite and infinite trail games are now recorded. We define the events $E_-[-j, k] = \{X(\tau^{j, k}) = -j - 1\}$ and $E_+[j, k] = \{X(\tau^{j, k}) = k + 1\}$.

LEMMA 3.11. *We have that*

$$E_- = \bigcup_{k=1}^{\infty} \bigcap_{j=1}^{\infty} E_-[j, k] \quad \text{and} \quad E_+ = \bigcup_{j=1}^{\infty} \bigcap_{k=1}^{\infty} E_+[j, k].$$

Proof. These follow from the definitions of the events E_- and E_+ . \square

LEMMA 3.12. *We have that*

$$\lim_{k \rightarrow \infty} \mathbb{P}_{S_-, S_+}^i \left(E_- \setminus \left\{ \lim_{j \rightarrow \infty} m(X_{\tau^{j, k}}) = m_{-\infty} \right\} \right) = 0. \quad (30)$$

and

$$\lim_{j \rightarrow \infty} \mathbb{P}_{S_-, S_+}^i \left(E_+ \setminus \left\{ \lim_{k \rightarrow \infty} m(X_{\tau^{j, k}}) = m_{\infty} \right\} \right) = 0.$$

These statements are also valid if we replace all instances of m by n .

Proof. By Lemma 3.11 for E_- , we see that, on this event, there exists a random value $K \in \mathbb{N}_+$ such that, for all $j \in \mathbb{N}_+$, $X(\tau^{j, K}) = -j - 1$. Since $m_{-i} \rightarrow m_{-\infty}$ as $i \rightarrow \infty$, we see that, on E_- , $\lim_j m(X(\tau^{j, K})) = m_{-\infty}$. Thus, we obtain (30). The other three assertions made by the lemma have the same proof up to evident notational changes. \square

The structure of time-invariant Nash equilibria

The aim of this chapter is to prove Theorem 2.6, our result that relates time-invariant Nash equilibria and positive ABMN solutions. On the way to this result, we will establish some basic properties of time-invariant Nash equilibria. In the first section, we prove Theorem 2.6(1) alongside some simple properties of strategy pairs. The second proves Theorem 2.6(2).

4.1. Time-invariant Nash equilibria result in positive ABMN solutions

Here, we prove Theorem 2.6(1). Our style of argument is hands on: we build up inferences on the behaviour of a time-invariant Nash equilibrium step-by-step. With one exception: to close out the proof, we will invoke the unanimity Theorem 2.15(2,3), which is argued independently by explicit solution of the ABMN system in Chapter 5.

Recall the mean payoff notation (5) and (9). A strategy pair $(S_-, S_+) \in \mathcal{S}^2$ has *finite mean costs* if neither $\mathbb{E}_{S_-, S_+}^k[P_-]$ nor $\mathbb{E}_{S_-, S_+}^k[P_+]$ equals minus infinity, for any $k \in \mathbb{Z}$.

Let $(S_-, S_+) \in \mathcal{S}_0^2$. We adopt our standard convention of writing $b_i = S_-(i, j)$ and $a_i = S_+(i, j)$ for $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$. The *idle zone* $\mathcal{I} \subset \mathbb{Z}$ is given by $\mathcal{I} = \{j \in \mathbb{Z} : a_j = b_j = 0\}$.

LEMMA 4.1. *Let $(S_-, S_+) \in \mathcal{S}_0^2$ be such that $\mathcal{I} \neq \emptyset$. For $k \in \mathbb{Z}$, consider the gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$ under \mathbb{P}_{S_-, S_+}^k . For $i \in \mathbb{N}$ given, condition on the event that X_i is a given element of \mathcal{I} . (If $i = 0$, suppose that $k \in \mathcal{I}$.) Let $j \geq i$, $j = \inf \{m \in \mathbb{N} : X_m \notin \mathcal{I}\}$. Then the conditional law of $X : \llbracket i, j \rrbracket : \mathbb{N} \rightarrow \mathbb{Z}$ is equal to simple random walk given the value X_i stopped on leaving \mathcal{I} .*

Proof. At each turn with index in $\llbracket i, j - 1 \rrbracket$, neither Mina nor Maxine offers a positive stake, since the b and a values vanish in the idle zone. The gameplay increments $X(h + 1) - X(h)$ for $h \in \llbracket i, j - 1 \rrbracket$ are thus unbiased ± 1 steps as determined by the $0/0 = 1/2$ rule that was specified in Section 2.1. \square

Recall the escape event E from (6). An element of \mathcal{S}_0^2 is non-zero when at least one of its components is not identically zero.

PROPOSITION 4.2. *Let $(S_-, S_+) \in \mathcal{S}_0^2$ be non-zero, with finite mean costs. Then escape occurs almost surely: $\mathbb{P}_{S_-, S_+}^k(E) = 1$ for $k \in \mathbb{Z}$.*

Proof. Let $k \in \mathbb{Z}$ and suppose that $\mathbb{P}_{S_-, S_+}^k(E^c) > 0$. We may find $\ell \in \mathbb{Z}$ such that it is with positive probability that the process X , under the law \mathbb{P}_{S_-, S_+}^k , visits ℓ infinitely often. Since (S_-, S_+) is time-invariant, the strong Markov property implies that

$$\mathbb{P}_{S_-, S_+}^k \left(X \text{ visits } \ell \text{ infinitely often} \mid X \text{ visits } \ell \text{ at least once} \right) = 1. \quad (31)$$

Let $i \in \mathbb{Z} \cup \{-\infty\}$, $j \in \mathbb{Z} \cup \{\infty\}$, $i \leq j$, be such that at least one of i and j is finite; $\ell \in \llbracket i, j \rrbracket$; $a_m = b_m = 0$ for $m \in \mathbb{Z} \cap (i, j)$; and at least one of a_m and b_m is positive for any endpoint $m \in \{i, j\}$ that is finite. (It may be that $i = j = \ell$; in this case, some of these conditions are vacuous. In the other event, $i < \ell < j$.)

Suppose that $i < \ell < j$. Note that $\llbracket i + 1, j - 1 \rrbracket \subset \mathcal{I}$. We now consider the conditional law of X under \mathbb{P}_{S_-, S_+}^k given that X visits ℓ infinitely often. We invoke (31) to note that the conditioning disappears at the first visit of X to ℓ . Lemma 4.1 thus implies that, on each occasion that X visits ℓ , X pursues a simple random walk until it reaches i or j . Suppose, without loss of generality, that the index i is finite, and that $a_i > 0$. It is with probability at least $2^{-(\ell-i)}$ that X proceeds from a visit to ℓ by means of a string of leftward steps to reach i . Later, the conditioned walk X inevitably returns to ℓ , and a further opportunity to reach i directly ensues. Thus, X will infinitely often visit i , a location to which a assigns positive value. (Note that this conclusion also holds trivially in the opposing case, where $i = j = \ell$.) The cost $\sum_{t \geq 1} C_+(t)$ incurred by Maxine (which is specified in Section 2.1) is thus seen to be almost surely infinite on the \mathbb{P}_{S_-, S_+}^k -positive probability event that X visits ℓ infinitely often. (Were b_i instead supposed positive, then it would be Mina's cost $\sum_{t \geq 1} C_-(t)$ that is found to be infinite.) This is contrary to our assumption that (S_-, S_+) has finite mean costs. We conclude, as desired, that $\mathbb{P}_{S_-, S_+}^k(E) = 1$. \square

For $S \in \mathcal{S}_0$, let $\text{Left}(S) \in \mathbb{Z} \cup \{-\infty\} \cup \{\infty\}$ denote $\inf\{i \in \mathbb{Z} : S(i, 1) > 0\}$; and let $\text{Right}(S) \in \mathbb{Z} \cup \{-\infty\} \cup \{\infty\}$ denote $\sup\{i \in \mathbb{Z} : S(i, 1) > 0\}$. We say that S is *wide* if $\text{Left}(S) = -\infty$ and $\text{Right}(S) = \infty$; if S is not wide, it is *narrow*.

The right rocket $\eta \cdot \text{Rocket}^{i \rightarrow}$ at $i \in \mathbb{Z}$ of strength $\eta \in (0, \infty)$ is the element of \mathcal{S}_0 given

$$\eta \cdot \text{Rocket}_j^{i \rightarrow} = \eta \cdot 2^{-(j-i)-1} \mathbf{1}_{j \geq i}, \quad j \in \mathbb{Z}.$$

The counterpart left rocket $\eta \cdot \text{Rocket}^{\leftarrow i} \in \mathcal{S}_0$ is

$$\eta \cdot \text{Rocket}_j^{\leftarrow i} = \eta \cdot 2^{-(i-j)-1} \mathbf{1}_{j \leq i}, \quad j \in \mathbb{Z}.$$

The right boost at $i \in \mathbb{Z}$ of strength η is the map $\text{Boost}_\eta^{i \rightarrow} : \mathcal{S}_0 \rightarrow \mathcal{S}_0$ that sends $q = (q_i : i \in \mathbb{Z}) \in \mathcal{S}_0$ to $q + \eta \cdot \text{Rocket}^{i \rightarrow}$. The corresponding left boost $\text{Boost}_\eta^{i \leftarrow} : \mathcal{S}_0 \rightarrow \mathcal{S}_0$ sends q to $q + \eta \cdot \text{Rocket}^{i \leftarrow}$.

The right drag at $i \in \mathbb{Z}$ is the map $\text{Drag}^{i \rightarrow} : \mathcal{S}_0 \rightarrow \mathcal{S}_0$ that sends $q \in \mathcal{S}_0$ to the map

$$\mathbb{Z} \rightarrow (0, \infty) : j \rightarrow \begin{cases} q_j/2 & \text{if } j \geq i \\ q_j & \text{if } j < i. \end{cases}$$

The counterpart left drag $\text{Drag}^{i \leftarrow} : \mathcal{S}_0 \rightarrow \mathcal{S}_0$ sends $q \in \mathcal{S}_0$ to

$$\mathbb{Z} \rightarrow (0, \infty) : j \rightarrow \begin{cases} q_j/2 & \text{if } j \leq i \\ q_j & \text{if } j > i. \end{cases}$$

LEMMA 4.3. *Let $(S_-, S_+) \in \mathcal{S}_0^2$.*

- (1) *Suppose that the quantities $\text{Right}(S_-)$ and $\text{Right}(S_+)$ are finite. Let $i \in \mathbb{Z}$ exceed their maximum. There $\mathbb{E}_{S_-, \text{Boost}_\eta^{i \rightarrow}(S_+)}^i[P_+] > \mathbb{E}_{S_-, S_+}^i[P_+]$ for $\eta \in (0, m_\infty - m_{-\infty})$.*
- (2) *Suppose that $\text{Right}(S_+) = \infty$ and $\text{Right}(S_-) < \infty$. Let $i \in \mathbb{Z}$, $i > \text{Right}(S_-)$, satisfy $S_+(i, 1) > 0$. Then $\mathbb{E}_{S_-, \text{Drag}^{i \rightarrow}(S_+)}^i[P_+] > \mathbb{E}_{S_-, S_+}^i[P_+]$.*

- (3) If $\text{Left}(S_-)$ and $\text{Left}(S_+)$ exceed $-\infty$ and $i \in \mathbb{Z}$ is less than their minimum, then, provided that $\eta \in (0, n_{-\infty} - n_{\infty})$, we have that $\mathbb{E}_{\text{Boost}_{\eta}^{\leftarrow i}(S_-), S_+}^i[P_-] > \mathbb{E}_{S_-, S_+}^i[P_-]$.
- (4) If $\text{Left}(S_-) = -\infty$ and $\text{Left}(S_+) > -\infty$ and $i \in \mathbb{Z}$, $i < \text{Left}(S_+)$, satisfies $S_-(i, 1) > 0$, then $\mathbb{E}_{\text{Drag}^{\leftarrow i}(S_-), S_+}^i[P_-] > \mathbb{E}_{S_-, S_+}^i[P_-]$.

Proof: (1). The idle zone \mathcal{I} determined by (S_-, S_+) includes $\llbracket i, \infty$. By Lemma 4.1, X under \mathbb{P}_{S_-, S_+}^i thus behaves as a simple random walk when it visits $\llbracket i, \infty$. Right escape E_+ is thus impossible, so Maxine's mean terminal payoff $\mathbb{E}_{S_-, S_+}^i[T_+]$ is at most $m_{-\infty}$ because it is a weighted average of m_* and $m_{-\infty}$. Since $P_+ \leq T_+$ in view of running costs C_+ in (9) being non-negative, we find that $\mathbb{E}_{S_-, S_+}^i[P_+] \leq m_{-\infty}$. Now consider $\mathbb{P}_{S_-, \text{Boost}_{\eta}^i(S_+)}^i$. Under this law, X proceeds non-randomly by rightward steps, so that E_+ occurs almost surely. Since E_+ occurs, we have $T_+ = m_{\infty}$ almost surely. By the non-random rightward movement, we further have that $\sum_{t=1}^{\infty} C_+(t)$ equals $\sum_{t=1}^{\infty} \eta \cdot 2^{-t} = \eta$. By (9), we see then that $\mathbb{E}_{S_-, \text{Boost}_{\eta}^i(S_+)}^i[P_+] = m_{\infty} - \eta$. This confirms Lemma 4.3(1).

(2). Under gameplay governed by the law \mathbb{P}_{S_-, S_+}^i , Maxine offers a positive stake at i , and at infinitely many locations to its right, while Mina offers no stake at or to the right of i . Thus $X : \mathbb{N} \rightarrow \mathbb{Z}$, $X_0 = i$, remains always to the right of X , and tends to infinity. If Maxine switches from S_+ to $\text{Drag}^{i \rightarrow}(S_+)$, the law of gameplay X is unaffected, because the original and altered gameplays may be coupled so that Maxine's altered stake process is one-half of her original one, while Mina's remains identically zero—with the result that Maxine wins exactly the same turns in the altered gameplay as she did in the original one. The value of T_+ is thus almost surely equal to m_{∞} under $\mathbb{P}_{S_-, \text{Drag}^{i \rightarrow}(S_+)}^i$ as well under \mathbb{P}_{S_-, S_+}^i . But $\mathbb{E}_{S_-, \text{Drag}^{i \rightarrow}(S_+)}^i \sum_{t=1}^{\infty} C_+(t) = \frac{1}{2} \mathbb{E}_{S_-, S_+}^i \sum_{t=1}^{\infty} C_+(t)$ and $\mathbb{E}_{S_-, S_+}^i \sum_{t=1}^{\infty} C_+(t) \geq \mathbb{E}_{S_-, S_+}^i[C_+(1)] > 0$, so that $\mathbb{E}_{S_-, \text{Drag}^{i \rightarrow}(S_+)}^i \sum_{t=1}^{\infty} C_+(t) < \mathbb{E}_{S_-, S_+}^i \sum_{t=1}^{\infty} C_+(t)$. In summary, the switch to the altered strategy has maintained Maxine's terminal receipt but has reduced her running costs, so that Lemma 4.3(2) holds by (9).

(3,4). The preceding proofs may be readily adapted to prove these statements. \square

LEMMA 4.4.

- (1) Any element of \mathcal{N} has finite mean costs.¹
- (2) If $(S_-, S_+) \in \mathcal{S}_0^2$ is an element of \mathcal{N} then S_- and S_+ are wide.

In the ensuing proof and later, we denote the identically zero strategy by 0.

Proof of Lemma 4.4(1). Let $(S_-, S_+) \in \mathcal{N}$, and let $i \in \mathbb{Z}$. Note that $\mathbb{E}_{S_-, S_+}^i[P_+] \geq \mathbb{E}_{S_-, 0}^i[P_+]$. In evaluating the latter term, note that no running costs to Maxine have been incurred, so that the quantity is an average of terminal receipts m_{∞} , $m_{-\infty}$ and m_* to Maxine in the events E_+ , E_- and E^c . We see that $\mathbb{E}_{S_-, 0}^i[P_+] \geq \min\{m_{-\infty}, m_{\infty}, m_*\} = m_* > -\infty$, the latter inequality by assumption. Likewise, $\mathbb{E}_{S_-, S_+}^i[P_-] > -\infty$.

¹When the value of $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ is clear—and it is usually a generic quadruple satisfying (4)—we will often omit to record this notation when we denote \mathcal{N} . This includes the present case, where such a generic value is specified by the result, Theorem 2.6(1), that we are seeking to prove.

(2). We argue by contradiction and suppose, without loss of generality—for the other case is similar—that S_- is narrow. (Lemma 4.3(4) is not used in the ensuing proof. It is needed for the case whose proof we omit). Either $\text{Left}(S_-) > -\infty$ or $\text{Right}(S_-) < \infty$.

Suppose that $\text{Right}(S_-) < \infty$. If $\text{Right}(S_+) < \infty$, then Lemma 4.3(1) provides a strategy \hat{S}_+ to Maxine along with a value of $i \in \mathbb{Z}$ such that $\mathbb{E}_{S_-, \hat{S}_+}^i [P_+] > \mathbb{E}_{S_-, S_+}^i [P_+]$. But this is contrary to $(S_-, S_+) \in \mathcal{N}$. If $\text{Right}(S_+) = \infty$, then it is Lemma 4.3(2) that provides such $\hat{S}_+ \in \mathcal{S}_0$ and $i \in \mathbb{Z}$. A contradiction has thus been found in the case that $\text{Right}(S_-) < \infty$.

Suppose now that $\text{Left}(S_-) > -\infty$. If $\text{Left}(S_+) > -\infty$, then Lemma 4.3(3) furnishes a strategy \hat{S}_- for Mina and an index $i \in \mathbb{Z}$ for which $\mathbb{E}_{\hat{S}_-, S_+}^i [P_-] > \mathbb{E}_{S_-, S_+}^i [P_-]$ holds, contrary to $(S_-, S_+) \in \mathcal{N}$.

The case that $\text{Left}(S_-) > -\infty$ and $\text{Left}(S_+) = -\infty$ remains. The pair $(S_-, S_+) \in \mathcal{S}_0^2 \cap \mathcal{N}$ is non-zero, because S_+ is; it has finite mean costs by Lemma 4.4(1). Thus $\mathbb{P}_{S_-, S_+}^i (E^c) = 0$ by Proposition 4.2. Select $i \in \mathbb{Z}$ for which $S_+(i, 1) > 0$ and $S_-(j, 1) = 0$ for $j \in (-\infty, i]$. Note that $\mathbb{P}_{S_-, S_+}^i (E_-^c) = 0$ because gameplay X is at least i almost surely. Thus, $\mathbb{P}_{S_-, S_+}^i (E_+) = 1$, so that $T_+ = m_\infty$ almost surely. If Maxine drags down her strategy S_+ by replacing the stake she offers at i to be one-half of its value, the resulting strategy \hat{S}_+ is such that gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$ is equal under the laws \mathbb{P}_{S_-, S_+}^i and $\mathbb{P}_{S_-, \hat{S}_+}^i$; $T_+ = m_\infty$ almost surely under each of them; but $\sum_{t=1}^\infty C_+(t)$ is almost surely less under $\mathbb{P}_{S_-, \hat{S}_+}^i$ than it is under \mathbb{P}_{S_-, S_+}^i , because the value of $C_+(1)$ is lower. Thus (9) shows that $\mathbb{E}_{\hat{S}_-, S_+}^i [P_+] > \mathbb{E}_{S_-, S_+}^i [P_+]$. Again, we have a contradiction to $(S_-, S_+) \in \mathcal{N}$. This completes the proof of Lemma 4.4(2). \square

COROLLARY 4.5. *For $(S_-, S_+) \in \mathcal{N} \cap \mathcal{S}_0^2$ and $i \in \mathbb{Z}$, $\mathbb{P}_{S_-, S_+}^i (E) = 1$.*

Proof. Due to Proposition 4.2 and Lemma 4.4(1,2). \square

Recall that an element $(S_-, S_+) \in \mathcal{S}_0^2$ may be identified as a sequence $\{(b_i, a_i) : i \in \mathbb{Z}\}$ to which Definition 2.1 associates a quadruple $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$.

LEMMA 4.6. *Suppose that $(S_-, S_+) \in \mathcal{N} \cap \mathcal{S}_0^2$. Then $m_i \leq m_{i+1}$ and $n_{i+1} \leq n_i$ for $i \in \mathbb{Z}$.*

Proof. Recall that \mathbb{P}_{S_-, S_+}^i denotes the law of gameplay when $X_0 = i$. Let $\sigma_{i+1} \in \mathbb{N}_+ \cup \{\infty\}$ denote the stopping time $\inf \{\ell \in \mathbb{N}_+ : X_\ell = i + 1\}$. Noting the non-negativity of running costs $C_-(t)$ in Lemma 3.9 with $k = i$ and $Q = \sigma_{i+1}$, we find that

$$\mathbb{E}_{S_-, S_+}^i [P_-] \leq \mathbb{E}_{S_-, S_+}^i [\mathbb{E}_{S_-, S_+}^{X(\sigma_{i+1})} [P_-]],$$

whose left-hand side equals m_i by definition and whose right-hand side takes the form

$$m_{i+1} \mathbb{P}_{S_-, S_+}^i (\sigma_{i+1} < \infty) + m_{-\infty} \mathbb{P}_{S_-, S_+}^i (\sigma_{i+1} = \infty, E) + m_* \mathbb{P}_{S_-, S_+}^i (\sigma_{i+1} = \infty, E^c).$$

However, the third term vanishes in view of Corollary 4.5. Thus, m_i is seen to be a weighted average of $m_{-\infty}$ and m_{i+1} . To conclude, as we seek to do, that $m_i \leq m_{i+1}$, it is thus enough to argue that $m_{-\infty} \leq m_{i+1}$. To obtain this bound, we first *claim* that $\mathbb{E}_{S_-, 0}^{i+1} [P_+] = m_{-\infty}$. To check this, we invoke Lemma 4.4(2) to say that S_- is wide. Thus, E_- , and $T_- = m_{-\infty}$, are $\mathbb{P}_{S_-, 0}^i$ -almost certain. The absence of running costs for Maxine means that $P_+ = T_+$ under $\mathbb{P}_{S_-, 0}^{i+1}$. This yields the claim. Using it, and $(S_-, S_+) \in \mathcal{N}$, we find that

$$m_{i+1} = \mathbb{E}_{S_-, S_+}^{i+1} [P_+] \geq \mathbb{E}_{S_-, 0}^{i+1} [P_+] = m_{-\infty}.$$

We have confirmed that $m_i \leq m_{i+1}$. We omit the similar proof that $n_{i+1} \leq n_i$. This completes the proof of Lemma 4.6. \square

LEMMA 4.7. *Let $\{(b_i, a_i) : i \in \mathbb{Z}\} \in \mathcal{N} \cap \mathcal{S}_0^2$. Recall from Definition 2.1 that m_i equals Maxine's mean receipt when the counter starts at $i \in \mathbb{Z}$. Suppose that $a_i + b_i > 0$. Then*

$$m_i = \frac{a_i}{a_i + b_i} m_{i+1} + \frac{b_i}{a_i + b_i} m_{i-1} - a_i. \quad (32)$$

Proof. Maxine will spend a_i at the first turn; she will win the turn with probability $\frac{a_i}{a_i + b_i}$; if she does so, the counter will reach $i + 1$, and her resulting conditional mean receipt will be m_{i+1} ; if she does not, this receipt will instead be m_{i-1} . Note that the two ratios on the right-hand side of (32) are well defined, because $a_i + b_i > 0$. \square

LEMMA 4.8. *Let $(S_-, S_+) \in \mathcal{N} \cap \mathcal{S}_0^2$, and let $i \in \mathbb{Z}$. Then $a_i > 0$ implies that $m_{i+1} > m_i$. And $b_i > 0$ implies that $n_{i-1} > n_i$.*

Proof. Lemma 4.7 and $a_i > 0$ imply that $m_i < \max\{m_{i-1}, m_{i+1}\}$. But the maximum is attained by m_{i+2} in view of Lemma 4.6. The second assertion in the lemma is similarly obtained. \square

PROPOSITION 4.9. *Let $(S_-, S_+) \in \mathcal{N} \cap \mathcal{S}_0^2$. Then $a_i > 0$, $b_i > 0$, $m_{i+1} > m_i$ and $n_i > n_{i+1}$ for all $i \in \mathbb{Z}$.*

Proof. By Lemma 4.4(2), S_- is wide. To show that every a -coefficient is positive, it is thus enough to argue that $a_i > 0$ implies $a_{i+1} > 0$ for $i \in \mathbb{Z}$, because every index $i \in \mathbb{Z}$ has a positive a -coefficient indexed somewhere to its left. Suppose to the contrary that $a_i > 0$ but $a_{i+1} = 0$. Applying (32) at index $i + 1$, we see that $b_{i+1} > 0$ implies that $m_{i+1} = m_i$. But Lemma 4.8 and $a_i > 0$ imply that $m_{i+1} > m_i$. Thus, $b_{i+1} = 0$. In view of $a_{i+1} = 0$, we see from (32) at index $i + 1$ (with use of the $0/0 = 1/2$ rule) that $m_{i+1} = \frac{m_i + m_{i+2}}{2}$. However: given that $b_{i+1} = 0$, the same equation shows that a sufficiently small positive choice of a_{i+1} would yield a value for m_{i+1} which is arbitrarily close to m_{i+2} , a quantity that exceeds $\frac{m_i + m_{i+2}}{2}$ because (in view of Lemma 4.8, $a_i > 0$ and $a_{i+1} > 0$) we have the bound $m_{i+2} > m_i$. Thus, $(S_-, S_+) \notin \mathcal{N}$, contrary to assumption. We have confirmed that $a_{i+1} > 0$, and thus that every a -coefficient is positive. The argument that $b_i > 0$ for $i \in \mathbb{Z}$ is no different. Lemma 4.8 then shows that each difference $m_{i+1} - m_i$ and $n_i - n_{i+1}$ is positive. This completes the proof of Proposition 4.9. \square

We may now prove the first part of Theorem 2.6.

