

A Game of Timing and Visibility*

Zvi Lotker
Ben Gurion University
zvilo@cse.bgu.ac.il

Boaz Patt-Shamir[†]
Tel Aviv University
boaz@eng.tau.ac.il

Mark R. Tuttle
Intel
tuttle@acm.org

February 11, 2007

Abstract

We consider the following abstraction of competing publications. There are n players in the game. Each player i chooses a point x_i in the interval $[0, 1]$, and a player's payoff is the distance from its point x_i to the next larger point, or to 1 if x_i is the largest. For this game, we give a complete characterization of the Nash equilibrium for the two-player game, and, more important, we give an efficient approximation algorithm to compute numerically the symmetric Nash equilibrium for the n -player game. The approximation is computed via a discrete version of the game. In both cases, we show that the (symmetric) equilibrium is unique. Our algorithmic approach to the n -player game is non-standard in that it does not involve solving a system of differential equations. We believe that our techniques can be useful in the analysis of other timing games.

*An extended abstract of this work appeared in 18th ACM Symposium on Parallelism in Algorithms and Architectures, July 2006.

[†]Contact author. Postal address: Dept. of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel.

1 Introduction

We consider the following non-cooperative, complete-information strategic game. The pure strategies are real numbers in $[0, 1]$. Each player i chooses a number x_i in $[0, 1]$, and player i 's payoff is the distance to the next larger point. More precisely, the payoff is $x' - x_i$, where $x' = \min \{x_j \mid x_j \geq x_i \text{ and } i \neq j\}$, or $x' = 1$ if x_i is larger than the number chosen by any other player. This game formalizes a simplified version of the situation where players seek to maximize their visibility time.

In this paper we completely characterize the Nash equilibrium with a closed-form solution for the two-player version of the game and prove it is symmetric and unique. While the analysis in this case is standard, the equilibrium strategy turns out to be somewhat counter-intuitive. Our main result, however, is for the general n -player game. For this model we develop an algorithm that approximates the symmetric equilibrium strategy. Our approach is to use a discrete version of the game, in which a player must choose from a finite set of points in the unit interval. We develop an efficient numerical algorithm for the discretized game, and use it to approximate the symmetric equilibrium of the continuous game to any degree of accuracy. We remark that proving the algorithm correct entails an interesting methodology and some non-trivial analysis.

Related work. The game we consider is, to the best of our knowledge, a new variant of the family of “timing games” (see, e.g., [2]). More specifically, our game resembles the “War of Attrition” game, abbreviated henceforth “WoA.” In the two-player version of WoA, first formalized by Maynard Smith [5], the players are engaged in a costly competition and they need to choose a time to concede. More formally, the first player to concede (called “leader”) gets a smaller payoff than the other player (called “follower”). Furthermore, the payoff to the leader strictly decreases as time progresses, i.e., conceding early is better than conceding late. Hendricks *et al.* [3] axiomatize and analyze a general setting of complete information WoA. Our game violates one crucial axiom of [3]: in our game, the payoff to the leader does not decrease with time.

WoA and other 2-player continuous-time timing games were generalized by Baye *et al.* to a “general linear model of contests” [1]. Implicitly, the general formula presented in [1] covers our Lemma 3.2. Our other results (for the two-player game, and, of course, all results for the n -player game) are unrelated to the results of [1].

Another family of games that superficially resembles our game is the Hotelling location games [4], but these games either are zero-sum or they involve pricing, and hence they are fundamentally different.

Organization. The remainder of this paper is organized as follows. In Section 2 we formally define the game and review some relevant facts from game theory. In Section 3 we thoroughly analyze the two-player case in the continuous model. In Section 4 we consider the n -player case, and present an algorithm to compute the equilibrium in the discrete model. Section 4.3 presents some experimental results. We discuss some applications motivating the game in Section 5, and give concluding remarks in Section 6. Some standard proofs are presented in the appendix.

2 Preliminaries

Definition of the visibility game. The visibility game is a symmetric game of n players, whose actions are (in the continuous case) real numbers in the unit interval $[0, 1]$. The payoff to each player i is defined as follows. Given the choices (x_1, \dots, x_n) of the players, define

$$L(i) \stackrel{\text{def}}{=} \{x_j \mid x_j \geq x_i \text{ and } j \neq i\} .$$

$L(i)$ is the set of all values at least x_i excluding x_i . With this definition, the payoff function u_i for player i is defined by

$$u_i(x_1, x_2, \dots, x_n) \stackrel{\text{def}}{=} \begin{cases} \min(L(i)) - x_i , & \text{if } L(i) \neq \emptyset \\ 1 - x_i , & \text{otherwise} \end{cases}$$

In words, u_i is the distance from x_i to either the next value up, or to 1 if x_i is the unique maximum. Note that in our definition, if two players happen to choose the same value, the payoff to both of them is 0. We call this definition *non-conserving*. In a *conserving* variant of the game, colliding players somehow share the interval to their right, so that only the interval to the left of the smallest x_i is not claimed by anyone. Unless otherwise stated, we will mostly study the non-conserving variant as defined above, which is mathematically more convenient.¹

Game theory fact. We will make frequent use of the following standard property of Nash equilibria, which we state in the n -player continuous case of our game (see, e.g., [6]).

Theorem 2.1 *Let (f_1, \dots, f_n) be a Nash equilibrium point, with expected payoff v_i to player i at the equilibrium point. Let $\pi_i(x)$ denote the expected payoff for player i when he plays the pure strategy x and all other players play their equilibrium mixed strategy. Then $\pi_i(x) \leq v_i$ for all $x \in [0, 1]$, and furthermore, there exists a set \mathcal{Z} of measure 0 such that $\pi_i(x) = v_i$ for all $x \in \text{support}(f_i) \setminus \mathcal{Z}$.*

3 The 2-player continuous game

In this section we study the two player game. We start with the simple observation that this game does not admit any pure-strategy equilibrium. The theorem is proved for the non-conserving game variant, but the result holds (using arguments of the same type) in the conserving case as well.

