Reinforced Random Processes in Competitive Systems*

Paul Kattuman† Weiye Yang‡

September 20, 2014

Abstract

We study dominance in competitive systems with positive feedbacks. The conventional view is that relatively insignificant and randomly occurring imbalances between the shares of competitors can, in the presence of positive feedbacks, lead the system to a path that decisively favours the one that gained early advantage. This is based on non-linear feedback. In this paper we consider the firm targetting some specific market share (e.g., 60%), with sufficiently high probability (e.g., 0.9) and using Pólya's linear urn process, examine the path to such intermediate levels of dominance over finite or infinite horizons. We determine the trade-off between *initial market presence* and *feedback strength* in locking-in to any specified degree of dominance .

Keywords: Market Dominance, Stochastic Dominance, Positive Feedback, Polya's Urn, Initial Asymmetry, Feedback strength

*Draft. We should be grateful for comments. Please do not quote without permission.

†University of Cambridge, Judge Business School, Trumpington Street, Cambridge CB2 1AG, UK. Email: p.kattuman@jbs.cam.ac.uk; Tel: 44(0)1223764136

‡University of Cambridge, DPMMS, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK. Email: weiye.yang@cantab.net

1 Introduction

The competitive diffusion of innovative products, technologies, ideas and practices is of great importance in a world abounds in them. There are many reasons why positive feedback characterises consumer choice among competing alternatives in markets for innovative goods. they include word-of-mouth publicity, social learning, cost-advantages of larger scales of operation, and network externalities – by which each customer gains more value the more others have made the same choice, and therefore tends to make her choice to accord with the choice made by others. In influential theoretical models of markets with positive feedback, these and similar factors make the future adoption rate highly elastic with respect to present market share, and monopoly is the certain eventual market outcome. The process whereby intermediate degrees of dominance emerges in markets is of great practical and academic interest and remains underexplored.

In the extant literature the certain eventual monopoly outcome is generally explained in terms of non-linear positive feedback driving an adoption process (Arthur, 1989). Under these types of self-reinforcement, the market outcome is history-dependent in the sense that small share differences early on can be decisive in picking out the winner in the long run. The battle between QWERTY and DVORAK keyboard formats has been explained as a case of an *inefficient* QWERTY being picked out by accidents of history, and emerging due to positive feedbacks from the economies of learning-by-doing and learning-by-using (David 1985). Other often mentioned examples include the battle between CP/M, DOS and Macintosh among operating systems, and between VHS and Betamax among video-recorder formats, which we examine in section 8. The focus in the literature has been on the potential for market failure, in the sense of the potential for inefficient firms, technologies, products to come to dominate the market, even when accounts of the dominance of the inefficient have been contested (Liebowitz and Margolis, 1999).

In this paper we consider the firm targetting some specific market share (e.g., 60%), with sufficiently high probability (e.g., 0.9) and using Pólya's linear urn process, examine the path to such intermediate levels of dominance over finite or infinite horizons. We also consider the fact that in feedback strength is a variable. Feedback may be positive, but not very strong in some markets. Intuition suggests that if an competitor has a large lead in its initial market share, then it should not take as strong a positive feedback to eventually attain a specified degree of dominance, compared to the situation where this agent has a smaller lead in initial share. How much does a greater initial advantage compensate for smaller feedback strength in leading to a specific equilibrium? The explicit consideration of the trade-off involved between initial market presence and feedback strength, for chosen dominance target focusses attention on the pay-off to competing for various extents of initial market share, conditional on the strength of positive feedback.

2 The Pólya urn Model

Thus the urn can represent a market with two competing firms. The initial numbers of balls of each colour represent the initial sizes of the two firms. The strength of feedback, S, is common to both competing agents, i.e. it is independent of the colour of the sampled ball. The proportions of balls of the two different colours in the urn evolve stochastically depending on the sequence of sampled ball types, and the replacement rule. This history dependence of the process is reflected in the way the distribution of the proportions of balls of different colours in the urn change over time.

The limiting equilibrium of a Pólya process with a linear replacement rule has a continuous distribution. More general urn models, with non-linear replacement rules which reflect nonlinear (positive) feedback, can potentially have discrete distributions as limiting equilibria.

2.1 Notation

The basic notation and ideas can be introduced using a model with two competing agents, represented by two colours: black and white. Without loss of generality we will focus on the evolution of the proportion of black balls in the urn. We now proceed formally:

Notation 2.1. Let $\mathbb{Z}_+ := \{x \in \mathbb{Z} : x \geq 0\}$. Note that this differs from $\mathbb{N} = \{x \in \mathbb{Z} : x > 0\}$.

Let $n \in \mathbb{Z}_+$ index time, i.e. the rounds over which sampling (and replacement) occur.

Let B_0 be the initial number of black balls in our two-colour urn (i.e. at time $n = 0$) and let W_0 likewise be the initial number of white balls. Let S be the number of additional balls replaced according to the replacement rule.

For modelling positive feedback, $S > 0$. Let B_n and W_n be the number of black and white balls respectively in the urn after n rounds. Denote by Y_n the proportion of black balls in the urn at round n, that is,

$$Y_n := \frac{B_n}{B_n + W_n}. \tag{2.1}$$

The limiting proportion, whose existence we will shortly prove, is then defined by $Y_\infty := \lim_{n\to\infty} Y_n$.

For $n \geq 0$ we inductively define

$$\begin{aligned} B_{n+1} &:= B_n + S1_{\{U_{n+1} \leq Y_n\}} \\ W_{n+1} &:= W_n + S1_{\{U_{n+1} > Y_n\}}, \end{aligned} \tag{2.2}$$

where 1_A is the indicator function of the event A, and the random variables $(U_n)_{n\geq 1}$ are independent and identically distributed with distribution Unif$[0, 1]$. $\{U_{n+1} \leq Y_n\}$ is thus the event of drawing a black ball in round n. These uniform draws correspond to drawing a black ball with probability Y_n, independently of past draws. The probability comes from the random variable U_{n+1}, which is the only source of randomness in going from round n to round $n + 1$. Notice the probability of drawing a black ball at any time is equal to the proportion of black balls in the urn - it is a *linear* urn model.

2.2 Feedback strength and initial asymmetry in the urn process

The initial proportion of black balls Y_0 measures the degree of *initial asymmetry* in the urn.

The number of balls returned each period after sampling from the urn, according to the replacement policy, represents the strength of feedback. The different values of S differentiate dynamic processes and induce different limiting distributions: $S > 0$ models positive feedback in a growth process; $S = 0$ is a degenerate process ($B_n = B_0$ and $W_n = W_0$ for all n); and $S < 0$ is a model of negative feedback. The large family of urn models is reviewed in Pemantle (2006).

Normalizing S with the total initial number of balls in the urn, we define the *feedback strength* of the urn as the expression

$$\frac{S}{B_0 + W_0}. \tag{2.3}$$

This of course is the growth rate of the market in the initial period.

2.3 Pólya's Result

Proposition 2.2. *The random variables Y_n converge almost surely as $n \to \infty$ to a limit Y_∞ where*

$$Y_\infty \sim \text{Beta}\left(\frac{B_0}{S}, \frac{W_0}{S}\right). \tag{2.4}$$

In particular, when $B_0 = W_0 = S$, the limit variable Y_∞ is uniform on the interval $[0, 1]$.

It is an interesting and non-intuitive result that Pólya's urn has a random limit. The proof of the above proposition goes back to Pólya (1930); see also Freedman (1965).

2.4 Multivariate generalisation of the Pólya urn model

Athreya (1969) showed that the two colour result above generalises to the case with any number $d \geq 2$ of colours, with the shares of colours in this generalised Pólya urn following a generalised Beta distribution (i.e. a Dirichlet distribution) of order d over the unit $(d-1)$-simplex. We can specify a general d-variate Pólya urn process by the parameters (x_0, S), where the constant $x_0 \in (0, \infty)^d$ is the vector of initial numbers of balls of each colour in the urn and the constant $S > 0$ is the feedback. Notice that we have now generalised to an abstract urn in which the numbers of balls of each colour do not have to be integers; they must merely be positive.

Notation 2.3. If $X = (X_n^1, \ldots, X_n^d)_{n \geq 0}$ is a d-variate Pólya urn process with parameters (x_0, S), then we use the notation

$$X \sim \text{PU}(d; x_0, S). \tag{2.5}$$

The random variable X_n^i then denotes the amount of colour i in the urn at time n.