Proof of Theorem 2.6(1). Suppose that $(S_-, S_+) \in \mathcal{S}_0^2$ is a time-invariant Nash equilibrium for $\text{Trail}(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$. We abusively identify (S_-, S_+) with the sequence $\{(b_i, a_i) : i \in \mathbb{Z}\} \in \mathcal{S}_0^2$ as usual (and, by doing so, we conform notation with the theorem's statement). We note at the outset that, in view of Proposition 4.9, each a_i and b_i , and each difference $m_{i+1} - m_i$ and $n_i - n_{i+1}$, is positive.

Equation ABMN(1) results from rearranging the formula in Lemma 4.7. Equation ABMN(2) is similarly derived.

Next we derive ABMN(3, 4). Recall $S_-(i, j) = b_i$ and $S_+(i, j) = a_i$ for each $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$. For given $i \in \mathbb{Z}$, we will consider a perturbed strategy $\hat{S}_+ \in \mathcal{S}$ for Maxine in which only her first-turn stake is altered, and only then if the counter is at i . In this way, $\hat{S}_+(j, k) = a_j$ for $j \in \mathbb{Z}$ and

$k \geq 2$; and also for $k = 1$ and $j \in \mathbb{Z}$, $j \neq i$. We let $\eta > -a_i$ be small in absolute value, and set $\hat{S}_+(1, i) = a_i + \eta$.

The *original* scenario refers to the law \mathbb{P}_{S_-, S_+}^i , which records counter evolution $X : \mathbb{N} \rightarrow \mathbb{Z}$ given the initial condition $X_0 = i$ under the strategy pair (S_-, S_+) . The *altered* scenario refers to the same law, instead governed by the pair (S_-, \hat{S}_+) . Let O_+ and A_+ denote the mean payoff to Maxine in the original and altered scenarios: that is, $O_+ = \mathbb{E}_{S_-, S_+}^i[P_+]$ and $A_+ = \mathbb{E}_{S_-, \hat{S}_+}^i[P_+]$. Then

$$O_+ = \frac{a_i}{a_i + b_i} m_{i+1} + \frac{b_i}{a_i + b_i} m_{i-1} - a_i \text{ and } A_+ = \frac{a_i + \eta}{a_i + \eta + b_i} m_{i+1} + \frac{b_i}{a_i + \eta + b_i} m_{i-1} - a_i - \eta,$$

so that

$$A_+ - O_+ = \left(\frac{b_i}{(a_i + b_i)^2} (m_{i+1} - m_{i-1}) - 1 \right) \cdot \eta \cdot (1 + o(1)), \quad (33)$$

where the $o(1)$ term is small in the sense of $|\eta| \rightarrow 0$. Since $(S_-, S_+) \in \mathcal{N}$, A_+ is at most O_+ , whatever the value of $\eta > -a_i$. The derivative in η of $A_+ - O_+$ thus vanishes at zero, so that $\frac{b_i}{(a_i + b_i)^2} (m_{i+1} - m_{i-1}) - 1 = 0$ or equivalently

$$b_i (m_{i+1} - m_{i-1}) = (a_i + b_i)^2. \quad (34)$$

We now consider the same original scenario alongside a new altered scenario in which it is Mina who adopts a perturbed strategy \hat{S}_- (as a function of a given choice of $i \in \mathbb{Z}$). Analogously to what we have done, we choose $\eta > -b_i$, and set $\hat{S}_-(j, k) = b_j$ for $j \in \mathbb{Z}$ and $k \geq 2$ or when $k = 1$ and $j \in \mathbb{Z}$, $j \neq i$; and then we set $\hat{S}_-(1, i) = b_i + \eta$. We denote by O_- and A_- Mina's mean payoff in the original and in the newly altered scenarios; to wit, $O_- = \mathbb{E}_{S_-, S_+}^i[P_-]$ and $A_- = \mathbb{E}_{\hat{S}_-, S_+}^i[P_-]$. We find then that

$$O_- = \frac{b_i}{a_i + b_i} n_{i+1} + \frac{a_i}{b_i + a_i} n_{i-1} - b_i \text{ and } A_- = \frac{b_i + \eta}{b_i + \eta + a_i} n_{i+1} + \frac{b_i}{a_i + \eta + b_i} n_{i-1} - b_i - \eta;$$

and, analogously to (33),

$$A_- - O_- = \left(\frac{a_i}{(a_i + b_i)^2} (n_{i-1} - n_{i+1}) - 1 \right) \cdot \eta \cdot (1 + o(1)).$$

The condition that $(S_-, S_+) \in \mathcal{N}$ ensures that $O_- \geq A_-$, whatever the value of $\eta > -b_i$. Thus,

$$a_i (n_{i-1} - n_{i+1}) = (a_i + b_i)^2. \quad (35)$$

The derived equations (34) and (35) are ABMN(3, 4) with index i .

We have established that $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ solves the ABMN equations. To complete the proof of Theorem 2.6(1), it remains to confirm that the boundary values (11) are achieved. We will argue that $\lim_{i \rightarrow \infty} m_{-i} = m_{-\infty}$; the three other limits are similarly shown. The sequence $\{m_{-i} : i \in \mathbb{N}\}$ decreases by Proposition 4.9 to a limiting value that we call $\mathbf{m}_{-\infty}$. Since $m_i = \mathbb{P}_{S_-, S_+}^i[P_+] \geq \mathbb{P}_{S_-, 0}^i[P_+] = m_{-\infty}$, we have that $\mathbf{m}_{-\infty} \geq m_{-\infty}$; we wish to obtain the opposite inequality. By removing non-negative running costs from the right-side of the expression for m_i in Lemma 3.10(2), we see that $m_i \leq \mathbb{P}_{S_-, S_+}^i(E_-) \cdot m_{-\infty} + \mathbb{P}_{S_-, S_+}^i(E_+) \cdot m_{\infty}$ where we invoked Corollary 4.5. Thus $\mathbf{m}_{-\infty} \leq m_{-\infty}$ provided that we argue that $\lim_{i \rightarrow -\infty} \mathbb{P}_{S_-, S_+}^i(E_+) = 0$: far to the left is the domain of Mina's likely victory. It would be of interest to argue directly; and to do so would be more in keeping with the style of this section. It is quicker however to simply invoke the eventual gameplay unanimity Theorem 2.15(3), which will be proved by independent arguments when we find an explicit solution of the ABMN system in Chapter 5. (Theorem 2.15(2) is invoked in the corresponding place in two of the three omitted limit derivations.) We have thus obtained Theorem 2.6(1). \square

4.2. The reverse implication

Here we prove Theorem 2.6(2). It is here that the infinite-turn nature of the game has to be tamed by comparison with finite-trail counterparts.

We begin by developing definitions and results that will lead to the proof of the desired result at the end of the section. As such, we now enforce the notation in the hypothesis of Theorem 2.6(2), so that, from now on,

$\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ denotes a positive solution of the ABMN equations

with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$ that satisfies (4). Let $S_-, S_+ \in \mathcal{S}$ satisfy

$$S_-(i, j) = b_i \quad \text{and} \quad S_+(i, j) = a_i \quad \text{for } (i, j) \in \mathbb{Z} \times \mathbb{N}_+. \quad (36)$$

Recall that $S_+ = a$ with the usual notational abuse.

DEFINITION 4.10. Let $i \in \mathbb{Z}$. The forward play-cone F_i of i is given by

$$F_i = \left\{ (k, \ell) \in \mathbb{Z} \times \mathbb{N}_+ : |k - i| \leq \ell, |k - i| + \ell \in 2\mathbb{N} \right\}.$$

This is the set of space-time sites that are in principle accessible for gameplay $X : \mathbb{N} \rightarrow \mathbb{Z}$ under \mathbb{P}_{S_1, S_2}^i for some strategy pair $(S_1, S_2) \in \mathcal{S}^2$.

Let $S \in \mathcal{S}$. An element $(q, \ell) \in F_i$ such that $S(q, \ell + 1) \neq b_q$ is called a *Mina deviation point*. The Mina deviation set $D_-(S, i) \subseteq F_i$ is the collection of Mina deviation points. The strategy S is called *deviating for Mina* if $D_-(S, i)$ is non-empty. A *Maxine deviation point* $(q, \ell) \in F_i$ satisfies $S(q, \ell + 1) \neq a_q$. The set $D_+(S, i)$ of such points is the Maxine deviation set; if $D_+(S, i) \neq \emptyset$, then S is deviating for Maxine.

When gameplay under \mathbb{P}_{S_-, S_+}^i runs through a Mina deviation point—when $X_\ell = q$ for $(q, \ell) \in D_-(S, i)$ —her stake according to strategy S —namely, $S(q, \ell + 1)$ —may be viewed as a mistake when her opponent plays her element S_+ of the putative Nash equilibrium (S_-, S_+) . The next result, which is fundamental to proving Theorem 2.6(2), validates this notion. It measures the magnitude of the mistakes that result from a player's deviation in the sense of decrease in mean payoff in finite trail games. It finds the mistakes to be uniformly costly as the finite trails vary.

PROPOSITION 4.11. *Let $i \in \mathbb{Z}$ be given.*

(1) *Let $S_-^{dev} \in \mathcal{S}$ be deviating for Mina. Suppose that $\mathbb{P}_{S_-^{dev}, S_+}^i(E) = 1$. Then*

$$\sup \mathbb{E}_{S_-^{dev}, S_+}^i [P_-^{j,k}] < \mathbb{E}_{S_-, S_+}^i [P_-],$$

where the supremum is taken over $j, k \in \mathbb{N}_+$ such that $i \in \llbracket -j, k \rrbracket$ and for which there exists an element (u, ℓ) of $D_-(S_-^{dev}, i)$ with $u \in \llbracket -j, k \rrbracket$.

(2) *Now suppose that $S_+^{dev} \in \mathcal{S}$ is deviating for Maxine, and $\mathbb{P}_{S_-, S_+^{dev}}^i(E) = 1$. Then*

$$\sup \mathbb{E}_{S_-, S_+^{dev}}^i [P_+^{j,k}] < \mathbb{E}_{S_-, S_+}^i [P_+],$$

where now the supremum is taken over $j, k \in \mathbb{N}_+$ with $i \in \llbracket -j, k \rrbracket$ and for which there exists $(u, \ell) \in D_+(S_+^{dev}, i)$ such that $u \in \llbracket -j, k \rrbracket$.

It is a short step from the just stated result to the next conclusion, which asserts that a player's deviation will cost her in the infinite trail game. This is in essence what it means for (S_-, S_+) to be a Nash equilibrium. Indeed, we next close out the proof of Theorem 2.6 by first deriving Proposition 4.12 from Proposition 4.11; and second showing how the latter result leads to the desired conclusion. These tasks done, we will turn to the remaining and more substantial one: to prove Proposition 4.11.

PROPOSITION 4.12. *Let $i \in \mathbb{Z}$.*

(1) *Let $S_-^{dev} \in \mathcal{S}$ be deviating for Mina. Then*

$$\mathbb{E}_{S_-^{dev}, S_+}^i [P_-] < \mathbb{E}_{S_-, S_+}^i [P_-].$$

(2) *Now let $S_+^{dev} \in \mathcal{S}$ be deviating for Maxine. Then*

$$\mathbb{E}_{S_-, S_+^{dev}}^i [P_+] < \mathbb{E}_{S_-, S_+}^i [P_+].$$

Proof: (1). Suppose first that $\mathbb{P}_{S_-^{dev}, S_+}^i (E^c) > 0$. Lemma 3.6 implies that $\mathbb{E}_{(S_-^{dev}, S_+)}^i [P_-] = -\infty$. But $\mathbb{E}_{S_-, S_+}^i [P_-] = m_i$ by Lemma 3.10(1). We have that $m_i \geq m_{-\infty}$ since the sequence $\{m_i : i \in \mathbb{Z}\}$ increases for any positive ABMN solution by Theorem 2.3(1). And we know that $m_{-\infty} > -\infty$ by hypothesis. Thus we see that $\mathbb{E}_{S_-, S_+}^i [P_-] > -\infty$, so that Proposition 4.12(1) has been established in this case.

Now we suppose instead that $\mathbb{P}_{S_-^{dev}, S_+}^i (E) = 1$. Let $\eta > 0$ be arbitrary. Note that $T_-^{j,k} = n(X(\tau^{j,k}))$ for $j, k \in \mathbb{N}_+$; and that T_- equals $n_{-\infty}$ on E_- , and n_∞ on E_+ . By Lemma 3.12 and $\mathbb{P}_{S_-^{dev}, S_+}^i (E) = 1$, we may thus find $j_0, k_0 \in \mathbb{N}_+$ such that, when $j \geq j_0$ and $k \geq k_0$,

$$\mathbb{P}_{S_-^{dev}, S_+}^i (|T_- - T_-^{j,k}| \geq \eta) \leq \eta.$$

By Lemma 3.8(2), we see that

$$\mathbb{P}_{S_-^{dev}, S_+}^i (P_- \leq P_-^{j,k} + \eta) \geq 1 - \eta.$$

By Lemma 3.8(3),

$$\mathbb{E}_{S_-^{dev}, S_+}^i [P_-] \leq \mathbb{E}_{S_-^{dev}, S_+}^i [P_-^{j,k}] + (1 + n_{-\infty} - n_\infty)\eta.$$

By taking $\eta > 0$ to be one-half of the difference of the two sides in the conclusion of Proposition 4.11(1), the latter result is seen to imply Proposition 4.12(1).

(2). We omit this similar argument. □

Proof of Theorem 2.6(2). Recall that $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ is a positive ABMN solution with boundary data $(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty) \in \mathbb{R}^4$. Further recall that $(S_-, S_+) = (b, a)$, with the usual notational abuse. Let $S \in \mathcal{S}$. If S is not deviating for Mina, then $\mathbb{E}_{S, S_+}^i [P_-] = \mathbb{E}_{S_-, S_+}^i [P_-]$ since the laws \mathbb{P}_{S, S_+}^i and \mathbb{P}_{S_-, S_+}^i are equal. Otherwise, $\mathbb{E}_{S, S_+}^i [P_-] < \mathbb{E}_{S_-, S_+}^i [P_-]$ by Proposition 4.12(1). (We recall that implicit in the notation \mathbb{P}_{S_1, S_2}^i and $\mathbb{E}_{S_1, S_2}^i [\cdot]$ are the values $(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$, because these values appear as terminal receipts.)

By Proposition 4.12(2), it follows similarly that $\mathbb{E}_{S_-, S}^i [P_+] < \mathbb{E}_{S_-, S_+}^i [P_+]$ if S is deviating for Maxine. Further, $\mathbb{E}_{S_-, S}^i [P_+] = \mathbb{E}_{S_-, S_+}^i [P_+]$ if Maxine's S is not deviating. We have confirmed that $(S_-, S_+) \in \mathcal{N}(m_{-\infty}, m_\infty, n_{-\infty}, n_\infty)$ and thus obtain Theorem 2.6(2). □

We now prepare to prove Proposition 4.11(1). (The proof of Proposition 4.11(2) is essentially the same.) Henceforth, Proposition 4.11(1)'s hypotheses are understood to be in force: S_- and S_+ are the non-deviating strategies given by (36); $i \in \mathbb{Z}$ is given; and $S_-^{\text{dev}} \in \mathcal{S}$ is deviating for Mina, with $\mathbb{P}_{S_-^{\text{dev}}, S_+}^i(E) = 1$.

Let $j, k \in \mathbb{N}$ satisfy $i \in \llbracket -j, k \rrbracket$. Developing Definition 4.10, we set

$$D_-^{j,k}(S, i) = \left\{ (q, \ell) \in D_-(S, i) : q \in \llbracket -j, k \rrbracket \right\}$$

for $S \in \mathcal{S}$. It may be that $D_-^{j,k}(S_-^{\text{dev}}, i)$ is infinite. It serves our purpose to approximate S_-^{dev} by strategies for which the counterpart set is finite. We now specify these strategies. Enumerate $D_-^{j,k}(S_-^{\text{dev}}, i)$ in increasing order of the vertical component, using an arbitrary rule to break the ties that arise when elements share the same height. For $v \in \mathbb{N}_+$, let $D_{-,v}^{j,k}(S_-^{\text{dev}}, i)$ denote the set whose elements are the first v members of $D_-^{j,k}(S_-^{\text{dev}}, i)$. Let $S_-^{\text{dev}}[v]$ denote the strategy that equals S_-^{dev} on $D_{-,v}^{j,k}(S_-^{\text{dev}}, i)$ and S_- otherwise; note that $D_-^{j,k}(S_-^{\text{dev}}[v], i)$ equals $D_{-,v}^{j,k}(S_-^{\text{dev}}, i)$. We make another basic comparison in terms of the next definition.

DEFINITION 4.13. For $S \in \mathcal{S}$, let $\text{ground}^{j,k}(S, i) \in \mathbb{N}$ denote the minimum vertical coordinate assumed by an element of $D_-^{j,k}(S, i)$.

Note then that $\text{ground}^{j,k}(S_-^{\text{dev}}[v], i)$ is independent of $v \in \mathbb{N}_+$.

We wish to argue that Mina's deviant play under the strategies $S_-^{\text{dev}}[v]$, $v \in \mathbb{N}_+$, and S_-^{dev} , is suitably penalized in the trail game on $\llbracket -j-1, k+1 \rrbracket$. In the notation of the next definition, Lemma 4.15 establishes such a conclusion for the finitely deviating strategies $S_-^{\text{dev}}[v]$: there is a penalty incurred by use of these strategies; and, in a suitable sense, the penalty is uniform among them, and is governed by the limiting strategy S_-^{dev} . After we prove Lemma 4.15, it will remain to address the penalty suffered by using S_-^{dev} itself. Definition 4.14 speaks of a 'strong' penalty as a contrast with a modified definition that will be used to treat the perhaps infinitely deviating S_-^{dev} , this appearing after the proof of Lemma 4.15.

DEFINITION 4.14. Let $S_1, S_2 \in \mathcal{S}$. Consider the following conditions:

- (1) We have that $\mathbb{E}_{S_1, S_+}^{u, \ell} [P_-^{j,k}] \leq n_u$ for all $\ell \in \mathbb{N}_+$ and $u \in \llbracket -j, k \rrbracket$.
- (2) Writing $g = \text{ground}^{j,k}(S_1, i)$, consider any $u \in \llbracket -j, k \rrbracket$ for which $(u, g) \in D_-^{j,k}(S_1, i)$. Then the value $n_u - \mathbb{E}_{S_1, S_+}^{u, g} [P_-^{j,k}]$ is positive, and indeed is bounded below by a positive quantity that is determined solely by $S_2(u, g+1)$.

If these conditions are met, we say that S_1 receives the strong (i, j, k) -penalty merited by S_2 . Let $S \in \mathcal{S}$. If S receives the strong (i, j, k) -penalty merited by S , we say that S justly receives a strong (i, j, k) -penalty.

(Although it is omitted from the notation of the strong (i, j, k) -penalty, it is the strategy $S_+ = a$ that Mina is facing when she plays S_1 or S_2 . The above definition and the next result are intended to capture the sense of Mina's mistake when she declines to stake at the b -level dictated by S_- against Maxine's a -stake offered by S_+ .)

LEMMA 4.15. *Let $j, k \in \mathbb{N}$ satisfy $i \in \llbracket -j, k \rrbracket$. Let v be at least the number of elements of $\mathcal{D}_-^{j,k}(S_-^{dev}, i)$ of minimum height. Then $S_-^{dev}[v]$ receives the strong (i, j, k) -penalty merited by S_-^{dev} .*

(The value of g implicit in Lemma 4.15 does not depend on the value of $v \in \mathbb{N}_+$ used in $S_-^{dev}[v]$, because $\text{ground}^{j,k}(S_-^{dev}[v], i)$ is independent of $v \in \mathbb{N}_+$.)

The finite-error strategies $S_-^{dev}[v]$ have been introduced because they may be analysed using the fundamental game-theoretic technique of backwards induction. When Mina uses $S_-^{dev}[v]$ for any given $v \in \mathbb{N}_+$, she never deviates at late enough time. Lemma 3.10(2) then serves to show that she incurs no penalty by doing so. As turn index retreats in backwards induction, Mina will make deviating moves. At the heart of the analysis of the inductive step is the consideration of one turn when Mina deviates. What is being played here is a game of Penny Forfeit, treated in Section 3.2. The next result gathers what we need to know about one step in the game.

LEMMA 4.16.

- (1) *Let $j, k \in \mathbb{Z}$ satisfy $i \in \llbracket -j, k \rrbracket$. For $\ell \in \mathbb{N}_+$, let $S_1, S_2 \in \mathcal{S}$ be such that, if $(u, h) \in \mathbb{Z} \times \mathbb{N}_+$ satisfies $S_1(u, h) \neq S_2(u, h)$, then $h \leq \ell$. Then $\mathbb{E}_{S_1, S_+}^{u, h}[P_-^{j, k}] = \mathbb{E}_{S_2, S_+}^{u, h}[P_-^{j, k}]$ for any $(u, h) \in \llbracket -j, k \rrbracket \times \llbracket \ell, \infty \rrbracket$.*

Let $S \in \mathcal{S}$ and $(u, \ell) \in \llbracket -j, k \rrbracket \times \mathbb{N}$. Suppose that $\mathbb{E}_{S, S_+}^{v, \ell+1}[P_-^{j, k}] \leq n_j$ for $v \in \{u-1, u+1\}$.

- (2) *We have that $\mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}] \leq n_u$.*
- (3) *Suppose further that $(u, \ell) \in \mathcal{D}_-^{j, k}(S, i)$. Then $n_u - \mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}]$ is bounded below by a positive quantity that is determined solely by the value of $S(u, \ell+1) \neq b_u$.*

Proof: (1). The laws $\mathbb{P}_{S_1, S_+}^{u, h}$ and $\mathbb{P}_{S_2, S_+}^{u, h}$ are identical because S_1 and S_2 coincide at any point (u, ℓ) with $\ell \geq h+1$.

(2). Note that

$$\mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}] = \frac{S(u, \ell)}{a_i + S(u, \ell)} \mathbb{E}_{S, S_+}^{u-1, \ell+1}[P_-^{j, k}] + \frac{a_i}{a_i + S(u, \ell)} \mathbb{E}_{S, S_+}^{u+1, \ell+1}[P_-^{j, k}] - S(u, \ell).$$

Since $\mathbb{E}_{S, S_+}^{u-1, \ell+1}[P_-^{j, k}] \leq n_{u-1}$ and $\mathbb{E}_{S, S_+}^{u+1, \ell+1}[P_-^{j, k}] \leq n_{u+1}$, we see that

$$\mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}] \leq \frac{S(u, \ell)}{a_u + S(u, \ell)} n_{u-1} + \frac{a_u}{a_u + S(u, \ell)} n_{u+1} - S(u, \ell).$$

By Lemma 3.5, this right-hand side has a unique maximum in b at $b = b_u$, when it assumes the value n_u .

(3). Since $S(u, \ell+1)$ is not equal to b_u , we see that the above right-hand side, and thus $\mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}]$, is less than n_u . The difference $n_u - \mathbb{E}_{S, S_+}^{u, \ell}[P_-^{j, k}]$ is determined solely by $S(u, \ell+1)$. \square

The next result leads quickly to Lemma 4.15. Indeed, its proof (in a perhaps slightly disguised form) is the backwards inductive argument that underlies Lemma 4.15.

LEMMA 4.17. *Suppose that $S \in \mathcal{S}$ is such that $\mathcal{D}_-^{j, k}(S, i)$ is finite. Then S justly receives a strong (i, j, k) -penalty.*

Proof. We will induct on the cardinality of $\mathcal{D}_-^{j, k}(S, i)$.

Let $S \in \mathcal{S}$. Set $g = \text{ground}^{j,k}(S, i)$. For $\ell \in \mathbb{N}$, $\ell \neq g$, let $\text{IH}(S, \ell)$ denote the assertion that

$$\mathbb{E}_{S, S_+}^{u, \ell} [P_-^{j,k}] \leq n_u \text{ for } u \in \llbracket -j, k \rrbracket. \quad (37)$$

For $\ell = g$, let $\text{IH}(S, \ell)$ denote the assertion that the preceding display holds and so does the following.

Consider any $u \in \llbracket -j, k \rrbracket$ for which $(u, g) \in \text{D}_-^{j,k}(S, i)$.

Then the value $n_u - \mathbb{E}_{S, S_+}^{u, g} [P_-^{j,k}]$ is positive,

and indeed is bounded below by a positive quantity that is determined solely by $S(u, g + 1)$.

We take the inductive hypothesis indexed by $q \in \mathbb{N}_+$ to be the assertion that the statements $\text{IH}(S, \ell)$, $\ell \in \mathbb{N}$, are true for each $S \in \mathcal{S}$ such that $\#\text{D}_-^{j,k}(S, i) \leq q$. The base case will be $q = 0$. This is the assertion that (37) holds for $\ell \in \mathbb{N}_+$, when $S \in \mathcal{S}$ is such that $\text{D}_-^{j,k}(S, i)$ is empty.

The base case holds by Lemma 3.10(3).

Let $q \in \mathbb{N}$ and assume the inductive hypothesis indexed by q . Let $S \in \mathcal{S}_0$ be such that $\#\text{D}_-^{j,k}(i) = q + 1$. Again set $g = \text{ground}^{j,k}(S, i)$. Let $\hat{S} \in \mathcal{S}_0$ be given by

$$\hat{S}(i, \ell) = \begin{cases} S_-(i, g + 1) = b_i & \text{if } \ell = g + 1 \\ S(i, \ell) & \text{if } \ell \in \mathbb{N}_+, \ell \neq g + 1. \end{cases}$$

for $i \in \mathbb{Z}$. The set $\text{D}_-^{j,k}(\hat{S}, i)$ is formed from $\text{D}_-^{j,k}(S, i)$ by the removal of the elements of minimum height—which is height g . Hence, $\#\text{D}_-^{j,k}(\hat{S}, i) < \#\text{D}_-^{j,k}(S, i)$; the hypotheses $\text{IH}(\hat{S}, \ell)$, $\ell \in \mathbb{N}_+$, are thus available. By Lemma 4.16(1) with $S_1 = \hat{S}$, $S_2 = S$ and $\ell = g + 1$, we find that $\text{IH}(S, \ell)$ holds for $\ell \geq g + 1$.