Theorem 3.1 *There is no equilibrium of pure strategies for the game.*

Proof: By contradiction. Let (x, y) be such an equilibrium. First we note that $x \neq y$, because otherwise the payoff for the players is 0 and each player can increase his payoff by changing his strategy. It follows that at least one of the players does not play $\frac{1}{2}$. Assume w.l.o.g. that $x \neq \frac{1}{2}$. Consider first the case that $x < \frac{1}{2}$. Then player 2 can improve his payoff by playing $x + \epsilon$ for some arbitrarily small $\epsilon > 0$. It follows that there is no equilibrium where a player plays less than $\frac{1}{2}$. But

¹The non-conserving model captures some real situations. For example, suppose the action is interpreted as writing to a shared register. In some implementations, concurrent writes to the same register yield unpredictable results, and may simply fail.

there could be no equilibrium if both players play at least $\frac{1}{2}$: If $x > \frac{1}{2}$, then the best strategy for player 2 is to play 0, contradiction. ■

In the remainder of this section we analyze the mixed-strategy Nash equilibrium for two players. It turns out that there is only one equilibrium point, which is symmetric.

3.1 Mixed Strategy Equilibrium

Let us start by assuming the existence of an equilibrium point (existence is not immediately guaranteed because the game is infinite and the payoff functions are not continuous, but as it will turn out, the equilibrium does exist). So fix a Nash equilibrium point. Let (f_1, f_2) be the probability density functions (pdf's) of players 1 and 2, respectively, at the equilibrium point. The following lemma characterizes the density functions in the equilibrium point on nearly all the support set. We remark that this result is implicit in [1]. All proofs of this section are given in the appendix.

Lemma 3.2 *There exists a set \mathcal{Z} of measure zero, such that for all $x \in \text{support}(f_1) \setminus \mathcal{Z}$, $f_2(x) = \frac{1}{1-x}$.*

We note that the density function of Lemma 3.2 remains invariant under affine transformations of the payoff functions (possibly different transformations for the two players). More precisely, if for some $a_1 > 0$ and any real b_1 , the payoff function of Player 1 is defined by

$$u_1(x, y) = \begin{cases} a(y - x) + b, & \text{if } y > x \\ a(1 - x) + b, & \text{if } x > y, \text{ and} \\ 0, & \text{otherwise,} \end{cases}$$

and the payoff of player 2 is defined similarly using $a_2 > 0$ and any real b_2 , then the proof of Lemma 3.2 can be extended to show that $f_2(x) = \frac{1}{1-x}$ for $x \in \text{support}(f_1) \setminus \mathcal{Z}$.

Next, we determine the support sets. First we note that the supports of f_1 and f_2 are essentially the same.

Lemma 3.3 *With the possible exception of a set of measure zero, $\text{support}(f_1) = \text{support}(f_2)$.*

We now set to determine the support of f_1 . To avoid dealing with pathological cases, we make a simplifying assumption that the equilibrium strategies satisfy $f_1(x) = \frac{1}{1-x}$ for all $x \in \text{support}(f_1)$ except for at most a finite number of points. We consider an equilibrium point (f_1, f_2) with expected payoffs v_1 and v_2 to Players 1 and 2, respectively.

Lemma 3.4 $\inf(\text{support}(f_1)) = 0$, and $\sup(\text{support}(f_1)) = 1 - v_1$.

Lemma 3.5 *For all intervals $[x_1, x_2]$ with $0 < x_1 < x_2 < 1 - v_1$ we have that $\int_{x_1}^{x_2} f_1(x) dx > 0$.*

Thus we know how the pdf looks like “almost everywhere,” and we know that it spans the interval $[0, 1 - v]$, where v is the value of the game. In the theorem below, we show that there are no “atoms” (discrete points with positive probability) in the support of the equilibrium strategy. This completes the characterization of the equilibrium strategy for the two-player game.

Theorem 3.6 *Up to a set of measure zero, there is a unique Nash equilibrium point to the two-player visibility game defined by $f_1(x) = \frac{1}{1-x}$ for $0 \leq x \leq 1 - \frac{1}{e}$ and 0 otherwise. Furthermore, the expected*




Figure 1: *Nash equilibrium strategy for 2-player game. The pdf is on the left, and the cdf is on the right. The cutoff point is $1 - \frac{1}{e} \approx 0.632$, and the game value is about 0.368.*

payoff for each player is $\frac{1}{e}$.

We remark that the actual equilibrium strategy, as depicted in Fig. 1, was surprising to us. It may be particularly interesting to find some natural phenomena that adhere to this distribution.

4 The n -player game

In this section we present our main results. We consider a discrete, symmetric n -player version of our game. The game is defined by two parameters: the number of players n , and the resolution of the actions k . Specifically, in our version, players can only choose from the $k + 1$ points $\{0, \frac{1}{k}, \frac{2}{k}, \dots, 1\}$. We refer to the game as the (n, k) -game. Since this is a finite symmetric game, it admits a symmetric mixed Nash equilibrium. In this section we consider only symmetric mixed Nash equilibria.

4.1 Elementary properties of the equilibrium

It is convenient to define $t_i = 1 - \frac{i}{k}$ for $i = 0, \dots, k$ so that $\frac{i}{k}$ is the distance from t_i to 1, meaning that $t_k = 0$ and $t_0 = 1$ (see Figure 2). For any given equilibrium, we denote by p_k, p_{k-1}, \dots, p_0 the probabilities of playing the pure strategies t_k, t_{k-1}, \dots, t_0 in the equilibrium strategy, and denote by π_{t_i} the payoff to a player playing the pure strategy t_i when all other players follow the equilibrium strategy. We start with a general property that holds also for the continuous case, and then restrict attention to the discrete case.

Lemma 4.1 *Consider the n -player game (continuous or discrete, conserving or not), and let v a symmetric equilibrium value. Then $v < 1/n$. If the game is conserving, then we also have $v > 1/(n+1)$.*




Figure 2: *A schematic representation of the discrete game.*

Proof: Let $\pi^1, \pi^2, \dots, \pi^n$ be the random variables denoting the payoffs to the players $1, 2, \dots, n$ following the equilibrium strategy. Obviously, $\sum_{i=1}^n \pi^i \leq 1$, so by linearity of expectation $nv = \sum_{i=1}^n E[\pi^i] \leq 1$ i.e., $v \leq 1/n$. To show that $v \neq 1/n$, note that $v = 1/n$ only if at least one player chooses 0 with probability 1. However, in a symmetric equilibrium this happens only if all players choose 0 with probability 1, and this is clearly not an equilibrium strategy (obviously for non-conserving games, and by considering the action $1/k$ for conserving games). This proves the upper bound on v for the general case. In the case of the conserving n -player game, let x_1, \dots, x_n denote the actions taken by the players, and let $\hat{x} \stackrel{\text{def}}{=} \min\{x_1, \dots, x_n\}$. Since the game is conserving, we have that $\sum_{i=1}^n \pi^i = 1 - \hat{x}$, since only the leftmost interval $[0, \hat{x}]$ is not claimed by any player. It follows from linearity of expectation that $nv = \sum_{i=1}^n E[\pi^i] = 1 - E[\hat{x}]$. Suppose, for contradiction, that $v \leq 1/(n+1)$. If player n plays the pure strategy 0 when all remaining players follow the equilibrium strategy, then his payoff is $\min\{x_1, \dots, x_{n-1}\}$ and therefore his expected payoff is

$$E[\min\{x_1, \dots, x_{n-1}\}] > E[\min\{x_1, \dots, x_n\}] = 1 - nv \geq 1 - \frac{n}{n+1} = \frac{1}{n+1} \geq v,$$

contradicting the fact that v is the equilibrium value, so $v > 1/(n+1)$. ■

Next, we calculate the expected payoff of playing the pure strategy t_i against the equilibrium strategy. We will use the following notation extensively.

