

Balls and Bins Models with Feedback

Eleni Drinea*
Harvard University
edrinea@deas.harvard.edu

Alan Frieze†
Carnegie Mellon University
alan@math.cmu.edu

Michael Mitzenmacher‡
Harvard University
michaelm@eecs.harvard.edu

Abstract

We examine generalizations of the classical balls and bins models, where the probability a ball lands in a bin is proportional to the number of balls already in the bin raised to some exponent p . Such systems exhibit positive or negative feedback, depending on the exponent p , with a phase transition occurring at $p = 1$. Similar models have proven useful in economics and chemistry; for example, systems with positive feedback ($p > 1$) tend naturally toward monopoly. We provide several results and useful heuristics for these models, including showing a bound on the time to achieve monopoly with high probability.

1 Introduction

There have been several recent instances in technology where a small number of companies compete in a market until one obtains a non-negligible advantage in the market share, at which point its share rapidly grows to a monopoly or near-monopoly. Economists have described this tendency toward monopoly in terms of *positive feedback* [12]. One loose explanation for this principle, commonly referred to as Metcalfe's Law, is that the inherent potential value of a system grows super-linearly in the number of existing users. For example, in a system with n users, there are $\binom{n}{2}$ possible pairwise connections, and this may more accurately reflect the value of the system than the number of users.

The video recording battle between VHS and Beta formats is often cited as a classic example of the power of positive feedback; VHS won out thanks to early user adoption, even though Beta was argued to be technically superior. In the Microsoft anti-trust trial, economists argued about the relationship of positive feedback and Microsoft's operating systems monopoly. Even the long-lasting dominance of the QWERTY keyboard has been ascribed to positive feedback [1, 12].

Similarly, there are situations where *negative feedback* occurs, so that a competitor with a larger market share has difficulty keeping its advantage. When large competitors suffer from inefficiencies, negative feedback is likely to occur [12].

In this paper, we provide a simple mathematical model that elucidates the power of positive and negative feedback. Our model is a non-linear generalization of classical balls and bins

*Supported in part by NSF Grant CCR-9983832 and NSF Grant ITR/SY-0121154.

†Supported in part by NSF Grant CCR-9818411.

‡Supported in part by an Alfred P. Sloan Research Fellowship, NSF Grant CCR-9983832, and NSF Grant ITR/SY-0121154.

models. While we developed this model independently, we have found since that these variations on standard balls and bins models have been known and applied by economists [1]. For example, it has previously been shown that under certain conditions positive feedback provably leads to monopoly in the limit [2]. From these limiting results, however, it is unclear how quickly monopoly will occur.

In this paper, we follow a more concrete approach, examining a specific family of models: the probability that a ball lands in a bin with x balls is proportional to x^p . We call p the *exponent* of the model. In the case where $p = 0$, this is just the standard model of throwing balls into bins independently and uniformly at random. In the case where $p = 1$, this is equivalent to the Pólya–Eggenberger model [8].

In economic terms, the model captures the effect of positive feedback in competitive situations. For example, suppose that there are two instant messaging services that do not interoperate well from which to choose. There is strong incentive to choose the service with more current users, as it offers more potential interactions. Of course this does not necessarily mean that all new users flock to a single system. We model the effect here as a probabilistic one, where new users are more likely to sign up to more popular services, and the strength of this feedback is governed by the exponent p .

We focus our analysis on the case of two bins. This is most interesting in practice; generally two companies are fighting to attract users for their competing systems [12]. Moreover, a simple union bound argument in Section 4 demonstrates that the problem of two bins encapsulates the significant behaviors.

It is well known in the case of $p = 1$ that if we start with two bins, each with one ball, the resulting distribution when there are n balls in the system is uniform; the probability of ending with k balls in the first bin is $1/(n - 1)$. More generally, it is clear that if one bin has a fraction q of the balls, it tends to maintain a fraction q of the balls in the future [7]. Positive feedback occurs when the exponent p is greater than 1. To see the difference in behavior when $p > 1$, note that if we start with one ball in each bin, the probability that a specific bin obtains *all* the balls is

$$\prod_{x=1}^{\infty} \left(1 - \frac{1}{1 + x^p} \right),$$

which for $p > 1$ is a constant depending on p . We demonstrate that for any constant exponent $p > 1$, any constant $\epsilon > 0$, and a sufficiently large number n of balls thrown, the probability that neither of the bins obtains a $1 - \epsilon$ fraction of the balls is inversely polynomial in n . The exact polynomial depends on ϵ and p . An interpretation of this statement is that monopoly occurs quickly with high probability. Similarly, negative feedback occurs when the exponent p is less than 1. For any constant $p < 1$, any constant $\epsilon > 0$, and a sufficiently large number of balls thrown n , the probability a bin obtains more than a $1/2 + \epsilon$ fraction of the balls is inversely polynomial in n . This result emphasizes the phase transition in this model at $p = 1$.