Now consider $u \in \llbracket -j, k \rrbracket$ such that $(u, g) \in \text{D}_-^{j,k}(S, i)$. Lemma 4.16(3) (and Lemma 4.16(2) for other $u \in \llbracket -j, k \rrbracket$) implies $\text{IH}(S, g)$.

To complete the inductive step, it remains to verify $\text{IH}(S, \ell)$ for $\ell \in \llbracket 0, g - 1 \rrbracket$. We do so iteratively in decreasing ℓ . It is Lemma 4.16(2) that demonstrates the generic step in this iteration. This completes the proof of Lemma 4.17. \square

Proof of Lemma 4.15. Apply Lemma 4.17 for $S = S_-^{\text{dev}}[v]$ for $v \in \mathbb{N}_+$. We learn that Definition 4.14 holds with $S_1 = S_2 = S_-^{\text{dev}}[v]$. Thus, the positive quantity in Definition 4.14(2) is determined by $S_-^{\text{dev}}[v](u, g + 1)$. When v satisfies the bound in Lemma 4.15, we have that $S_-^{\text{dev}}[v](u, g + 1)$ equals $S_-^{\text{dev}}(u, g + 1)$. As a result, Definition 4.14 holds with $S_1 = S_-^{\text{dev}}[v]$ and $S_2 = S_-^{\text{dev}}$. This is what Lemma 4.15 asserts. \square

Lemma 4.15 is a stepping stone to a counterpart that describes the penalty incurred by use of the perhaps infinitely deviating strategy $S_-^{\text{dev}} \in \mathcal{S}$. The counterpart, Lemma 4.19, depends on a variation of Definition 4.14.

DEFINITION 4.18. Let $S \in \mathcal{S}$. An element $(q, \ell) \in F_i$ is said to be (S, S_+) -accessible from $(i, 0)$ if $\mathbb{P}_{S, S_+}^i(X_\ell = q) > 0$. Let $\text{A}(S, i)$ denote the set of elements of F_i that are (S, S_+) -accessible from $(i, 0)$.

Alter Definition 4.14 by taking S_1 and S_2 equal to S ; the first part to include the condition that the point (u, ℓ) belongs to $\text{A}(S, i)$; and the second to include the condition that $(u, g) \in \text{A}(S, i)$. Thus, no requirement is imposed by a given part when (u, ℓ) or (u, g) is not (S, S_+) -accessible from $(i, 0)$. When the altered set of conditions is satisfied, we say that S *justly receives a weak (i, j, k) -penalty*.

LEMMA 4.19. *Let $j, k \in \mathbb{N}$ satisfy $i \in \llbracket -j, k \rrbracket$. The strategy S_-^{dev} justly receives a weak (i, j, k) -penalty.*

To prove this result, we intend to make use of Proposition 4.11(1)'s hypothesis that $\mathbb{P}_{S_-^{\text{dev}}, S_+}^i(E) = 1$. Since escape is certain under $\mathbb{P}_{S_-^{\text{dev}}, S_+}^i$, gameplay will exit $\llbracket -j, k \rrbracket$ in finite time, so that Mina's choice between S and $S[v]$, for high v , will typically leave gameplay unaffected. Thus we aim to reduce the proof of the new result to quoting Lemma 4.15. To do this, it is useful to state a consequence of $\mathbb{P}_{S_-^{\text{dev}}, S_+}^i(E) = 1$.

LEMMA 4.20. *Let $(u, \ell) \in \mathbf{A}(S_-^{\text{dev}}, i)$. Then $\mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}(E) = 1$.*

Proof. We have that

$$1 = \mathbb{P}_{S_-^{\text{dev}}, S_+}^i(E) = \sum_{u \in \mathbb{Z}} \mathbb{P}_{S_-^{\text{dev}}, S_+}^i(X_\ell = u) \cdot \mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}(E).$$

Since $\mathbb{P}_{S_-^{\text{dev}}, S_+}^i(X_\ell = u) > 0$ if and only if $(u, \ell) \in \mathbf{A}(S_-^{\text{dev}}, i)$, we see that $\mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}(E)$ equals one when this condition is satisfied. \square

Proof of Lemma 4.19. Let $h(v)$ be the vertical coordinate of the v^{th} element of $\mathbf{D}_-^{j, k}(S_-^{\text{dev}}, i)$. Let $(u, \ell) \in \llbracket -j, k \rrbracket \times \mathbb{N}$. Under $\mathbb{P}_{S_-^{\text{dev}}[v], S_+}^{u, \ell}$ given $\tau^{j, k} \geq h(v)$, Mina does not deviate after time $\tau^{j, k}$. By Lemma 3.10(2) and Theorem 2.3(1), the conditional mean of $P_-^{j, k}$ under $\mathbb{P}_{S_-^{\text{dev}}[v], S_+}^{u, \ell}$ given that $\tau^{j, k} > h(v)$ is thus seen to be at least n_{k+1} . Now consider (25) with $\pm = -1$ and $(P, S_-) \rightarrow (P^{j, k}, S_-^{\text{dev}})$; note that running costs here are non-negative, and that terminal receipt is at most n_{-j-1} by Theorem 2.3(1). We see then that the conditional mean of $P_-^{j, k}$ under $\mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}$ given that $\tau^{j, k} > h(v)$ is at most n_{-j-1} .

We find then that

$$\mathbb{E}_{S_-^{\text{dev}}, S_+}^{u, \ell}[P_-^{j, k}] - \mathbb{E}_{S_-^{\text{dev}}[v], S_+}^{u, \ell}[P_-^{j, k}] \leq \mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}(\tau^{j, k} \geq h(v)) \cdot (n_{-j-1} - n_{k+1}). \quad (38)$$

Lemma 4.19 will follow from Lemma 4.15 provided that we show that the right-hand side of this display vanishes in high v whenever $(u, \ell) \in \mathbf{A}(S_-^{\text{dev}}, i)$. By Lemma 4.20, and the hypothesis of Proposition 4.11(1), we know that $\mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}(E) = 1$. Thus, $\tau^{j, k}$ is finite, $\mathbb{P}_{S_-^{\text{dev}}, S_+}^{u, \ell}$ -almost surely. The right-hand side of (38) thus indeed tends to zero in the limit of high v . Lemma 4.15 implies Lemma 4.19, as we sought to show. \square

We are ready for the following proof.

Proof of Proposition 4.11(1). For $j, k \in \mathbb{N}_+$ such that $i \in \llbracket -j, k \rrbracket$, let g denote the minimum vertical coordinate among elements of $\mathbf{D}_-^{j, k}(S_-^{\text{dev}}, i)$. Any element $(u, g) \in F_i$ belongs to $\mathbf{A}(S_-^{\text{dev}}, i)$ because, under (S_-^{dev}, S_+) , gameplay is governed before the g^{th} turn by the positive-element pair (S_-, S_+) . Lemma 4.19 thus implies that, when $(u, g) \in F_i$,

$$\mathbb{E}_{S_-^{\text{dev}}, S_+}^{u, g}[P_-^{j, k}] \leq n_u \quad (39)$$

and

$$\mathbb{E}_{S_-^{\text{dev}}, S_+}^{u, g}[P_-^{j, k}] < n_u \text{ if } (u, g) \in \mathbf{D}_-^{j, k}(S_-^{\text{dev}}, i). \quad (40)$$

Now note that

$$\mathbb{E}_{S_{\text{dev}}, S_+}^i [P_-^{j,k}] = -\mathbb{E}_{S_{\text{dev}}, S_+}^i \sum_{t=1}^{g-1} C_-(t) \mathbf{1}_{t < \tau^{j,k}} + \sum_{\substack{u \in \llbracket -j, k \rrbracket: \\ (u,g) \in F_i}} \mathbb{P}_{S_{\text{dev}}, S_+}^i (X_g^{j,k} = u) \cdot \mathbb{E}_{S_{\text{dev}}, S_+}^{u,g} [P_-^{j,k}].$$

The joint law of $C_-(t)$, $t \in \llbracket 1, g-1 \rrbracket$, is equal under $\mathbb{P}_{S_{\text{dev}}, S_+}^i$ and \mathbb{P}_{S_-, S_+}^i , because S and S_- coincide on $\mathbb{Z} \times \llbracket 1, g-1 \rrbracket$. The costs $C_-(t)$ are non-negative, and upper bounds on the conditional mean payoffs in the preceding display are offered by (39) and (40).

By way of comparison,

$$\mathbb{E}_{S_-, S_+}^i [P_-^{j,k}] = -\mathbb{E}_{S_-, S_+}^i \sum_{t=1}^{g-1} C_-(t) \mathbf{1}_{t < \tau^{j,k}} + \sum_{\substack{u \in \llbracket -j, k \rrbracket: \\ (u,g) \in F_i}} \mathbb{P}_{S_-, S_+}^i (X_g^{j,k} = u) \cdot \mathbb{E}_{S_-, S_+}^{u,g} [P_-^{j,k}],$$

with

$$\mathbb{E}_{S_-, S_+}^{u,g} [P_-^{j,k}] = n_u \text{ for } u \in \llbracket -j, k \rrbracket$$

by Lemma 3.10(3). Consider a pair (j, k) over which the supremum in Proposition 4.11(1) is taken. Since $D_-^{j,k}(S_{\text{dev}}, i)$ is non-empty, we may find $q \in \llbracket -j, k \rrbracket$ such that $(q, g) \in D_-^{j,k}(S_{\text{dev}}, i)$. Since $X : \llbracket 0, g \rrbracket \rightarrow \mathbb{Z}$ coincides under \mathbb{P}_{S_-, S_+}^i and $\mathbb{P}_{S_{\text{dev}}, S_+}^i$, we see that

$$\mathbb{E}_{S_-, S_+}^i [P_-^{j,k}] - \mathbb{E}_{S_{\text{dev}}, S_+}^i [P_-^{j,k}] = \sum_{\substack{u \in \llbracket -j, k \rrbracket: \\ (u,g) \in F_i}} \mathbb{P}_{S_-, S_+}^i (X_g^{j,k} = u) \cdot \left(n_u - \mathbb{E}_{S_{\text{dev}}, S_+}^{u,g} [P_-^{j,k}] \right),$$

where the term in parentheses on the right-hand side is strictly positive if $u = q$ (by (40)), and is non-negative if $u \in \llbracket -j, k \rrbracket$, $u \neq q$ (by (39)). This implies that

$$\mathbb{E}_{S_-, S_+}^i [P_-^{j,k}] - \mathbb{E}_{S_{\text{dev}}, S_+}^i [P_-^{j,k}] \geq \mathbb{P}_{S_-, S_+}^i (X_g^{j,k} = q) \cdot \left(n_q - \mathbb{E}_{S_{\text{dev}}, S_+}^{q,g} [P_-^{j,k}] \right).$$

We claim that $\mathbb{P}_{S_-, S_+}^i (X_g^{j,k} = q) > 0$. Indeed, it is enough to find any access route for X from $(i, 0)$ to (q, g) that never leaves $\llbracket -j, k \rrbracket \times \llbracket 0, g \rrbracket$, because the strategies in the pair $(S_-, S_+) = (b, a)$ have positive coefficients; that such a route exists is due to $(q, g) \in F_i$, $i, q \in \llbracket -j, k \rrbracket$ and $k > -j$. For example, $\mathbb{P}_{S_-, S_+}^i (X_g^{j,k} = q)$ is at least η^g , where $\eta = \min \{a_i \wedge b_i : i \in \llbracket -j, k \rrbracket\}$. That is,

$$\mathbb{E}_{S_-, S_+}^i [P_-^{j,k}] - \mathbb{E}_{S_{\text{dev}}, S_+}^i [P_-^{j,k}] \geq \eta^g \cdot \left(n_q - \mathbb{E}_{S_{\text{dev}}, S_+}^{q,g} [P_-^{j,k}] \right).$$

Since the positive right-hand side is independent of the choice of the pair $j, k \in \mathbb{N}_+$ over which the supremum is taken in Proposition 4.11(1), we have obtained this result. \square

Proof of Proposition 4.11(2). The essentially identical argument is omitted. \square

Explicit ABMN solutions and their consequences

Here we explicitly solve the ABMN system, proving Theorem 2.21, and its softer cousin Proposition 2.16. Then we analyse the asymptotic decay in high index values of ABMN solutions, proving Theorem 2.14. Two consequences of this decay—finiteness of boundary data in Theorem 2.3(3), and the almost sure eventual unanimity of gameplay in Theorem 2.15—are derived.

5.1. Explicit ABMN solutions

Fundamental to deriving Theorem 2.21 is an alternative representation of the ABMN system that we offer first, in Proposition 5.1. The real-valued variables $\{m_i, n_i : i \in \mathbb{Z}\}$ satisfy the MN system on \mathbb{Z} if

$$(m_i - m_{i-1})(m_{i+1} - m_{i-1} + n_{i-1} - n_{i+1})^2 = (m_{i+1} - m_{i-1})^3 \quad \text{MN(1)}$$

$$(n_i - n_{i+1})(m_{i+1} - m_{i-1} + n_{i-1} - n_{i+1})^2 = (n_{i-1} - n_{i+1})^3 \quad \text{MN(2)},$$

for $i \in \mathbb{Z}$. As for ABMN(1, 2, 3, 4) from Definition 2.2, we refer to the above equations as MN(1) and MN(2) rather than by the usual convention of numbered equations.

PROPOSITION 5.1. *A positive solution of the ABMN system on \mathbb{Z} solves the MN system on \mathbb{Z} .*

Proof. For $i \in \mathbb{Z}$, set $M_i = m_{i+1} - m_{i-1}$ and $N_i = n_{i-1} - n_{i+1}$. We claim that

$$a_i = \frac{M_i^2 N_i}{(M_i + N_i)^2}, \quad b_i = \frac{M_i N_i^2}{(M_i + N_i)^2} \quad \text{and} \quad \frac{a_i}{a_i + b_i} = \frac{M_i}{M_i + N_i}. \quad (41)$$

These follow from ABMN(3, 4). Expressing ABMN(1) in the form (32), we find from (41) that

$$m_i = m_{i-1} + \frac{M_i^2}{(M_i + N_i)^2} - \frac{M_i^2 N_i}{(M_i + N_i)^2},$$

whence MN(1) holds. Equation MN(2) is obtained similarly, from ABMN(2). \square

Recall $c, d, s : (0, \infty) \rightarrow (0, \infty)$ from Definition 2.18.

DEFINITION 5.2. Let $\gamma, \delta : (0, \infty) \rightarrow (0, \infty)$ be given by $\gamma(x) = c(x)^{-1}$ and $\delta(x) = d(x)^{-1}$. Set $\beta : (0, \infty) \rightarrow (0, \infty)$, $\beta(x) = \frac{\omega-1}{4}$, where recall that $\omega = \sqrt{8x+1}$ for $x \in (0, \infty)$.

LEMMA 5.3.

- (1) *The functions $c, d, s : (0, \infty) \rightarrow (0, \infty)$ are increasing.¹*
- (2) *We have that $s(x) = x^2/2 + O(x^3)$ as $x \searrow 0$.*
- (3) *For $x \in (0, \infty)$, $s(x) = \frac{\beta(x)^2}{\beta(x)+2}$.*

¹Let $* \in \{c, d, s\}$. By ‘Lemma 5.3(1:*)’ will be meant ‘* is increasing’.

(4) For $x \in (0, \infty)$, $\beta(x) \leq x$.

(5) For $x \in (0, \infty)$, $s(x) < x$.

Proof: (1). The expressions for $c(x)$, $d(x)$ and $s(x)$ in Definition 2.18 are readily seen to be increasing in the variable $\omega \in (1, \infty)$; since $\omega = \sqrt{8x+1}$, they are also increasing in $x \in (0, \infty)$.

(2). We have that $\omega = \sqrt{8x+1} = 1 + 4x + O(x^2)$, whence

$$s(x) = \frac{(\omega-1)^2}{4(\omega+7)} = \frac{16x^2 + O(x^3)}{4(8+O(x))} = x^2/2 + O(x^3).$$

(3). This is due to $s(x) = \frac{(\omega-1)^2}{4(\omega+7)}$ and $\beta(x) = (\omega-1)/4$.

(4). Since $\omega(x) = \sqrt{8x+1} \leq 4x+1$, $\beta(x) \leq x$.

(5). Lemma 5.3(3), $\beta > 0$ and Lemma 5.3(4) imply that

$$s(x) = \frac{\beta(x)^2}{\beta(x)+2} < \beta(x) \leq x$$

as desired. \square

Recall Definition 2.12.

PROPOSITION 5.4. For $i \in \mathbb{Z}$, we have that²

$$\gamma(\phi_i) = \frac{m_i - m_{i-1}}{m_{i+1} - m_{i-1}}, \quad \delta(\phi_i) = \frac{n_{i-1} - n_i}{n_{i-1} - n_{i+1}} \quad \text{and} \quad s(\phi_i) = \phi_{i+1}.$$

Notation to be used only in the proof of this proposition³ makes the task to show that $*(\phi_i)$ equals $*_i$ for $* \in \{\gamma, \delta, s\}$.

DEFINITION 5.5. For $i \in \mathbb{Z}$, set $\gamma_i = \frac{m_i - m_{i-1}}{m_{i+1} - m_{i-1}}$, $\delta_i = \frac{n_{i-1} - n_i}{n_{i-1} - n_{i+1}}$ and $s_i = \phi_{i+1}$. We also set $\beta_i = \frac{n_{i-1} - n_{i+1}}{m_{i+1} - m_{i-1}}$, and write $\omega_i = \omega(\phi_i) = \sqrt{8\phi_i + 1}$.

LEMMA 5.6. We have that

$$(1 + \beta_i)^2 \gamma_i = 1, \quad 1 - \delta_i = \frac{\beta_i^2}{(1 + \beta_i)^2}, \quad \phi_i = \delta_i \beta_i / \gamma_i, \quad \phi_{i+1} = \frac{\beta_i(1 - \delta_i)}{1 - \gamma_i}.$$

Proof. Equation MN(1) implies that $(1 + \beta_i)^2 \gamma_i = 1$. Equation MN(2) implies $1 - \delta_i = \frac{\beta_i^2}{(1 + \beta_i)^2}$. That $\phi_i = \delta_i \beta_i / \gamma_i$ follows by the definitions of the concerned quantities. Noting that

$$1 - \delta_i = \frac{n_i - n_{i+1}}{n_{i-1} - n_{i+1}} \quad \text{and} \quad 1 - \gamma_i = \frac{m_{i+1} - m_i}{m_{i+1} - m_{i-1}},$$

we find from the definitions of β_i and ϕ_{i+1} that $\phi_{i+1} = \frac{\beta_i(1 - \delta_i)}{1 - \gamma_i}$ holds. \square

LEMMA 5.7. For $i \in \mathbb{Z}$,

$$\gamma_i^{-1} = \frac{1}{16}(\omega_i + 3)^2, \quad \delta_i^{-1} = \frac{(\omega_i + 3)^2}{8(\omega_i + 1)}, \quad s_i = \frac{(\omega_i - 1)^2}{4(\omega_i + 7)} \quad \text{and} \quad \beta_i = \frac{1}{4}(\omega_i - 1).$$

²Let $* \in \{\gamma, \delta, s\}$. By ‘Proposition 5.4(*)’, we will mean the statement made concerning the labelled quantity.

³In particular, the temporary usage of s_i introduced in Definition 5.5 is an abuse, because the denoted quantity is not the function s_i ; nor is it the value $s_i(x)$ for $x = \phi_0$. Indeed, $s_i(x)$ equals ϕ_i , while s_i with the temporary usage equals ϕ_{i+1} .

Proof. Omitting i subscripts, consider the four equations stated in Lemma 5.6 when we take $\phi \in (0, \infty)$ given. The first and third equations imply that $\delta\beta(1 + \beta)^2 = \phi$. Using the second equation, we find that $(2\beta + 1)\beta = \phi$; since β is positive, we confirm that $\beta = (\omega - 1)/4$. From the first equation, we then obtain $\gamma = 16(\omega + 3)^{-2}$. The third equation $\delta = \phi\gamma/\beta$ then yields $\delta = \frac{16\phi}{(\omega+3)^2} \cdot \frac{4}{\omega-1}$ which equals $\frac{8(\omega+1)}{(\omega+3)^2}$ in view of $\omega^2 - 1 = 8\phi$. Finally, $s = \phi_{i+1}$ by definition, so that the fourth equation implies that $s = \frac{\omega-1}{4} \cdot \frac{(\omega+3)^2 - 8\omega - 8}{(\omega+3)^2 - 16}$ whose right-hand side is seen to equal $\frac{(\omega-1)^2}{4(\omega+7)}$ after cancellation of $\omega - 1 > 0$ from numerator and denominator. \square

LEMMA 5.8.

(1) We have that $\gamma_i^{-1} - 1 = 2\phi_i + O(\phi_i^2)$.

(2) And that $\beta_i = \phi_i + O(\phi_i^2)$.

Proof: (1). From Lemma 5.7, note that $\gamma_i^{-1} = \frac{1}{16}(\omega_i + 3)^2 = (1 + \phi_i + O(\phi_i^2))^2$.

(2). By the same result, $\beta_i = \frac{1}{4}(\omega_i - 1) = \phi_i + O(\phi_i^2)$. \square

Proof of Proposition 5.4. By Lemma 5.7 and Definitions 2.18, 5.2 and 5.5,

$$\gamma_i = 16(\omega_i + 3)^{-2} = c(\phi_i)^{-1} = \gamma(\phi_i) ; \quad \delta_i = \frac{8(\omega_i+1)}{(\omega_i+3)^2} = d(\phi_i)^{-1} = \delta(\phi_i) ;$$

and $s_i = \frac{(\omega_i-1)^2}{4(\omega_i+7)} = s(\phi_i)$. \square

Proofs of Proposition 2.16 and Theorem 2.21. For given $x \in (0, \infty)$, let (a, b, m, n) be an ABMN solution with $\frac{n-1-n_0}{m_0-m_{-1}} = x$.

Since $c_i(x) = c(s_i(x)) = c(\phi_i)$, Definition 5.2 and Proposition 5.4(γ) imply that

$$c_i(x) - 1 = \frac{1 - \gamma(\phi_i)}{\gamma(\phi_i)} = \frac{m_{i+1} - m_i}{m_i - m_{i-1}}. \quad (42)$$

Adopting the notation in Definition 2.20, we find that

$$\frac{m_{j+1} - m_j}{m_0 - m_{-1}} = \prod_{i=0}^j (c_i(x) - 1) \quad (43)$$

for any $j \in \mathbb{Z}$. Since a default solution has $m_0 - m_{-1} = 1$ by definition, we deduce that the formula for $m_{k+1}^{\text{def}} - m_k^{\text{def}}$ in Definition 2.20 holds. Similarly to (42), we find via Proposition 5.4(δ) that

$$d_i(x) - 1 = \frac{1 - \delta(\phi_i)}{\delta(\phi_i)} = \frac{n_i - n_{i+1}}{n_{i-1} - n_i},$$

whence

$$\frac{n_j - n_{j+1}}{n_{-1} - n_0} = \prod_{i=0}^j (d_i(x) - 1)$$

for $j \in \mathbb{Z}$. Since $n_{-1} - n_0 = x(m_0 - m_{-1}) = x$ for any default solution, we find that the formula for $n_k^{\text{def}} - n_{k+1}^{\text{def}}$ in Definition 2.20 is valid. Proposition 5.1 implies that the sought formulas for a_i^{def} and b_i^{def} for $i \in \mathbb{Z}$ hold. The exhibited solution exists and is unique. This completes the proof of Theorem 2.21. The noted existence and uniqueness also prove Proposition 2.16. \square

5.2. Asymptotic decay of solutions

Here we prove Theorem 2.14.

LEMMA 5.9.

(1) For $\phi_i \in (0, 1)$, we have that

$$\phi_i^2/2 - O(\phi_i^3) \leq \phi_{i+1} \leq \phi_i^2/2,$$

where the positive constant implied by the O -notation is bounded above in terms of $h \in (0, 1)$, where $\phi_i \in (0, 1 - h)$.

(2) For any $i \in \mathbb{Z}$, $\phi_{i+1} < \phi_i$.

Proof: (1). From Proposition 5.4(s) and Lemma 5.3(3,4), we see that

$$\phi_{i+1} \leq \beta(\phi_i)^2/2 \leq \phi_i^2/2.$$

By Proposition 5.4(s) and Lemma 5.3(2), $\phi_{i+1} = s(\phi_i) = \phi_i^2/2 + O(\phi_i^3)$.

(2). By Lemma 5.3(5), and Proposition 5.4(s), $\phi_{i+1} = s(\phi_i) < \phi_i$. □

We are about to prove Theorem 2.14. Since this result uses the notion of the battlefield index specified in Definition 2.13, we now offer a proof that this index is well-defined.

LEMMA 5.10. Let $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ be a positive ABMN solution on \mathbb{Z} . There is a unique value of $k \in \mathbb{Z}$ for which $\phi_k \in (1/3, 3]$.