Notation 4.1 *Given a symmetric strategy for the (n, k) game:*

- p_i denotes the probability of choosing t_i ,
- $F_{i,\ell} \stackrel{\text{def}}{=} p_i + p_{i-1} + \dots + p_{i-\ell+1}$ denotes the probability of choosing one of the ℓ strategies $t_i, t_{i-1}, \dots, t_{i-\ell+1}$ to the right of t_i , and
- π_i denotes the payoff to a player playing t_i .

Lemma 4.2 *Given a symmetric strategy for the (n, k) -game, $E[\pi_i] = \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1}$.*

Proof: Since the value of π_i is always of the form $\frac{\ell}{k}$ for some integer ℓ satisfying $0 \leq \ell \leq k$, the expected value of π_i is

$$\begin{aligned} E[\pi_i] &= \sum_{\ell=1}^i \frac{\ell}{k} \mathbf{P} \left[\pi_i = \frac{\ell}{k} \right] = \frac{1}{k} \sum_{\ell=1}^i \mathbf{P} \left[\pi_i \geq \frac{\ell}{k} \right] \\ &= \frac{1}{k} \sum_{\ell=1}^i \mathbf{P} [\text{no one plays } t_i, \dots, t_{i-\ell+1}] = \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1}. \quad \blacksquare \end{aligned}$$

One nice property of the (n, k) game is that the support for an equilibrium forms a segment at the left of the unit interval, as in the continuous case for two players (cf. Lemmas 3.4 and 3.5).

Lemma 4.3 *For the (n, k) -game, if v is an equilibrium value, then t_k, t_{k-1}, \dots, t_c is the equilibrium support, where $c = \lfloor vk \rfloor + 1$.*

Proof: Let t_i be the largest strategy in the support, meaning that t_i is in the support and t_{i-1}, \dots, t_0 are not. Consider π_i . Notice that $F_{i,\ell} = p_i$ since t_{i-1}, \dots, t_0 are not in the support, so Theorem 2.1

and Lemma 4.2 imply that $v = E[\pi_i] = \frac{i}{k}(1 - p_i)^{n-1} < \frac{i}{k}$, so $i > vk$ and $i \geq \lfloor vk \rfloor + 1 = c$. This shows that t_{c-1}, \dots, t_0 are not in the support.

Suppose that t_c is not in the support. The expected payoff for playing the pure strategy t_c against the equilibrium strategy is $\frac{c}{k}$ since no other player plays t_c, \dots, t_0 since they are not in the support. This value is at most the equilibrium value, so $\frac{c}{k} \leq v$ or $c \leq vk$, but this contradicts the fact that $c = \lfloor vk \rfloor + 1 > vk$.

Finally, suppose t_i is in the support and t_{i+1} is not. Then by Lemma 4.2,

$$\begin{aligned}
E[\pi_{i+1}] &= \frac{1}{k} \sum_{\ell=1}^{i+1} (1 - F_{i+1,\ell})^{n-1} \\
&= \frac{1}{k} (1 - F_{i+1,1})^{n-1} + \frac{1}{k} \sum_{\ell=2}^{i+1} (1 - F_{i+1,\ell})^{n-1} \\
&= \frac{1}{k} + \frac{1}{k} \sum_{\ell=2}^{i+1} (1 - F_{i,\ell-1})^{n-1} \\
&= \frac{1}{k} + \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1} \\
&= \frac{1}{k} + E[\pi_i] > E[\pi_i] = v.
\end{aligned}$$

The third equality holds since t_{i+1} is not in the support, meaning none of the players following the equilibrium strategy play t_{i+1} , so $F_{i+1,1} = 0$ and $F_{i+1,\ell} = F_{i,\ell-1}$. But this means that playing t_{i+1} beats an equilibrium strategy, contradiction. This concludes the proof that t_k, t_{k-1}, \dots, t_c are the equilibrium support. \blacksquare

4.2 Symmetric equilibrium: uniqueness and algorithm

In this section we prove the following two theorems.

Theorem 4.4 *There is a unique symmetric equilibrium to the (n, k) game.*

Theorem 4.5 *Algorithm EQUILIBRIUM (n, k, ϵ) of Figure 3 computes the equilibrium value of the (n, k) game to within ϵ .*

We remark that in fact, we compute the equilibrium strategy and not only the equilibrium value.

Overview. The high-level idea in the algorithm is as follows. We guess the game value v , which implies the support set of the equilibrium strategy. Then we start computing the probabilities associated with each point in the support set, *starting with the rightmost point and advancing to the left*. This can be done since the payoff to a player who plays x depends only on the actions taken to its right, i.e., actions with value larger than x . Since by induction the probability of these actions have already been computed, we can proceed by solving the polynomials suggested by Lemma 4.2. To make this idea work, we need to analyze the case of wrongly guessing the game value. It requires some non-trivial

Subroutine FEASIBLE(n, k, v) =

- (1) Let $c = \lfloor vk \rfloor + 1$.
- (2) Let $p_c = 1 - \left(v \frac{k}{c}\right)^{\frac{1}{n-1}}$
- (3) For $i \leftarrow c + 1, \dots, k$ do
 - (3.1) Let p_i be the least positive solution to $v = \mathcal{E}_i(p_i, \dots, p_0)$
 - (3.2) Return “failure” if no positive solutions exist or $F_i > 1$.
- (4) Return p_k, \dots, p_c .

EQUILIBRIUM(n, k, ϵ) =

- (1) Let $l \leftarrow \frac{1}{n}$ and $r \leftarrow 0$.
- (2) While $l - r > \epsilon$ do
 - (2.1) Let $m \leftarrow (l + r)/2$.
 - (2.2) If FEASIBLE(n, k, m) succeeds then $l \leftarrow m$ else $r \leftarrow m$
- (3) Return l .