For example, our previous two-colour urn $(B_n, W_n)_{n\geq 0}$ can now be expressed as the distribution $\mathrm{PU}(2; (B_0, W_0), S)$. The version of Athreya's result that we will be using can be stated as follows:

Proposition 2.4. *Let $X \sim \mathrm{PU}(d; x_0, S)$ and let the process $Y = (Y_n^1, \ldots, Y_n^d)_{n\geq 0}$ be defined for each $i \in \{1, \ldots, d\}$ and $n \geq 0$ by*

$$Y_n^i = \frac{X_n^i}{\sum_{j=1}^d X_n^j}, \tag{2.6}$$

that is, Y_n^i represents the proportion of balls in the urn with colour i at time n. Then the process (Y_n) converges almost surely as $n \to \infty$ to a limit Y_∞ with a Dirichlet distribution, in particular

$$Y_\infty \sim \mathrm{Dir}\left(\frac{x_0}{S}\right). \tag{2.7}$$

2.5 Aggregation property

One of the most useful features of the Dirichlet distribution is its *aggregation property*, proven by Frigyik *et al.* (2010):

Proposition 2.5. *Suppose $Z \equiv (Z^1, \ldots, Z^d) \sim \mathrm{Dir}(\alpha)$, where $\alpha = (\alpha^1, \ldots, \alpha^d)$. Let $\{A_1, \ldots, A_r\}$ be a partition of $\{1, \ldots, d\}$. Then*

$$\left(\sum_{j\in A_1} Z^j, \ldots, \sum_{j\in A_r} Z^j\right) \sim \mathrm{Dir}\left(\sum_{j\in A_1} \alpha^j, \ldots, \sum_{j\in A_r} \alpha^j\right). \tag{2.8}$$

In our competing firms analogy, this would be consistent with the idea of two (or more) firms merging in an attempt to capture a larger share of the market.

The aggregation property can be used to prove that the marginals of a Dirichlet distribution are Beta distributions. Specifically, if $Z \sim \mathrm{Dir}(\alpha)$ is d-variate, then for $i \in \{1, \ldots, d\}$,

$$Z^i \sim \mathrm{Beta}\left(\alpha^i, \sum_{j\neq i}^d \alpha^j\right). \tag{2.9}$$

It follows that if we are only concerned with the limiting market share of a single specific firm, then it is enough to consider a market with only two competitors: the firm we are concerned with, and a corporate group consisting of all of the other firms. This we do in section 3.

2.6 Critique

The linear Pólya urn process is a highly structured model with a specific form of history dependence. The probability of adding S balls of a particular colour is linear with respect to the proportion of that colour.

Generalisations where the linear urn function is varied to allow the probability of an addition to a colour to be an arbitrary function of the proportion of all colours, and in addition the urn function is allowed to vary in a structured way with time, have been explored by Hill, Lane and Sudderth (1980) and by Arthur, Ermoliev, Kaniovski (1983, 1984). In this class of non-linear Pólya processes, the equilibrium proportion of each agent is dynamically selected from among the fixed points of the mapping from proportions to probabilities. When the stable fixed points of the non-linear urn function occur only at the unit vectors, the equilibrium will yield monopoly to one of the agents almost surely. Lock-in to dominance is swift in the non-linear Pólya urn process but the choice between agents will depend on history.

This is unlike the linear process where each equilibrium is a continuous distribution over the unit interval - the full range of shares. The parameters of the process, initial asymmetry and feedback strength, can weight the equilibrium probabilistically towards any degree of dominance, ranging from monopoly in the limit to one colour, all the way to symmetry between colours. Depending on "history", the linear Pólya process could converge to any proportion of balls. But as the number of balls in the urn increases, the current proportions will grow more stable, and balls will continue to be selected in the same proportions. Initially, each round of addition of balls to the urn has a large influence on the probability of choice of colour of the next batch of S balls and the positive feedback will dominate. Eventually however the number of balls will have grown so large that the next batch of S balls has negligible effect on the proportion of colours in the urn. This also implies that the market growth rate (and thus feedback strength) declines continuously over time, converging to 0. The market movement is ever towards saturation.

Finally, it is worth emphasizing that the linear model is Markov. That is, for the outcome in round n, the precise sequence in which the balls accumulated by round $n-1$ have been drawn does not matter. The long-run equilibrium, looking forward from any date, depends only upon the set of occurrences that have happened beforehand, and not upon their chronological order. Independence of equilibrium from this order makes it easier to obtain predictions of the future looking forward from any chosen date.

3 Stochastic dominance

Definition 3.1. Let Z_1 and Z_2 be real-valued random variables. Let F_1 and F_2 be their respective cumulative distribution functions.

1. Z_1 has *first-order stochastic dominance* over Z_2 if $F_1(x) \leq F_2(x)$ for all $x \in \mathbb{R}$, with strict inequality at some x.

2. Z_1 has *second-order stochastic dominance* over Z_2 if $\int_{-\infty}^{x} F_1(u)du \leq \int_{-\infty}^{x} F_2(u)du$ for all $x \in \mathbb{R}$, with strict inequality at some x.

In this section we obtain stochastic dominance results on the limiting distributions of markets that can be modelled using linear Pólya processes. We will only be considering the limiting share of a single competitor in the market, so by the aggregation property we may use the bivariate process $\mathrm{PU}(2;(B_0,W_0),S)$ without loss of generality. Let $\alpha = \frac{B_0}{S}$ and $\beta = \frac{W_0}{S}$. The limiting market share of the firm B is then, as before, the random variable $Y_\infty \sim \mathrm{Beta}(\alpha,\beta)$. Our objective is a partial ordering on the set of parameter pairs $(\alpha,\beta) \in (0,\infty)^2$ for this limiting random variable, in the sense of stochastic dominance. In other words, given two parameter pairs for Y_∞, which would be preferred by firm B?

Since the Beta distribution is supported only in $[0,1]$, in this case we can replace "$x \in \mathbb{R}$" in the above definitions of stochastic dominance with "$x \in [0,1]$", and likewise replace "$\int_{-\infty}^{x}$" with "$\int_{0}^{x}$". Throughout this section, let $\mathrm{Beta}(\alpha_1,\beta_1)$ and $\mathrm{Beta}(\alpha_2,\beta_2)$ be two Beta distributions with $(\alpha_1,\beta_1) \neq (\alpha_2,\beta_2)$. Let $I(x;\alpha_1,\beta_1)$ and $I(x;\alpha_2,\beta_2)$ be their respective cumulative distribution functions, i.e.

$$I(x;\alpha,\beta) = \frac{1}{B(\alpha,\beta)} \int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt. \tag{3.1}$$

We look for conditions under which stochastic dominance arises between these two distributions. Define the difference function δ by

$$\delta(x) = I(x; \alpha_1, \beta_1) - I(x; \alpha_2, \beta_2). \tag{3.2}$$

We now state and prove an extremely important theorem:

Theorem 3.2. *If either $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, or $\alpha_1 < \alpha_2$ and $\beta_1 < \beta_2$, then there exists a unique $x_* \in (0,1)$ such that $I(x_*; \alpha_1, \beta_1) = I(x_*; \alpha_2, \beta_2)$. Otherwise, there exists no $x \in (0,1)$ such that $I(x; \alpha_1, \beta_1) = I(x; \alpha_2, \beta_2)$.*

Proof in Appendix A.1.

3.1 First-order stochastic dominance

Proposition 3.3. *Suppose $\alpha_1 \geq \alpha_2$ and $\beta_1 \leq \beta_2$. Then* $\mathrm{Beta}(\alpha_1, \beta_1)$ *has first-order stochastic dominance over* $\mathrm{Beta}(\alpha_2, \beta_2)$.

Proof in Appendix A.2.

This Proposition is intuitive, since in the standard two-colour Pólya urn, adding black balls or removing white balls at the start should give the black balls an advantage. Note also that first-order dominance always implies second-order dominance.

3.2 Second-order stochastic dominance

The indefinite integral of the CDF of the Beta distribution is given by

$$\int_0^x I(u; \alpha, \beta) du = \left(x - \frac{\alpha}{\alpha + \beta} \right) I(x; \alpha, \beta) + \frac{x^\alpha (1-x)^\beta}{(\alpha+\beta) B(\alpha, \beta)} \tag{3.3}$$

for $x \in [0,1]$, and can be checked simply by differentiation. Notice in particular that

$$\int_0^1 I(u; \alpha, \beta) du = \frac{\beta}{\alpha + \beta}. \tag{3.4}$$

Lemma 3.4. *Suppose $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, and let $x_* \in (0,1)$ be the unique intersection point of $I(x; \alpha_1, \beta_1)$ and $I(x; \alpha_2, \beta_2)$. Then $\delta'(x_*) > 0$.*

Proof in Appendix A.3.