Proof. By Lemma 5.9(2), the sequence $\{\phi_i : i \in \mathbb{Z}\}$ is decreasing. Taking $s(0) = 0$, the value $\lim_{i \rightarrow \infty} \phi_i$ is a fixed point of $s : [0, \infty) \rightarrow [0, \infty)$ because s is continuous and $s(\phi_i) = \phi_{i+1}$ (the latter by Proposition 5.4(s)). But $s(x) < x$ for $x > 0$ by Lemma 5.3(5). Thus, $\phi_i \searrow 0$ as $i \rightarrow \infty$. The opposite limiting value $\lim_{i \rightarrow \infty} \phi_{-i}$ would also be a fixed point for $s : [0, \infty) \rightarrow [0, \infty)$ were it to be finite; we see then that $\lim_{i \rightarrow \infty} \phi_{-i}$ is infinite. We may thus set $k \in \mathbb{Z}$ so that $k = \inf \{i \in \mathbb{Z} : \phi_i \leq 3\}$ and be assured that k is well-defined. Now, $\phi_j > 3$ for $j \leq k - 1$, while ϕ_k , being $s(\phi_{k-1})$, exceeds $s(3) = 1/3$ by Lemma 5.3(1:s). On the other hand, if $j \geq k + 1$, then $\phi_j \leq \phi_{k+1} = s(\phi_k) \leq s(3) = 1/3$. Thus, $k \in \mathbb{Z}$ is the unique index whose ϕ -value exceeds one-third and is at most three. □

Proof of Theorem 2.14(1). For $i \in \mathbb{N}$, set $\epsilon_i = \phi_{k+i}/2$ and $g_i = -\log \epsilon_i$. By $s(3) = 1/3$ and Lemma 5.3(1:s), we have that $s(x) \leq 1/3$ for $x \in (0, 3]$. Definition 2.13 and $s(\phi_i) = \phi_{i+1}$ (from Proposition 5.4(s)) thus imply that $\phi_{k+j} \leq 1/3$ for $j \geq 1$. We may then apply Lemma 5.9(1) to find that $\epsilon_i^2(1 - O(\epsilon_i)) \leq \epsilon_{i+1} \leq \epsilon_i^2$, where the positive constant implicit in the O -notation may be chosen independently of the ABMN solution $\{(a_j, b_j, m_j, n_j) : j \in \mathbb{Z}\}$ and the value of the index $i \geq 1$. (We say that a positive constant is universal, or is bounded universally, if it may be so chosen.) We learn that

$$2g_i \leq g_{i+1} \leq 2g_i + O(e^{-g_i}), \tag{44}$$

where the implicit positive constant is again universal. Thus, $g_i > \log 6$ for $i \geq 1$, and we may write $g_i = 2^{\ell_i}$ for a real-valued sequence $\{\ell_i : i \in \mathbb{N}_+\}$ whose terms are bounded below by $\frac{\log \log 6}{\log 2} > 0$. From (44), we find that

$$0 \leq g_{i+1} - 2g_i = (2^{\ell_{i+1} - \ell_i - 1} - 1)2^{\ell_i + 1} = O(\exp\{-2^{\ell_i}\});$$

using $\ell_i > 0$, we readily obtain

$$0 \leq \ell_{i+1} - \ell_i - 1 = O(\exp\{-2^{\ell_i}\}).$$

Since $\ell_1 > 0$ and $\ell_{i+1} \geq \ell_i + 1$, we have that $\ell_i > i - 1$ for $i \geq 1$. Thus,

$$0 \leq \ell_{i+1} - \ell_i - 1 = O(\exp\{-2^{i-1}\}).$$

We may find $B \in \mathbb{R}$ so that $\ell_i = B + i + O(\exp\{-2^{i-1}\})$ for $i \in \mathbb{N}_+$. The universal form of O and the fact that ℓ_1 is bounded (since $\epsilon_1 \in (1/6, 3/2]$) implies that B is bounded in a universal sense. Set $A = 2^B$ (so that A is bounded away from zero and infinity in a universal sense), and exponentiate with base two to obtain

$$g_i = A \cdot 2^{i+O(\exp(-2^{i-1}))}$$

for $i \geq 1$. Since $\phi_{k+i} = 2e^{-g_i}$, we see then that, for $i \geq k + 1$,

$$\phi_i = 2 \exp \left\{ -A \cdot 2^{i-k+O(\exp(-2^{i-k-1}))} \right\}. \quad (45)$$

Similarly as we derived (43), we find that

$$m_j - m_{j-1} = (m_k - m_{k-1}) \prod_{i=k}^{j-1} (\gamma_i^{-1} - 1)$$

for $j \geq k + 1$. By Lemma 5.8(1),

$$\begin{aligned} m_j - m_{j-1} &= (m_k - m_{k-1}) \prod_{i=k}^{j-1} \left(4 \exp \left\{ -A \cdot 2^{i-k+\kappa_i \exp(-2^{i-k-1})} \right\} + O(1)e^{-A \cdot 2^{i-k+1/2}} \right) \\ &= (m_k - m_{k-1}) 4^{j-k} E_{k,j} \prod_{i=k}^{j-1} \exp \left\{ -A \cdot 2^{i-k+\kappa_i \exp(-2^{i-k-1})} \right\}, \end{aligned}$$

where the values of κ_i are bounded above in absolute value (in a universal sense), and where

$$\begin{aligned} E_{k,j} &= \prod_{i=k}^{j-1} \left(1 + O(1) \exp \left\{ -A 2^{i-k} (2^{1/2} - 2^{\kappa_i \exp(-2^{i-k-1})}) \right\} \right) \\ &= \prod_{i=k}^{j-1} \left(1 + \exp \left\{ -O(1)A \cdot 2^{i-k} \right\} \right) \end{aligned}$$

satisfies $E_{k,j} = E(1 + e^{-O(1)A 2^{j-k}})$ with

$$E = \prod_{i=k}^{\infty} \left(1 + \exp \left\{ -O(1)A \cdot 2^{i-k} \right\} \right).$$

The quantity E is positive and bounded away from zero and infinity universally. Note that

$$\begin{aligned} \sum_{i=k}^{j-1} 2^{i-k+\kappa_i \exp(-2^{i-k-1})} &= 2^{j-k} - 1 + \sum_{i=k}^{j-1} 2^{i-k} (2^{\kappa_i \exp(-2^{i-k-1})} - 1) \\ &= 2^{j-k} - 1 + \rho - \sum_{i=j}^{\infty} 2^{i-k} (2^{\kappa_i \exp(-2^{i-k-1})} - 1) \\ &= 2^{j-k} - 1 + \rho + O(1)e^{-2^{j-k}O(1)}, \end{aligned}$$

where $\rho = \sum_{i=k}^{\infty} 2^{i-k} (2^{\kappa_i \exp\{-2^{i-k-1}\}} - 1)$. Thus, $m_j - m_{j-1}$ equals

$$\begin{aligned} & (m_k - m_{k-1}) 4^{j-k} \exp\{-2^{j-k} A\} E \exp\{A(1-\rho)\} (1 + e^{-O(1)A2^{j-k}}) \\ & \quad \times \left(1 + e^{-O(1)A2^{j-k}}\right) \exp\left\{A \cdot O(1)e^{-2^{j-k}O(1)}\right\} \\ & = (m_k - m_{k-1}) 4^{j-k} \exp\{-2^{j-k} A\} E \exp\{A(1-\rho)\} (1 + e^{-O(1)2^{j-k}}), \end{aligned}$$

where we used that $A = \Theta(1)$ —namely, A is bounded away from zero and infinity in a universal sense—for the displayed equality. Set F equal to $E \exp\{A(1-\rho)\}$, and note that this positive expression is bounded away from zero and infinity universally. We find that

$$m_j - m_{j-1} = (m_k - m_{k-1}) \cdot F \cdot 2^{2(j-k)} \exp\{-2^{j-k} A\} (1 + e^{-O(1)2^{j-k}}), \quad (46)$$

which is the inference that Theorem 2.14 makes for the sequence of m -differences. With $M = m_{j+1} - m_{j-1}$ and $N = n_{j-1} - n_{j+1}$, we have that $a_j = \frac{M^2 N}{(M+N)^2}$ and $b_j = \frac{MN^2}{(M+N)^2}$ from (41). Using the definition of β_j in the guise $N = \beta_j M$, and Lemma 5.8(2) with $i = j$, we find that

$$a_j = (m_{j+1} - m_{j-1}) \frac{\beta_j}{(1+\beta_j)^2} = (m_{j+1} - m_{j-1}) (\phi_j + O(\phi_j^2))$$

and

$$b_j = (m_{j+1} - m_{j-1}) \frac{\beta_j^2}{(1+\beta_j)^2} = (m_{j+1} - m_{j-1}) (\phi_j^2 + O(\phi_j^3)).$$

We may use (46) to replace the quantity $m_{j+1} - m_{j-1}$ in these expressions. The expressions in terms of ϕ_i may be bounded by means of (45):

$$\begin{aligned} \phi_j & = 2 \exp\{-A \cdot 2^{j-k} (1 + O(\exp\{-2^{j-k-1}\}))\} \\ & = 2 \exp\{-A \cdot 2^{j-k}\} \exp\{O(e^{-2^{j-k}c})\} = 2 \exp\{-A \cdot 2^{j-k}\} (1 + O(e^{-2^{j-k}c})). \end{aligned}$$

Here, the value of c is positive (and universal) in the second line. We thus obtain the expressions for a_j and b_j in Theorem 2.14(1). It remains to derive the asymptotic expression for the quantity $n_j - n_{j-1}$. Here, we use $n_{j-1} - n_j = \phi_j(m_j - m_{j-1})$, (46) and the preceding display.

(2). According to Definition 2.13, the battlefield index $k \in \mathbb{Z}$ is the unique solution of $\phi_k \in (1/3, 3]$. Consider the role-reversal transformation that replaces index i by $2k - i$, and (a, b, m, n) by (b, a, n, m) . The resulting system is also a solution of the ABMN system by a minor variation of Proposition 3.2. Write $\hat{\phi}_i$ for the value of ϕ_i in the transformed solution. Then $\hat{\phi}_i = 1/\phi_{2k+1-i}$ for $i \in \mathbb{Z}$. We see then that $\hat{\phi}_{k+1} \in [1/3, 3)$ (so that $k+1$ is the battlefield index of the transformed system except when $\phi_k = 1/3$). Theorem 2.14(2) thus reduces to Theorem 2.14(1), because the proof of the latter operates as well as when $\phi_k = 1/3$ as when $\phi_k \in (1/3, 3]$. \square

5.3. Consequences of asymptotic decay

We may now complete the proof of Theorem 2.3.

Proof of Theorem 2.3(3). By Theorem 2.3(2), we know that $m_\infty, m_{-\infty}, n_\infty$ and $n_{-\infty}$ exist as elements of $\mathbb{R} \cup \{\infty\} \cup \{-\infty\}$. Since we know that m_0 and n_0 belong to \mathbb{R} , it is enough, in order to exclude the possibility that one of the four quantities is infinite, to argue that $\lim_{i \rightarrow \infty} (m_i - m_0) < \infty$, $\lim_{i \rightarrow \infty} (m_{-i} - m_0) > -\infty$, $\lim_{i \rightarrow \infty} (n_i - n_0) > -\infty$ and $\lim_{i \rightarrow \infty} (n_{-i} - n_0) < \infty$. These results follow from the asymptotic expressions for $m_j - m_{j-1}$ and $n_{j-1} - n_j$ in Theorem 2.14(1,2).

The almost sure occurrence of the unanimity event U is a consequence of Theorem 2.14, and we prove it now.

Proof of Theorem 2.15(4). By Theorem 2.6(1), this reduces to Theorem 2.15(1,2,3).

(1,2,3). We abusively write $(S_-, S_+) = (b, a)$ as usual. Theorem 2.6 and Theorem 2.14(1) imply that, for $i \geq k$, $\frac{a_i}{a_i + b_i} = 1 - 2 \exp\{-2^{i-k} A\} (1 + e^{-O(1)2^{i-k}})$. Thus, the \mathbb{P}_{S_-, S_+}^i -probability that every move is won by Maxine equals

$$\prod_{j=i}^{\infty} \frac{a_j}{a_j + b_j} = \prod_{j=i}^{\infty} \left(1 - 2 \exp\{-2^{j-k} A\} (1 + e^{-O(1)2^{j-k}}) \right) = 1 - 2 \exp\{-2^{i-k} A\} (1 + e^{-O(1)2^{i-k}}).$$

This bound proves Theorem 2.15(2). The corresponding bound for $i \leq k - 1$, and the proof of Theorem 2.15(3), are similar. It remains then to derive Theorem 2.15(1). The displayed and omitted bounds permit us to choose $L \in \mathbb{N}$ such that

$$\text{if } |i - k| > L, \text{ then } \mathbb{P}_{S_-, S_+}^i(U) \geq 1/2. \quad (47)$$

The *status report* $\text{Stat} : \mathbb{N} \rightarrow \{I, O, F\}$ is a random process defined under the law \mathbb{P}_{S_-, S_+}^i that we will use to prove Theorem 2.15(1). This process takes values in a three-point set whose labels denote ‘inner’, ‘outer’ and ‘final’. To record the status report, we will iteratively specify an increasing sequence $\{\tau_i : i \in \mathbb{N}\}$ of times valued in $\mathbb{N} \cup \{\infty\}$. We set $\tau_0 = 0$. We check whether $|X_0 - k| \leq L$, where the value of L was specified in the preceding paragraph. If this condition is met then we set $\text{Stat}(0) = I$. If the condition is not met, we set $\text{Stat}(0) = O$.

Let $i \in \mathbb{N}_+$. Suppose that an initial status report $\text{Stat}(j) \in \{I, O, F\}$, $j \in \llbracket 0, i-1 \rrbracket$, and an increasing sequence $\tau_j \in \mathbb{N} \cup \{\infty\}$, $j \in \llbracket 0, i-1 \rrbracket$, has been recorded.

If $\text{Stat}(i-1) = F$, we set $\tau_i = \infty$ and $\text{Stat}(i) = F$.

If $\text{Stat}(i-1) = I$, we set $\tau_i = \tau_{i-1} + L$. We set

$$\text{Stat}(i) = \begin{cases} I & \text{if } |X_{\tau_i} - k| \leq L \\ O & \text{in the other case.} \end{cases}$$

If $\text{Stat}(i-1) = O$, we begin to view the process X run forward from time τ_{i-1} . We watch for the first occasion $F \geq \tau_{i-1} + 2$ at which the sequence of observed differences $X_{j+1} - X_j$, $F - 1 \geq j \geq \tau_{i-1}$, has assumed both values -1 and 1 . If this occasion never occurs, so that $F = \infty$, we set $\text{Stat}(i) = F$ and $\tau_i = \infty$. If the occasion does occur, we set $\tau_i = F$. The last display is used to set $\text{Stat}(i)$. This completes the description of the iterative scheme for the generic later step indexed by $i \geq 1$.

The status report $\text{Stat} : \mathbb{N} \rightarrow \{I, O, F\}$ is not a Markov process, but it has simple properties that serve to prove that unanimity U is an almost sure event under \mathbb{P}_{S_-, S_+}^i for any $i \in \mathbb{Z}$. Consider then the process Stat under the just mentioned law. By construction, Stat arrives, and is absorbed, in F precisely when the event U occurs. To prove Theorem 2.15(1), our task is thus to show that Stat almost surely reaches F . Two properties suffice to show this.

Property I. Let $j \in \mathbb{N}_+$. Suppose given a status report history Stat_i , $i \in \llbracket 0, j-1 \rrbracket$, for which $\text{Stat}_{j-1} = I$. There exists a constant $c > 0$ that does not depend on this history such that the conditional probability that $\text{Stat}(j) = O$ is at least c .

Property O. Let $j \in \mathbb{N}_+$. Suppose given a history Stat_i , $i \in \llbracket 0, j-1 \rrbracket$, for which $\text{Stat}_{j-1} = O$. The conditional probability that $\text{Stat}(j) = F$ is at least one-half.

Properties *I* and *O* show that, whatever the status report history up to a given moment, there is probability at least $c/2$ that one of the next two entries in the report is F . Thus, it is inevitable

that the report will eventually contain an entry in F . The proof of Theorem 2.15(1) has thus been reduced to the task of deriving the two properties.

The proofs of Properties I and O depend on a *claim*. This states that all the information in any report history Stat_i , $i \in \llbracket 0, j-1 \rrbracket$, in which F is not recorded, is contained in the gameplay history X_i , $i \in \llbracket 0, \tau_{j-1} \rrbracket$. The claim may be proved by induction on j . The specifications of τ_j above are stopping times for the process X that are finite when $\tau_i \in \{I, O\}$. This proves the claim.

We now prove Property I . The coefficients a_i and b_i are positive by Theorem 2.6; and they are bounded by Theorem 2.14. Consider then the event that X makes L rightward jumps from time τ_{j-1} . To find a lower bound on the conditional probability of this event given the circumstance of Property I , note that the claim permits us to further condition on X until time τ_{j-1} . Since $|X_{\tau_{j-1}} - k| \leq L$, a lower bound is offered by the minimum over $\ell \in \llbracket k-L, k+L \rrbracket$ of the product $\prod_{i=0}^{L-1} \frac{a_{\ell+i}}{a_{\ell+i} + b_{\ell+i}}$. This minimum is positive because the positive and bounded quantities a and b that are involved are finite in number.

And now we prove Property O . Again, by the claim, we may condition on X until time τ_{j-1} . Since $|X_{\tau_{j-1}} - k| > L$, we may invoke (47) to show the sought property.

This completes the proof of Theorem 2.15(1). □

The Mina margin map

Here we prove our results concerning the Mina margin map in three sections. Finite-trail counterparts to the map are defined and estimated in Section 6.1, and Theorem 2.23 and several consequences are derived. In Section 6.2, the θ^{-1} - and Θ -transforms of the map are compared, and Theorem 2.28 is proved. In Section 6.3, the $\lambda \leq 0.999904$ bound Theorem 2.4(3) is derived by a scheme of explicit approximation for a well-chosen value of a suitable finite-trail counterpart for \mathcal{M} .

6.1. Approximating the Mina margin by its finite trail counterpart

Here we prove Theorem 2.23, the third part contingent on Theorem 2.4(3). At the end of the section, we prove the consequent Theorems 2.4(1,2); 2.8; 2.10; and 2.24(1,2).

6.1.1. An explicit form for the finite-trail Mina margin map.

LEMMA 6.1. *Let $x \in (0, \infty)$. For $k, \ell \in \mathbb{Z} \cup \{\infty\} \cup \{-\infty\}$, $k < \ell$, the value of $\frac{n_k - n_\ell}{m_\ell - m_k} \in (0, \infty)$ is a constant function of the element $\{(a_i, b_i, m_i, n_i) \in (0, \infty)^2 \times \mathbb{R}^2 : i \in \mathbb{Z}\}$ in the equivalence class $\mathcal{C}(x)$.*

Remark. When we write expressions $\frac{n_k - n_\ell}{m_\ell - m_k}$ in this section, we refer to the quantities $\frac{n_k - n_\ell}{m_\ell - m_k}(x)$ that the above lemma identifies; the value of $x \in (0, \infty)$ is often understood.

Proof of Lemma 6.1. That each expression $\frac{n_k - n_\ell}{m_\ell - m_k}$ is a finite number follows from Theorem 2.3(1,3). Each expression $\frac{n_k - n_\ell}{m_\ell - m_k}$ is invariant under the translation $\chi_{u,v}$, $u, v \in \mathbb{R}$ and dilation τ_x , $x \in \mathbb{R}$, maps that must be used to interpolate any two elements of $\mathcal{C}(x)$. \square

Recall the functions $s_j, c_j, d_j : (0, \infty) \rightarrow (0, \infty)$, $j \in \mathbb{Z}$, from Definition 2.19.

Set $P_0 = S_0 = 1$. For $k \in \mathbb{N}$, we iteratively specify

$$P_{k+1}(x) - P_k(x) = \prod_{i=0}^k (c_i(x) - 1), \quad (48)$$

and

$$S_{k+1}(x) - S_k(x) = \prod_{i=0}^k (d_i(x) - 1). \quad (49)$$

Set $Q_1 = T_1 = 0$. For $k \in \mathbb{N}_+$, we then set

$$Q_{k+1}(x) - Q_k(x) = \prod_{i=1}^k (c_{-i}(x) - 1)^{-1}, \quad (50)$$

and

$$T_{k+1}(x) - T_k(x) = \prod_{i=1}^k (d_{-i}(x) - 1)^{-1}. \quad (51)$$

LEMMA 6.2. *Let x equal ϕ_0 from Definition 2.12. For $k \in \mathbb{N}$,*

$$P_k(x) = \frac{m_k - m_{-1}}{m_0 - m_{-1}} \quad \text{and} \quad S_k(x) = \frac{n_{-1} - n_k}{n_{-1} - n_0}.$$

For $\ell \in \mathbb{N}_+$,

$$Q_\ell(x) = \frac{m_{-1} - m_{-\ell}}{m_0 - m_{-1}} \quad \text{and} \quad T_\ell(x) = \frac{n_{-\ell} - n_{-1}}{n_{-1} - n_0}.$$

Proof. The claimed formula for $P_k(x)$ is trivial when $k = 0$. To prove the general formula for $P_k(x)$, it suffices to argue that $\prod_{i=0}^k (c_i(x) - 1)$ equals $\frac{m_{k+1} - m_k}{m_0 - m_{-1}}$ for $k \in \mathbb{N}$, and we do this by induction on k . The generic step in the induction is enabled by showing that $c_k(x) - 1 = \frac{m_{k+1} - m_k}{m_k - m_{k-1}}$, which we obtain as follows:

$$c_k(x) - 1 = \frac{1 - \gamma(s_k(x))}{\gamma(s_k(x))} = \frac{1 - \gamma(\phi_k)}{\gamma(\phi_k)} = \frac{m_{k+1} - m_k}{m_k - m_{k-1}},$$

the respective equalities by Definition 5.2; by iterating Proposition 5.4(s); and by Proposition 5.4(γ).

Likewise, the claimed formula for $Q_\ell(x)$ is trivial when $\ell = 1$. Establishing the formula in the general case is a matter of showing that $\prod_{i=0}^\ell (c_{-i}(x) - 1)^{-1}$ equals $\frac{m_{-\ell} - m_{-\ell-1}}{m_0 - m_{-1}}$ for $\ell \geq 2$. The generic inductive step here amounts to showing that $(c_{-\ell}(x) - 1)^{-1} = \frac{m_{-\ell} - m_{-\ell-1}}{m_{-\ell+1} - m_{-\ell}}$ for such ℓ , and follows, similarly as above, from $(c_{-\ell}(x) - 1)^{-1} = \frac{\gamma(s_{-\ell}(x))}{1 - \gamma(s_{-\ell}(x))}$.

The formulas for S and T follow when the changes

$$P \rightarrow S, Q \rightarrow T, k \rightarrow \ell, c \rightarrow d, \gamma \rightarrow \delta \quad \text{and} \quad m_i \rightarrow -n_{-i}$$

are made. □

Recall from (14) that the finite-trail Mina margin map $\mathcal{M}_{\ell,k} : (0, \infty) \rightarrow (0, \infty)$ satisfies $\mathcal{M}_{\ell,k}(x) = \frac{n_{-\ell} - n_k}{m_k - m_{-\ell}}$ for $k \in \mathbb{N}$ and $\ell \in \mathbb{N}_+$, where $x = \phi_0$.

LEMMA 6.3. *We have that*

$$\mathcal{M}_{\ell,k}(x) = \frac{x(S_k + T_\ell)}{P_k + Q_\ell}$$

for $k \in \mathbb{N}$ and $\ell \in \mathbb{N}_+$.

In reading the proof of this result, recall the notation explained in the remark that follows Lemma 6.1.

Proof of Lemma 6.3. By Lemma 6.2,

$$m_k - m_{-\ell} = (P_k + Q_\ell)(m_0 - m_{-1}) \quad \text{and} \quad n_{-\ell} - n_k = (S_k + T_\ell)(n_{-1} - n_0). \quad (52)$$

But $x = \phi_0$, which is to say, $x = \frac{n_{-1} - n_0}{m_0 - m_{-1}}$. We find then that

$$\mathcal{M}_{\ell,k}(x) = \frac{(S_k + T_\ell)(n_{-1} - n_0)}{(P_k + Q_\ell)(m_0 - m_{-1})} = \frac{x(S_k + T_\ell)}{P_k + Q_\ell},$$

as we sought to do. □

6.1.2. Estimates for the finite trail Mina margin map. In this subsection, we derive the following compact-uniform Cauchy sequence property of the finite-trail Mina margin maps.

PROPOSITION 6.4. *For $k \geq 0$, $\ell \geq 2$ and $1/3 \leq x \leq 3$,*

$$\sup_{\substack{i \geq k+1 \\ j \geq \ell+1}} |\mathcal{M}_{i,j}(x) - \mathcal{M}_{\ell,k}(x)| \leq 3^5 2^{2k-2} 6^{1-2^k} + 3^3 2^{\ell-2} 6^{\ell-2^{\ell-1}}.$$

The next lemma assembles key elements for the proof of Proposition 6.4. We omit to denote the argument ' (x) ' of $\mathcal{M}_{\cdot,\cdot}$, P , Q , S and T as we derive this proposition.

LEMMA 6.5. *Let $k \in \mathbb{N}$ and $x \in \mathbb{R}$.*

- (1) *For $k \geq 0$ and $x \leq 3$, $P_{k+1} - P_k \leq 2^{2k} 6^{1-2^k}$.*
- (2) *For $k \geq 1$ and $x \geq 1/3$, $Q_{k+1} - Q_k \leq 2^{2k} 6^{1-2^k}$.*
- (3) *For $k \geq 0$ and $x \leq 3$, $S_{k+1} - S_k \leq 2^{2k+1} 6^{1-2^{k+1}}$.*
- (4) *For $\ell \geq 2$ and $x \geq 1/3$, $T_{\ell+1} - T_\ell \leq 3(12)^{\ell-1} 6^{1-2^{\ell-1}}$.*

Two simple lemmas gather estimates needed to prove Lemma 6.5.

LEMMA 6.6.