Figure 3: Algorithm EQUILIBRIUM(n, k, ϵ) computes the equilibrium value for the (n, k) -game to within ϵ . $F_i = \sum_{j=0}^i p_j$ and $\mathcal{E}_i(p_i, \dots, p_0)$ is defined in Eq. (1).

analysis to show that if our guess of the game value is too large, then the sum of the probabilities over the support set is smaller than 1, and if the guess is too small, then the algorithm will fail at some point (indicated by the nonexistence of a positive real probability that solves the polynomial for that point). Thus, we can conduct binary search on the game value, approximating it to any desired degree. The uniqueness of the symmetric equilibrium is a by-product of our analysis of the algorithm.

We start by applying Lemma 4.2 as follows. Suppose v is an equilibrium value for the unique (n, k) -game, and suppose t_k, \dots, t_c is the equilibrium support with equilibrium distribution p_k, \dots, p_c . We know from Lemma 4.3 that $c = \lfloor vk \rfloor + 1$. Since by Theorem 2.1 $E[\pi_i]$ is equal to the equilibrium value, we have $v = E[\pi_i] = \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1}$ for each $i = c, \dots, k$. Equivalently, if we define

$$\mathcal{E}_i(p_i, \dots, p_0) \stackrel{\text{def}}{=} \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1} \quad (1)$$

then we have $v = \mathcal{E}_i(p_i, \dots, p_0)$ for each $i = c, \dots, k$. Since p_{c-1}, \dots, p_0 are all 0, \mathcal{E}_c is a polynomial in p_c , \mathcal{E}_{c+1} is a polynomial in p_c and p_{c+1} , and so on. This suggests that we can iteratively solve these equations for p_c , then p_{c+1} , and so on. At each step, having computed the values of p_c, \dots, p_{i-1} , we just have to solve a polynomial of degree $n - 1$ in the single variable p_i , which can be efficiently solved to any desired accuracy. Let us state the simplest case for later use.

Lemma 4.6 $p_c = 1 - \left(v \cdot \frac{k}{c}\right)^{\frac{1}{n-1}}$

Proof: Solve $v = \mathcal{E}_c(p_c, 0, \dots, 0) = \frac{1}{k} \sum_{\ell=1}^c (1 - F_{c,\ell})^{n-1} = \frac{c}{k} (1 - p_c)^{n-1}$ for p_c . ■

The main problem we face now is to show that we can find a set of values p_i that will indeed be probabilities. To this end, we define a *feasible solution* for a potential equilibrium value v as follows.

Let $F_i \stackrel{\text{def}}{=} p_i + \dots + p_0$, namely F_i is the probability of choosing at least $1 - i/k$.

Definition 4.2 We say that p_k, p_{k-1}, \dots, p_c is a feasible solution for v if $c \stackrel{\text{def}}{=} \lfloor vk \rfloor + 1$ and the following conditions hold for each $i = c, \dots, k$:

$$C_1(i): v = \mathcal{E}_i(p_i, \dots, p_0).$$

$$C_2(i): p_i \geq 0.$$

$$C_3(i): F_i \leq 1.$$

It is obvious from this discussion that an equilibrium value has a feasible solution:

Lemma 4.7 If there is an equilibrium with value v and with support t_k, \dots, t_c and probabilities p_k, \dots, p_c (where $p_{c-1} = p_{c-2} = \dots = p_0 = 0$), then p_k, \dots, p_c is a feasible solution for v and $F_k = 1$.

It is not so obvious that the feasible solution for a given value v is unique. To do that, we first rewrite the \mathcal{E}_i polynomials and change their variables, as defined in the following lemma.

Lemma 4.8 $\mathcal{E}_i(p_i, \dots, p_0) = \mathcal{P}_i(1 - F_i)$ where

$$\mathcal{P}_i(x) = a_{i,0} + a_{i,1}x + \dots + a_{i,n-1}x^{n-1}, \quad \text{and} \quad a_{i,j} = \frac{1}{k} \binom{n-1}{j} \sum_{\ell=1}^i (F_{i-\ell})^{n-1-j}.$$

Proof: Follows immediately from

$$\begin{aligned} \mathcal{E}_i(p_i, \dots, p_0) &= \sum_{\ell=1}^i \frac{1}{k} (1 - F_{i,\ell})^{n-1} \\ &= \sum_{\ell=1}^i \frac{1}{k} (1 - F_i + F_{i-\ell})^{n-1} \\ &= \sum_{\ell=1}^i \frac{1}{k} \sum_{j=0}^{n-1} \binom{n-1}{j} (1 - F_i)^j (F_{i-\ell})^{n-1-j} \\ &= \sum_{j=0}^{n-1} \frac{1}{k} \binom{n-1}{j} \left(\sum_{\ell=1}^i (F_{i-\ell})^{n-1-j} \right) (1 - F_i)^j. \quad \blacksquare \end{aligned}$$

Using the \mathcal{P}_i representation, and applying Descartes's Rule of Signs, we can prove that solving for p_i (after solving for p_0, \dots, p_{i-1}) results in at most one solution that makes sense, as formalized in the following key lemma.

Lemma 4.9 Suppose $v = \mathcal{E}_{i-1}(p_{i-1}, \dots, p_0)$ and $F_{i-1} \leq 1$ for some nonnegative p_{i-1}, \dots, p_0 , and suppose $i - 1 > vk$. If $v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0)$ and $F_i \leq 1$, then p_i is the least positive solution to the equation $v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0)$.

Proof: First we prove that p_i is positive, and then we prove that p_i is the only positive solution to both $v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0)$ and $F_i \leq 1$. It follows that p_i is the least positive solution to $v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0)$, since any smaller positive solution would also satisfy $F_i \leq 1$, contradicting the uniqueness of p_i .

To see that p_i is positive, suppose $p_i \leq 0$. Since p_{i-1}, \dots, p_0 are nonnegative and $F_{i-1} \leq 1$, we have $F_{i,\ell} \leq F_{i-1,\ell-1} \leq F_{i-1} \leq 1$, so $1 - F_{i,\ell} \geq 1 - F_{i-1,\ell-1} \geq 0$, and

$$\begin{aligned}
v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0) &= \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i,\ell})^{n-1} \\
&\geq \frac{1}{k} \sum_{\ell=1}^i (1 - F_{i-1,\ell-1})^{n-1} \\
&= \frac{1}{k} + \frac{1}{k} \sum_{\ell=2}^i (1 - F_{i-1,\ell-1})^{n-1} \\
&= \frac{1}{k} + \frac{1}{k} \sum_{\ell=1}^{i-1} (1 - F_{i-1,\ell})^{n-1} \\
&= \frac{1}{k} + \mathcal{E}_{i-1}(p_{i-1}, \dots, p_0) = \frac{1}{k} + v > v
\end{aligned}$$

which is a contradiction, so $p_i > 0$.