Remark 3.5. It follows that δ is negative in a lower neighbourhood of x_*, and positive in an upper neighbourhood of it. Thus in this case there is no first-order stochastic dominance.

Proposition 3.6. *Suppose $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$. Then*

- $\mathrm{Beta}(\alpha_2, \beta_2)$ *does not have second-order stochastic dominance over* $\mathrm{Beta}(\alpha_1, \beta_1)$.
- $\mathrm{Beta}(\alpha_1, \beta_1)$ *has second-order stochastic dominance over* $\mathrm{Beta}(\alpha_2, \beta_2)$ *if and only if*

$$\frac{\alpha_1}{\alpha_1 + \beta_1} \geq \frac{\alpha_2}{\alpha_2 + \beta_2}. \tag{3.5}$$

Proof in Appendix A.4.

Remark 3.7. The expression $\frac{\alpha}{\alpha+\beta}$ is exactly the mean of the $\mathrm{Beta}(\alpha, \beta)$ distribution.

Remark 3.8. By a simple rearrangement of Proposition 3.6, if $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, we have that

$$\frac{\alpha_1}{\alpha_1 + \beta_1} \leq \frac{\alpha_2}{\alpha_2 + \beta_2} \tag{3.6}$$

if and only if $\mathrm{Beta}(\beta_1, \alpha_1)$ has second-order stochastic dominance over $\mathrm{Beta}(\beta_2, \alpha_2)$. The interpretation of this in our two-colour Pólya urn model is that increasing the initial amounts of both colours in the urn will always benefit at least one of the colours, and could potentially benefit both! This is perhaps slightly counterintuitive, but explainable in the following way: the variance of the $\mathrm{Beta}(\alpha, \beta)$ distribution is given by

$$\frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}. \tag{3.7}$$

In particular, observe that if we set $\beta = k\alpha$ for a fixed k, and increase α, the variance will decrease. Therefore if both firms (colours) are risk-averse, then increasing the parameters of the limiting distribution while keeping its mean fixed will decrease the variance, and thus benefit both firms (colours).

3.3 Further remarks

Notation 3.9. Since first- and second-order stochastic dominance both define partial orders on the set of parameter pairs $(0, \infty)^2$ of the Beta distribution, we will use the notation

$$(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2) \tag{3.8}$$

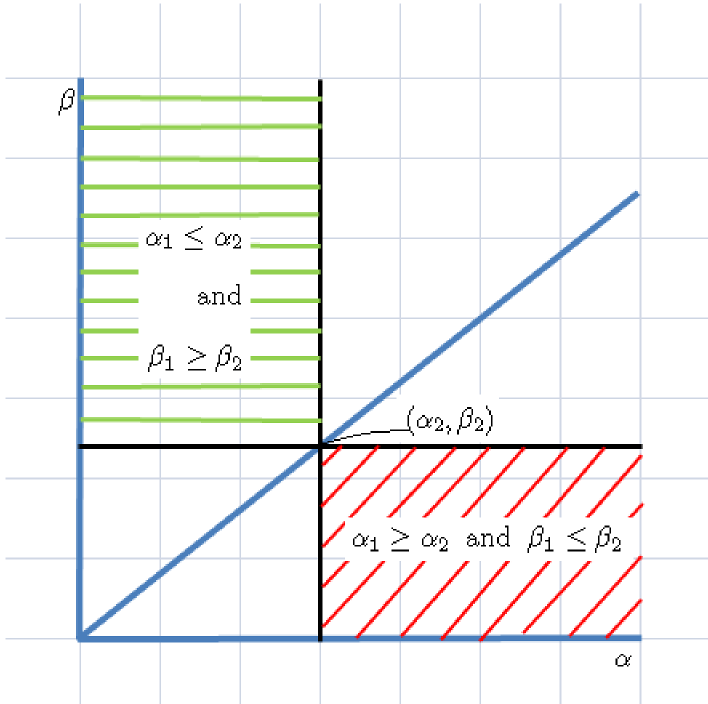

Figure 1: Regions of (α_1, β_1) for which $(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)$ or $(\alpha_1, \beta_1) <_1 (\alpha_2, \beta_2)$. The boundary points of the regions are contained in them. To be precise, the point (α_2, β_2) should not actually be a member of either region.

to mean that Beta(α_1, β_1) has first-order stochastic dominance over Beta(α_2, β_2), and likewise we will use

$$(\alpha_1, \beta_1) >_2 (\alpha_2, \beta_2) \tag{3.9}$$

for second-order stochastic dominance.

We may be able to better visualise the orderings we have defined in this section by *fixing* a pair $(\alpha_2, \beta_2) \in (0, \infty)^2$ and looking at the *sets* of parameter pairs (α_1, β_1) for which, for example, $(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)$.

Figure 1 displays the sets of parameter pairs (α_1, β_1) for which first-order stochastic dominance exists in some direction with respect to some fixed pair (α_2, β_2). The lower-right region is the set $\{(\alpha_1, \beta_1) : (\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)\}$, whereas the upper-left region is $\{(\alpha_1, \beta_1) : (\alpha_1, \beta_1) <_1 (\alpha_2, \beta_2)\}$.

Likewise, Figure 2 displays the sets of parameter pairs (α_1, β_1) for which second-order

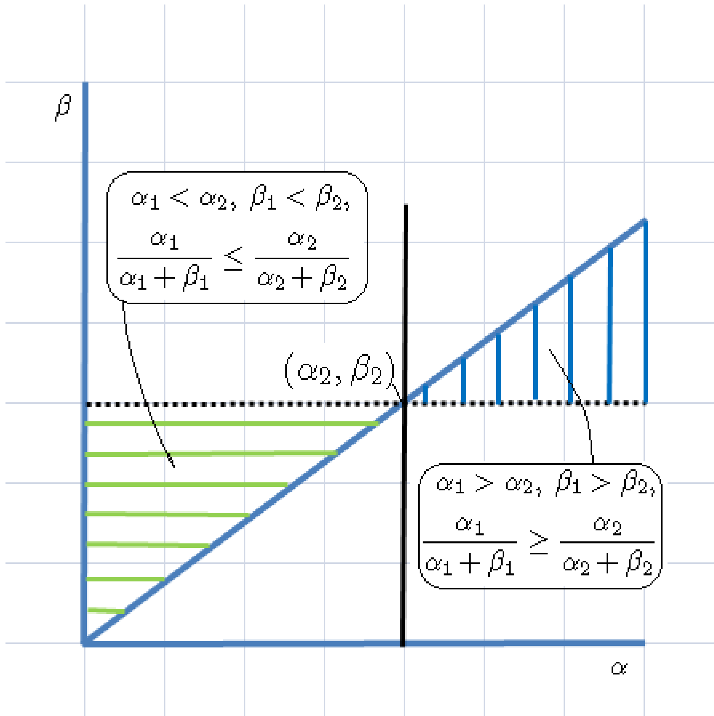

Figure 2: Additional regions of (α_1, β_1) for which $(\alpha_1, \beta_1) >_2 (\alpha_2, \beta_2)$ or $(\alpha_1, \beta_1) <_2 (\alpha_2, \beta_2)$, but for which no first-order dominance exists. The regions contain the points on their diagonal boundaries, but not the points on their horizontal boundaries.

stochastic dominance exists in some direction with respect to (α_2, β_2), but first-order dominance does not (recall that first-order dominance implies second-order dominance). The upper-right region is $\{(\alpha_1, \beta_1) : (\alpha_1, \beta_1) >_2 (\alpha_2, \beta_2)\}$, whereas the lower-left region is $\{(\alpha_1, \beta_1) : (\alpha_1, \beta_1) <_2 (\alpha_2, \beta_2)\}$.

Remark 3.10. Using first- and second-order stochastic dominance, the only case in which two Beta distributions are not ordered is when one has higher parameters but a lower mean than the other. Which one is preferred will come down to the details of the specific utility function used.