- (1) *For $x \in (0, \infty)$, $s(x) \leq x^2/2$.*
- (2) *For $x \in (0, \infty)$, $c(x) \leq 1 + 2x$.*
- (3) *For $x \in (0, \infty)$, $c(x) \geq 1 + x/2$.*
- (4) *For $x \in (0, 3]$, $d(x) - 1 \leq 1/3$.*
- (5) *For $x \in (0, \infty)$, $d(x) \geq 2^{-3/2} x^{1/2}$.*

Proof: (1). Since $\omega \geq 1$, $\beta(x) \geq 0$. Thus, Lemma 5.3(3) implies that $s(x) \leq \beta(x)^2/2$. So the result reduces to Lemma 5.3(4).

(2). By Definition 2.18, $c(x) = \frac{(\omega+3)^2}{16} = \frac{8x+10+6\omega}{16} \leq 1 + 2x$ where the inequality is due to $\omega = \sqrt{1+8x} \leq 1 + 4x$ for $x \geq 0$.

(3). We have that $c(x) = \frac{8x+10+6\omega}{16} \geq 1 + x/2$ from $\omega \geq 1$.

(4). By Lemma 5.3(1:d), $d(x) - 1 \leq d(3) - 1 = 1/3$.

(5). Recall that $d(x) = \frac{(\omega+3)^2}{8(\omega+1)}$ where $\omega = \sqrt{8x+1}$. Thus, $d(x) \geq (\omega+3)/8 \geq 2^{-3/2} x^{1/2}$. □

LEMMA 6.7. *Let $j \in \mathbb{N}_+$.*

- (1) *For $x \leq 3$, $s_j(x) \leq 2 \cdot 6^{-2^{j-1}}$.*
- (2) *For $x \geq 1/3$, $s_{-j}(x) \geq 2^{-1} 6^{2^{j-1}}$.*
- (3) *For $i \geq 1$ and $x \leq 3$, $c_i(x) \leq 1 + 2^2 6^{-2^{i-1}}$.*
- (4) *For $i \in \mathbb{Z}$ and $x \in (0, \infty)$, $d_i(x) - 1 \leq s_i(x)^2$.*
- (5) *For $i \geq 2$ and $x \geq 1/3$, $d_{-i}(x) - 1 \geq 2^{-1} 6^{2^{i-2}-1}$.*

Proof: (1). Note that $s(3) = 1/3$ since $\omega(3) = 5$. We may thus use Lemma 6.6(1) to prove the desired statement by induction.

(2). Due to the preceding and $s_{-j}(x) = 1/s_j(1/x)$.

(3). By Lemma 5.3(1:c), Lemma 6.7(1) and Lemma 6.6(2),

$$c_i(x) = c(s_i(x)) \leq c(2 \cdot 6^{-2^{i-1}}) \leq 1 + 2^2 6^{-2^{i-1}}$$

for $i \geq 1$ and $x \leq 3$.

(4). It is enough to show that $d(x) \leq 1 + x^2$. To see this, note that $d(x) - 1 = \delta(x)^{-1} - 1$. From $1 - \delta(x) = (1 - \frac{1}{\beta(x)+1})^2$ and Lemma 5.3(4), we find that $\delta(x) \geq 1 - \frac{x^2}{(1+x)^2}$, so that $d(x) - 1 \leq \frac{x^2}{1+2x} \leq x^2$.

(5). Note that $d_{-i}(x) = d(s_{-i}(x)) \geq d(2^{-1}6^{2^{i-1}}) \geq 2^{-2}6^{2^{i-2}}$, where the first inequality is due to Lemma 5.3(1:d) and Lemma 6.7(2), and the second to Lemma 6.6(5). From this, the sought result follows. \square

Proof of Lemma 6.5: (1). Note that $c(x) \leq 4$ for $x \in (0, 3]$ by Lemma 5.3(1:c) and $c(3) = 4$. Thus we bound the first term in the product in (48). Bounding the latter terms by Lemma 6.7(3), we find that

$$P_{k+1} - P_k = \prod_{i=0}^k (c_i(x) - 1) \leq 3 \prod_{i=1}^k 2^2 6^{-2^{i-1}},$$

whence the sought result.

(2). Note that

$$c_{-i}(x) - 1 = c(s_{-i}(x)) - 1 \geq c(2^{-1}6^{2^{i-1}}) - 1 \geq 2^{-2}6^{2^{i-1}},$$

where the first inequality holds when $x \geq 1/3$ in view of Lemma 5.3(1:c) and Lemma 6.7(2); the second is due to Lemma 6.6(3). By (50), $Q_{k+1} - Q_k \leq \prod_{i=1}^k 2^2 6^{-2^{i-1}} = 2^{2k} 6^{1-2^k}$, whence Lemma 6.5(2).

(3). Note that

$$S_{k+1} - S_k \leq 3^{-1} \prod_{i=1}^k 2^2 6^{-2^i} = 2^{2k+1} 6^{1-2^{k+1}},$$

where, in the first inequality, the first term in the product expression in (49) is bounded by use of Lemma 6.6(4), and the later terms are taken care of by the bounds $d_i(x) - 1 \leq s_i(x)^2 \leq 2^2 6^{-2^i}$, which are valid for $x \leq 3$ and $i \geq 1$ in view of Lemma 6.7(1,4).

(4). Since $s_{-1}(1/3) = 3$, Proposition 3.4 and Lemma 5.3(1:s) imply that $s_{-1}(x) \geq 3$ for $x \geq 1/3$. And since $d(3) = 4/3$, the same result implies that $(d_{-1}(x) - 1)^{-1} \leq 3$ for such x . Applying these bounds alongside Lemma 6.7(5) to (51), we see that

$$T_{\ell+1} - T_\ell \leq 3 \cdot \prod_{i=2}^{\ell} 2 \cdot 6^{1-2^{i-2}} = 3(12)^{\ell-1} 6^{1-2^{\ell-1}}$$

for $x \geq 1/3$ and $\ell \geq 2$. Whence Lemma 6.5(2). \square

Two further lemmas will permit the derivation of Proposition 6.4 from Lemma 6.5.

LEMMA 6.8. *We have that*

$$x^{-1} |\mathcal{M}_{k+1,\ell} - \mathcal{M}_{\ell,k}| \leq \max \left\{ S_{k+1} - S_k, (S_k + T_\ell)(P_{k+1} - P_k) \right\},$$

and

$$x^{-1}|\mathcal{M}_{\ell,k+1} - \mathcal{M}_{\ell,k}| \leq \max \left\{ T_{\ell+1} - T_{\ell}, (S_k + T_{\ell})(Q_{\ell+1} - Q_{\ell}) \right\}.$$

Proof. Since P_j and Q_j are at least one whenever j is at least one, it is enough to show that

$$x^{-1}|\mathcal{M}_{k+1,\ell} - \mathcal{M}_{\ell,k}| \leq \max \left\{ \frac{S_{k+1} - S_k}{P_{k+1} + Q_{\ell}}, \frac{(S_k + T_{\ell})(P_{k+1} - P_k)}{(P_{k+1} + Q_{\ell})(P_k + Q_{\ell})} \right\}, \quad (53)$$

and

$$x^{-1}|\mathcal{M}_{\ell+1,k} - \mathcal{M}_{\ell,k}| \leq \max \left\{ \frac{T_{k+1} - T_k}{P_{k+1} + Q_{\ell}}, \frac{(S_k + T_{\ell})(Q_{k+1} - Q_k)}{(P_{k+1} + Q_{\ell})(P_k + Q_{\ell})} \right\}. \quad (54)$$

Note that

$$\begin{aligned} x^{-1}(\mathcal{M}_{k+1,\ell} - \mathcal{M}_{\ell,k}) &= \frac{S_{k+1} + T_{\ell}}{P_{k+1} + Q_{\ell}} - \frac{S_k + T_{\ell}}{P_k + Q_{\ell}} \\ &= \frac{(S_{k+1} + T_{\ell})(P_k + Q_{\ell}) - (S_k + T_{\ell})(P_{k+1} + Q_{\ell})}{(P_{k+1} + Q_{\ell})(P_k + Q_{\ell})}. \end{aligned}$$

The numerator in the latter term equals $(P_k + Q_{\ell})(S_{k+1} - S_k) - (S_k + T_{\ell})(P_{k+1} - P_k)$. Since this is a difference of positive terms, and the right-hand denominator above is positive, we obtain (53).

Note further that

$$\begin{aligned} x^{-1}(\mathcal{M}_{\ell+1,k} - \mathcal{M}_{\ell,k}) &= \frac{S_k + T_{\ell+1}}{P_k + Q_{\ell+1}} - \frac{S_k + T_{\ell}}{P_k + Q_{\ell}} \\ &= \frac{(S_k + T_{\ell+1})(P_k + Q_{\ell}) - (S_k + T_{\ell})(P_k + Q_{\ell+1})}{(P_k + Q_{\ell+1})(P_k + Q_{\ell})}. \end{aligned}$$

In this case, the numerator in the last line is $(P_k + Q_{\ell})(T_{\ell+1} - T_{\ell}) - (S_k + T_{\ell})(Q_{\ell+1} - Q_{\ell})$. By reasoning as we did above, we obtain (54). This completes the proof of Lemma 6.8. \square

LEMMA 6.9.

- (1) For $x \leq 3$, $\sup_{k \geq 1} S_k \leq 3/2$.
- (2) For $x \geq 1/3$, $\sup_{k \geq 1} T_k \leq 12$.

Proof: (1). By $S_0 = 1$ and Lemma 6.5(2), $\sup_{k \geq 1} S_k \leq 1 + \sum_{k=0}^{\infty} 2^{2k+1}6^{1-2^{k+1}} = 1 + 3^{-1} + 2^{-4}3^{-7} + 2^{-10}3^{-15} + \dots = 1.333361 \dots \leq 3/2$.

(2). Recall that, by definition, $T_1 = 0$. The quantity T_2 equals $(d_{-1}(x) - 1)^{-1}$ which is at most 3 when $x \geq 1/3$, as we noted in the proof of Lemma 6.5(4). Using these alongside Lemma 6.5(4), we find that $\sup_{k \geq 1} T_k \leq 3 + \sum_{\ell=2}^{\infty} 3(12)^{\ell-1}6^{1-2^{\ell-1}} = 3 + \sum_{k=0}^{\infty} 2^k 6^{k+3-2^{k+1}} \leq 12$. \square

Proof of Proposition 6.4. By Lemma 6.8(1), Lemma 6.9(1,2) and Lemma 6.5(1,3), we have that

$$x^{-1}|\mathcal{M}_{k+1,\ell} - \mathcal{M}_{\ell,k}| \leq \max \left\{ 2^{2k+1}6^{1-2^{k+1}}, 2^{2k-4}6^{4-2^k} \right\}$$

for $k \geq 0$ and $\ell \geq 1$, where here we used $\frac{27}{2} \cdot 2^{2k}6^{1-2^k} = 2^{2k-4}6^{4-2^k}$. And by Lemma 6.8(1), Lemma 6.9(1,2) and Lemma 6.5(2,4),

$$x^{-1}|\mathcal{M}_{\ell+1,k} - \mathcal{M}_{\ell,k}| \leq \max \left\{ 3(12)^{\ell-1}6^{1-2^{\ell-1}}, \frac{27}{2}2^{2\ell}6^{1-2^{\ell}} \right\}$$

for $k \geq 0$ and $\ell \geq 2$. In each of the two displayed maximums, it is the first expression which is the greater for the stated ranges of k and ℓ . Set $g(i) = 2^{2i-1}6^{1-2^i}$ and $h(j) = (12)^{j-1}6^{1-2^{j-1}}$, and note

that $g(i+1)/g(i) \leq 1/3$ and $h(j+1)/h(j) \leq 1/3$ provided that $i \leq k$ and $j \geq \ell$ where $k \geq 0$ and $\ell \geq 2$. What we learn is that

$$x^{-1} \sup_{\substack{i \geq k+1 \\ j \geq \ell+1}} |\mathcal{M}_{i,j} - \mathcal{M}_{\ell,k}| \leq \frac{3}{2} \cdot \left(3^3 2^{2k-1} 6^{1-2k} + 3(12)^{\ell-1} 6^{1-2\ell-1} \right)$$

for $k \leq 0$, $\ell \geq 2$ and $1/3 \leq x \leq 3$. Using $x \leq 3$, and rewriting the product of three and the right-hand side of this display, we obtain Proposition 6.4. \square

6.1.3. Proofs via the finite trail Mina margin map.

Proof of Theorem 2.23(1). Note that

$$\mathcal{M}_{k,k}(x) = \frac{n_{-k} - n_k}{m_k - m_{-k}} \rightarrow \frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} = \mathcal{M}(x) \quad (55)$$

where the convergence, which is in the limit $k \rightarrow \infty$, is explained by the proof of Theorem 2.3(3); the latter equality here is due to the specification of $\mathcal{M}(x)$ in Definition 2.22 and to Lemma 6.1. Note next that, for $i \in \mathbb{Z}$, the standard element in $\mathcal{C}(s_i(x))$ is equal to the left shift by i places of the standard element in $\mathcal{C}(x)$. Thus,

$$\mathcal{M}_{k,k}(s_i(x)) = \frac{n_{-k+i} - n_{k+i}}{m_{k+i} - m_{-k+i}}.$$

The left-hand side converges to $\mathcal{M}(s_i(x))$ in the limit of high k , by (55). The right-hand side converges to $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} = \mathcal{M}(x)$ since m - and n -differences vanish asymptotically at high values of the index by Theorem 2.14. Thus we find that $\mathcal{M}(s_i(x)) = \mathcal{M}(x)$ for $i \in \mathbb{Z}$ and $x \in [1/3, 3]$. Since $\cup_{i \in \mathbb{Z}} s_i[1/3, 3] = (0, \infty)$, we see that $\mathcal{M}(x)$ exists for all $x \in (0, \infty)$, and that in fact $\mathcal{M}(s_i(x)) = \mathcal{M}(x)$ for $i \in \mathbb{Z}$ and $x \in (0, \infty)$. This completes the proof of Theorem 2.23(1).

Proof of Theorem 2.23: (2). We first show that \mathcal{M} is continuous on $(0, \infty)$. Proposition 6.4 shows that $\mathcal{M}_{k,k}$ converges uniformly as $k \rightarrow \infty$ on $[1/3, 3]$. By (55), the limiting function is the restriction of \mathcal{M} to $[1/3, 3]$. Since the constituent functions $c_i, d_i : (0, \infty) \rightarrow (0, \infty)$, $i \in \mathbb{Z}$, are continuous, we see that the map $\mathcal{M}_{k,k} : [1/3, 3] \rightarrow (0, \infty)$ is continuous for any $k \in \mathbb{N}_+$. Thus, \mathcal{M} is continuous on this interval. But $\mathcal{M}(x) = \mathcal{M}(s(x))$ for $x \in (0, \infty)$ by Theorem 2.23(1), and $\mathcal{M}(3) = \mathcal{M}(1/3)$ since $s(3) = 1/3$. Since $s : (0, \infty) \rightarrow (0, \infty)$ is seen to be continuous from its specification in Definition 2.18, we confirm that \mathcal{M} is continuous on $(0, \infty)$.

To derive the formula for $\mathcal{M}(x)$ claimed in Theorem 2.23(2), note, by decoding the notation for products in Definition 2.20, that this formula may expressed in our present notation in the form $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} = \frac{x(S_{\infty} + T_{\infty})}{P_{\infty} + Q_{\infty}}$, where in fact we have extended this notation to write $*_{\infty}$ for $\lim_{k \rightarrow \infty} *_{k}$ with $* \in \{P, Q, S, T\}$. Since $x = \frac{m_0 - m_{-1}}{n_{-1} - n_0}$, the sought formula is a consequence of

$$m_{\infty} - m_{-\infty} = (P_{\infty} + Q_{\infty})(m_0 - m_{-1}) \quad \text{and} \quad n_{-\infty} - n_{\infty} = (S_{\infty} + T_{\infty})(n_{-1} - n_0).$$

To obtain these identities, we take the limit in high k and ℓ of the two formulas in (52), using Theorem 2.3(3) to justify that the limiting expressions are finite real numbers. This completes the proof of Theorem 2.23(2). \square

6.1.4. Some further consequences. In order to prove Theorem 2.4 and Theorem 2.23(3), we now offer a definition of the quantity $\lambda \in (0, 1]$ to which these results refer.

DEFINITION 6.10. We set $\lambda = \inf\{\mathcal{M}(x) : x \in [1/3, 3]\}$.

LEMMA 6.11. *There exists $x_0 \in [1/3, 3]$ such that $\mathcal{M}(x_0) = \lambda$. We have that*

$$\lambda = \inf\{\mathcal{M}(x) : x \in (0, \infty)\}.$$

Proof. Since $\mathcal{M} : [1/3, 3] \rightarrow (0, \infty)$ is continuous by Theorem 2.23(2), the infimum is attained on $[1/3, 3]$, and we may find $x_0 \in [1/3, 3]$ so that $\mathcal{M}(x_0) = \lambda \in (0, \infty)$. The proof of Lemma 5.10 shows that $\cup_{i \in \mathbb{Z}} s_i[1/3, 3] = (0, \infty)$. By Theorem 2.23(3), we see thus that $\lambda = \inf\{\mathcal{M}(x) : x \in (0, \infty)\}$. \square

LEMMA 6.12. *For $x \in (0, \infty)$, $\mathcal{M}(x^{-1}) = \mathcal{M}(x)^{-1}$.*

Proof. Recall from Definition 2.22 that $\mathcal{M}(x) = n_{-\infty}^{\text{st}}(x)$ for $x \in (0, \infty)$ is the Mina margin of the standard solution $(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z})$. The ϕ_0 -value of this solution is equal to $\frac{n_{-1}^{\text{st}}(x) - n_0^{\text{st}}(x)}{m_0^{\text{st}}(x) - m_{-1}^{\text{st}}(x)} = x$. By Proposition 3.2 and dilation, the quadruple

$$n_{-\infty}^{\text{st}}(x)^{-1} \cdot \left(b_{-i}^{\text{st}}(x), a_{-i}^{\text{st}}(x), n_{-i}^{\text{st}}(x), m_{-i}^{\text{st}}(x) : i \in \mathbb{Z} \right)$$

is also a standard ABMN solution. Its ϕ_1 -value equals $\frac{m_0^{\text{st}}(x) - m_{-1}^{\text{st}}(x)}{n_{-1}^{\text{st}}(x) - n_0^{\text{st}}(x)} = x^{-1}$. The left shift by one place of the displayed quadruple is thus a standard ABMN solution whose ϕ_0 -value equals x^{-1} . The quantity $\mathcal{M}(x^{-1})$, which by definition equals $n_{-\infty}^{\text{st}}(x^{-1})$, is thus found to be equal to $n_{-\infty}^{\text{st}}(x)^{-1} \cdot m_{\infty}^{\text{st}}(x) = n_{-\infty}^{\text{st}}(x)^{-1} = \mathcal{M}(x)^{-1}$. Here, we used that $m_{\infty}^{\text{st}}(x) = 1$ since the solution $(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z})$ is standard. \square

LEMMA 6.13. *We have that $\lambda^{-1} = \sup\{\mathcal{M}(x) : x \in [1, 3, 3]\} = \sup\{\mathcal{M}(x) : x \in (0, \infty)\}$. Further, there exists $y_0 \in [1, 3, 3]$ such that $\mathcal{M}(y_0) = \lambda^{-1}$.*

Proof. The ranges $\mathcal{M}[1/3, 3]$ and $\mathcal{M}(0, \infty)$ are invariant under the transformation $z \rightarrow z^{-1}$ in view of Lemma 6.12. The supremum of the continuous function \mathcal{M} is attained on $[1/3, 3]$. \square

Proof of Theorem 2.23(3). The range $\mathcal{M}[1/3, 3]$ has maximum λ^{-1} and minimum λ , by Lemmas 6.11 and 6.13. By the continuity of \mathcal{M} on $[1/3, 3]$, $\mathcal{M}[1/3, 3]$ is seen to equal $[\lambda, \lambda^{-1}]$. Since $\cup_{i \in \mathbb{Z}} s_i[1/3, 3] = (0, \infty)$, $\mathcal{M}(0, \infty)$ equals $\mathcal{M}[1/3, 3]$. Note that $\lambda \in (0, 1]$ since $\mathcal{M}[1/3, 3] = [\lambda, \lambda^{-1}]$ and \mathcal{M} is continuous. The remaining assertion that we need to validate, which is that λ is at most 0.999904, is Theorem 2.4(3), whose proof will appear in Section 6.3. \square

Proof of Theorem 2.4(1). Theorem 2.23(3) shows that the set of values of the Mina margins of standard positive ABMN solutions is equal to $[\lambda, \lambda^{-1}]$. Now consider an arbitrary positive ABMN solution. The value of the Mina margin is shared between this solution and the equivalent standard solution. Thus, no new values for the Mina margin emerge as the solution set is enlarged from standard to general.

(2). Consider a positive ABMN solution (a, b, m, n) with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$. By Theorem 2.3(1), (a, b, m, n) is strict. Thus, $m_{-\infty} < m_{\infty}$ and $n_{\infty} < n_{-\infty}$. By Theorem 2.4(1), $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} \in [\lambda, \lambda^{-1}]$.

Conversely, suppose that $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) \in \mathbb{R}^4$ satisfies $m_{-\infty} < m_{\infty}$, $n_{\infty} < n_{-\infty}$ and $\frac{n_{-\infty} - n_{\infty}}{m_{\infty} - m_{-\infty}} \in [\lambda, \lambda^{-1}]$. Set x equal to the latter quantity. In the notation of Section 2.3, the image of the standard solution $(a_i^{\text{st}}(x), b_i^{\text{st}}(x), m_i^{\text{st}}(x), n_i^{\text{st}}(x) : i \in \mathbb{Z})$ under the transformation $\chi_{x,y} \circ \tau_u$, where $x = m_{-\infty}$, $n = n_{-\infty}$ and $u = m_{\infty} - m_{-\infty}$, is a positive ABMN solution with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty})$. This completes the proof of Theorem 2.4(2). \square

Proof of Theorem 2.8. Let $x \in (0, \infty)$. By Theorem 2.6, the game $\text{Standard}(x)$ has a time-invariant Nash equilibrium if and only if there exists a positive ABMN solution whose Mina margin equals x . The latter condition is equivalent to $x \in [\lambda, \lambda^{-1}]$ by Theorem 2.4(1). \square

Proof of Theorem 2.10. Let $y \in [1/3, 3]$, $\mathcal{M}(y) = \lambda^{-1}$, be the value y_0 assured by Lemma 6.13. Let $w \leq y$ be maximal such that $\mathcal{M}(w) = \lambda$, where Theorem 2.23 assures the existence of this quantity. We have that $w < y$ because \mathcal{M} is continuous and assigns different values to these two points. Set $z = s_{-1}(w)$. By Proposition 3.4 and Lemma 5.3(5), $z > w$. Since $\mathcal{M}(z) = \mathcal{M}(w) = \lambda$ by Theorem 2.23(1), we thus have that $z > y$.

Now let $x \in (\lambda, \lambda^{-1})$. By the continuity of \mathcal{M} , we may find $u \in (w, y)$ and $v \in (y, z)$ such that $\mathcal{M}(u) = \mathcal{M}(v) = x$. Note that $w < u < v < z = s_{-1}(w)$. The quadruples

$$(a_i^{\text{st}}(u), b_i^{\text{st}}(u), m_i^{\text{st}}(u), n_i^{\text{st}}(u) : i \in \mathbb{Z}) \quad \text{and} \quad (a_i^{\text{st}}(v), b_i^{\text{st}}(v), m_i^{\text{st}}(v), n_i^{\text{st}}(v) : i \in \mathbb{Z})$$

are standard ABMN solutions of Mina margin x . They are shift inequivalent because u is not equal to $s_i(v)$ for any $i \in \mathbb{Z}$. Indeed, the condition $s_i(v) \in [w, s_{-1}(w))$ implies that $i = 0$, but $s_0(v) = v \neq u$. This pair of solutions demonstrates that $Q(x) \geq 2$, as required to obtain Theorem 2.10. \square

We end this section by proving Theorem 2.24, in part as a consequence of Theorem 2.23(1). This makes now a convenient moment to derive the next result, thus rendering rigorous a verbal argument in the first paragraph of Section 2.3.

Proof of Proposition 2.17. To prove the two parts of this result, it is enough to argue that there is a unique standard solution, and a unique default solution, to which any positive ABMN solution is equivalent. Suppose then that $(a, b, m, n) = \{(a_i, b_i, m_i, n_i) \in (0, \infty)^2 \times \mathbb{R}^2 : i \in \mathbb{Z}\}$ is a positive ABMN solution. The boundary values $m_{-\infty}$ and n_{∞} exist as real numbers by Theorem 2.3(3). Note then that the translation $(a', b', m', n') = \tau_{-m_{-\infty}, -n_{\infty}}$ has $m'_{-\infty} = n'_{\infty} = 0$. Write $x = \frac{n'_{-1} - n'_0}{m'_0 - m'_{-1}}$ and $y = m'_{\infty} - m'_{-\infty}$. By applying the dilation τ_u to (a', b', m', n') , we obtain a default solution if $u = x^{-1}$ and a standard solution if $u = y^{-1}$. We have seen that $\tau_{x^{-1}} \circ \chi_{-m_{-\infty}, -n_{\infty}}(a, b, m, n)$ is a default ABMN solution. It is clear that any variation of the parameters $(x^{-1}, m_{-\infty}, n_{\infty})$ will result in an ABMN solution that fails to be default. Likewise, $\tau_{y^{-1}} \circ \chi_{-m_{-\infty}, -n_{\infty}}(a, b, m, n)$ has been shown to be a standard ABMN solution. Any variation of $(y^{-1}, m_{-\infty}, n_{\infty})$ will result in an ABMN solution that fails to be standard. Thus we complete the proof of Proposition 2.17. \square

Proof of Theorem 2.24(1). By Theorem 2.6, a time-invariant Nash equilibrium in $\text{Standard}(x)$ is the reverse-ordered (a, b) -component of a standard ABMN solution whose Mina margin equals x . Since Proposition 2.17 implies that standard ABMN solutions are indexed by the value $z \in (0, \infty)$ of their CenRatio, we obtain Theorem 2.24(1).