To see that p_i is unique, Lemma 4.8 says that each choice of p_i satisfying $v = \mathcal{E}_i(p_i, p_{i-1}, \dots, p_0)$ and $F_i \leq 1$ yields a distinct nonnegative root $1 - F_i$ of the polynomial

$$(a_{i,0} - v) + a_{i,1}x + \dots + a_{i,n-1}x^{n-1}$$

so we need only count the number of nonnegative roots of this polynomial. The coefficients of the nonconstant terms are nonnegative since the $F_{i-\ell}$ defining the $a_{i,j}$ are nonnegative. In fact, some of them must be positive since some of the nonnegative p_{i-1}, \dots, p_0 must be positive: If all of the p_{i-1}, \dots, p_0 are zero, then $i-1 > vk$ implies $v = \mathcal{E}_{i-1}(p_{i-1}, \dots, p_0) = \frac{i-1}{k} > v$, which is a contradiction.

Suppose zero is a root. Then the coefficient of the constant term ($a_{i,0} - v$) must be zero, and since the other coefficients are nonnegative and some are positive, there can be no positive roots. Thus, there is a unique nonnegative root 0, so $1 - F_i = 0$ and $p_i = 1 - F_{i-1}$ is uniquely determined.

Suppose zero is not a root. We know that there is a nonnegative root $1 - F_i$, so there must be at least one positive root. Descartes's Rule of Signs says that the number of positive real roots (counting multiplicities) of a polynomial is equal to the number of alternations in the signs of its nonzero coefficients minus an even number. Since all coefficients of the non-constant terms are nonnegative and some are positive, and since there is at least one positive root, the coefficient of the constant term must be negative or the number of sign alternations would be zero, contradicting the existence of a positive root. Thus, the number of sign alternations is one and there is a unique nonnegative root r , so $1 - F_i = r$ and $p_i = r - 1 - F_{i-1}$ is uniquely determined. \blacksquare

Corollary 4.10 *If there is a feasible solution for v , then it is unique.*

Proof: The definition of a feasible solution says that p_{c-1}, \dots, p_0 must be 0, where $c = \lfloor vk \rfloor + 1$. Since $c > vk$, Lemma 4.6 says that p_c is uniquely determined and positive, and, by induction on $i > c$, Lemma 4.9 shows that p_i is uniquely determined since $i - 1 \geq c > vk$. \blacksquare

We now prove the central property of the algorithm: if there exists a feasible solution for v , there exist solutions for all $v' \geq v$. The main results of this section follow directly from Lemma 4.11 below.

Lemma 4.11 *If there is a feasible solution p_k, \dots, p_c for v , then there is a feasible solution $p'_k, \dots, p'_{c'}$ for each $v' > v$. In this case, $F'_k < F_k$ where $F'_i = p'_i + \dots + p'_{c'}$ and $F_i = p_i + \dots + p_c$.*

Proof: We show that there is a feasible solution for each $v' > v$ satisfying $\frac{c}{k} \geq v' > \frac{c-1}{k}$. The same proof shows that if there is a feasible solution for $\frac{c'}{k}$, then there is a feasible solution for each v' satisfying $\frac{c'+1}{k} \geq v' \geq \frac{c'}{k}$. Thus, by induction on c' , there is a feasible solution for each $v' > v$.

The values p_k, \dots, p_c and v satisfy $C_1(i)$, $C_2(i)$, and $C_3(i)$ for all $i \geq c$. It is enough to construct $p'_k, \dots, p'_{c'}$ so that the values $p'_k, \dots, p'_{c'}$ and v' satisfy $C_1(i)$, $C_2(i)$, and $C_3(i)$ for all $i \geq c$. To see why, consider two cases. If $\frac{c}{k} > v'$, then since $\frac{c}{k} > v' > v \geq \frac{c-1}{k}$ we have $c' = \lfloor v'k \rfloor + 1 = \lfloor vk \rfloor + 1 = c$, and p'_k, \dots, p'_c is a feasible solution for v' . If $\frac{c}{k} = v'$, then $c' = c + 1$, but Lemma 4.6 says that $p'_c = 1 - (v' \cdot \frac{k}{c})^{\frac{1}{n-1}} = 0$. Since the values p'_k, \dots, p'_c and v' satisfy $C_1(i)$, $C_2(i)$, and $C_3(i)$ for $i \geq c$, they certainly do for $i \geq c + 1$, and p'_k, \dots, p'_{c+1} is a feasible solution for v' .

We proceed by induction on $i \geq c$ to construct values p'_i, \dots, p'_0 satisfying the properties (1) $v' = \mathcal{E}_j(p'_j, \dots, p'_0)$, (2) $p'_j \geq 0$, and (3) $F'_j < F_j \leq 1$ for all $c \leq j \leq i$.

Suppose $i = c$. Lemma 4.6 and $v' > v$ imply that

$$p'_c = 1 - \left(\frac{v'k}{c}\right)^{\frac{1}{n-1}} < 1 - \left(\frac{vk}{c}\right)^{\frac{1}{n-1}} = p_c$$

and properties (1–3) clearly hold. In particular, $p'_c \geq 0$ since $\frac{c}{k} \geq v'$, and $F'_c = p'_c < p_c = F_c \leq 1$.