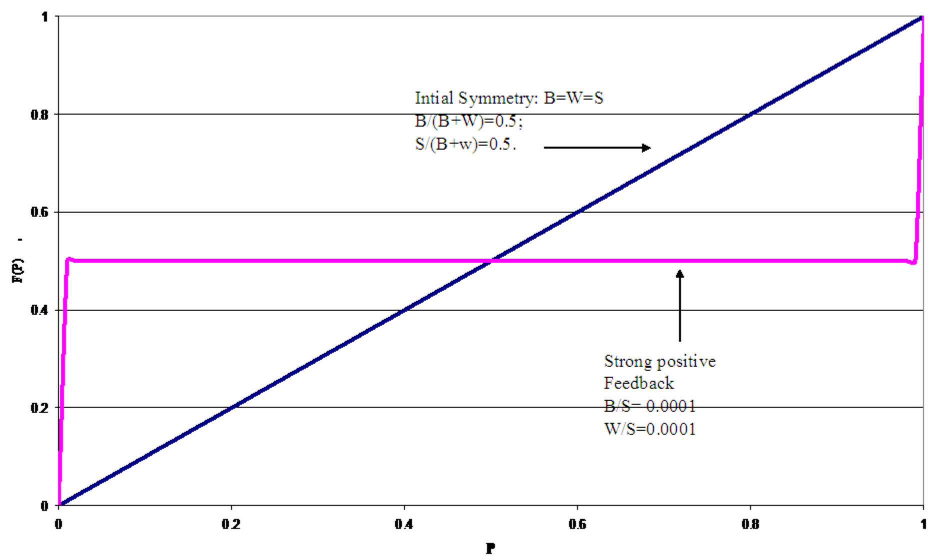

Figure 3: Polya Process: Asymptotic CDFs of proportion of black balls - Extreme cases.

4 Probabilistic market dominance

Competition law is concerned with the abuse of 'dominant market position', which is conventionally defined in terms of market shares; for example, by the US Department of Justice, and likewise in Article 82 of the EC treaty. From the point of view of an evolving market, a valid view that can be taken is of the outcome in the indefinite future. If we define dominance as "agent i eventually monopolises the market", then the corresponding event is

$$\left\{Y^i_\infty = 1\right\}. \tag{4.1}$$

As noted in section 1, this can occur with positive probability in, for example, a non-linear bivariate Pólya process when the stable fixed points of the urn function are at 0 and/or 1. In a linear bivariate Pólya process the limiting random variable has a continuous distribution so the above event is a null set. However, taking S very large relative to B_0 and W_0 results in arbitrarily high probabilities of near-monopoly outcomes, as illustrated in Figure 3.

A more general and applicable notion of dominance should permit us characterise the evolving market, looking to the future, as heading for dominance if an agent is on track to exceed some specified (high) share, with some stipulated probability. There must of course be some basis specifying what amounts to a "high" market share and the related probability, in order to pin down this notion of dominance.

Remark 4.1. We stress the difference between the notion of dominance to be introduced in this section and the notion of stochastic dominance in section 3. The two concepts are unrelated.

We define dominance as having occurred if the share of a given colour is sufficiently likely to reach a (high) value. This probability value can be specified depending on how conservative we wish to be in defining dominance; the values $10\%, 5\%$ and 1% are conventional significance levels, but other values may be more appropriate in different contexts. For example, if a parameter pair in the bivariate Pólya process $\mathrm{PU}(2; (B_0, W_0), S)$ gives at least 95% probability that the limiting share of colour B (black) is at least 60% then we may conclude that the system is heading towards 60% dominance by B at the 5% level. More generally, denoting the significance level for defining dominance by $p \in (0, 1)$ and the target share by $x \in (0, 1)$, we have dominance for colour B if

$$\mathbb{P}\left[Y_\infty \geq x\right] \geq 1 - p, \tag{4.2}$$

where Y_∞ is as before the limiting proportion of black balls.

5 The iso-dominance function

The probabilistic definition of dominance requires us to specify a *pair* of numbers $(x, p) \in (0, 1)^2$. We say that dominance occurs if, with probability at least $1 - p$, a firm will have a limiting share of the market of at least x.

Let us consider a firm's preferences over (and potential efforts to control) initial parameters of the market (by the aggregation property we can assume again that the market only contains two firms without loss of generality). This corresponds to preferences over the parameters of the distribution of its limiting market share. Suppose that the firm would like to attain a limiting share of at least $x \in (0,1)$ with probability $1-p \in (0,1)$. Our question is then: given this constraint, what choice of parameters does the firm have?

Mathematically, our question is equivalent to the following: given the pair (x, p), what is the set of Beta distribution parameters $(\alpha, \beta) \in (0, \infty)^2$ such that $I(x; \alpha, \beta) = p$? The aim of this section is to derive a few of the properties of this set.

In Özçağ *et al.* (2008) a proof is given that, for $x \in (0,1)$ and all (α, β), all of the partial derivatives of the incomplete Beta function

$$B_x(\alpha, \beta) := \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt \tag{5.1}$$

with respect to (α, β) exist. In fact, all we really need from this result is continuity. Since the Beta function itself is well known to be continuous for $(\alpha, \beta) \in (0, \infty)^2$, it follows that the CDF of the Beta distribution

$$I(x; \alpha, \beta) \equiv \frac{B_x(\alpha, \beta)}{B(\alpha, \beta)} \tag{5.2}$$

is continuous in (α, β) for all $x \in (0,1)$ and all $(\alpha, \beta) \in (0, \infty)^2$. We now prove a lemma concerning exactly how $I(x; \alpha, \beta)$ varies with (α, β):

Lemma 5.1. *Fix $x \in (0,1)$ and $\alpha \in (0, \infty)$. Then*

- $\lim_{\beta \to 0} I(x; \alpha, \beta) = 0,$
- $\lim_{\beta \to \infty} I(x; \alpha, \beta) = 1.$

Proof in Appendix A.5.

Remark 5.2. The intuition behind this lemma is simple. Consider our two-colour urn. Reducing the initial number of white balls should push the CDF of Black's limiting proportion to the right (i.e. to higher values) and increasing the initial number of white balls should do the opposite. Now when we talk about CDFs, which are increasing functions, moving to the right is the same as moving downwards. Thus we would expect I to decrease when we decrease β.

We can now state and prove the following important theorem:

Theorem 5.3. *Let $(x, p) \in (0, 1)^2$ and $\alpha \in (0, \infty)$. Then there exists a unique $\beta \in (0, \infty)$ such that*

$$I(x; \alpha, \beta) = p. \tag{5.3}$$

Proof in Appendix A.6.

Definition 5.4. We will denote the unique β in the above theorem by $\beta_{x,p}(\alpha)$, and view it as a family of functions of α parametrised by (x, p). Formally, we thus have the function

$$\beta_{x,p} : (0, \infty) \to (0, \infty) \tag{5.4}$$

for each $(x, p) \in (0, 1)^2$. We will call $\beta_{x,p}$ an *iso-dominance function* and we will call its graph in the (α, β)-plane an *iso-dominance curve*.

Finding an analytical expression for $\beta_{x,p}(\alpha)$ in terms of α appears to be non-trivial in general, and the only obvious case is $(x, p) = (\frac{1}{2}, \frac{1}{2})$, for which $\beta_{\frac{1}{2},\frac{1}{2}}(\alpha) = \alpha$.

Proposition 5.5 (Properties of $\beta_{x,p}$)**.** *Let $(x, p) \in (0, 1)^2$. Then the function $\beta_{x,p}$ is strictly increasing and continuous.*

Proof in Appendix A.7.

Remark 5.6 (Digression into duality)**.** Suppose we define *point sets* $\mathcal{D} = (0, 1)^2$ and $\mathcal{P} = (0, \infty)^2$, and define *line sets*

$$\begin{aligned} L(\mathcal{D}) &= \{\{(x, I(x; \alpha, \beta)) : x \in (0, 1)\} \subseteq \mathcal{D} : (\alpha, \beta) \in \mathcal{P}\}, \\ L(\mathcal{P}) &= \{\{(\alpha, \beta_{x,p}(\alpha)) : \alpha \in (0, \infty)\} \subseteq \mathcal{P} : (x, p) \in \mathcal{D}\}. \end{aligned} \tag{5.5}$$

That is, "lines" are simply the graphs of functions of the form $I(\cdot; \alpha, \beta)$ or $\beta_{x,p}$. Then by construction, the *incidence structures* $(\mathcal{D}, L(\mathcal{D}), \in)$ and $(\mathcal{P}, L(\mathcal{P}), \in)$ are isomorphic to each other's duals.

5.1 Limiting properties

In this section we seek to prove a number of results about the asymptotic properties of iso-dominance functions. We must first however prove a few results about the convergence properties of the Beta distribution. Fix $(x, p) \in (0,1)^2$ and consider the family of functions $I(\cdot; \alpha, \beta_{x,p}(\alpha))$ parametrised by α. In other words, we are considering the family of Beta CDFs that pass through the point (x, p).