Our belief is that these non-linear balls and bins models, which naturally arise in economic, chemical, and biological systems, may also be useful for describing phenomena in computer science. As an example, we suggest how we may generalize random Web graph models using similar non-linear models. We also provide heuristics and calculation methods that may prove useful for analyzing such systems.

We wish to note that after submitting this paper, we learned of other work being done on this problem by Spencer and Wormald. They provide an elegant framework for the problem that gives many additional insights into the behavior of these types of systems, particularly in the case of many bins [13].

2 The case $p > 1$

We begin with some useful definitions.

Definition 1 *If there are n balls divided among m bins, we say that one bin has an ϵ -advantage if it has at least a $1/m + \epsilon$ fraction of the balls. Similarly, a bin is all-but- ϵ -dominant if it has at least a $1 - \epsilon$ fraction of the balls.*

Consider a fixed $p > 1$. In this section we cover the case of two bins. We will prove that when a ball lands in a bin with x balls with probability proportional to x^p , and we start with one ball in each bin, one bin becomes all-but- δ -dominant with probability q after n balls, where n is polynomial in q and $1/\delta$. We note that the starting point is chosen for convenience, and in Section 4 we use a simple union bound argument to extend the result to $m > 2$ bins.

Our proof follows a sequence of steps. We first show that one bin obtains an ϵ_0 -advantage for some ϵ_0 . From here, we show that the separation grows. Roughly, if we double the number of balls in the system, we increase the advantage by a constant factor (with high probability). We then show that if one bin becomes all-but- ϵ_1 -dominant for a sufficiently small ϵ_1 , the dominance improves (that is ϵ_1 shrinks) by a constant factor when we double the number of balls in the system. Putting it all together gives our result. We note that in what follows, we make no efforts to optimize the various constants used in the theorems.

2.1 Initial separation

We first show that if $p > 1$, the probability that neither of the bins gains an ϵ_0 -advantage is inversely polynomial in the number of balls thrown for some constant ϵ_0 . While this can be proven regardless of the initial state, for convenience we start with one ball in each bin.

Theorem 1 *Consider a system with exponent p and two bins B_0 and B_1 that begin with one ball each. Then there exist constants $\epsilon_0 > 0$ and $\gamma > 0$ such that after n steps, the probability that the two bins fail to ϵ_0 -separate is at most $O(n^{-\gamma})$.*

Proof: See the Appendix. □

2.2 Increasing advantage

Assume that B_0 (w.l.o.g.) has a constant ϵ -advantage over B_1 after n balls have been thrown into the system. Let $x(t)$ and $y(t)$ be the loads of B_0 and B_1 respectively when there are t balls in the system. We would like to say that as we continue throwing balls into the system, the probability of a ball going into B_0 is

$$\frac{x(n)^p}{x(n)^p + y(n)^p},$$

and use this to show that the advantage grows. This is not quite the case, however, since a new ball may go into B_1 , in which case the probability the next ball falls into B_0 sinks below $\frac{x(n)^p}{x(n)^p + y(n)^p}$.

To circumvent this issue, we consider throwing balls in waves of $\epsilon n/k$, for some $k \geq 1$. If we throw in $\epsilon n/k$ balls k/ϵ times, then the number of balls in the system doubles. Consider the first wave. Let X be the number of new balls that land in B_0 and Y the number of new balls than

land in B_1 . We underestimate the probability that a new ball lands in B_0 by assuming that all previous balls in the wave have landed in B_1 . Even in this worst case situation,

$$\frac{x(t)}{y(t)} \geq \left(\frac{\frac{1}{2} + \epsilon}{\frac{1}{2} - \frac{k-1}{k}\epsilon} \right)^p$$

for all t in $[n, n + \epsilon n/k]$. Hence, by use of Chernoff bounds, we find that with all but exponentially small probability,

$$\frac{X}{Y} \geq \left(\frac{\frac{1}{2} + \epsilon}{\frac{1}{2} - \frac{k-1}{k}\epsilon} \right)^p - \epsilon'$$

for some constant ϵ' . For n sufficiently large, we may take k large enough and ϵ' small enough so that the difference between $\frac{X}{Y}$ and $\left(\frac{\frac{1}{2} + \epsilon}{\frac{1}{2} - \epsilon} \right)^p$ is an arbitrarily small constant. Note that this implies that B_0 will continue to have an ϵ advantage over the next wave.