(2). By Theorem 2.23(1) and Proposition 3.4, $\mathcal{M}(s_k(x))$ equals $\mathcal{M}(x)$ for all $x \in (0, \infty)$ and $k \in \mathbb{Z}$. This, the set X is the disjoint union of $s_k(Y)$ as k ranges over \mathbb{Z} . Proposition 2.26 then yields Theorem 2.24(2). \square

6.2. The Mina margin map after domain coordinate change

In this section, we prove Theorem 2.28. The map $\Theta : \mathbb{R} \rightarrow (0, \infty)$ is an increasing surjection, so we may set $\Psi = \Theta^{-1} : (0, \infty) \rightarrow \mathbb{R}$. The proof of the theorem will harness the next result.

LEMMA 6.14. *There exists a constant $C > 0$ such that $|\theta(x) - \Psi(x)| \leq C$ for $x \geq 1/3$.*

Two further results will serve to prove Lemma 6.14.

LEMMA 6.15. *We have that*

$$\Psi(x) = \begin{cases} \log_2(\log_2(x) + 1) & \text{if } x \in [1, \infty), \\ -\log_2(-\log_2(x) + 1) & \text{if } x \in (0, 1). \end{cases}$$

Proof. The formulas follow from the expressions 2^{2^x-1} and $2^{-(2^{-x}-1)}$ for $\Theta(x)$ that are respectively valid when $x \geq 0$ and $x < 0$. \square

For $x \in [1/3, 3]$ and $i \in \mathbb{N}$, we write $s_{-i}(x)$ in the form $2^{2^i c_i(x)-1}$.

LEMMA 6.16. *There exists a constant $C > 0$ and a function $c : [1/3, 3] \rightarrow (0, \infty)$ such that the bound $|c_i(x) - c(x)| \leq C2^{-i}$ holds for $x \in [1/3, 3]$ and $i \in \mathbb{N}$. Further, $\inf \{c(x) : x \in [1/3, 3]\} > 0$.*

Proof. From the relation

$$s_i(x) = 2s_{-(i-1)}(x)^2 + O(s_{-(i-1)}(x))$$

and the form $s_{-j}(x) = 2^{2^j c_j - 1}$ (where we write $c_j = c_j(x)$), we find that

$$2^{2^j c_j - 1} = 2^{2^j c_{j-1}(x) - 1} + O(2^{2^{j-1} c_{j-1}(x) - 1}) = 2^{2^j c_{j-1}(x) - 1} (1 + 2^{-2^{j-1} c_{j-1}(x)})$$

so that

$$2^j c_j = 2^j c_{j-1} - 1 + \Theta(1)2^{-2^{j-1} c_{j-1}}$$

and

$$c_j = c_{j-1} - 2^{-j} + \Theta(1)2^{-j-2^{j-1} c_{j-1}} = c_{j-1} + \Theta(1)2^{-j}.$$

We thus learn that there exists $c = c(x) \in [0, \infty)$ such that $|c_j - c| \leq O(1)2^{-j}$. We may exclude the possibility that c equals zero because, in this case, we would have that $c_j \leq O(1)2^{-j}$, which would imply the false assertion that $s_{-j}(x) = 2^{2^j c_j - 1} = O(1)$ is bounded above independently of $j \in \mathbb{N}$.

We now argue that $\inf \{c(x) : x \in [1/3, 3]\} > 0$. This follows from $c(1/3) > 0$ and the fact that $s_{-i}(x)$ is increasing in $x \in [1/3, 3]$ for each $i \in \mathbb{N}$. This completes the proof of Lemma 6.16. \square

Proof of Lemma 6.14. Note that $s_{-i}(x) \geq 3$ for $i \in \mathbb{N}_+$ and $x \in [1/3, 3]$. By Lemma 6.15 and $s_{-i}(x) = 2^{2^i c_i(x) - 1}$, we find that $\Psi(s_{-i}(x)) = i + \log_2 c_i$. Indeed, we find that

$$|\Psi(s_{-i}(x)) - i| = \log_2(c_i/c) = \log(1 + \frac{c_i - c}{c}) \leq O(1)2^{-i},$$

the inequality due to Lemma 6.16.

On the other hand, $\theta(s_{-i}(x))$ is equal to the unique value $J \in \mathbb{Z}$ such that $s_J(s_{-i}(x)) \in [1/3, 3]$. When $x \in [1/3, 3)$, we see then that $J = i$.

We find then that $|\Psi(x) - \theta(x)| = O(1)$ for $x \geq 1/3$, as we sought to do in proving Lemma 6.14. \square

Proof of Theorem 2.28(1,2). We first claim that

$$s_k(\theta^{-1}(x+k)) = \theta^{-1}(x). \tag{56}$$

To check this, note that $\theta(s_k(z)) = \theta(z) - k$, so that

$$\theta(s_k(\theta^{-1}(x+k))) = \theta(\theta^{-1}(x+k)) - k = (k+x) - x = x,$$

as desired.

Note then that

$$\psi(x) = \mathcal{M}(\theta^{-1}(x)) = \mathcal{M}(s(\theta^{-1}(x+1))) = \mathcal{M}(\theta^{-1}(x+1)) = \psi(x+1)$$

where the respective equalities are due to the definition of ψ ; the above claim with $k = 1$; Theorem 2.23(1); and the definition of ψ once more. We have obtained Theorem 2.28(1).

Note next that $\mathcal{S}_k \text{StSol}(x+k)$ equals

$$\begin{aligned} & \mathcal{S}_k \left(a^{\text{st}}(\theta^{-1}(x+k)), b^{\text{st}}(\theta^{-1}(x+k)), m^{\text{st}}(\theta^{-1}(x+k)), n^{\text{st}}(\theta^{-1}(x+k)) \right) \\ &= \left(a^{\text{st}} \left(s_k(\theta^{-1}(x+k)) \right), b^{\text{st}} \left(s_k(\theta^{-1}(x+k)) \right), m^{\text{st}} \left(s_k(\theta^{-1}(x+k)) \right), n^{\text{st}} \left(s_k(\theta^{-1}(x+k)) \right) \right), \end{aligned}$$

the latter equality by Proposition 2.26. Applying (56), we find that

$$\mathcal{S}_k \text{StSol}(x+k) = \left(a^{\text{st}}(\theta^{-1}(x)), b^{\text{st}}(\theta^{-1}(x)), m^{\text{st}}(\theta^{-1}(x)), n^{\text{st}}(\theta^{-1}(x)) \right) = \text{StSol}(x).$$

This implies that $\text{StSol}(x+k) = \mathcal{S}_{-k} \text{StSol}(x)$, which is what Theorem 2.28(3) asserts.

(3). Let $z \in \mathbb{R}$ and set $\Theta(z) = x$. Since $|\Psi(x) - \theta(x)| = O(1)$ by Lemma 6.14, and Ψ and θ are increasing, we have that

$$\Theta(z - O(1)) \leq \theta^{-1}(z) \leq \Theta(z + O(1)).$$

Substituting the expressions for 2^{2^z-1} and $2^{-(2^{-z}-1)}$ for $\Theta(z)$, valid when $z \geq 0$ and $z < 0$, we obtain Theorem 2.28(3). \square

6.3. The Mina margin map is not identically equal to one

Here we prove Theorem 2.4(3). We will obtain evidence for Conjecture 2.5 as we do so.

PROPOSITION 6.17. *The value of $\mathcal{M}_{5,4}(0.58)$ lies in the interval $[0.9999032032, 0.9999032038]$.*

Proof of Theorem 2.4(3). Since λ is equal to the infimum of $\mathcal{M}(x)$ over $x \in (0, \infty)$, we have that $\lambda \leq \mathcal{M}(0.58)$.

Note that Proposition 6.4 with $(\ell, k) = (5, 4)$ implies that the value $\mathcal{M}(z) = \lim_j \mathcal{M}_{j,j}(z)$ (where $z = 0.58 \in [1/3, 3]$) satisfies

$$|\mathcal{M}(z) - \mathcal{M}_{5,4}(z)| \leq 3.3 \times 10^{-8} + 5.95 \times 10^{-7} \leq 6.3 \times 10^{-7}.$$

Applying the upper bound on $\mathcal{M}_{5,4}(z)$ in Proposition 6.17, we find that

$$\mathcal{M}(0.58) \leq 0.9999032038 + 6.3 \times 10^{-7} = 0.9999038338.$$

We confirm then that λ , being at most $\mathcal{M}(0.58)$, is bounded above by 0.999904. This completes the proof of Theorem 2.4(3). \square

Numerical work with Mathematica indicates that $\mathcal{M}_{5,4}(0.5809)$ equals 0.999903202726 to twelve decimal places; that $\mathcal{M}_{5,4}(0.5809)$ equals $\min \{ \mathcal{M}_{5,4}(x) : x \in [1/3, 3] \cap 10^{-4}\mathbb{Z} \}$; and that the error between $\inf \{ \mathcal{M}_{5,4}(x) : x \in [1/3, 3] \}$ and this minimum may jeopardise only the final digit of the twelve. If this evidence is admitted, then the preceding proof yields that $\inf \{ \mathcal{M}(x) : x \in [1/3, 3] \}$ is at least $\mathcal{M}_{5,4}(0.5809) - 10^{-11} - 6.3 \times 10^{-7} \geq 0.99990257 \geq 0.999902$; whence Conjecture 2.5. In fact, the conjecture is cautious: λ is likely to exceed 0.999903, as an estimate on a higher indexed $\mathcal{M}_{\ell,k}$ might show.

The formula for $(0, \infty) \rightarrow \mathbb{R} : x \rightarrow \mathcal{M}_{5,4}(x)$ in Lemma 6.3 may be recorded explicitly—it involves several applications of such operations as inverse and square-root—but it is messy, and would occupy several pages of standard print. Arguably a claim that mathematical software evaluates this function at 0.58 to be within the range claimed by Proposition 6.17 may be admitted as a proof of

this result. But a diligent reader who is given this information has no practical way to confirm it. In the following proof, we provide an approximation scheme, from above and below, for computing $\mathcal{M}_{5,4}(0.58)$. All quantities in the scheme are values in $10^{-10}\mathbb{Z}$, and the proof is reduced to verifying about fifty explicit statements of the form ‘if $x = u$, then $f(x) = v$ ’, where u and v are given elements of $10^{-10}\mathbb{Z}$, and f is the application of a function such as s , c and d from Definition 2.18 followed by a rounding down or up on to the lattice $10^{-10}\mathbb{Z}$. In this way, the diligent reader has a mundane but manageable task to verify every detail of the derivation of Proposition 6.17.

We note that, were $\mathcal{M}_{5,4}$ shown to be differentiable, and a suitable bound on its derivative found, then a similarly explicit record of the values of $\mathcal{M}_{5,4}$ on a fine enough mesh of points in $[1/3, 3]$ would furnish a proof of Conjecture 2.5. If the number of points in the mesh were large, then a manual check on the explicit bounds would be impracticable, so that in such a case the proof would be at least modestly computer-assisted.

We now turn to introducing and implementing the approximation scheme. Let $k \in \mathbb{N}$, and set

$$\lfloor x \rfloor_k = 10^{-k} \lfloor 10^k x \rfloor \quad \text{and} \quad \lceil x \rceil_k = 10^{-k} \lceil 10^k x \rceil + 10^{-k} \in \mathbb{R}.$$

Namely, the real line is partitioned

$$\mathbb{R} = \bigcup_{j \in \mathbb{Z}} 10^{-k} \cdot [j, j + 1)$$

into intervals whose endpoints are consecutive elements in the lattice $10^{-k}\mathbb{Z}$; $[\lfloor x \rfloor_k, \lceil x \rceil_k)$ is the unique interval in the partition that contains x . From the outset, we set the parameter k equal to ten, and omit to denote it. It should thus be understood that $\lfloor x \rfloor$ and $\lceil x \rceil$ denote $\lfloor x \rfloor_{10}$ and $\lceil x \rceil_{10}$, rather than the usual integer roundings of x .

Recall Definition 2.18. We specify $s^\uparrow, s^\downarrow, c^\uparrow, c^\downarrow, d^\uparrow, d^\downarrow : (0, \infty) \rightarrow (0, \infty)$ by setting

$$*^\uparrow(x) = \lceil *(x) \rceil \quad \text{and} \quad *^\downarrow(x) = \lfloor *(x) \rfloor \quad \text{for } * \in \{s, c, d\}.$$

For $x \in (0, \infty)$, we specify $\{s_i^\uparrow(x) : i \in \mathbb{Z}\}$ and $\{s_i^\downarrow(x) : i \in \mathbb{Z}\}$, the upper and lower s -sequences evaluated at x . Indeed, we set $s_0^\uparrow(x) = s_0^\downarrow(x) = x$. For $i \geq 1$, we then iteratively set $s_i^\uparrow(x) = s^\uparrow(s_{i-1}^\uparrow(x))$ and $s_i^\downarrow(x) = s^\downarrow(s_{i-1}^\downarrow(x))$. We further set $s_{-1}^\uparrow(x) = \lceil s_{-1}(x) \rceil$ and $s_{-1}^\downarrow(x) = \lfloor s_{-1}(x) \rfloor$. For $i \leq -2$, we iteratively set $s_{-i}^\uparrow(x) = s_{-1}^\uparrow(s_{-i+1}^\uparrow(x))$ and $s_{-i}^\downarrow(x) = s_{-1}^\downarrow(s_{-i+1}^\downarrow(x))$.

Set $z = 0.58$.

We will write $s_i^\uparrow = s_i^\uparrow(z)$ and $s_i^\downarrow = s_i^\downarrow(z)$ for $i \in \mathbb{Z}$. In this way, the value of $z = 0.58$ is understood.

We further define the upper and lower c - and d -sequences, $\{c_i^\uparrow, c_i^\downarrow, d_i^\uparrow, d_i^\downarrow : i \in \mathbb{Z}\}$, where again the value of z is understood. We set $c_i^\uparrow = c^\uparrow(s_i^\uparrow)$, $c_i^\downarrow = c^\downarrow(s_i^\downarrow)$, $d_i^\uparrow = d^\uparrow(s_i^\uparrow)$ and $d_i^\downarrow = d^\downarrow(s_i^\downarrow)$.

The data $\{s_i^\uparrow, s_i^\downarrow, c_i^\uparrow, c_i^\downarrow, d_i^\uparrow, d_i^\downarrow\}$, $i \in \llbracket -4, 3 \rrbracket$, are forty-eight elements of $10^{-10}\mathbb{Z}$. These values are presented in Tables 1 and 2. Two of the values are known without computation: $s_0^\uparrow = s_0^\downarrow = 0.58$. The remaining values may be computed, one at a time, where each step is a computation INPUT \rightarrow OUTPUT of one element of the lattice $10^{-10}\mathbb{Z}$ from another. Each step takes one of the forms $s_i^\uparrow \rightarrow s_{i+1}^\uparrow$ for $i \in \llbracket 0, 2 \rrbracket$; $s_i^\uparrow \rightarrow s_{i-1}^\uparrow$ for $i \in \llbracket -3, 0 \rrbracket$; $s_i^\uparrow \rightarrow c_i^\uparrow$ or $s_i^\uparrow \rightarrow d_i^\uparrow$ for $i \in \llbracket -4, 3 \rrbracket$; or it is formed by replacing \uparrow by \downarrow in one of these steps. Forty-six such steps lead to the completion of the two tables, given the two initial entries.

i	s_i^\uparrow	s_i^\downarrow
-4	954911606.03	954911605.92
-3	21848.5122538904	21848.5122525938
-2	102.3071054647	102.3071054616
-1	5.3556473847	5.3556473846
0	0.5800000000	0.5800000000
1	0.0504077253	0.0504077252
2	0.0010408205	0.0010408204
3	0.0000005392	0.0000005391

TABLE 1. Values of s_i^\uparrow and s_i^\downarrow for $i \in \llbracket -4, 3 \rrbracket$

i	c_i^\uparrow	c_i^\downarrow	d_i^\uparrow	d_i^\downarrow
-4	477488579.78	477488579.73	10926.0060411432	10926.0060404948
-3	11081.6603248978	11081.6603242447	52.8859257466	52.8859257450
-2	62.5133614707	62.5133614689	4.2201465577	4.2201465576
-1	5.7859121540	5.7859121538	1.5182994418	1.5182994417
0	1.8055756566	1.8055756565	1.0700124766	1.0700124765
1	1.0944264319	1.0944264316	1.0019497202	1.0019497201
2	1.0020784046	1.0020784043	1.0000010767	1.0000010766
3	1.0000010785	1.0000010782	1.0000000001	1.0000000000

TABLE 2. Values of c_i^\uparrow , c_i^\downarrow , d_i^\uparrow and d_i^\downarrow for $i \in \llbracket -4, 3 \rrbracket$

According to Lemma 6.3,

$$\mathcal{M}_{5,4}(z) = \frac{z(S_4 + T_5)}{P_4 + Q_5},$$

where

$$P_4 = 1 + (c_0 - 1) + (c_0 - 1)(c_1 - 1) + (c_0 - 1)(c_1 - 1)(c_2 - 1) + (c_0 - 1)(c_1 - 1)(c_2 - 1)(c_3 - 1);$$

$$Q_5 = (c_{-1} - 1)^{-1} + (c_{-1} - 1)^{-1}(c_{-2} - 1)^{-1} + (c_{-1} - 1)^{-1}(c_{-2} - 1)^{-1}(c_{-3} - 1)^{-1} \\ + (c_{-1} - 1)^{-1}(c_{-2} - 1)^{-1}(c_{-3} - 1)^{-1}(c_{-4} - 1)^{-1};$$

$$S_4 = 1 + (d_0 - 1) + (d_0 - 1)(d_1 - 1) + (d_0 - 1)(d_1 - 1)(d_2 - 1) + (d_0 - 1)(d_1 - 1)(d_2 - 1)(d_3 - 1);$$

and

$$T_5 = (d_{-1} - 1)^{-1} + (d_{-1} - 1)^{-1}(d_{-2} - 1)^{-1} + (d_{-1} - 1)^{-1}(d_{-2} - 1)^{-1}(d_{-3} - 1)^{-1} \\ + (d_{-1} - 1)^{-1}(d_{-2} - 1)^{-1}(d_{-3} - 1)^{-1}(d_{-4} - 1)^{-1}.$$

We further specify quantities $*^\uparrow$ and $*^\downarrow$, where $*$ $\in \{P_4, Q_5, S_4, T_5\}$. To do so, we record variable dependence in the form $P_4 = P_4(c_0, c_1, c_2, c_3)$, $S_4 = S_4(d_0, d_1, d_2, d_3)$, $Q_5 = Q_5(c_{-1}, c_{-2}, c_{-3}, c_{-4})$ and $T_5 = T_5(d_{-1}, d_{-2}, d_{-3}, d_{-4})$. We may then set

$$P_4^\uparrow = [P_4(c_0^\uparrow, c_1^\uparrow, c_2^\uparrow, c_3^\uparrow)] \quad \text{and} \quad P_4^\downarrow = [P_4(c_0^\downarrow, c_1^\downarrow, c_2^\downarrow, c_3^\downarrow)]; \\ S_4^\uparrow = [S_4(d_0^\uparrow, d_1^\uparrow, d_2^\uparrow, d_3^\uparrow)] \quad \text{and} \quad S_4^\downarrow = [S_4(d_0^\downarrow, d_1^\downarrow, d_2^\downarrow, d_3^\downarrow)]; \\ Q_5^\uparrow = [Q_5(c_{-1}^\uparrow, c_{-2}^\uparrow, c_{-3}^\uparrow, c_{-4}^\uparrow)] \quad \text{and} \quad Q_5^\downarrow = [Q_5(c_{-1}^\downarrow, c_{-2}^\downarrow, c_{-3}^\downarrow, c_{-4}^\downarrow)];$$

and

$$T_5^\uparrow = [T_5(c_{-1}^\downarrow, c_{-2}^\downarrow, c_{-3}^\downarrow, c_{-4}^\downarrow)] \quad \text{and} \quad T_5^\downarrow = [T_5(c_{-1}^\uparrow, c_{-2}^\uparrow, c_{-3}^\uparrow, c_{-4}^\uparrow)].$$

(Note the reversals of the uses of \downarrow and \uparrow in the replaced terms for Q_5 and T_5 .)

The tables then permit us to record the values (all of which are elements of the lattice $10^{-10}\mathbb{Z}$)

$$\begin{aligned} S_4^\uparrow &= 1.0701489815, & S_4^\downarrow &= 1.0701489813 \\ T_5^\uparrow &= 2.5400964392, & T_5^\downarrow &= 2.5400964386 \\ P_4^\uparrow &= 1.8818013910, & P_4^\downarrow &= 1.8818013906 \\ Q_5^\uparrow &= 0.2123436589, & Q_5^\downarrow &= 0.2123436587. \end{aligned} \tag{57}$$

Next we specify two further elements of $10^{-10}\mathbb{Z}$:

$$\mathcal{M}_{5,4}^\uparrow(z) = \left\lceil \frac{z(S_4^\uparrow + T_5^\uparrow)}{P_4^\downarrow + Q_5^\downarrow} \right\rceil \quad \text{and} \quad \mathcal{M}_{5,4}^\downarrow(z) = \left\lfloor \frac{z(S_4^\downarrow + T_5^\downarrow)}{P_4^\uparrow + Q_5^\uparrow} \right\rfloor. \tag{58}$$

LEMMA 6.18. *Let $i \in \mathbb{Z}$.*

- (1) $s_i^\downarrow \leq s_i \leq s_i^\uparrow$.
- (2) $c_i^\downarrow \leq c_i \leq c_i^\uparrow$.
- (3) $d_i^\downarrow \leq d_i \leq d_i^\uparrow$.
- (4) $P_4 \downarrow \leq P_4 \leq P_4^\uparrow$, $Q_4 \downarrow \leq Q_4 \leq Q_4^\uparrow$, $S_4 \downarrow \leq S_4 \leq S_4^\uparrow$ and $T_4 \downarrow \leq T_4 \leq T_4^\uparrow$.
- (5) $\mathcal{M}_{5,4}^\downarrow \leq \mathcal{M}_{5,4} \leq \mathcal{M}_{5,4}^\uparrow$.

Proof. Note that $s^\downarrow(x) \leq s(x) \leq s^\uparrow(x)$ for $x \in (0, \infty)$ by the definitions of s^\downarrow and s^\uparrow . By induction on $i \geq 1$, we will show that $s_i^\uparrow \geq s_i$. Indeed, note that $s_i^\uparrow = s^\uparrow(s_{i-1}^\uparrow) \geq s(s_{i-1}^\uparrow) \geq s(s_{i-1}) = s_i$, where the latter inequality is due to the inductive hypothesis at index $i-1$ and to Lemma 5.3(1:s). We also prove that $s_{-i}^\uparrow \geq s_{-i}$ for $i \geq 1$ by induction on i . In this regard, note that $s_{-i-1} = s_{-1}^\uparrow(s_{-i}^\uparrow) \geq s_{-1}(s_{-i}^\uparrow) \geq s_{-1}(s_{-i}) = s_{-1-i}$, where the first bound is due to $s_{-1}^\uparrow(x) \geq s_{-1}(x)$ for $x \in (0, \infty)$, which follows from the definition of s_{-1}^\uparrow ; the second is due to the inductive hypothesis at index i and $x \rightarrow s_{-1}(x)$ being increasing, which fact follows from Proposition 3.4 and Lemma 5.3(1:s).

Similar arguments prove that $s_i^\downarrow \leq s_i$ for $i \in \mathbb{Z}$.

(2). Note that $c^\uparrow(s_i^\uparrow) \geq c(s_i^\uparrow) \geq c(s_i) = c_i$, where the first bound is due to the definition of c^\uparrow and the second is due to $s_i^\uparrow \geq s_i$ and Lemma 5.3(1:c). Similarly may we show that $c_i^\downarrow \leq c_i$.

(3). This is similar to the preceding part.

(4). Note that P_4 is an increasing function of the variables c_i , $i \in \llbracket 0, 3 \rrbracket$; Q_5 is decreasing in c_{-i} , $i \in \llbracket 1, 4 \rrbracket$; S_4 is increasing in d_i , $i \in \llbracket 0, 3 \rrbracket$; and T_5 is decreasing in d_{-i} , $i \in \llbracket 1, 4 \rrbracket$. (The noted properties of Q_5 and T_5 are valid only insofar as the variables c_{-i} and d_{-i} remain greater than one. But this condition is always met in applications, including the present one.) Given these monotonicities, Lemma 6.18(2) shows that

$$P_4(c_0^\downarrow, c_1^\downarrow, c_2^\downarrow, c_3^\downarrow) \leq P_4 \leq P_4(c_0^\uparrow, c_1^\uparrow, c_2^\uparrow, c_3^\uparrow),$$

so that the monotonicities of $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ prove the assertions concerning P_4 . The derivation for Q_5 is similar. So are the others: Lemma 6.18(3) is used in regard to S_4 ; and Lemma 6.18(4) for T_5 .

(5). The expression $\mathcal{M}_{5,4}$ is an increasing function of S_4 and T_5 , and it is decreasing in P_4 and Q_5 —thus, we may use the preceding part to reach the desired conclusion. \square

Proof of Proposition 6.17. Using the data (57), note that the expressions (58) have the evaluations

$$\mathcal{M}_{5,4}^\uparrow = \left\lceil 0.58 \times \frac{1.0701489815 + 2.5400964392}{1.8818013906 + 0.2123436587} \right\rceil = 0.9999032038$$

and

$$\mathcal{M}_{5,4}^\downarrow = \left\lfloor 0.58 \times \frac{1.0701489813 + 2.5400964386}{1.8818013910 + 0.2123436589} \right\rfloor = 0.9999032032.$$

By Lemma 6.18(5), we learn that $\mathcal{M}_{5,4} \in [0.9999032032, 0.9999032038]$, as Proposition 6.17 states. \square

Trail prospects

We discuss five topics prompted by the article.