Suppose $i > c$ and we have constructed p'_{i-1}, \dots, p'_0 satisfying properties (1–3) for $c \leq j \leq i - 1$. Lemma 4.8 says that $1 - F_i$ is a nonnegative root of the polynomial $\mathcal{P}_i(x) - v$ where

$$\mathcal{P}_i(x) = a_{i,0} + a_{i,1}x + \dots + a_{i,n-1}x^{n-1}, \quad \text{and} \quad a_{i,j} = \frac{1}{k} \cdot \binom{n-1}{j} \cdot \sum_{\ell=1}^i (F_{i-\ell})^{n-1-j}.$$

We show that there is a positive root of the polynomial $\mathcal{P}'_i(x) - v'$ where

$$\mathcal{P}'_i(x) = a'_{i,0} + a'_{i,1}x + \dots + a'_{i,n-1}x^{n-1}, \quad \text{and} \quad a'_{i,j} = \frac{1}{k} \cdot \binom{n-1}{j} \cdot \sum_{\ell=1}^i (F'_{i-\ell})^{n-1-j},$$

and use this to construct the desired p'_i . These polynomials have several useful properties. First, $a'_{i,j} < a_{i,j}$ since $0 \leq F'_{i-\ell} < F_{i-\ell} \leq 1$ for all $\ell \geq 1$ by the induction hypothesis. Second, the $a_{i,j}$ are positive and the $a'_{i,j}$ are nonnegative for the same reason. Third, $a'_{i,n-1}$ is actually positive, since

$$a'_{i,n-1} = \frac{1}{k} \cdot \binom{n-1}{n-1} \cdot \sum_{\ell=1}^i (F'_{i-\ell})^0 = \frac{i}{k} > 0$$

because $i > c \geq 0$.

Descartes's Rule of Signs says that the number of positive real roots (counting multiplicities) of a polynomial is equal to the number of alternations in the signs of its nonzero coefficients minus an even number. We know that $1 - F_i$ is a nonnegative root of $\mathcal{P}_i(x) - v$. If $1 - F_i$ is zero, then the constant term $a_{i,0} - v$ of $\mathcal{P}_i(x) - v$ must be zero. If $1 - F_i$ is positive, then the constant term $a_{i,0} - v$ of $\mathcal{P}_i(x) - v$ must be negative, for if not then all nonzero coefficients of $\mathcal{P}_i(x) - v$ are positive and the number of

sign alternations is zero, contradicting the fact that $\mathcal{P}_i(x) - v$ has a positive root. Thus, in either case, $a_{i,0} - v \leq 0$, and $a'_{i,0} < a_{i,0}$ implies $a'_{i,0} - v < a_{i,0} - v = 0$, so the constant term of $\mathcal{P}'_i(x) - v$ must be negative. The nonzero coefficients of $\mathcal{P}'_i(x) - v$ therefore consist of positive coefficients followed by a negative constant term, so the number of sign alternations is one, and there is a positive root R'_i of $\mathcal{P}'_i(x) - v$. Let $p'_i = (1 - R'_i) - F'_{i-1}$ so that $1 - F'_i = R'_i$ is this root.

(1) We have $v' = \mathcal{P}'_i(1 - F'_i) = \mathcal{E}_i(p'_i, \dots, p'_0)$ by Lemma 4.8, since $1 - F'_i$ is a root of $\mathcal{P}'_i(x) - v'$.

(2) We have $p'_i > 0$ by Lemma 4.9.

(3) Suppose $F'_i \geq F_i$. Since $1 - F'_i$ is a positive root of $\mathcal{P}'_i(x) - v$, we have $0 \leq F'_i, F_i \leq 1$, and hence $0 \leq 1 - F'_i \leq 1 - F_i \leq 1$. Notice that $\mathcal{P}'_i(x)$ and $\mathcal{P}_i(x)$ are nondecreasing between 0 and 1 since the coefficients $a'_{i,j}$ and $a_{i,j}$ are nonnegative. Notice also that $\mathcal{P}'_i(x) \leq \mathcal{P}_i(x)$ between 0 and 1 since the coefficients satisfy $a'_{i,j} \leq a_{i,j}$. It follows that $v' = \mathcal{P}'_i(1 - F'_i) \leq \mathcal{P}_i(1 - F'_i) \leq \mathcal{P}_i(1 - F_i) = v$, contradicting $v' > v$, and we are done. ■

Lemma 4.11 immediately implies Theorem 4.4:

Proof of Theorem 4.4: If $v' > v$ are distinct equilibria, then $1 = F'_k < F_k = 1$, contradiction. ■


Lemma 4.11 also implies that we can use binary search to find the equilibrium value. The algorithm EQUILIBRIUM(n, k, ϵ) uses binary search to compute the equilibrium value of the (n, k) -game to within ϵ . It repeatedly calls FEASIBLE(n, k, v) to test whether there is a feasible solution for v . To prove correctness, we need only check that FEASIBLE(n, k, v) computes the feasible solution for v .

Lemma 4.12 FEASIBLE(n, k, v) returns p_k, \dots, p_c iff p_k, \dots, p_c is a feasible solution for v .


Proof: If the algorithm returns p_k, \dots, p_c , then $v = \mathcal{E}_i(p_i, \dots, p_0)$ and $p_i \geq 0$ and $F_i \leq 1$ for each $i = c, \dots, k$, since the algorithm would have returned “failure” if any of these conditions were false, so p_k, \dots, p_c is a feasible solution for v . Conversely, suppose p_k, \dots, p_c is a feasible solution for v . On each iteration i of the loop on line 3, line 3a computes the least positive solution to $v = \mathcal{E}_i(p_i, \dots, p_0)$, and Lemma 4.9 says this is precisely p_i ; and since p_i is part of a feasible solution and thus passes the test on line 3b, the algorithm returns the feasible solution p_k, \dots, p_c . ■

We can now prove the correctness of the algorithm.


Proof of Theorem 4.5: Let v be the equilibrium value. We will show that the algorithm preserves the invariant that (1) $l \leq v < r$, (2) there is a feasible solution for l , and (3) there is no feasible solution for r . Since the algorithm terminates when $l - r \leq \epsilon$, Part (1) of the invariant will imply the result. The base case for Part (1) follows from Lemma 4.1 and Line 1 of Algorithm EQUILIBRIUM(n, k, ϵ). Part (2) follows from Lemma 4.11. Part (3) follows from the fact that if there were a feasible solution for $r > v$ then there were a feasible solution for v with $F_k < 1$, contradiction to Lemma 4.7. For the induction step, suppose first that FEASIBLE(n, k, m) succeeds. Then by Lemma 4.12 there exists a feasible solution for $1 - m/k$ and hence, by Lemma 4.11, $v \geq 1 - m/k$, proving Part (1). Part (2) in this case follows from Lemma 4.12, and Part (3) follows from the induction hypothesis. If FEASIBLE(n, k, m) fails, then $v < 1 - m/k$, since otherwise, there exists a feasible solution to v with $F_k < 1$, contradiction to Lemma 4.7. This proves Part (1). Part (2) follows from the induction hypothesis and Part (3) from Lemma 4.12. In any case, the assignment made in Line 2b of the algorithm guarantees that the



The 3-player equilibrium strategy with game value 0.27.