Suppose we show that $\left(\frac{1/2 + \epsilon}{1/2 - \epsilon} \right)^p > \frac{1/2 + \beta\epsilon}{1/2 - \beta\epsilon}$, for some $\beta > 1$. Then

$$\frac{x(2n)}{y(2n)} > \frac{\left(\frac{1}{2} + \epsilon \right) n + \left(\frac{1}{2} + \beta\epsilon \right) n}{\left(\frac{1}{2} - \epsilon \right) n + \left(\frac{1}{2} - \beta\epsilon \right) n} = \frac{\frac{1}{2} + \frac{1+\beta}{2}\epsilon}{\frac{1}{2} - \frac{1+\beta}{2}\epsilon}.$$

(Note that the arbitrarily small constant between $\frac{X}{Y}$ and $\left(\frac{1/2 + \epsilon}{1/2 - \epsilon} \right)^p$ will get absorbed.) Hence our ϵ -advantage increases to $\frac{1+\beta}{2}\epsilon$ once we double the number of balls, with high probability.

Theorem 2 *Suppose that B_0 has an $\epsilon \geq \epsilon_0$ advantage over B_1 when n balls are in the system. If we throw n more balls into the system then with high probability B_0 's advantage increases by a factor of $1 + \frac{(p-1)(1-2\epsilon)}{1+2\epsilon(p-1)}$.*

Proof:

$$\left(\frac{1/2 + \epsilon}{1/2 - \epsilon} \right)^p = \left(1 + \frac{4\epsilon}{1-2\epsilon} \right)^p > 1 + \frac{4p\epsilon}{1-2\epsilon} = \frac{1/2 + \alpha}{1/2 - \alpha}$$

where

$$\alpha = \left(1 + \frac{(p-1)(1-2\epsilon)}{1+2\epsilon(p-1)} \right) \epsilon.$$

□

So if for example, $\epsilon < 0.4$, the advantage, with high probability, increases by a factor of at least $1 + \frac{p-1}{5+4(p-1)}$.

2.3 To Complete Dominance

By Theorem 2, the advantage increases until one bin is all-but-0.1-dominant. At this point, a similar argument shows the dominance improves (that is, the initial 0.1 shrinks) geometrically.

Theorem 3 *If B_0 is all-but- ϵ_1 -dominant for $\epsilon_1 \leq 0.2$, then when we double the number of balls in the system, B_0 becomes all-but- $\frac{p+1}{2p}\epsilon_1$ -dominant with high probability.*

Proof: As before, by breaking the next group of balls into suitable blocks, we obtain that $\frac{X}{Y}$ can be made arbitrarily close to $\left(\frac{1-\epsilon_1}{\epsilon_1} \right)^p$ with high probability. Now if $\left(\frac{1-\epsilon_1}{\epsilon_1} \right)^p - \frac{1-\epsilon_1/p}{\epsilon_1/p} > 0$, then with high probability

$$\frac{x(2n)}{y(2n)} > \frac{(1-\epsilon_1)n + (1-\epsilon_1/p)n}{\epsilon_1 n + \epsilon_1 n/p} = \frac{1-\epsilon_1 \frac{p+1}{2p}}{\epsilon_1 \frac{p+1}{2p}},$$

proving the lemma.

Let $g(\epsilon, p) = \left(\frac{1-\epsilon}{\epsilon}\right)^p - \frac{1-\epsilon/p}{\epsilon/p}$. To show $g(\epsilon, p) > 0$ for $p > 1$ and $0 < \epsilon \leq 0.2$, we consider the function

$$\phi_p(\epsilon) = \frac{\epsilon(1-\epsilon)^p}{(p-\epsilon)\epsilon^p}.$$

We need to show that $\phi_p(\epsilon) > 1$ for $\epsilon < 0.2$. Now

$$\phi_p(0.2) = \frac{4^p}{5p-1} > 1$$

and taking logarithms and differentiating gives

$$\frac{\phi'_p(\epsilon)}{\phi_p(\epsilon)} = -\frac{p-1}{\epsilon} - \frac{p}{1-\epsilon} + \frac{1}{p-\epsilon} < 0.$$

Hence $\phi_p(\epsilon) > 1$ for $0 < \epsilon \leq 0.2$. □

2.4 Wrapping up

The following lemma estimates the number of balls in the system when B_0 's advantage is arbitrarily close to 1, or in other words, when B_0 is all-but- δ -dominant for an arbitrarily small constant δ . Suppose we start with B_0 having an ϵ_0 -advantage and n_0 balls in the system, as given in Theorem 1.