Let $(X_\alpha)_{\alpha \in (0,\infty)}$ be a family of random variables such that, for each $\alpha \in (0, \infty)$,

$$X_\alpha \sim \text{Beta}(\alpha, \beta_{x,p}(\alpha)). \tag{5.6}$$

Proposition 5.7 (Small-α degenerate limit)**.** *If $x' \in (0,1)$, then*

$$\lim_{\alpha \to 0} I(x'; \alpha, \beta_{x,p}(\alpha)) = p. \tag{5.7}$$

(Equivalently: As $\alpha \to 0$, X_α converges in distribution to a Bernoulli random variable X_0 with $\mathbb{P}[X_0 = 0] = p$.)

Proof in Appendix A.8.

Proposition 5.8 (Large-α degenerate limit)**.** *If $x' \in (0, x)$, then*

$$\lim_{\alpha \to \infty} I(x'; \alpha, \beta_{x,p}(\alpha)) = 0. \tag{5.8}$$

If $x' \in (x, 1)$, then

$$\lim_{\alpha \to \infty} I(x'; \alpha, \beta_{x,p}(\alpha)) = 1. \tag{5.9}$$

(Equivalently: As $\alpha \to \infty$, X_α converges in distribution to a degenerate random variable X_∞ satisfying $X_\infty = x$ almost surely.)

Proof in Appendix A.9.

We have now done enough preliminary work to prove the result that we've been aiming for:

Corollary 5.9 (First-order asymptotics of the iso-dominance function)**.**

$$\begin{aligned} \lim_{\alpha \to 0} \frac{\beta_{x,p}(\alpha)}{\alpha} &= \frac{p}{1-p}, \\ \lim_{\alpha \to \infty} \frac{\beta_{x,p}(\alpha)}{\alpha} &= \frac{1-x}{x}. \end{aligned} \tag{5.10}$$

Proof in Appendix A.10.

Remark 5.10. Note that the previous corollary implies in particular that $\lim_{\alpha \to 0} \beta_{x,p}(\alpha) = 0$ and $\lim_{\alpha \to \infty} \beta_{x,p}(\alpha) = \infty$. It follows that the iso-dominance function is a bijection. In other words, the iso-dominance curve represents a one-to-one relationship between α and β: for each $\alpha \in (0, \infty)$ there exists a unique $\beta \in (0, \infty)$ such that (α, β) lies on the curve, and for each β there exists a unique α such that (α, β) lies on the curve.

5.2 Pairs of Dominance conditions

Let us return to our competing firms. Suppose a firm decides it would like to attain a limiting share of x_1 with probability $1 - p_1$ as before, but it now specifies *another* pair (x_2, p_2) corresponding to an additional dominance condition. Is it possible to satisfy both conditions at once? Mathematically, given two points $(x_1, p_1), (x_2, p_2) \in (0, 1)^2$, is there a Beta distribution whose CDF passes through both of these points? The requirements for this to occur will seem rather familiar.

Proposition 5.11. *Let $(x_1, p_1), (x_2, p_2) \in (0, 1)^2$ with $(x_1, p_1) \neq (x_2, p_2)$. If either $x_1 > x_2$ and $p_1 > p_2$, or $x_1 < x_2$ and $p_1 < p_2$, then there exists a unique pair $(\alpha_*, \beta_*) \in (0, \infty)^2$ such that $I(x_1; \alpha_*, \beta_*) = p_1$ and $I(x_2; \alpha_*, \beta_*) = p_2$. Otherwise, there exists no pair $(\alpha, \beta) \in (0, \infty)^2$ such that $I(x_1; \alpha, \beta) = p_1$ and $I(x_2; \alpha, \beta) = p_2$.*

Proof in Appendix A.11.

Remark 5.12. Proposition 5.11 can be stated in the following equivalent way in terms of iso-dominance functions:

Let $(x_1, p_1), (x_2, p_2) \in (0, 1)^2$ with $(x_1, p_1) \neq (x_2, p_2)$. If either $x_1 > x_2$ and $p_1 > p_2$, or $x_1 < x_2$ and $p_1 < p_2$, then there exists a unique $\alpha_* \in (0, \infty)$ such that $\beta_{x_1,p_1}(\alpha_*) = \beta_{x_2,p_2}(\alpha_*)$. Otherwise, there exists no $\alpha \in (0, \infty)$ such that $\beta_{x_1,p_1}(\alpha) = \beta_{x_2,p_2}(\alpha)$.

Notice the similarities between this and Theorem 3.2. This is an example of a symmetry between the dual incidence structures $(\mathcal{D}, L(\mathcal{D}), \in)$ and $(\mathcal{P}, L(\mathcal{P}), \in)$ defined in Remark 5.6.

5.3 Asymmetry-feedback trade-off

Again we look at the bivariate case $\mathrm{PU}(2; (B_0, W_0), S)$ without loss of generality. Let $\alpha = \frac{B_0}{S}$ and $\beta = \frac{W_0}{S}$, the parameters associated with the Beta distribution of Y_∞, which is as before the limiting market share of firm B. Recall the definitions of *initial asymmetry* $\frac{B_0}{B_0+W_0}$ and *feedback strength* $\frac{S}{B_0+W_0}$. There is then a homeomorphism between the set of pairs $(\alpha, \beta) \in (0, \infty)^2$ and the set of pairs $(\frac{B_0}{B_0+W_0}, \frac{S}{B_0+W_0}) \in (0, 1) \times (0, \infty)$ given by

$$\begin{aligned} \frac{B_0}{B_0 + W_0} &= \frac{\alpha}{\alpha + \beta}, \\ \frac{S}{B_0 + W_0} &= \frac{1}{\alpha + \beta}, \end{aligned} \tag{5.11}$$

and

$$\begin{aligned} \alpha &= \left(\frac{B_0}{B_0 + W_0} \right) \cdot \left(\frac{S}{B_0 + W_0} \right)^{-1}, \\ \beta &= \left(1 - \frac{B_0}{B_0 + W_0} \right) \cdot \left(\frac{S}{B_0 + W_0} \right)^{-1}. \end{aligned} \tag{5.12}$$

It follows that we may instead plot our iso-dominance curves on a graph of $\frac{B_0}{B_0+W_0}$ against $\frac{S}{B_0+W_0}$ without losing any information - it is simply a change of variables. This provides us with the trade-off between initial asymmetry and feedback strength.

Remark 5.13. The set of points described by an iso-dominance curve in the $(\frac{B_0}{B_0+W_0}, \frac{S}{B_0+W_0})$-plane may not in general represent a bijection between $\frac{B_0}{B_0+W_0}$ and $\frac{S}{B_0+W_0}$. For example, if $(x, p) = (\frac{1}{2}, \frac{1}{2})$ then every point of the corresponding iso-dominance curve satisfies $\frac{B_0}{B_0+W_0} = \frac{1}{2}$. However, it is the case that for any $(x, p) \in (0, 1)^2$, the function

$$\alpha \mapsto \frac{1}{\alpha + \beta_{x,p}(\alpha)} \tag{5.13}$$

is a strictly decreasing bijection that maps $(0, \infty) \to (0, \infty)$. Now suppose that we have an iso-dominance curve in the $(\frac{B_0}{B_0+W_0}, \frac{S}{B_0+W_0})$-plane. Recalling the expression of $\frac{S}{B_0+W_0}$ in terms of α and β, it follows that for every value of $\frac{S}{B_0+W_0} \in (0, \infty)$, there exists a *unique* value of $\frac{B_0}{B_0+W_0} \in (0, 1)$ such that this pair of values lies on the curve.

Illustration is provided in Figures 4 and 5.

6 Finite-time case: Dirichlet-multinomial distribution

Useful real-life applications of dominance analysis cannot rely on the luxury of infinite time. We hence seek to derive finite-time analogues to the limiting results of section 1.