The 4-player equilibrium strategy with game value 0.21.



The 5-player equilibrium strategy with game value 0.17.

Figure 4: *Equilibrium strategies and game values for the 3-, 4-, and 5-player games. On the left is the probability function, and on the right we zoom in dramatically (notice the y-axis scale) to show minute probability fluctuations for strategies near zero. Game values are approximations correct to within $\epsilon \approx .01$.*

invariant holds, and we are done. ■

4.3 Experimental results

Figure 4 on the next page gives the equilibrium strategies and approximate game values for 3, 4, and 5-player games with $k \approx 100$ and $\epsilon \approx 0.01$ as computed by an implementation of our algorithm in Mathematica. Notice that the game value decreases as the number of players increases, as predicted by Lemma 4.1. Notice also how the probabilities initially decrease and then increase for the 3-player game, in contrast to the 2-player game where the probabilities increase monotonically up to the cutoff point (recall Figure 1). Most interesting, however, are the minute oscillations in the probabilities for strategies near zero. On the right side of Figure 4 we have zoomed in and increased the resolution of the y -axis scale by a factor of 100 to show the oscillations. Some of this oscillation may be due to numerical errors in our Mathematica implementation, but moving from 2- to 3- to 4-player games we find 0, 1, and 2 extreme points in the probability functions. While we have no mathematical

explanation for this phenomenon at this point, we speculate that there is a subtle interaction among the players' strategies that would be interesting to explore in the future.

5 Applications

The visibility game captures an abstract notion of visibility: Each player chooses a time to perform an action that maximizes the time the action is visible to others.

In the commercial world, the action might be launching a new product in the stores. It is well known that the public has a short attention span. In the case of advertising, this means that one launch can be overshadowed by a following launch: The impact of the advertising campaign for the first launch can be lost in the noise and excitement surrounding the second launch. An advertiser might be inclined to minimize the negative impact of subsequent launches by timing its launch to maximize the time between launches.

In computer science, applications abound. As one example, a valuable resource is usually protected by having a processor write a value into a control register or acquire a lock. If the register can be overwritten or the lock can be preempted by another processor, then one processor might want to maximize its access to the resource by carefully choosing the time to acquire it. As another example, consider load balancing. Suppose servers compete for tasks (because they can charge customers for performing tasks) by updating a pointer directing a task to a server. In any symmetric equilibrium to our game, all players would see the same expected payoff, which would translate into all servers getting the same number of tasks.

In other settings, people might actually want to minimize visibility. With smoking frowned upon so widely in the United States, a community leader might want to choose a time to enter a smoking room that maximizes the time he or she will be alone in the room.

Our simple formulation of the visibility game does restrict its application in several ways.

First, players choose times in advance before the game starts. This means the choice is made "off-line" with regard to the evolution of the game once it has started. In the real world, the inability to respond immediately (which is inherent to any physical process) forces at least a part of the game to be played in an off-line fashion. As another example, consider a product whose release is a major event, meaning its release time must be planned far in advance.

Second, players choose times with little or no knowledge of the times chosen by other players, which is related to having players choose times in advance. Incomplete knowledge is a fact of life in computer science. In asynchronous distributed systems, processors are typically modeled as completely independent agents that make progress at unpredictable rates that are totally unrelated to each other. In reactive systems, processors respond to input from users potentially sitting in different room with no knowledge of the other users. Timing actions with incomplete knowledge occurs in other settings, too: In the smoking scenario, there may be many doors into the smoking room, and who knows who will enter or when?

Finally, two players choosing the same time are awarded a payoff of zero. We do consider other payoffs in the paper, but in computer science the payoff often is zero. In a wireless network, computers

transmitting at the same time simply generate noise in the network, and no communication happens. In some models of concurrent computation, two processors writing to the same register collide and both writes fail. Even in advertising, since the attention of the audience in a certain domain is typically dominated by a very small set of recent events, the payoff for a collision is very small.

6 Conclusion

Our results represent a first step toward understanding the effect of delayed actions on the outcome of timely games. These games arise naturally in many situations such as recommendation systems and other economic systems. To do that, we have defined and analyzed a simple game we called the visibility game. To the best of our knowledge, this is the first time this game is explicitly addressed. We have a fairly good understanding of the game for two players, and a general method to solve the game for n players. For example, consider the following extension: At every point in time, all previous actions are ordered from newest to oldest, and the payoff is the integral of a decreasing function of the ranking (the function for the basic game is 1 for the last player and 0 to everyone else). Our algorithm can be extended to this case. Other interesting directions include the following.

- Analyzing the numerical stability of our algorithm.
- Analyzing an on-line problem with delay considerations.
- Understanding the repeated game version.
- Analyzing the case where each player can write k times in a single time unit, for some $k > 1$.

References

- [1] M. R. Baye, D. Kovenock, and C. de Vries. A general linear model of contests. Mimeo, 1998. Unpublished manuscript available from <http://www.nash-equilibrium.com/baye/Contests.pdf>.
- [2] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [3] K. Hendricks, A. Weiss, and C. Wilson. The war of attrition in continuous time with complete information. *International Economic Review*, 29(4):663–680, Nov. 1988.
- [4] H. Hotelling. Stability in competition. *Economic Journal*, 39:41–57, 1929.
- [5] J. Maynard Smith. The theory of games and the evolution in animal conflicts. *Journal of Theoretical Biology*, 47:209–221, 1974.
- [6] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

Acknowledgment: We thank Ariel Rubinstein for very helpful comments.

Appendix: Additional Proofs

Proof of Lemma 3.2: Let F_1, F_2 denote the cumulative density functions (cdf's) of the density functions f_1, f_2 , respectively. Let v_1 be the expected payoff for player 1 at the equilibrium point. By Theorem 2.1, there exists a set \mathcal{Z} of measure 0 such that for all $x \in \text{support}(f_1) \setminus \mathcal{Z}$, we have

$$\begin{aligned}
 v_1 &= \int_0^1 u_1(x, y) f_2(y) dy \\
 &= \int_0^x (1-x) f_2(y) dy + \int_x^1 (y-x) f_2(y) dy \\
 &= (1-x) F_2(x) + \int_x^1 y f_2(y) dy - x(1-F_2(x)) \\
 &= F_2(x) - x + \int_x^1 y f_2(y) dy .
 \end{aligned}$$