Theorem 4 *Assume that we throw balls into the system until B_0 is all-but- δ -dominant for some $\delta > 0$. Then, if $p > 1$, with probability $1 - e^{-\Omega(n_0)}$, B_0 is all-but- δ -dominant when the system has $2^{x+z} \cdot n_0$ balls, where $x = \log_{1+\frac{p-1}{5+4(p-1)}} \frac{0.4}{\epsilon_0}$ and $z = \log_{\frac{2p}{p+1}} \frac{0.1}{\delta}$.*

Proof: Recall that in each doubling stage, we succeed with all but exponentially small probability in the number of balls in the system, which is greater than n_0 . Each time we double the number of balls in the system, the initial advantage ϵ_0 increases by a factor of at least $1 + \frac{p-1}{5+4(p-1)}$ until it becomes 0.4; this requires $x = \log_{1+\frac{p-1}{5+4(p-1)}} \frac{0.4}{\epsilon_0}$ doubling stages. From then on, B_0 goes from all-but-0.1-dominant to all-but- δ -dominant, shrinking the fraction of balls not in B_0 by a factor of $\frac{p+1}{2p}$ with each doubling stage. Hence, we need $z = \log_{\frac{2p}{p+1}} \frac{0.1}{\delta}$ doubling stages until B_0 is all-but- δ -dominant. □

Essentially, our argument shows that once we achieve a little separation, the separation continues to grow with very high probability. In fact, the only reason our probability bounds are polynomial in the number of balls is because of the need to establish an initial gap in Theorem 1.

3 The case $p < 1$

In the case where $p < 1$, we have similar results, except that in this case the system tends to converge toward an equal number of balls in each bin. That is, we have negative feedback. For convenience, we consider only the case where $0 < p < 1$. (The case where $p \leq 0$ is trivial.)

Consider a fixed exponent $p < 1$ in a system with two bins, B_0 and B_1 . Suppose that n balls are in the system and B_0 (w.l.o.g.) has an ϵ_0 -advantage. We show that the advantage shrinks. We first show that if ϵ_0 is at least $1/\sqrt{2(p+1)(p+2)}$, the corresponding all-but- δ -dominance for B_0 increases. Once its advantage becomes sufficiently small, it decreases by a constant factor by throwing n more balls in the system.

Theorem 5 *Suppose that B_0 has an ϵ -advantage. If we throw n more balls in the system and $\epsilon \leq 1/\sqrt{2(p+1)(p+2)}$, B_0 's advantage decreases by a factor of $(3+p)/4$ with high probability. Otherwise, suppose B_0 is an all-but- ϵ -dominant, where $0 < \epsilon \leq \frac{1}{2} - \frac{1}{\sqrt{2(p+1)(p+2)}}$. If we throw n more balls in the system then B_0 becomes all-but- $\frac{p+1}{2p}\epsilon$ -dominant with high probability.*

Proof: The proof is similar to Theorem 2. We first consider when the advantage shrinks by the constant factor $(3+p)/4$. Using the idea of throwing balls in waves and Chernoff bounds as in Theorem 2, we see that the argument boils down to showing that the probability a ball lands in the most full bin, or $\left(\frac{1/2+\epsilon}{1/2-\epsilon}\right)^p$, is bounded above by $\frac{1/2+(1+p)\epsilon/2}{1/2-(1+p)\epsilon/2}$. Therefore it suffices to determine where $q(\epsilon, p) = \frac{1-(1+p)\epsilon}{1+(1+p)\epsilon} \cdot \left(\frac{1/2+\epsilon}{1/2-\epsilon}\right)^p < 1$. Note $q(0, p) = 1$.

We first show $q_p(\epsilon)$ is decreasing in ϵ . It is easier to look at $\log q_p(\epsilon)$, which decreases when $q_p(\epsilon)$ does. The derivative of $\log q_p(\epsilon)$ with respect to ϵ is $\frac{4p}{1-4\epsilon^2} - \frac{2(1+p)}{1-(1+p)^2\epsilon^2}$. It is straightforward to check that $q_p(\epsilon)$ is decreasing for $\epsilon < 1/\sqrt{2(p+1)(p+2)}$ and increasing past that point. Hence $q_p(\epsilon) < 1$ in the range $(0, 1/\sqrt{2(p+1)(p+2)}]$, and the advantage shrinks by a constant factor when we double the number of balls in the system for $\epsilon \leq 1/\sqrt{2(p+1)(p+2)}$.