7.1. Properties of the Mina margin map and prospective routes to conjectures

Conjecture 2.11 concerns the level sets of the Mina margin map, and via (15), Conjecture 2.30 concerns the level sets of the finite trail counterparts to this map. Consider the Θ -transformed finite-trail Mina margin maps $\mathcal{M}_{j+1,k+1} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$ depicted for several pairs $(j+1, k+1)$ in Figure 2.3. In these sketches, there are a total of $2(j+k) - 5$ elements in any level set through which every swerve of the function passes; such level sets are indexed by $[\lambda, \lambda^{-1}]$ for $\lambda = 0.999903 \dots$ up to an error that vanishes in high j and k ; the functions converge to a limit $\mathcal{M} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$, and this limit has level sets with two elements in each period (such as in $\Theta^{-1}(1/3, 3]$) for heights in (λ, λ^{-1}) . These claims constitute the content of the two conjectures and they can be said to be visually more-or-less evident. But can they be proved? In regard to Conjecture 2.11 at least, control on derivatives and explicit evaluation on a suitably fine mesh may be a tractable approach: see the discussion regarding Conjecture 2.5 in Section 6.3.

7.2. The possible existence of further Nash equilibria

We have studied time-invariant Nash equilibria. It is natural to ask whether further Nash equilibria exist. We discuss two directions.

7.2.1. Time-invariant random Nash equilibria. Our formulation of the notion of Nash equilibrium in Section 2.1 is deterministic. What if time-invariant random play is permitted? A strategy would consist of a set of laws on the non-negative reals indexed by \mathbb{Z} . When such a strategy is played, the stake offered would be sampled from the law indexed by the present counter location, the sampling being independent of other randomness. To avoid extra notation, we have not formulated this notion in the main part of the article. We do not believe that non-trivial random time-invariant Nash equilibria exist. Indeed, we remarked after the Penny Forfeit Lemma 3.5 that random play is suboptimal for the one-step game. By iterating this result and invoking the monotonicity in Penny Forfeit argued in the proof of Lemma 4.16(2), the possibility of a non-trivial role for randomness of the form we have discussed may be excluded.

7.2.2. Nash equilibria that are not time-invariant. A deterministic strategy pair that may not be time-invariant takes the form $(b, a) : \mathbb{Z} \times \mathbb{N}_+ \rightarrow (0, \infty)^2$. We may anticipate that, were such a pair a Nash equilibrium, the naturally associated dynamical quadruple (a, b, m, n) , specified by suitably modifying Definition 2.1, would satisfy a dynamical form dABMN of the ABMN system on \mathbb{Z} . For simplicity, we describe these equations on a finite trail $\llbracket -K - 1, K + 1 \rrbracket$ and for a finite time interval $\llbracket 0, T \rrbracket$ (so that $K + 1, T \in \mathbb{N}_+$). Boundary data is a quadruple $(m_{-K}, m_K, n_K, n_{-K}) \in \mathbb{R}^4$ which equals $(0, 1, 1, 0)$ in the simple symmetric case; and two terminal functions $m_{\text{ter}}, n_{\text{ter}} : \llbracket -K - 1, K + 1 \rrbracket \rightarrow \mathbb{R}$. If write $*_i(j) = *(i, j)$ for $* \in \{a, b, m, n\}$, so that, for

example, $a_i(j)$ is the stake offered by Maxine at the j^{th} turn in the event that $X_{j-1} = i$, the revised equations are

$$(a_i(j) + b_i(j))(m_i(j) + a_i(j)) = a_i(j)m_{i+1}(j+1) + b_i(j)m_{i-1}(j+1) \quad \text{dABMN(1)}$$

$$(a_i(j) + b_i(j))(n_i(j) + b_i(j)) = a_i(j)n_{i+1}(j+1) + b_i(j)n_{i-1}(j+1) \quad \text{dABMN(2)}$$

$$(a_i(j) + b_i(j))^2 = b_i(j)(m_{i+1}(j+1) - m_{i-1}(j+1)) \quad \text{dABMN(3)}$$

$$(a_i(j) + b_i(j))^2 = a_i(j)(n_{i-1}(j+1) - n_{i+1}(j+1)) \quad \text{dABMN(4)},$$

where (i, j) ranges over $\llbracket -K, K \rrbracket \times \llbracket 0, T-1 \rrbracket$. Boundary conditions enter via

$$\begin{aligned} m_{\pm(K+1)}(j) &= m_{\pm(K+1)}, \quad n_{\pm(K+1)}(j) = n_{\pm(K+1)} \text{ for } j \in \llbracket 1, K \rrbracket; \text{ and} \\ m_i(T) &= n_{\text{term}}(i), \quad n_i(T) = n_{\text{term}}(i) \text{ for } i \in \llbracket -K, K \rrbracket. \end{aligned}$$

Of course, any ABMN solution (a, b, m, n) solves dABMN if we extend notation to set $*_i(j) = *_i$ for $(i, j) \in \mathbb{Z} \times \mathbb{N}_+$ and $* \in \{a, b, m, n\}$ (and then restrict the domain suitably). Do other solutions of dABMN exist?

Certainly there are some such. Points (i, j) in $\mathbb{Z} \times \mathbb{N}_+$ are odd or even according to whether $i + j$ is odd or even. The parity of $j + X_i(j)$ never changes from its initial $j = 0$ value in any instance of the trail game. If we select two solutions (a', b', m', n') and $(\hat{a}, \hat{b}, \hat{m}, \hat{n})$ of the ABMN equations, and set

$$(a, b, m, n)(i, j) = \begin{cases} (a'_i, b'_i, m'_i, n'_i) & \text{when } i + j \text{ is even,} \\ (\hat{a}_i, \hat{b}_i, \hat{m}_i, \hat{n}_i) & \text{when } i + j \text{ is odd,} \end{cases}$$

then (a, b, m, n) solves dABMN (when suitably restricted in the domain) and Theorem 2.6 directly implies that (b, a) is a Nash equilibrium in the trail game. Conceptually, this is not really a new solution, however. Gameplay resides on the odd or even lattice and use of any of these new Nash equilibria will coincide with that of a time-invariant Nash equilibrium in any given game.

For $k \in \mathbb{N}_+$, the system dABMN can be solved on $\llbracket -K, K \rrbracket \times \llbracket 1, k \rrbracket$ by choosing a given terminal condition $\{m_i(k), n_i(k) : i \in \llbracket -K-1, K+1 \rrbracket\}$ and iteratively solving dABMN for decreasing values of j . In searching for a solution that is not invariant in time, we seek a terminal condition such that, if this condition is imposed even for a very high value of k , the solved solution stabilises for bounded values of j to a form that is not time-invariant. We have tested a few terminal conditions; Figure 7.1 depicts a solution of dABMN on the trail $\llbracket -8, 8 \rrbracket$. The m -component of the terminal condition, defined on this trail, rises sharply at both ends, and otherwise has the form of a rough plateau of height one-half; the n -component is the reflection of the m -component in the vertical axis. The region of each sharp rise for the m -component represents a battlefield which may be rather stable as time evolves. As the two plots in Figure 7.1 show, the two battlefields rapidly approach one another by a short distance, and then remain in a rough stasis in which a gradual movement towards the origin is discernible, before rapidly breaking towards a single central battlefield. A total of 4200 time-steps are involved in the simulation, with the $j = 140$ black curve in each plot (which is the penultimate depicted in the backwards-in-time evolution) showing an eruption towards the centre, and the final black curve in each plot adopting the central location which is in essence the fixed point of the evolution. Battlefield pairs with greater separation may endure much longer, and may present a metastable effect for dABMN that causes these equations to converge very slowly to fixed points. There is no evidence however of non-convergence: our limited investigation has not produced examples that support the notion that further time-invariant Nash equilibria exist beyond the simple parity-based class discussed above.

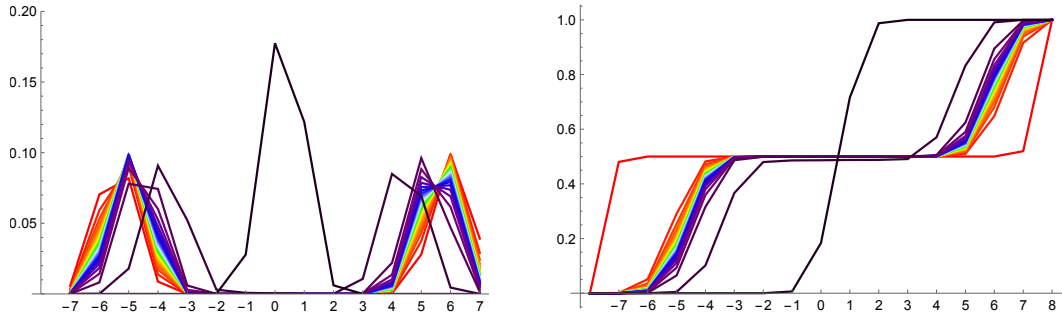


FIGURE 7.1. The dynamic ABMN equations dABMN on trail $\llbracket -8, 8 \rrbracket$ and time interval $\llbracket 0, 4200 \rrbracket$ are solved with $m_{\text{ter}} : \llbracket -8, 8 \rrbracket \rightarrow [0, 1]$, $m_{\text{ter}}(-8) = 0$, $m_{\text{ter}}(8) = 1$, rising sharply from zero to run along a rough plateau at height one-half, and ending with a further sharp rise to height one. We set $n_{\text{ter}}(\cdot) = m_{\text{ter}}(\cdot)$ and work with a standard symmetric boundary quadruple $(0, 1, 1, 0)$. The red curve on the right plot, which is most exposed both on the left and the right, is m_{ter} . The solutions of the equations are depicted at values of j in $\llbracket 0, 4200 \rrbracket$ that are multiples of 140, so that thirty curves excepting the final condition are depicted in each plot. On the left, the a -components of dABMN on the open-play set $\llbracket -7, 7 \rrbracket$ for the j -values in question are shown (with linear interpolation between integers); on the right, the m -components on the trail $\llbracket -8, 8 \rrbracket$ with such interpolation are shown. The curves are coloured on a spectrum leading from red to black as time passes backwards. These curves make a staccato advance (with this flow of time) from the sides to the centre, with the final black curve in each plot, indexed by $j = 0$, representing a single battlefield around the origin. The b - and n -components are formed by reflecting the a - and m -components in the vertical axis.

7.3. Gameboards beyond \mathbb{Z}

By use of a setup involving directed graphs, self-funded stake-governed random-turn games derived from games such as Hex or Go may be considered. It would be of interest to determine for a suitable class of games whether some of the features of the Trail of Lost Pennies of \mathbb{Z} are present more generally. The central ratio $\frac{n_{-1}-n_0}{m_0-m_{-1}}$ is the ratio of changes in mean payoff for Mina and Maxine arising from Mina's victory at the first turn. This or similar quantities may be considered in suitable infinite games, permitting us to ask whether Theorem 2.8 generalizes to these games: do Nash equilibria exist precisely when the quantity lies in an interval of the form $[\lambda, \lambda^{-1}]$? Do these game-determined λ -values differ from one by a notably small but positive quantity, as this value for the trail game on \mathbb{Z} appears to differ by about 10^{-4} ? Do more general games exemplify as ours does the Battlefield Cyl Fog premise from Section 1.1.1, having one (or perhaps several) bounded battlefield regions of the space of configurations on the gameboard, specified by any given Nash equilibrium, in which players concentrate their stake expenditure, with the outcome of turns occurring therein being highly influential in the overall game? As we recalled in Section 1.2, Peres and Šunić [51] found a path decomposition that leads to rapid computation of λ -biased infinity harmonic functions. To address the last question, it may be natural at first to work with trees with suitable boundary data, and to seek to iteratively determine Nash equilibria for the player-funded stake-governed tug-of-war along elements in a suitable path decomposition. In this way, solutions

for this game in a broader class of graphs may decompose into components, these being trail games as we have analysed.

7.4. The trail game and “Incentive Inch, Outcome Mile”

We have suggested that the Trail of Lost Pennies advances the premise “Incentive Inch, Outcome Mile” presented in Section 1.1.1. Here we examine this point more closely, for the game on finite intervals and on \mathbb{Z} . The premise for the trail game on these graphs may be summarised as follows.

- (1) When the Mina margin exceeds $\lambda^{-1} + o(1)$, Mina is the dominant player, and is able to secure victory at little cost from any starting position (except perhaps a few near the right boundary).
- (2) When the Mina margin is less than $\lambda - o(1)$, it is Maxine who is dominant, and is likewise able (except near the left boundary).
- (3) In a short interval of Mina margin values of length of order 10^{-4} around one, namely $[\lambda + o(1), \lambda^{-1} - o(1)]$, there are many equilibria, and the dominance of one or other player as above arises from any starting vertex (except a few near the boundary) under a suitable equilibrium.

Here, $o(1)$ refers to a positive expression that vanishes in the limit of long finite intervals, and is zero in the case of \mathbb{Z} . The qualifications in parentheses should be omitted in the \mathbb{Z} -case. By ‘many’ in the third part, we mean $2n + \Theta(1)$ for a finite interval of length $n \geq 6$, and countably many for \mathbb{Z} .

To evaluate the premise as just presented in the case of finite trails and \mathbb{Z} , we begin by focusing on a Mina margin value that lies marginally in the regime of Mina’s claimed dominance for the gameboard $\llbracket -6, 6 \rrbracket$ (whose set of open play is $\llbracket -5, 5 \rrbracket$). The associated Θ -transformed Mina margin map $\mathcal{M}_{6,6} \circ \Theta : \mathbb{R} \rightarrow (0, \infty)$ is depicted in Figure 2.3(left). Select $z = 1 + 10^{-4}$, a value that lies slightly above λ^{-1} according to Conjecture 2.5 (so that the purple curve in Figure 2.3(left) has turned left off the highway to cross this height). Indeed, the equation $\mathcal{M}_{6,6}(x) = z$ is found (by some trail-and-error work in Mathematica) to have a unique solution in $x \in \mathbb{R}$, with this solution taking the form $x = 4.04493$ up to five decimal places. The corresponding standard solution $(a, b, m, n) : \llbracket -5, 5 \rrbracket \rightarrow (0, \infty)^2 \times \mathbb{R}^2$ is depicted in Figure 7.2.

The solution has central ratio $\frac{n-1-n_0}{m_0-m-1} = \Theta(x)$ equal to 46538 up to rounding error: the origin is comfortably to the left of the battlefield index. Indeed, this index lies at four, since $\phi_4 = 0.719 \dots \in [1/3, 3)$. Mina’s stake at vertex three is the greater, and she dominates staking, and mean payoffs, at vertices two and below. Were we to consider the analogous solution on longer gameboards $\llbracket -\ell, \ell \rrbracket$, $\ell > 6$, we would see that its battlefield index rises in ℓ , and that the region around the origin falls progressively further into the territory that Mina controls. In this territory, she vastly outbids Maxine, even though all stakes are tiny. The weak limit in high ℓ of the gameplay starting at the origin that is governed by the reverse-ordered (a, b) -component of the solution is likely to be a deterministic left-moving walk on \mathbb{Z} . The conclusion may seem to be that, when x exceeds λ^{-1} , the Trail of Lost Pennies on \mathbb{Z} has become uncompetitive because Mina’s position is too strong: she should, it appears, win without expenditure. And as usual likewise for Maxine in the opposing case, when x is less than λ . But care is needed in this interpretation. After all, the limit in high ℓ of the stakes offered near the origin is zero for both players, and the double-zero strategy will not gratify Mina’s ambition to win without cost. Overall, then, the limit from finite gameboards creates a sense of utter dominance for the player with a favourable value of the Mina margin, but our formal results are agnostic: as Theorem 2.8 shows, there are no time-invariant Nash equilibria in




FIGURE 7.2. The unique standard solution $(a, b, m, n) : \llbracket -5, 5 \rrbracket \rightarrow (0, \infty)^2 \times \mathbb{R}^2$ for the trail game on $\llbracket -6, 6 \rrbracket$ with Mina margin equal to $1 + 10^{-4}$ is depicted. The stakes (a_i, b_i) and expected mean payoffs (m_i, n_i) are displayed for $i \in \llbracket 0, 5 \rrbracket$ in the left and right charts. Note that the leftmost displayed data is indexed by the origin: the data to the left, indexed by $\llbracket -5, -1 \rrbracket$, is visually indistinguishable from the zero-indexed data in the two displays.

the infinite game; since this is the solution concept we have studied, our results offer no guidance to Mina as she prepares to play the trail game on \mathbb{Z} with boundary data specifying a value of the Mina margin that exceeds λ^{-1} . We expect “Incentive Inch, Outcome Mile” as presented above to be validated by lengthy finite trails; and we expect the third property, concerning the brief middle phase, to be valid for \mathbb{Z} ; but the properties concerning the phases of Mina’s or Maxine’s supposed dominance have no clear mathematical meaning for \mathbb{Z} .

7.5. The game of chicken in the Trail of Lost Pennies

Theorems 2.8 and 2.10 show that, for any $x \in (\lambda, \lambda^{-1})$, the trail game with boundary data $(m_{-\infty}, m_{\infty}, n_{-\infty}, n_{\infty}) = (0, 1, x, 0)$ has at least two distinct time-invariant Nash equilibria of any given integral battlefield index; if $x \in \{-\lambda, \lambda\}$, then there is at least one such. For any $x \in [-\lambda, \lambda]$, we may thus find an element $\mathcal{S}_0^2 \cap \mathcal{N}$ of battlefield index zero. For $k \in \mathbb{Z}$, we denote by $(S_-(k), S_+(k))$ the right-shift by k places of (S_-, S_+) . This is an element of $\mathcal{S}_0^2 \cap \mathcal{N}$ of battlefield index k . Suppose that the counter begins at the origin. Game outcomes under the strategy pairs $(S_-(k), S_+(k))$ become more favourable to Mina, and less favourable to Maxine, as the index k increases; as Theorem 2.15 indicates, the probability of victory for Maxine decays rapidly as k becomes positive. Suppose the game is about to begin, and players must commit to strategies. Mina may consider playing one of the strategies $S_-(k)$ for $k \in \mathbb{Z}$. If her opponent elects to play $S_+(k)$, then Mina would much prefer that the shared value of k be positive; Maxine would naturally prefer a negative choice. But the players must consider the case that the opposing player elects a different value of k . What then?

For simplicity, consider the symmetric game where $x = 1$. Suppose that $S_+ = a$ and $S_- = b$ in the usual notational abuse, where $a_i = b_{-i}$ for $i \in \mathbb{Z}$. (The choice $(a, b) = (a^{\text{st}}(3), b^{\text{st}}(3))$ meets this condition, as we will see in Proposition A.1(1).) Let $k \in \mathbb{N}_+$. Suppose that Mina chooses between the soft $S_-(-k)$ and the tough $S_-(k)$, while Maxine elects to play either the soft $S_+(k)$ or the tough $S_+(-k)$. By this restriction, we consider a two-person game where each player has two alternatives, and in Table 1, we depict mean payoffs in a two-by-two array whose rows index Mina’s choice, whose columns index Maxine’s, and each of whose coordinates contains a list of

	Maxine Soft : $S_+(k)$	Maxine Tough : $S_+(-k)$
Mina Soft : $S_-(-k)$	M, M	B, G
Mina Tough : $S_-(k)$	G, B	C, C

TABLE 1. Mina and Maxine choose between their components in two given Nash equilibria with battlefield indices $-k$ and k , for some $k \in \mathbb{N}_+$. The respective mean payoffs for Mina and Maxine for each of the four strategy pairs are recorded in each entry of the 2×2 array. The possible outcomes are $G = \text{Good}$, $M = \text{Medium}$, $B = \text{Bad}$ and $C = \text{Catastrophic}$.

Mina's and Maxine's mean payoffs when the indexing strategy pair is played. The good outcome G has value $1 - \exp\{-2^k O(1)\}$. The medium outcome M takes the form $1/2 - \exp\{-2^k O(1)\}$. The bad outcome B has value $\exp\{-2^k O(1)\}$. And the value C of the catastrophic outcome is ... minus infinity! We will illustrate how to obtain these assertions rather than present formal derivations.

The outcomes of G and B arise in the off-diagonal cases, where Nash equilibria are played, so that the claimed forms for G and B arise from Theorem 2.14 in the sense of the paragraph that follows this theorem. Consider the strategy pair (Soft = $S_-(-k)$, Soft = $S_+(k)$). At the first turn, Mina is playing k units to the right of her presumed location of the battlefield vertex, as if she has as good as lost already. But Maxine is playing k units to the left of where she is claiming the battlefield index to be, and also in effect nearly admits defeat. Maxine's and Mina's stakes are a_{-k} and b_k : both very small, but equal in our special case. So the first turn victor is chosen by the outcome of a fair coin toss. And this early winner will lose even one later move only with probability $\exp\{-2^k O(1)\}$ as the estimates in Theorem 2.14 show, because the victor's stakes rise and her opponent's fall as the counter moves closer to the victor's presumed battlefield location. In the case of (Tough = $S_-(k)$, Tough = $S_+(-k)$) play, a phenomenon opposite to the eventual unanimity of gameplay in Theorem 2.15 occurs. The counter location at late time has law approaching an equilibrium which heavily charges the origin and a few nearby sites. When the counter moves slightly to the left of the origin, it comes closer to Maxine's presumed battlefield index at $-k$ than it does to Mina's at k , so that Maxine's stake rises far higher than Mina's, and the counter is restored towards the origin. An opposing leftward force naturally acts on the counter when it is to the right of the origin. The implicit consensus against lengthy play in a bounded region discussed around Theorem 2.15 has been broken with double-tough play, and the players are trapped in an unending mutually destructive cycle.

In the classic game of chicken [52, Chapter 10], two players choose between soft and tough options of swerving or driving straight. When one player drives straight and the other swerves, their payoffs are the pleasure G of winning and the annoyance B of showing weakness. When both swerve, both receive an intermediate value M . When both drive straight, the shared outcome is a highly negative C as the cars crash. We see then that the Trail of Lost Pennies on \mathbb{Z} embeds the game of chicken. The translation symmetry of \mathbb{Z} makes the selection of which Nash equilibrium to play a difficult choice for players who may be infinitely punished for a perhaps unintentionally tough decision. The counterpart embedding of chicken occurs in the finite trail game, where the value of C , while often highly negative, is finite.

7.6. Play between people and algorithms

The finite trail $\llbracket -j, k \rrbracket$ —perhaps for values of j and k somewhere between one and five—may provide an attractive context for investigating how people or algorithms play the trail game. Given the smallness of $1 - \lambda$ and the multiplicity of Nash equilibria for many of the games with longer trails, it seems fanciful to believe that two people who play the same game repeatedly will typically adhere to such an equilibrium (at least when $j + k$ is high enough). Other strategies may seem natural.

7.6.1. Cooperative behaviour. Trust could be established during iterated play. If two players each stake zero throughout a standard symmetric trail game on $\llbracket -k, k \rrbracket$, $k \in \mathbb{N}_+$, whose counter starts at zero, their running costs are zero, and their mean payoffs are one-half (this is because play ends in finite time on a finite trail; we use the $0/0 = 1/2$ rule from Section 2.1).

7.6.2. Tit-for-tat. A consistent zero strategy has the flaw of being vulnerable to exploitation, and a player in an iterated game may prefer a tit-for-tat approach: stake zero in every game until the opponent makes a positive stake; in the next game, play more aggressively; revert to playing zero if the opponent reacts modestly to the aggressive play. Of course, there are degrees of aggression that may be adopted. The iterated prisoner's dilemma is a classic example where the Nash equilibrium (which proposes uncooperative play) often predicts wrongly how people will play, and where tit-for-tat and variants thereof are commonly adopted strategies for humans [17] that have been found in computer-against-computer tournaments to be effective [10].

7.6.3. The loadsamoney bully. On a finite trail, the loadsamoney bully chooses $\epsilon > 0$ small and consistently stakes ϵ against an opponent who stakes zero. He plays aggressively when the opponent makes a positive stake: he may react to a stake $a > 0$ by staking $2a$ at the next turn, for example. This player wins games against a zero-staking opponent while incurring almost no running cost. He seeks to rein financial terror on the opponent who deviates from a zero strategy, by seeming prepared to win the concerned game no matter what the cost. Hoping to create a sense of formidable financial resources, his long-term plan for the iterated game is to cow the opponent into a submissive zero strategy.

APPENDIX A

Proofs of the further formulas

A.1. Deriving alternative formulas for standard solutions and Mina margins

Proof of Theorem 2.31(1). In the notation of Definition 2.20, $m_1^{\text{def}}(x) - m_0^{\text{def}}(x) = c(x) - 1$ and $m_0^{\text{def}}(x) - m_{-1}^{\text{def}}(x) = 1$, so that $M_0(x) = c(x)$. And $n_0^{\text{def}}(x) - n_1^{\text{def}}(x) = x(d(x) - 1)$, while $n_{-1}^{\text{def}}(x) - n_0^{\text{def}}(x) = x$, which leads to $N_0(x) = xd(x)$. The formulas for $a_0^{\text{def}}(x)$ and $b_0^{\text{def}}(x)$ in terms of $M_0(x)$ and $N_0(x)$ yield that $a_0^{\text{def}}(x) = c(x)f(x)$ and $b_0^{\text{def}}(x) = xd(x)f(x)$. By (12) and Proposition 2.26, we obtain Theorem 2.31(1).

(2). By (12), $m_0^{\text{def}}(x) - m_{-1}^{\text{def}}(x) = 1$ and $n_{-1}^{\text{def}}(x) - n_0^{\text{def}}(x) = x$, we see that $m_0^{\text{st}}(x) - m_{-1}^{\text{st}}(x) = Z(x)^{-1}$ and $n_{-1}^{\text{st}}(x) - n_0^{\text{st}}(x) = xZ(x)^{-1}$. Proposition 2.26 then implies the special cases of Theorem 2.31(2) noted after the general cases, which are then derived by summation.