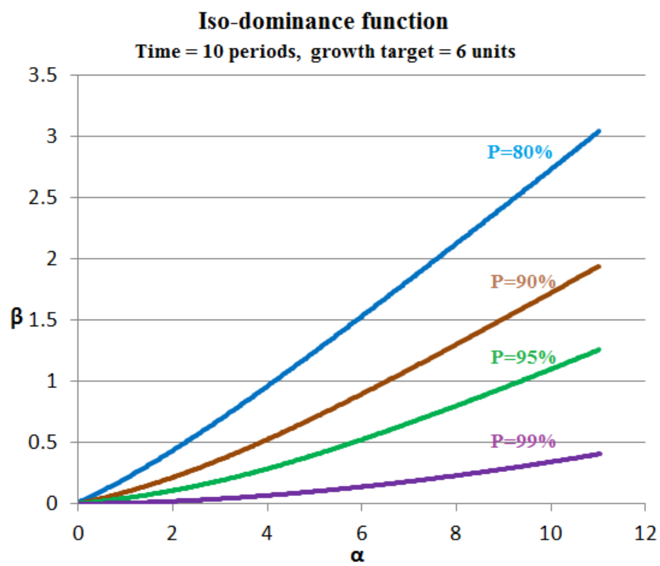

Figure 4: Iso-dominance map: Finite time case

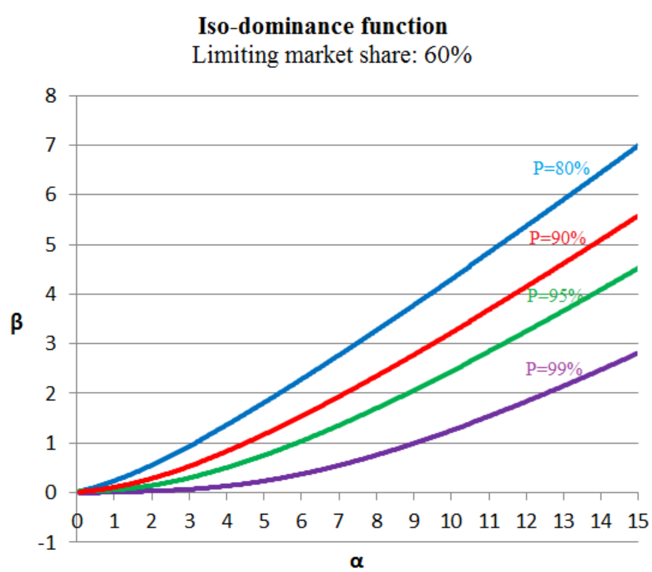

Figure 5: Iso-dominance map: Limiting Market share of 60%

For this, we turn to the Dirichlet-multinomial (DM; also known as the multivariate Pólya) distribution. This is a family parametrised by $\{(n, \alpha) : n \in \mathbb{N}, d \in \mathbb{N}, \alpha \in (0, \infty)^d\}$. It can be interpreted as a d-variate multinomial distribution $\mathrm{Multin}(n, P)$ which has a d-variate random parameter $P \sim \mathrm{Dir}(\alpha)$, and is the multivariate generalisation of the perhaps more well-known Beta-binomial distribution. Its probability mass function (see Mosimann (1962)) is given as follows: let $Z \sim \mathrm{DM}(n, \alpha)$, and let $z \in \mathbb{Z}_+^d$ such that $\sum_{j=1}^d z^j = n$. Then

$$\mathbb{P}[Z = z] = \frac{n!}{\prod_{j=1}^d (z^j!)} \frac{\Gamma(A)}{\Gamma(n + A)} \prod_{j=1}^d \frac{\Gamma(z^j + \alpha^j)}{\Gamma(\alpha^j)} \tag{6.1}$$

where $A = \sum_{j=1}^d \alpha^j$. For any other value of z the probability is 0.

Proposition 6.1. *Let $X \sim \mathrm{PU}(d; x_0, 1)$ be a Pólya urn process with unit feedback. Then for all $n \geq 1$,*

$$X_n - x_0 \sim \mathrm{DM}(n, x_0). \tag{6.2}$$

Proof in Appendix A.12.

It is fairly simple to generalise this proposition to all positive feedbacks S with a scaling argument:

Proposition 6.2. *Let $X \sim \mathrm{PU}(d; x_0, S)$. Let $k > 0$ be a constant. Then*

$$kX \sim \mathrm{PU}(d; kx_0, kS). \tag{6.3}$$

Proof in Appendix A.13.

Our desired result immediately follows:

Corollary 6.3. *Let $X \sim \mathrm{PU}(d; x_0, S)$. Then for all $n \geq 1$,*

$$\frac{X_n - x_0}{S} \sim \mathrm{DM}\left(n, \frac{x_0}{S}\right). \tag{6.4}$$

In addition to the fact that it deals directly with the fixed-time distributions of a Pólya process, an advantage of DM over Dirichlet is that it is a discrete distribution with a known probability mass function. It may thus be easier to implement numerically than the Dirichlet distribution. Finally, the DM distribution inherits the aggregation property (section 2.5) from the Dirichlet distribution:

Proposition 6.4 (Aggregation)**.** *Suppose* $Z \equiv (Z^1, \ldots, Z^d) \sim \mathrm{DM}(n, \alpha)$*, where* $\alpha = (\alpha^1, \ldots, \alpha^d)$*. Let* $\{A_1, \ldots, A_r\}$ *be a partition of* $\{1, \ldots, d\}$*. Then*

$$\left(\sum_{j \in A_1} Z^j, \ldots, \sum_{j \in A_r} Z^j \right) \sim \mathrm{DM}\left(n, \left(\sum_{j \in A_1} \alpha^j, \ldots, \sum_{j \in A_r} \alpha^j \right) \right). \tag{6.5}$$

Proof in Appendix A.14.

As with the Dirichlet distribution, the marginals of the Dirichlet-multinomial distribution follow from the aggregation property, and are, in this case, Beta-binomial (BB) distributions. If $Z \sim \mathrm{DM}(n, \alpha)$ is d-variate, then for $i \in \{1, \ldots, d\}$,

$$Z^i \sim \mathrm{BB}\left(n, \alpha^i, \sum_{j \neq i}^{d} \alpha^j \right). \tag{6.6}$$

Thus if we are concerned with the finite-time distributions of a single firm, we need only look at the properties of the much simpler Beta-binomial distribution.

6.1 Iso-dominance in finite time

After seeing that the DM distribution inherits the aggregation property from the Dirichlet distribution, it is natural to ask what other properties it might inherit. In particular, can we construct a family of iso-dominance curves in the finite-time case as we did for the limiting case in section 5? It turns out that we can. Since we will we looking at the size of a single firm, it is enough to use the Beta-binomial distribution to conduct our analysis, by the aggregation property.

Let $n \in \mathbb{N}$, $(\alpha, \beta) \in (0, \infty)^2$, and let $Z \sim \mathrm{BB}(n, \alpha, \beta)$. We will denote the CDF of this distribution by $I^n(k; \alpha, \beta)$, a function of $k \in \{0, 1, \ldots, n\}$. The probability mass function of Z is given by

$$\mathbb{P}[Z = k] = \binom{n}{k} \frac{B(k + \alpha, n - k + \beta)}{B(\alpha, \beta)} \tag{6.7}$$

for $k = 0, 1, \ldots, n$. As with the Dirichlet-multinomial distribution, we can view the Beta-binomial as a compound distribution: Z can be interpreted as a binomial distribution $\mathrm{Bin}(n, P)$ with random parameter $P \sim \mathrm{Beta}(\alpha, \beta)$. This gives rise to a conditional version of the probability mass function:

$$\begin{aligned}\mathbb{P}[Z = k] &= \mathbb{E}\left[\mathbb{P}[Z = k|P]\right] \\ &= \mathbb{E}\left[\binom{n}{k} P^k (1-P)^{n-k}\right].\end{aligned} \tag{6.8}$$

The CDF of Z can be expressed in a similar way, conditioning on P:

$$\begin{aligned}I^n(k;\alpha,\beta) &= \mathbb{E}\left[\mathbb{P}[Z \leq k|P]\right] \\ &= \mathbb{E}\left[I(1-P; n-k, 1+k)\right]\end{aligned} \tag{6.9}$$

where the familiar function I without a superscript is the regularised incomplete Beta function, i.e. the CDF of the Beta distribution.

Lamberson and Page investigated a much more general urn model than the Pólya model that this paper is concerned with, and one of their results was that strictly increasing the initial state of any one firm whilst keeping everything else the same would strictly increase the expected total sales of that firm up to any finite time. In our more restricted model we can do a bit better than this.

Proposition 6.5. *Let $n \in \mathbb{N}$, and let $(\alpha_1, \beta_1), (\alpha_2, \beta_2) \in (0, \infty)^2$. Suppose that $(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)$ (see section 3.3). Then for all $k \in \{0, 1, \ldots, n-1\}$,*

$$I^n(k; \alpha_1, \beta_1) < I^n(k; \alpha_2, \beta_2). \tag{6.10}$$

Proof in Appendix A.15.

Remark 6.6. This implies that, under the given conditions, $\mathrm{BB}(n, \alpha_1, \beta_1)$ has first-order stochastic dominance over $\mathrm{BB}(n, \alpha_2, \beta_2)$. This is analogous to some of the results that we derived in section 3, since it also implies that the CDFs of the two distributions "do not intersect" under certain conditions.