Now suppose B_0 is all-but- ϵ -dominant. Here we follow Theorem 3. Let $g(\epsilon, p) = \left(\frac{1-\epsilon}{\epsilon}\right)^p - \frac{1-\epsilon/p}{\epsilon/p}$. We study $g(\epsilon, p)$ for $0 < p < 1$ and $\epsilon \in (0, \frac{1}{2} - \frac{1}{\sqrt{2(p+1)(p+2)}}]$. It is easy to check that $g(\epsilon, p)$ is increasing in ϵ and (by use of Maple) that $g(\frac{1}{2} - \frac{1}{\sqrt{2(p+1)(p+2)}}, p) < 0$, for all $0 < p < 1$, so the lemma is proved. \square

Theorem 5 can be used to show that from any non-trivial starting state, even if one bin has a large advantage, when $p < 1$ the system will quickly return to a near-equal state.

4 From Two to Many

We use the results for the case of two bins to obtain similar results for the case of many bins using standard union bounds. A key point is that if we look at a pair of bins from a system with many bins, the evolution of this pair of bins is just that of a system with exponent p . That is because when we condition on a ball landing in the pair of bins, the probability that it falls into a bin with x balls is still proportional to x^p . The following simple proof avoids any conditioning issues, and applies when $p > 1$.

Lemma 1 *Suppose that when n balls are thrown into a pair of bins, the probability that neither is all-but- ϵ -dominant is upper bounded by $p(n, \epsilon)$. Here we assume $p(n, \epsilon)$ is non-increasing in n . Then when $1+mn/2$ balls are thrown into m bins, the probability that none is all-but- γ -dominant is at most $\binom{m}{2}p(n, \epsilon)$ for $\gamma = \epsilon/(\epsilon + (1-\epsilon)/(m-1))$.*

Proof: Consider the two bins with the most balls, B_0 and B_1 , with B_0 having more balls. The two bins together have at least n balls since $1+mn/2$ total balls are thrown. If B_0 is not all-but- γ -dominant over all the bins, then it has less than a $1-\gamma$ fraction of the balls and B_1 has at least a $\gamma/(m-1)$ fraction of the balls. For the value of γ stated,

$$\frac{1-\gamma}{\gamma/(m-1)} = \frac{1-\epsilon}{\epsilon}.$$

Hence in this case B_0 is not all-but- ϵ -dominant when considering the pair of bins B_0 and B_1 . But the probability that there is a pair of bins where neither is all-but- ϵ -dominant is bounded above by $\binom{m}{2}p(n, \epsilon)$. \square

Essentially this lemma says that going from two bins to m bins increases the number of balls thrown by a factor and the probability that all-but- ϵ -dominance does not occur by polynomial factors in m . Hence the probability one bin fails to all-but- ϵ -dominate is inversely proportional to a polynomial in the number of balls thrown, the number of bins, and $1/\epsilon$.

A similar lemma applies for the case $p < 1$.

Lemma 2 *Suppose that when n balls are thrown into a pair of bins, the probability that one obtains an ϵ -advantage is upper bounded by $p(n, \epsilon)$. Here we assume $p(n, \epsilon)$ is non-increasing in n . Then when $1 + mn$ balls are thrown into m bins, the probability that one bin has a γ -advantage is at most $\binom{m}{2}p(n, \epsilon)$ for $\gamma = 4\epsilon(m - 1)/(m(m - 2(m - 2)\epsilon))$.*

Proof: Consider the bin with the most balls, B_0 , and the bin with the fewest balls, B_1 . The bin B_0 has at least n balls since $1 + mn$ total balls are thrown. If B_0 has γ -advantage, then it has at least a $1/m + \gamma$ fraction of the balls, and B_1 has at most a $1/m - \gamma/(m - 1)$ fraction of the balls. For the value of γ stated,

$$\frac{1/m + \gamma}{1/m - \gamma/(m - 1)} = \frac{1/2 + \epsilon}{1/2 - \epsilon}.$$

Hence in this case B_0 has an ϵ -advantage when considering the pair of bins B_0 and B_1 . But the probability that there is a pair of bins where one bin has an ϵ -advantage over the other is bounded above by $\binom{m}{2}p(n, \epsilon)$. \square

5 Relation to Web models

Our original motivation for studying this problem arose when we considered related dynamic Web graph models. Several recently proposed dynamic Web models are similar to balls and bins models, with the pages being bins and the links being balls. The difference for Web graph models is that new pages and links both enter the system; hence, new bins arise as new balls are thrown. Proposed Web models have all been linear; for example, in most models the probability a new page links to an extant page is proportional to its indegree [3, 5, 10, 11].