Differentiating with respect to x , and applying the fundamental theorem of differential calculus, we get that for any $x \in \text{support}(f_1) \setminus \mathcal{Z}$ we have

$$0 = f_2(x) - 1 - x f_2(x) ,$$

and the lemma follows. ■

Proof of Lemma 3.3: First, note that $1 \notin \text{support}(f_1)$ since the pure strategy 1 is dominated by any other pure strategy. Next, we claim that the measure of the set $\text{support}(f_1) \setminus \text{support}(f_2)$ is zero: This is true since by Lemma 3.2 we have that with the exception of a set of measure 0, if $x \in \text{support}(f_1)$ then $f_2(x) > 0$ (since $0 \leq x < 1$) and hence $x \in \text{support}(f_2)$ as well. Repeating the symmetric argument with the players switched, we get that the measure of $\text{support}(f_2) \setminus \text{support}(f_1)$ is also zero, and the result follows. ■

Proof of Lemma 3.4: First we analyze the infimum. Suppose $\inf(\text{support}(f_1)) = \epsilon > 0$, and let v_1 be the expected payoff for Player 1 at the equilibrium point. By Lemma 3.3, $\inf(\text{support}(f_2)) = \epsilon$ too. By Theorem 2.1, we have that when Player 1 plays the pure strategy ϵ his expected payoff is

$$v_1 = \int_{\epsilon}^1 (y - \epsilon) f_2(y) dy$$

Hence, when Player 1 plays the pure strategy 0, his expected payoff is

$$\int_{\epsilon}^1 (y - 0) f_2(y) dy = v_1 + \epsilon \int_{\epsilon}^1 y f_2(y) dy > v_1 ,$$

contradiction to Theorem 2.1.

We now consider the supremum. Suppose first, for contradiction, that $\sup(\text{support}(f_1)) = x > 1 - v_1$. Then the expected payoff for Player 1 when playing x should be v_1 , but the payoff when playing x is never more than $1 - x < v_1$, contradiction. Next, suppose that $\sup(\text{support}(f_1)) = x < 1 - v_1$. Then the expected payoff for Player 1 when playing $\frac{x+1-v_1}{2}$ is always $1 - \frac{x+1-v_1}{2} > 1 - \frac{2(1-v_1)}{2} = v_1$, contradiction. ■

Proof of Lemma 3.5: Suppose, for contradiction, that there exist x_1 and x_2 such that $0 < x_1 < x_2 < 1 - v_1$ and $\int_{x_1}^{x_2} f_1(x)dx = 0$. Let $[x_1, x_2]$ be a maximal such interval, i.e., the integral of f_1 over any interval that contains $[x_1, x_2]$ is strictly positive. By Lemma 3.4, we may assume w.l.o.g. that $x_1, x_2 \in \text{support}(f_1)$. Hence, by Theorem 2.1, it must be the case that

$$\int_0^{1-v_1} u_1(x_1, y) f_2(y) dy = \int_0^{1-v_1} u_1(x_2, y) f_2(y) dy .$$

However, by definition of the payoff function, we have

$$\begin{aligned} \int_0^{1-v_1} u_1(x_2, y) f_2(y) dy & - \int_0^{1-v_1} u_1(x_1, y) f_2(y) dy \\ & = \int_0^{x_1} ((1-x_2) - (1-x_1)) f_2(y) dy + \int_{x_1}^{1-v_1} ((y-x_2) - (y-x_1)) f_2(y) dy \\ & = \int_0^{1-v_1} (x_1 - x_2) f_2(y) dy < 0 , \end{aligned}$$

contradiction. ■

Proof of Theorem 3.6: Fix a Nash equilibrium point. By Lemma 3.3, we may assume w.l.o.g. that both players employ the same mixed strategy f at that point, except perhaps for a set \mathcal{Z} of measure zero. By Lemmas 3.4 and 3.5, we have that $\text{support}(f) = [0, 1 - v]$, where v is the game value. We now argue that no point $x \in \text{support}(f)$ has positive probability. For suppose, towards contradiction, that there exists a point x_0 such that $\mathbf{P}[x_0] = \delta > 0$ under f . By our assumption that the number of such points is finite, there exists an $\epsilon' > 0$ such that $f(x) = \frac{1}{1-x}$ for $x \in \text{support}(f) \cap [x_0 - \epsilon, x_0 + \epsilon] \setminus \{x_0\}$. Let us assume that $x_0 < 1 - v$ first, and consider playing the pure strategies x_0 and $x_0 + \epsilon$ against f . Playing x_0 results in expected outcome

$$\begin{aligned} I_0 & = \int_0^{x_0^-} (1-x_0) f(y) dy + \int_{x_0^+}^1 (y-x_0) f(y) dy \\ & = \int_0^{x_0^-} (1-x_0) f(y) dy + \int_{x_0^+}^{x_0+\epsilon} (y-x_0) f(y) dy \\ & \quad + \int_{x_0+\epsilon}^1 (y-x_0) f(y) dy , \end{aligned} \tag{2}$$

where integrating up to x_0^- means that we are integrating over the half-open interval that does not include x_0 , and similarly when integrating from x_0^+ . On the other hand, when playing $x_0 + \epsilon$, we obtain

$$\begin{aligned} I_\epsilon & = \int_0^{x_0^-} (1-x_0-\epsilon) f(y) dy + \mathbf{P}[y=x_0] \cdot (1-x_0-\epsilon) \\ & \quad + \int_{x_0^+}^{x_0+\epsilon} (1-x_0-\epsilon) f(y) dy + \int_{x_0+\epsilon}^1 (y-x_0-\epsilon) f(y) dy \end{aligned} \tag{3}$$

It therefore follows that

$$\begin{aligned} I_\epsilon - I_0 & \geq \mathbf{P}[y=x_0] \cdot (1-x_0-\epsilon) - \epsilon \\ & \quad - \int_{x_0^+}^{x_0+\epsilon} (y-1+\epsilon) f(y) dy \\ & \geq \delta(1-x_0-\epsilon) - \epsilon , \end{aligned}$$

and hence, $\lim_{\epsilon \rightarrow 0^+} (I_\epsilon - I_0) = \delta(1 - x_0) > 0$. However, by Theorem 2.1, we must have $I_0 = I_\epsilon = v$, contradiction. If $x_0 = 1 - v$, we use the corresponding argument with $x_0 - \epsilon$. We can therefore conclude that $f(x) = \frac{1}{1-x}$ for *all* $x \in \text{support}(f)$, except perhaps for a set with f -measure is 0. Since $\int_0^{1-v} f(x)dx = 1$, we obtain that $\ln \frac{1}{1-v} = 1$, and hence $v = \frac{1}{e}$. ■