(3). This follows from (14) in the form $\mathcal{M}_{j+1,k+1}(x) = (n_{-j-1}^{\text{st}}(x) - n_{k+1}^{\text{st}}(x))(m_{k+1}^{\text{st}}(x) - m_{-j-1}^{\text{st}}(x))^{-1}$, and the preceding part.

(4). By Theorem 2.3(3), the numerator and denominator in the just obtained expression for $\mathcal{M}_{j,k}(x)$ converge in the limits $j, k \rightarrow \infty$ to $n_{-\infty}^{\text{st}}(x) - n_{\infty}^{\text{st}}(x) = n_{-\infty}^{\text{st}}(x)$ and $m_{\infty}^{\text{st}}(x) - m_{-\infty}^{\text{st}}(x) = 1$. Since the ratio of the latter quantities equals $\mathcal{M}(x)$, we obtain Theorem 2.31(4) from the limiting numerator. \square

Remark. By considering the limiting denominator in the argument for Theorem 2.31(4), we obtain the identity

$$\sum_{i \in \mathbb{Z}} Z(s_i(x))^{-1} = 1 \quad \text{for } x \in (0, \infty).$$

A.2. The symmetric game and the A equations

Here we prove Theorem 2.32.

A.2.1. The AM equations. Let $Q \in \mathbb{Z}$ be a parameter; we will consider $Q = 0$ and $Q = -1$. The AM system on \mathbb{Z} with parameter Q is a set of equations in the real variables a_i and m_i indexed by $i \in \mathbb{Z}$. These equations are

$$\begin{aligned} (a_i + a_{Q-i})(m_i + a_i) &= a_i m_{i+1} + a_{Q-i} m_{i-1} && \text{AM}(Q, 1) \\ (a_i + a_{Q-i})^2 &= a_{Q-i}(m_{i+1} - m_{i-1}) && \text{AM}(Q, 2) \end{aligned}$$

with the index i running over \mathbb{Z} .

PROPOSITION A.1. *Let $\{(a_i, b_i, m_i, n_i) : i \in \mathbb{Z}\}$ be an ABMN solution for which CenRatio is equal to a positive parameter x , and $m_{-\infty} = n_{\infty} = 0$.*

First let x equal to three.

- (1) *For $i \in \mathbb{Z}$, $a_i = b_{-i}$ and $m_i = n_{-i}$.*

- (2) The a - and m -variables solve the AM system on Z with parameter $Q = 0$.
- (3) The a -variables solve the A system on \mathbb{Z} (when we take $A_i = a_i$ for $i \in \mathbb{Z}$).

Now instead take x equal to one.

- (4) For $i \in \mathbb{Z}$, $a_i = b_{-1-i}$ and $m_i = n_{-1-i}$ for $i \in \mathbb{Z}$.
- (5) The a - and m -variables solve the AM system on Z with parameter $Q = -1$.
- (6) With a half-integer offset, the a -variables solve the A system on $\mathbb{Z} + 1/2$: that is, the variables $A_i = a_{i-1/2}$ for $i \in \mathbb{Z} + 1/2$ solve this system.

Proof: (1). We will argue that (a, b, m, n) is equal to $\mathcal{R}(a, b, m, n)$, where \mathcal{R} is the role-reversal map from Definition 3.1. This will suffice to prove Proposition A.1. For $i \in \mathbb{Z}$, let $\hat{\phi}_i$ denote the ϕ_i -value of $\mathcal{R}(a, b, m, n)$. Note that $\phi_0 = \frac{n_{-1}-n_0}{m_0-m_{-1}} = 3$ and $\hat{\phi}_1 = \frac{m_0-m_{-1}}{n_{-1}-n_0} = 1/3$. From the latter, and Proposition 3.4, $\hat{\phi}_0 = s_{-1}(1/3) = s^{-1}(1/3) = 3$. We have learnt that the value of CenRatio is shared between the ABMN solutions (a, b, m, n) and $\mathcal{R}(a, b, m, n)$ (the value equals three).

We write $\mathcal{R}(a, b, m, n) = (\hat{a}, \hat{b}, \hat{m}, \hat{n})$, and extend the hat-notation to boundary data. We next argue that this data satisfies

$$\hat{m}_{-\infty} = \hat{n}_{\infty} = 0 = m_{-\infty} = n_{\infty}. \quad (59)$$

That $m_{-\infty} = n_{\infty} = 0$ holds by assumption. But then we have $\hat{m}_{-\infty} = \lim_{k \rightarrow \infty} \hat{m}_{-k} = \lim_{k \rightarrow \infty} n_k = n_{\infty} = 0$. Likewise, $\hat{n}_{\infty} = 0$.

We further argue that

$$m_{\infty} = n_{-\infty} \quad \text{and} \quad \hat{m}_{\infty} = \hat{n}_{-\infty}. \quad (60)$$

To see these, note that $\mathcal{M}(3) = 1$, because this is a high k consequence of Corollary 3.3. But in the present case $\mathcal{M}(3)$ equals both $n_{-\infty}/m_{\infty}$ and $\hat{n}_{-\infty}/\hat{m}_{\infty}$; whence (60).

The two preceding displays imply that \hat{m}_{∞} , $n_{-\infty}$, $\hat{n}_{-\infty}$ and m_{∞} have a shared value. Let $u \in (0, \infty)$ denote the reciprocal of this value. The quadruples $\tau_u(a, b, m, n)$ and $\tau_u \mathcal{R}(a, b, m, n)$ are both the unique standard solution with CenRatio equal to three. The uniqueness of this solution asserted by Proposition 2.17(2) implies that these quadruples are equal. But then (a, b, m, n) and $\mathcal{R}(a, b, m, n)$ have been shown to be equal, as we sought to show in deriving Proposition A.1(1).

(4). We will argue that (a, b, m, n) and $\mathcal{A}(a, b, m, n)$ are equal, where $\mathcal{A} = \mathcal{S}_1 \circ \mathcal{R}$ is the left shift (by one place) of the role reversal map. This suffices to prove Proposition A.1(4). With a tilde notation for $\mathcal{A}(a, b, m, n)$ objects, we have that $\tilde{\phi}_0 = 1/\phi_0 = 1$. Thus, (a, b, m, n) and $\mathcal{A}(a, b, m, n)$ are ABMN solutions, each with CenRatio = 1. Since $\mathcal{M}(1) = 1$ by Corollary 3.3 in a high k limit, we find, similarly to above, that the values of \tilde{m}_{∞} , $n_{-\infty}$, $\tilde{n}_{-\infty}$ and m_{∞} coincide. The proof continues to mimic that of Proposition A.1(1).

(2). Given the first part, ABMN(1), or ABMN(2), implies AM(0, 1), while ABMN(3), or ABMN(4), implies AM(0, 2).

(5). Similarly, given the fourth part, and with AM(-1, 1) and AM(-1, 2) being implied.

(3,6). The AM system (for general Q) satisfies

$$(a_i + a_{Q-i})(m_i + a_i) = a_i \left(\frac{(a_i + a_{Q-i})^2}{a_{Q-i}} + m_{i-1} \right) + a_{Q-i} m_{i-1}.$$

This identity is seen to be equivalent to each of the following equalities in turn:

$$(a_i + a_{Q-i})(m_i + a_i - m_{i-1})a_{Q-i} = a_i(a_i + a_{Q-i})^2;$$

$a_i(a_i + a_{Q-i}) = a_{Q-i}(m_i + a_i - m_{i-1})$; and $m_i - m_{i-1} = a_i^2 a_{Q-i}^{-1}$. By summing the last of these over two consecutive indices, we find that

$$m_{i+1} - m_{i-1} = \frac{a_i^2}{a_{Q-i}} + \frac{a_{i+1}^2}{a_{Q-i-1}}.$$

Using AM(Q, 2) here, we obtain

$$\frac{(a_i + a_{Q-i})^2}{a_{Q-i}} = \frac{a_i^2}{a_{Q-i}} + \frac{a_{i+1}^2}{a_{Q-i-1}}.$$

Rearranging, and cancelling $a_{Q-i} \neq 0$, leads to

$$a_{Q-i-1}(2a_i + a_{Q-i}) = a_{i+1}^2$$

for $i \in \mathbb{Z}$. When $Q = 0$, the a -variables are thus seen to satisfy the A system on \mathbb{Z} , as Proposition A.1(3) asserts. When $Q = -1$, the A -variables $\{A_j : j \in \mathbb{Z} + 1/2\}$, $A_j = a_{j-1/2}$, solve the system $A_{-i-3/2}(2A_{i+1/2} + A_{i-1/2}) = A_{i+3/2}^2$ as i varies over \mathbb{Z} . Setting $k = i + 1/2$, we find that

$$A_{-k-1}(2A_k + A_{-k}) = A_{k+1}^2$$

for $k \in \mathbb{Z} + 1/2$. Thus, the variables $\{A_j : j \in \mathbb{Z} + 1/2\}$ solve the A system on $\mathbb{Z} + 1/2$, as we sought to show in proving Proposition A.1(6). This completes the proof of Proposition A.1. \square

A.2.2. Deriving Theorem 2.32. We are ready to derive the remaining result.

Proof of Theorem 2.32(1). If $\{a_i : i \in \mathbb{Z}\}$ solves the A system on \mathbb{Z} with $a_0 = 1$, then λa_i solves this system with $\lambda a_i = \lambda$. Thus it is without loss of generality that we take $\lambda = 1$.

Given that $a_0 = 1$, the 0- and -1 -indexed equations in the A system on \mathbb{Z} state that $3a_{-1} = a_1^2$ and $2a_{-1} + a_1 = 1$. These uniquely determine positive a_{-1} and a_1 . We may group the remaining A equations into pairs, indexed by $i = k$ and $i = -k - 1$ for $k \geq 1$, with the generic pair taking the form

$$O_-(2I_+ + I_-) = O_+^2 \quad \text{and} \quad I_+(2O_- + O_+) = I_-^2 \tag{61}$$

with input pair $(I_-, I_+) = (a_{-k}, a_k)$ and output pair $(O_-, O_+) = (a_{-k-1}, a_{k+1})$. These equation pairs may be iteratively solved, for $k \in \mathbb{N}_+$ rising, each step producing a unique pair of positive values in the output. The outcome is the unique solution $\{a_i : i \in \mathbb{Z}\}$ to the A system subject to $a_0 = 1$.

(2). Similarly to the first part, we may restrict to the case where $A_{-1/2} = 1$. The $-1/2$ -indexed equation then reads $2 + A_{1/2} = A_{1/2}^2$ and produces a unique positive value for $A_{1/2}$. Inductively we determine the A_i -values for $|i| \leq k - 1/2$. The equation pair indexed by $i = -k + 1/2$ and $i = k - 1/2$ takes the form (61) with $(I_-, I_+) = (A_{-k+1/2}, A_{k-1/2})$ and $(O_-, O_+) = (A_{-k-1/2}, A_{k+1/2})$, with a unique positive pair of outputs given the input pair. In this way, we construct the unique positive solution $\{A_i : i \in \mathbb{Z} + 1/2\}$ of the A system on $\mathbb{Z} + 1/2$ subject to $A_{-1/2} = 1$.

(3). Since $\mathcal{M}(3) = 1$, Theorem 2.6 implies that each strategy pair of the form

$$\lambda \cdot (b_i^{\text{st}}(3), a_i^{\text{st}}(3) : i \in \mathbb{Z}) \quad \text{for } \lambda \in (0, \infty) \tag{62}$$

is a time-invariant Nash equilibrium in the symmetric game. Note that $\{(b_i^{\text{st}}(3), a_i^{\text{st}}(3)) : i \in \mathbb{Z}\}$ equals $\{(a_{-i}^{\text{st}}(3), a_i^{\text{st}}(3)) : i \in \mathbb{Z}\}$ by Proposition A.1(1), and that $\{a_i^{\text{st}}(3) : i \in \mathbb{Z}\}$ solves the A

system on \mathbb{Z} by Proposition A.1(3). By Theorem 2.32(1), we see that the collection (62) takes the form

$$\left\{ \lambda \cdot (a_{-i}, a_i : i \in \mathbb{Z}) : \lambda \in (0, \infty) \right\},$$

where $\{a_i : i \in \mathbb{Z}\}$ is the unique solution of the A system on \mathbb{Z} subject to $a_0 = 1$. The sequence shift operation leaves invariant the space of time-invariant Nash equilibria for a given game, because it preserves boundary data. Since each of the last displayed maps is an equilibrium in the symmetric game, we have found that each element of S_1 is also such an equilibrium.

Since $\mathcal{M}(1) = 1$, similar reasoning, which invokes Proposition A.1(4,6), shows that every element of

$$\left\{ \lambda \cdot (A_{-i-1/2}, A_{i+1/2} : i \in \mathbb{Z}) : \lambda \in (0, \infty) \right\}$$

is a time-invariant Nash equilibrium for the symmetric game. All shifts of such equilibria are also equilibria; thus, we see that every element of S_2 is a time-invariant Nash equilibrium in the symmetric game.

It remains to argue that there is no coincidence between two elements specified in the definitions of the sets S_1 and S_2 . Our task is to extract from a strategy pair of the form $\lambda \cdot (b_{i+k}^{\text{st}}(z), a_{i+k}^{\text{st}}(z))$ that is specified by the data $(\lambda, k, z) \in (0, \infty) \times \mathbb{Z} \times \{1, 3\}$ the value that this triple takes in the indexing set. Proposition 2.26 permits us to write the pair in the form $\lambda \cdot (b_i^{\text{st}}(s_k(z)), a_i^{\text{st}}(s_k(z)))$. By Theorem 2.6, we may equally attempt to extract the data (λ, k, z) from the ABMN solution $\lambda \cdot (a_i^{\text{st}}(s_k(z)), b_i^{\text{st}}(s_k(z)), m_i^{\text{st}}(s_k(z)), n_i^{\text{st}}(s_k(z)))$. The value of CenRatio for the latter solution lies in the set $\{s_k(z) : k \in \mathbb{Z}, z \in \{1, 3\}\}$. We may extract k since it is the unique integer i for which $s_{-i}(\text{CenRatio}) \in (1/3, 3]$. The value of $s_{-k}(\text{CenRatio})$ lies in $\{1, 3\}$ and equals z . Since $m_\infty^{\text{st}} - m_{-\infty}^{\text{st}} = 1$, the value of λ may be identified as equal to $m_\infty^{\text{st}}(s_k(z)) - m_{-\infty}^{\text{st}}(s_k(z))$. We have confirmed that there is no coincidence between the elements in the union of S_1 and S_2 .

(4). We suppose that $Q(1) = 2$. Since $\mathcal{M}(1) = \mathcal{M}(3) = 1$, there are no further solutions to $\mathcal{M}(x) = 1$ among $x \in (1/3, 3]$. Theorem 2.24 implies that the set of time-invariant Nash equilibria in the game Standard(1) equals the set of pairs of the form $(b_{i+k}^{\text{st}}(z), a_{i+k}^{\text{st}}(z) : i \in \mathbb{Z})$ for $z \in \{1, 3\}$ and $k \in \mathbb{Z}$. The game Standard(1) is the symmetric game where $m_\infty = 1$. For the general cases, where $m_\infty = \lambda \in (0, \infty)$, the set of Nash equilibria is given by the set of pairs formed by multiplying the indicated expression by λ . In other words, the set in question is simply the union of S_1 and S_2 . \square

Bibliography

- [1] Murali K. Agastya and R. Preston McAfee. Continuing wars of attrition. *Working paper*, 2006.
- [2] Scott N. Armstrong and Charles K. Smart. A finite difference approach to the infinity Laplace equation and tug-of-war games. *Transactions of the American Mathematical Society*, 364(2):595–636, 2012.
- [3] Gunnar Aronsson. Extension of functions satisfying lipschitz conditions. *Arkiv för matematik*, 6(6):551–561, 1967.
- [4] Gunnar Aronsson. On the partial differential equation $u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy} = 0$. *Arkiv för Matematik*, 7(5):395 – 425, 1968.
- [5] G. Avni, I. Jecker, and D. Žikelić. Infinite-duration all-pay bidding games. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 617–636.
- [6] Guy Avni and Thomas A. Henzinger. An updated survey of bidding games on graphs. *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, pages No. 3: 3:1–3:6, 2022.
- [7] Guy Avni, Thomas A. Henzinger, and Ventsislav Chonev. Infinite-duration bidding games. *J. ACM*, 66(4), jul 2019.
- [8] Guy Avni, Thomas A. Henzinger, and Rasmus Ibsen-Jensen. Infinite-duration poorman-bidding games. In George Christodoulou and Tobias Harks, editors, *Web and Internet Economics*, pages 21–36, Cham, 2018. Springer International Publishing.
- [9] Guy Avni, Rasmus Ibsen-Jensen, and Josef Tkadlec. All-pay bidding games on graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:1798–1805, 04 2020.
- [10] Robert Axelrod and William D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, 1981.
- [11] Y. Azar, B. Birnbaum, L.E. Celis, N.R. Devanur, and Y. Peres. Convergence of local dynamics to balanced outcomes in exchange networks. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 293–302. IEEE, 2009.
- [12] Tamer Başar and Georges Zaccour, editors. *Handbook of Dynamic Game Theory*. Springer International Publishing, Cham, 2020.
- [13] Clayton Bjorland, Luis Caffarelli, and Alessio Figalli. Nonlocal tug-of-war and the infinity fractional laplacian. *Communications on pure and applied mathematics*, 65(3):337–380, 2012.
- [14] Pablo Blanc and J.D. Rossi. *Game Theory and Partial Differential Equations*. De Gruyter Series in Nonlinear Analysis and Applications 31. De Gruyter, Berlin, 2019.
- [15] Emile Borel. La théorie du jeu et les équations intégrales à noyau symétrique gauche. *C.R. Hebd. Seances Acad. Sci.*, 173:1304–1308, 1921.
- [16] M. G. Crandall, H. Ishii, and P.-L. Lions. User’s guide to viscosity solutions of second order partial differential equations. *Bulletin (new series) of the American Mathematical Society*, 27(1):1–67, 1992.
- [17] Pedro Dal Bo and Guillaume R. Frechette. Strategy choice in the infinitely repeated prisoner’s dilemma. *The American economic review*, 109(11):3929–3952, 2019.
- [18] Lawrence C. Evans and Ovidiu Savin. $C^{1,\alpha}$ regularity for infinity harmonic functions in two dimensions. *Calculus of variations and partial differential equations*, 32(3):325–347, 2008.
- [19] Christian Ewerhart and Julian Teichgräber. Multi-battle contests, finite automata, and the tug-of-war. *University of Zurich, Department of Economics, Working Paper No. 318. Econometrics: Econometric & Statistical Methods - Special Topics eJournal*, 2019.
- [20] Qiang Fu, Jingfeng Lu, and Yue Pan. Team contests with multiple pairwise battles. *American Economic Review*, 105(7):2120–40, July 2015.
- [21] Alan Gelder. From custer to thermopylae: Last stand behavior in multi-stage contests. *Games and Economic Behavior*, 87:442–466, 2014.
- [22] Samuel Häfner. A tug-of-war team contest. *Games and Economic Behavior*, 104:372–391, 2017.
- [23] Samuel Häfner and Kai Konrad. Eternal peace in the tug-of-war? Working Papers tax-mpg-rps-2016-09, Max Planck Institute for Tax Law and Public Finance, 2016.

- [24] Alan Hammond and Gábor Pete. Stake-governed tug-of-war and the biased infinity Laplacian. *arXiv:2206.08300*, 2022.
- [25] Christopher Harris and John Vickers. Racing with uncertainty. *mimeo*, Nuffield College, Oxford, 1986.
- [26] Christopher Harris and John Vickers. Racing with uncertainty. *The Review of Economic Studies*, 54(1):1–21, 1987.
- [27] Ryan B. Hayward. *Hex, inside and out : the full story*. CRC Press, Taylor and Francis Group, Boca Raton, FL, 2019.
- [28] Rufus Isaacs. *Differential games: a mathematical theory with applications to warfare and pursuit, control and optimization*. The SIAM series in applied mathematics. Wiley, New York, 1965.
- [29] R Jensen. Uniqueness of Lipschitz extensions: Minimizing the sup norm of the gradient. *Archive for rational mechanics and analysis*, 123(1):51–74, 1993.
- [30] Yashodhan Kanoria, Mohsen Bayati, Christian Borgs, Jennifer Chayes, and Andrea Montanari. Fast convergence of natural bargaining dynamics in exchange networks. In *Proceedings of the twenty-second annual ACM-SIAM symposium on discrete algorithms*, SODA '11, pages 1518–1537. Society for Industrial and Applied Mathematics, 2011.
- [31] Emin Karagözöglü, Çağrı Sağlam, and Agah R. Turan. Perseverance and suspense in tug-of-war. *Journal of Mathematical Economics*, 95:102469, 2021.
- [32] Jon Kleinberg and Éva Tardos. Balanced outcomes in social exchange networks. *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 295–304, 05 2008.
- [33] Tilman Klumpp. Colonel Blotto's tug of war. Working Papers 2021-3, University of Alberta, Department of Economics, 2021.
- [34] Tilman Klumpp, Kai A. Konrad, and Adam Solomon. The dynamics of majoritarian blotto games. *Games and Economic Behavior*, 117:402–419, 2019.
- [35] Tilman Klumpp and Mattias K. Polborn. Primaries and the New Hampshire effect. *Journal of Public Economics*, 90(6-7):1073–1114, 2006.
- [36] Kai Konrad and Dan Kovenock. Equilibrium and efficiency in the tug-of-war. Discussion Papers, Research Unit: Market Processes and Governance SP II 2005-14, WZB Berlin Social Science Center, 2005.
- [37] Kai A. Konrad. Dynamic contests and the discouragement effect. *Revue d'économie politique*, 122(2):233–256, 2012.
- [38] Kai A. Konrad and Dan Kovenock. Multi-battle contests. *Games and Economic Behavior*, 66(1):256–274, 2009.
- [39] Andrew J. Lazarus, Daniel E. Loeb, James G. Propp, Walter R. Stromquist, and Daniel H. Ullman. Combinatorial games under auction play. *Games and economic behavior*, 27(2):229–264, 1999.
- [40] Andrew J. Lazarus, Daniel E. Loeb, James G. Propp, and Daniel H. Ullman. Richman games. *Games of no chance*, *Math. Sci. Res. Inst. Publ.*, 29:439–449, 1994.
- [41] Tom Lee and Louis L. Wilde. Market structure and innovation: A reformulation. *The Quarterly Journal of Economics*, 94(2):429–436, 1980.
- [42] Marta Lewicka. *A Course on Tug-Of-War Games with Random Noise: Introduction and Basic Constructions*. Universitext. Springer International Publishing AG, Cham, 2020.
- [43] Marta Lewicka and Juan J. Manfredi. The obstacle problem for the p -laplacian via optimal stopping of tug-of-war games. *Probability theory and related fields*, 167(1-2):349–378, 2017.
- [44] Peter. Lindqvist. *Notes on the Stationary p -Laplace Equation*. SpringerBriefs in Mathematics. Springer International Publishing, Cham, 1st ed. 2019. edition, 2019.
- [45] David Malueg and Andrew Yates. Testing contest theory: Evidence from best-of-three tennis matches. *The Review of Economics and Statistics*, 92(3):689–692, 2010.
- [46] Juan J. Manfredi. *Introduction to Random Tug-of-War Games and PDEs*, pages 133–151. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
- [47] Yuval Peres, Gábor Pete, and Stephanie Somersille. Biased tug-of-war, the biased infinity Laplacian, and comparison with exponential cones. *Calc. Var. and PDE*, 38:541–564, 2010.
- [48] Yuval Peres, Oded Schramm, Scott Sheffield, and David Wilson. Random-turn hex and other selection games. *The American Mathematical Monthly*, 114(5):373–387, 2007.
- [49] Yuval Peres, Oded Schramm, Scott Sheffield, and David Wilson. Tug-of-war and the infinity Laplacian. *Journal of the American Mathematical Society*, 22(1):167–210, 2009.
- [50] Yuval Peres and Scott Sheffield. Tug-of-war with noise: A game-theoretic view of the p -Laplacian. *Duke mathematical journal*, 145(1):91–120, 2008.
- [51] Yuval Peres and Zoran Šunić. Biased infinity Laplacian Boundary Problem on finite graphs. *arXiv:1912.13394*, 2019.

- [52] William Poundstone. *Prisoner's Dilemma*. Knopf Doubleday Publishing Group, Westminster, 2011.
- [53] T. E. S. Raghavan, T. S. Ferguson, T. Parthasarathy, and O. J. Vrieze, editors. *Stochastic games and related topics: in honor of Professor L.S. Shapley*. Theory and decision library. Series C, Game theory, mathematical programming, and operations research; v. 7. Kluwer Academic Publishers, Dordrecht, 1991.
- [54] Ovidiu Savin. C^1 regularity for infinity harmonic functions in two dimensions. *Archive for rational mechanics and analysis*, 176(3):351–361, 2005.
- [55] Aner Sela. Best-of-three all-pay auctions. *Economics Letters*, 112(1):67–70, 2011.
- [56] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences - PNAS*, 39(10):1095–1100, 1953.
- [57] G. Tullock. Efficient rent seeking. In *Toward a Theory of Rent Seeking Society*, Texas A and M University Press, College Station, pages 97–112. 1980.
- [58] John Vickers. Patent races and market structure. *D.Phil thesis, Oxford 1985*.
- [59] Robert Wilson. Computing equilibria of N -person games. *SIAM Journal on Applied Mathematics*, 21:80–87, 1971.
- [60] Daniel John Zizzo. Racing with uncertainty: a patent race experiment. *International Journal of Industrial Organization*, 20(6):877–902, 2002.