Recent Web models, while capturing certain properties of the Web graph, do not appear completely accurate. For example, recent studies suggest that the Web has many long, stringy pieces [4]. Also, certain Web sites contain central pages, that everything links to. Let us consider a dynamic Web graph model where a new page with one outedge links to an extant page with probability proportional to the indegree to the p th power. The limiting cases for this model are interesting: when $p \rightarrow \infty$, essentially all edges point to a single node, and when $p \rightarrow -\infty$, the graph is essentially a single path. It is possible that some areas of the Web may be similar to this more general model with properly chosen parameters. Further discussion of this issue is given in [6, 9]; however, it suggests that non-linear systems provide interesting variations of Web graph models.

6 A Useful Heuristic

In this section, we consider a heuristic that may prove useful in applications. Suppose we have two bins, whose load we denote by $x(t)$ and $y(t)$, where the time t denotes the number of balls in the system. As before the probability that the new ball thrown at time t falls in the bins with

$x(t)$ balls is $\frac{(x(t))^p}{(x(t))^p + (y(t))^p}$. Then the expected change in $x(t)$, or $\Delta x(t)$, satisfies

$$\Delta x(t) = E[x(t+1) - x(t)] = \frac{(x(t))^p}{(x(t))^p + (y(t))^p},$$

and similarly for $y(t)$. Using the heuristic approximation

$$\frac{\Delta y(t)}{\Delta x(t)} = \frac{dy}{dx}$$

and dropping the t from the notation where the meaning is clear, we obtain the following approximation for the expected behavior of the system:

$$\frac{dy}{dx} = \frac{y^p}{x^p}.$$

This heuristic demonstrates the different types of behavior to be expected when $p < 1$, $p = 1$, and $p > 1$. When $p = 1$, the solution has the form $y = cx$. Otherwise, the solution has the form $y^{1-p} = x^{1-p} + c$. When $p < 1$, regardless of the initial values of x and y the limiting ratio of y/x goes to 1; in the long run, the two bins each contain roughly half of the balls. When $p > 1$, the limiting ratio of y/x goes to 0 or infinity.

This heuristic is appealing in that it allows us to approximate the behavior when $p > 1$ of the bins that are dominated. Specifically, let us consider more closely the case where the initial loads of the bins are $x(t_0)$ and $y(t_0)$ (with $y(t_0) > x(t_0)$) and $p > 1$. Then the solution has the form $y^{1-p} = x^{1-p} + y(t_0)^{1-p} - x(t_0)^{1-p}$. As $y \rightarrow \infty$, our heuristic suggests that x approaches $(x(t_0)^{1-p} - y(t_0)^{1-p})^{1/(1-p)}$. For example, consider the case where $x(300) = 100$, $y(300) = 200$, and $p = 2$. The heuristic suggests that even as the number of balls thrown grows to infinity, the expected value of $x(t)$ will only grow to about 200.

We point out that this heuristic is (at this point) just a heuristic. While in some cases differential equations can properly be used to determine the behavior of a system, the utility in this case is less clear. For example, from any starting point, there is some constant (though perhaps small) probability that the smaller bin will overtake the larger. From smaller starting points (say $x(3) = 1$ and $y(3) = 2$) there is more variation. Hence this heuristic is really valuable for determining the limiting behavior only when one bin dominates another sufficiently so that the probability that it is overtaken can essentially be dismissed.

We consider the performance of the heuristic with some examples. When $x(300) = 100$, $y(300) = 200$, and $p = 2$, the solution of the resulting differential equation is

$$\frac{1}{y} = \frac{1}{x} - \frac{1}{200}.$$

When there are 10,000 balls in the system, the differential equations predict $x(10,000) = 196$. Exact calculations show that the mean value of $x(10,000)$ is actually just above 197, although the mode is 190. More visually, Figure 1 shows the distribution for $x(10,000)$ is very concentrated; it looks close to a normal distribution, although it is asymmetric with a small probability of large values. Larger numbers of balls show similar behavior; for $x(100,000)$ and $x(1,000,000)$, which have essentially the same distribution, the mean is 201 though the distribution peaks at 195.

Figure 1 displays similar results for $p = 1.5$. The differential equations predict $x(10,000)$ should be about 637, which also is very accurate. They also predict that as the number of balls grows to infinity, $x(t)$ should converge to approximately 1,165, which is close to $x(1,000,000)$.