Recall that the PMF of the Beta-binomial distribution is given by the expression

$$\binom{n}{k} \frac{B(k+\alpha, n-k+\beta)}{B(\alpha, \beta)}. \tag{6.11}$$

This is clearly continuous in (α, β), and thus the CDF $I^n(k; \alpha, \beta)$ is also continuous in (α, β) for any fixed $n \in \mathbb{N}$ and $k \in \{0, 1, \ldots, n\}$. The next result is analogous to Lemma 5.1:

Lemma 6.7. *Fix $n \in \mathbb{N}$ and $\alpha \in (0, \infty)$. Then for $k \in \{0, 1, \ldots, n-1\}$,*

- $\lim_{\beta \to 0} I^n(k; \alpha, \beta) = 0$,
- $\lim_{\beta \to \infty} I^n(k; \alpha, \beta) = 1$.

Proof in Appendix A.16.

We now have enough machinery to prove our existence and uniqueness theorem.

Theorem 6.8. *Let $n \in \mathbb{N}$, $k \in \{0, 1, \ldots, n-1\}$, $p \in (0, 1)$ and $\alpha \in (0, \infty)$. Then there exists a unique $\beta \in (0, \infty)$ such that*

$$I^n(k; \alpha, \beta) = p. \tag{6.12}$$

Proof in Appendix A.17.

Definition 6.9. We will denote the unique β in the above theorem by $\beta^n_{k,p}(\alpha)$, and view it as a family of functions of α parametrised by (k, p) for each n. Formally, we thus have the finite-time iso-dominance function

$$\beta^n_{k,p} : (0, \infty) \to (0, \infty). \tag{6.13}$$

6.2 Properties of the finite-time iso-dominance function

Fix $n \in \mathbb{N}$, $k \in \{0, 1, \ldots, n-1\}$ and $p \in (0, 1)$. The iso-dominance function $\beta^n_{k,p}$ is strictly increasing, continuous and bijective, which are all properties that can be proven using methods not dissimilar to their infinite-time analogues in section 5. We may also, as before, plot finite-time iso-dominance curves in the $(\frac{B_0}{B_0+W_0}, \frac{S}{B_0+W_0})$-plane.

Again, finding an analytical expression for $\beta^n_{k,p}(\alpha)$ in terms of α seems to be tricky if not impossible in general. The only obvious case is when n is odd, $k = \frac{n-1}{2}$ and $p = \frac{1}{2}$, in which case $\beta^n_{\frac{n-1}{2}, \frac{1}{2}}(\alpha) = \alpha$.

An important point to note about finite-time iso-dominance functions is that they give the parameters for a firm to reach a threshold *size*, not a threshold market share as in the infinite-time case. Looking at market share in the finite-time case is slightly more complicated, since the total market size itself changes as we change the parameters α and β of the relevant Beta-binomial distribution. See the examples section below for a demonstration of this difficulty.

7 Examples

The theory developed so far is readily applicable . Suppose we have a market that evolves according to the positive feedback model $X \sim \mathrm{PU}(d; x_0, S)$. That is, there are d competing firms in the market and X_n^i denotes the "size" of firm i after n years (for some unspecified measure of size). It is currently year 0, and the initial sizes of the firms in the market are given by the vector x_0. At the end of every year, a single firm grows in size by S with probability equal to its current market share, and the other firms remain the same size.

7.1 Finite-time dominance

Suppose firm 1 would like to achieve a size of at least $x_0^1 + 4S$ with probability $\frac{3}{4}$, by the year $n = 5$. That is, it would like to grow in at least four out of the first five years. We assume that the firm has some influence over the parameters x_0 and S of the model, otherwise there is nothing to solve. By our knowledge of the finite-time marginal distributions of Pólya's urn, we know that the size of firm 1 at year 5 is X_5^1 where

$$\frac{X_5^1 - x_0^1}{S} \sim \mathrm{BB}\left(5, \frac{x_0^1}{S}, \frac{\sum_{j=2}^d x_0^j}{S}\right). \tag{7.1}$$

Let $Z = \frac{X_5^1 - x_0^1}{S}$, $\alpha = \frac{x_0^1}{S}$ and $\beta = \frac{\sum_{j=2}^d x_0^j}{S}$, so that $Z \sim \mathrm{BB}\,(5, \alpha, \beta)$. We would like to pick (α, β) such that

$$\mathbb{P}\left[X_5^1 \geq x_0^1 + 4S\right] = \frac{3}{4}. \tag{7.2}$$

This is equivalent to

$$\mathbb{P}\left[Z \leq 3\right] = \frac{1}{4}. \tag{7.3}$$

Thus $k = 3$, $p = \frac{1}{4}$ and we need to consider the iso-dominance function $\beta^5_{3,\frac{1}{4}}$. All pairs (α, β) such that $\beta = \beta^5_{3,\frac{1}{4}}(\alpha)$ would be acceptable parameters for our target market size. We should also consider parameter pairs that lie *below* the iso-dominance curve in the (α, β)-plane to be acceptable, since they dominate (in the sense of $>_1$) other parameter pairs that lie on the curve. These are parameter pairs for which the firm can reach its target size with a probability *greater than* its target probability.

Now we instead attempt to look at the market share of firm 1 in the finite time case. For example, suppose firm 1 would like to achieve at least 80% market share with probability $\frac{3}{4}$, by the year $n = 5$. If we define Z, α and β as before, we see that the market share of firm 1 in year 5 is given by

$$\begin{aligned} P_5^1 &= \frac{X_5^1}{\sum_{j=1}^d X_5^j} \\ &= \frac{x_0^1 + SZ}{\sum_{j=1}^d x_0^j + 5S} \\ &= \frac{\alpha + Z}{\alpha + \beta + 5}, \end{aligned} \tag{7.4}$$

so the set of acceptable parameter pairs is the set of pairs $(\alpha, \beta) \in (0, \infty)^2$ such that

$$\mathbb{P}\left[\frac{\alpha + Z}{\alpha + \beta + 5} \geq \frac{4}{5}\right] \geq \frac{3}{4} \tag{7.5}$$

and cannot be expressed by an iso-dominance function that we have defined.

7.2 Infinite-time dominance

In the theoretical limiting equilibrium of the market the sizes of all the firms tend to infinity so it is meaningless to discuss firm sizes. We thus restrict the discussion to market share. Suppose firm 1 would like to achieve at least 80% market share with probability $\frac{3}{4}$ in this theoretical long-run equillibrium. We know that the limiting market share of firm 1 has the distribution

$$Y_\infty^1 \sim \text{Beta}\left(\frac{x_0^1}{S}, \frac{\sum_{j=2}^d x_0^j}{S}\right). \tag{7.6}$$

Define α and β as in the previous example. We would like to find pairs (α, β) such that

$$\mathbb{P}\left[Y_\infty^1 \geq \frac{4}{5}\right] = \frac{3}{4}. \tag{7.7}$$

This is equivalent to

$$\mathbb{P}\left[Y_\infty^1 \leq \frac{4}{5}\right] = \frac{1}{4}. \tag{7.8}$$

Thus we should look at the iso-dominance function $\beta_{\frac{4}{5},\frac{1}{4}}$. All pairs (α, β) such that $\beta = \beta_{\frac{4}{5},\frac{1}{4}}(\alpha)$ would be acceptable, as would all pairs that lie below the iso-dominance curve in the (α, β)-plane, as mentioned in the previous example.

8 Application: VHS vs. Betamax

There are a number of detailed accounts of the fight for dominance among different VCR formats; see Cusumano, Mylonadis and Rosenbloom (1992), Grindely (1992) and Liebowitz and Margolis (1994). In brief, Sony pioneered the commercialization of home video recording technology in 1975 with the Betamax system. Eighteen months later the VHS standard was launched by a consortium consisting of Matsushita, JVC, and RCA. Tapes and machines were not compatible between the two standards and customers had to choose between the two. By 1979 VHS had gained a market share lead over Betamax. VHS continued to grow in the years that followed, while Betamax shrank. By 1988 VHS was so dominant that Sony abandoned the Betamax standard.

In the battle for the market, the attractiveness of the different formats to the consumer depended, as with any product, on a number of factors. Price, picture quality, play time, machine features such as programmability, ease of use, and size mattered. Absence of compatibility was a key factor. As the installed base of VHS format machines increased, so did the attractiveness of VHS format to potential buyers, and this in turn increased market share, boosting installed base further. The other main positive feedback came through complementary assets - rental stores chose to stock tapes in the more common format, and studios offered films in the format compatible with the more popular technology.

From its introduction, for 5 years till 1980, Betamax was the market leader, but were unable to deploy technological and business strategies to exploit the positive feedback potential. Cusumano *et al.* (1992) have pointed out that while Sony did not license Betamax to other firms, JVC and Matsushita aggressively sought partners and delayed VHS introduction till allies could agree on technical standards. Matsushita built and sold VCRs under the label of other firms, gaining access to their distribution channels, greater scale economies, as well as the potential to move down the learning curve.