Note this heuristic approach can easily be extended to the case of more than two bins. It would be interesting to develop a more formal statement in terms of probabilistic bounds based on this heuristic.

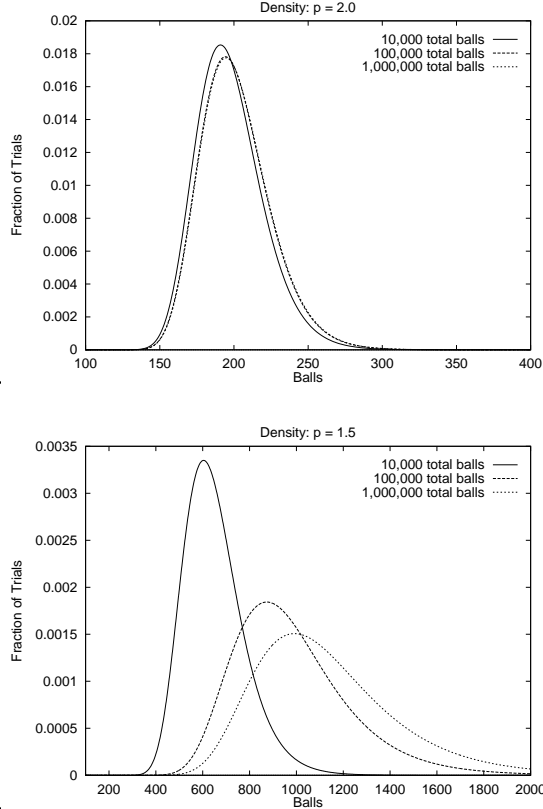


Figure 1: Density functions, starting with 100 balls in one bin and 200 in the other, $p = 2.0$ and $p = 1.5$.

7 Examples of Reaching Monopoly

We present some examples in order to demonstrate typical behavior for the $p > 1$ case, giving exact results determined by extensive numerical calculations with the appropriate recurrence. Specifically, if $w(x, y)$ is the probability of having x balls in B_0 and y balls in B_1 when there are $x + y$ balls in the system, then

$$w(x, y) = w(x - 1, y) \frac{(x - 1)^p}{(x - 1)^p + y^p} + w(x, y - 1) \frac{(y - 1)^p}{x^p + (y - 1)^p}.$$

The reason for showing these examples is to suggest that the number of balls necessary to converge to monopoly can be extremely large, especially for smaller values of p . This provides some evidence that the character of our result, namely that monopoly fails to happen with probability inversely polynomial in the number of balls in the system (and moreover with a small exponent), is correct. We point out that we do not currently have any bound that demonstrates that this probability could not fall exponentially with the number of balls; this remains an open question.

In Figure 2, we present the cumulative distribution for the number of balls in a bin when we begin with one ball in each bin, and place balls until 1,000,000 balls are in the system, using $p = 1.1$. While there is significant bias towards the periphery, there is still a reasonable




Figure 2: Cumulative distribution function, starting with 1 ball in each bin, $p = 1.1$ and 1,000,000 total balls.




Figure 3: Cumulative distribution function, starting with 1 ball in each bin, $p = 2.0$ and 1,000 total balls.

probability that one bin will not completely overwhelm the other. For example, the probability that one bin contains over 80% of the balls is less than 80%.

In contrast, consider the case of just 1,000 balls when $p = 2$ in Figure 3. Here almost all the weight lies in the area where one bin has almost all of the balls. The probability that one bin contains five or fewer balls is 0.864. This concentration, however, is a function of the dramatic effect of inequality at the beginning of the process; leading two or three balls to one is a huge advantage. If we begin with 1,000 balls in each bin, and place balls until there are 1,000,000 in the system, we see that while there is clear tendency toward monopoly, it appears more similar to the $p = 1.1$ case.

8 Conclusion

We have analyzed simple non-linear balls and bins models, where the probability of a new ball going to a bin with x balls is proportional to x^p . We have demonstrated a phase transition at $p = 1$; fast convergence to monopoly for $p > 1$; and fast convergence toward equality when $p < 1$.

We suggest a few problems worthy of future study that this framework introduces. First, it




Figure 4: Cumulative distribution function, starting with 1,000 balls in each bin, $p = 2.0$ and 1,000,000 total balls.

seems likely that our current arguments can be improved and simplified. In particular, a better understanding of the initial separation stage and a tighter argument for more than two bins might be helpful. Second, the impact of the initial conditions should be clarified. When two bins begin with nearly the same number of balls, how does the difference affect the probability that each will end up dominating the system? What is the distribution of the final state of the other bin? While we have heuristic approaches to this problem, rigorous bounds would be useful. Third, consideration of other natural families of functions besides x^p may be useful for real systems. In a similar vein, understanding systems where the function determining the probability that a ball goes in a bin may vary according to time may allow more realistic models.

References

- [1] B. Arthur. **Increasing Returns and Path Dependence in the Economy**. The University of Michigan Press, 1994.
- [2] B. Arthur, Y. Ermoliev, and Y. Kaniovski. On Generalized Urn Schemes of the Pólya Kind. English translation in *Cybernetics* 19:61-71, 1983.
- [3] A.-L. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks. *Physica A*, vol. 272, pages 173-189, 1999.
- [4] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graph structure in the Web: experiments and models. In *Proc. of the 9th World Wide Web Conference*, 2000. Available at <http://www9.org/w9cdrom/index.html>.
- [5] C. Cooper and A. Frieze. On a general model of Web graphs. To appear in ESA 2001.
- [6] E. Drinea, M. Enachescu, and M. Mitzenmacher. Variations on random graph models of the Web. Harvard Computer Science Technical Report TR-06-01.
- [7] B. Hill, D. Lane, and W. Sudderth. A Strong Law for Some Generalized Urn Processes. *Annals of Probability*, Vol. 8, Issue 2 (Apr.,1980), 214-226.
- [8] N. Johnson and S. Kotz. **Urn models and their application: an approach to modern discrete probability theory**. New York, Wiley, 1977.

- [9] P. L. Krapivsky and S. Redner. Organization of growing random networks. *Physical Review E*, 63, 066123-1 – 066123014, 2001.
- [10] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the Web graph. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pp. 57-65, 2000.
- [11] R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Extracting large scale knowledge bases from the Web. In *Proc. of the 25th VLDB Conference*, 1999.
- [12] C. Shapiro and H. R. Varian. **Information Rules**, Harvard Business School Press, Boston, MA, 1999.
- [13] J. Spencer and N. Wormald. Explosive processes. Draft manuscript.

Appendix

Theorem 1 Consider a system with exponent p and two bins B_0 and B_1 that begin with one ball each. Then there exist constants $\epsilon_0 > 0$ and $\gamma > 0$ such that after n steps, the probability that the two bins fail to ϵ_0 -separate is at most $O(n^{-\gamma})$.

Proof: We sketch the proof, which follows the same outline as Theorems 2 and 4. First, recall that when $p = 1$ and we begin with one ball in each bin, the resulting distribution after n balls are thrown is uniform. A simple coupling argument shows that when $p > 1$ the distribution of the number of balls in a bin has more weight at the extremes. Hence for any n_0 the probability that after n_0 balls are thrown neither bin has at least $1/2 + n_0^{3/4}$ balls is $O(n_0^{-1/4})$.

We build on this small advantage using a repeated doubling argument. Suppose that when we have $n_1 \geq n_0$ balls in the system and we throw n_1 more balls, the advantage grows by a constant factor with probability $e^{-n_1^a}$ for some constant $a > 0$. Then choose any suitable constant ϵ_0 , say $\epsilon_0 = 1/100p$. Then after $O(\log n_0)$ doublings, we obtain a constant ϵ_0 advantage with probability $O(n_0^{-1/4})$, and we have a polynomial in n_0 number of balls in the system.

We must take a bit more care in the Chernoff bounds to obtain the high probability result in the doubling stages. However, if we start a doubling phase with n_1 balls in the system and one bin having $\frac{1}{2}n_1 + x$ balls, where $x \geq n_1^{3/4}$, it suffices to throw the next n_1 balls in blocks of size $n_1^{5/8}$. The probability a ball lands in the bin with more balls is at least

$$z = \frac{\left(\frac{1}{2}n_1 + x\right)^p}{\left(\frac{1}{2}n_1 - x + n_1^{5/8}\right)^p + \left(\frac{1}{2}n_1 + x\right)^p}.$$

For n_0 sufficiently large, the $n_1^{5/8}$ term above does not affect that we expect the advantage to grow by a constant factor over the next n_1 balls, as in Theorem 2.

Moreover, by Chernoff's bounds, inductively over each block of size $n_1^{5/8}$, if X is the number of balls that go into the bin with more balls,

$$\Pr[X \leq zn_1^{5/8} - n_1^{1/2}] \leq \exp(-\Omega(n_1^{3/8})).$$

The $n_1^{1/2}$ term does not affect that the advantage grows by a constant factor with high probability for suitably large n_0 , and hence the theorem holds. \square