In this brief empirical analysis of the trade-off between history and feedback, we use sales data reported in Cusumano *et al.* (1992) to estimate, for each year between 1976 and 1988, the limiting probability distribution over market shares for both VHS and Betamax. The actual market shares and the overall growth of the market are shown below.

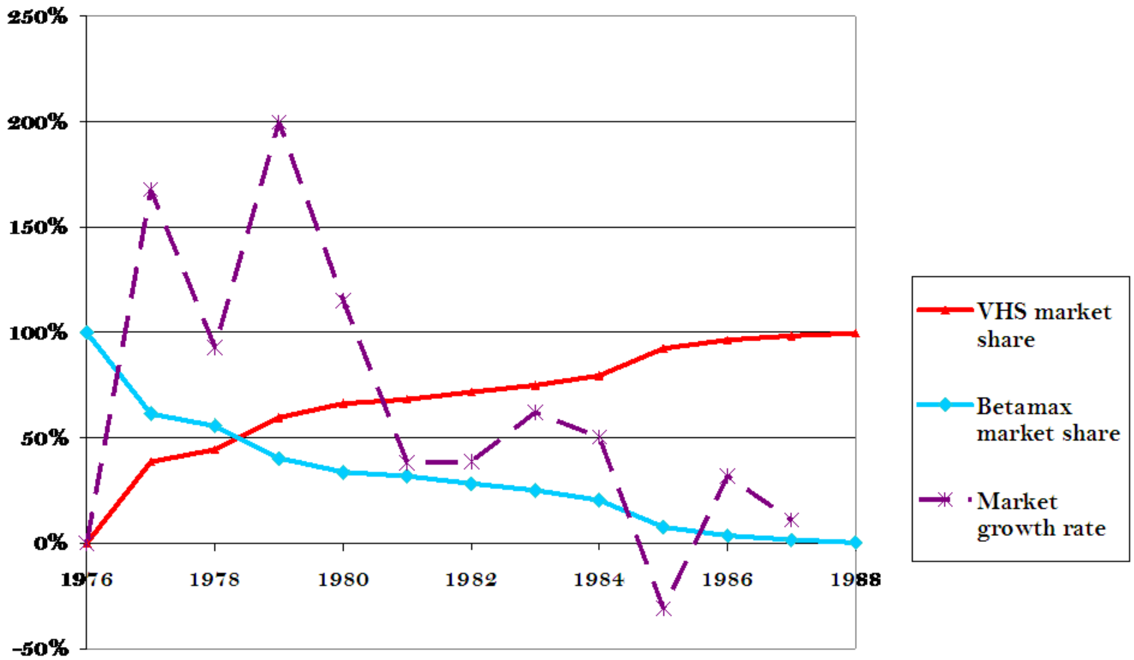

Figure 6: VHS and Betamax market shares and VCR market growth rate, 1975 to 1988.

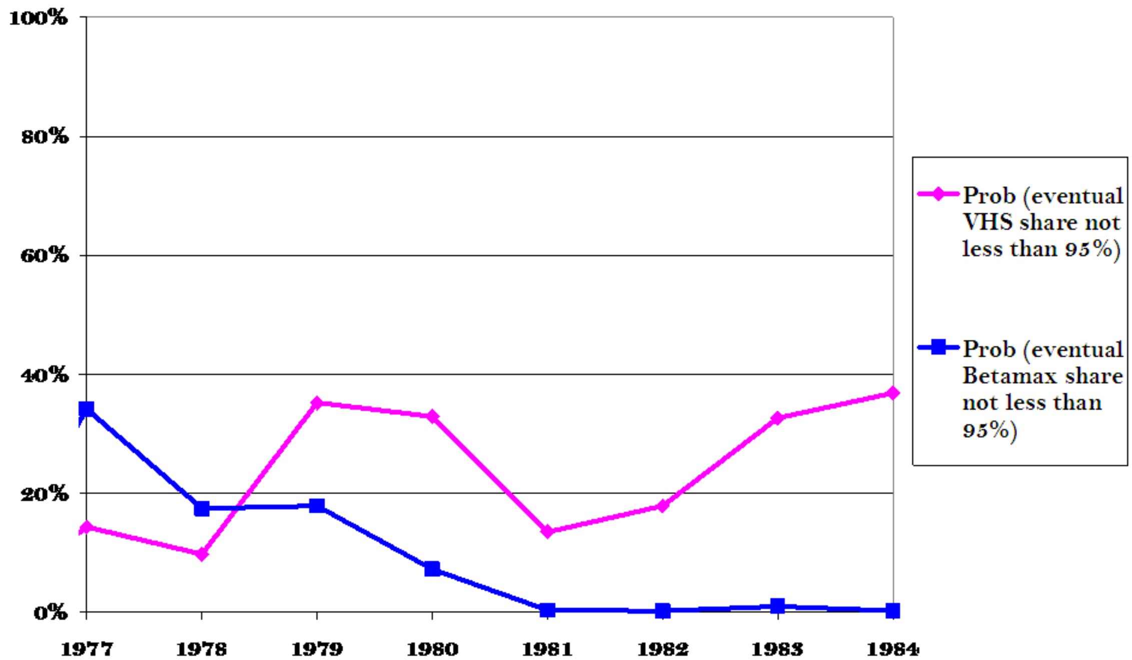

Figure 7: Probability of Dominance (limiting market share not less than 95%.

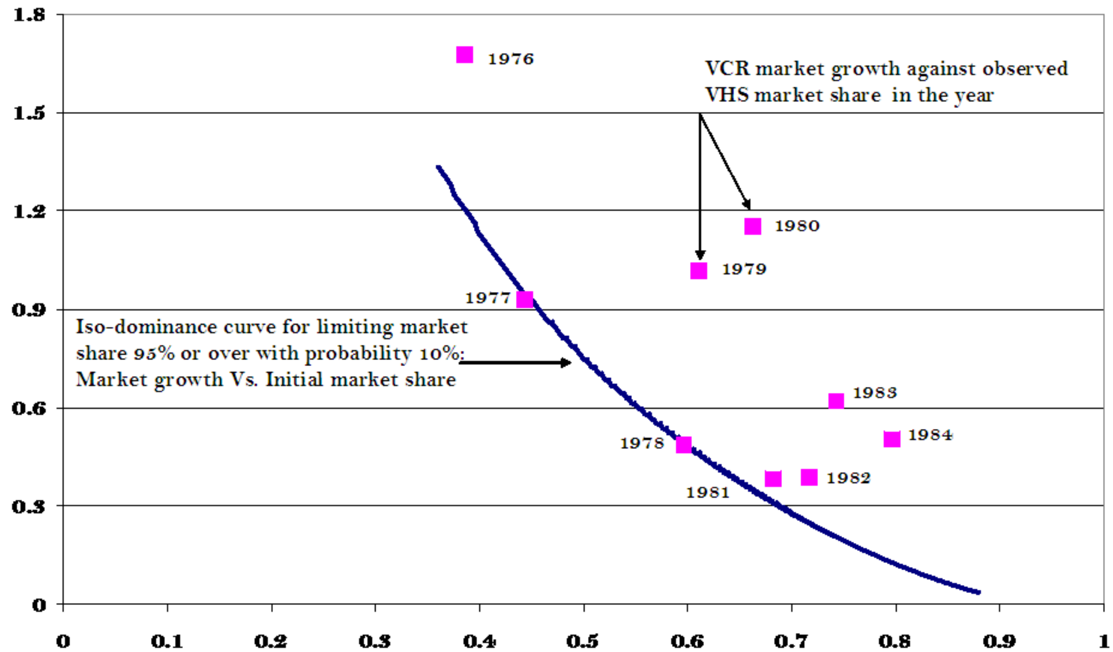

Figure 8: Iso-dominance curve: 10% probability that limiting market share is 95% or over

For each year between 1976 and 1988, we condition on the asymmetry between market shares and feedback strength (overall market growth rate) and estimate the probability that the long rum market share exceeds 95%. Between 1978 and 1979, VHS overtook Betamax, its actual market share rising from 44% to 60%. The probability of VHS dominance (limiting market share not less than 95%) rose from 10% to 35% between 1978 and 1979. The probability of Betamax having a limiting market share that is not less than 95% fell from 18% to 7% between 1979 and 1980. After 1980 the probability of Betamax dominance had fallen to 0. The decline in the VCR market between 1980 and 1981 reduced the probability of VHS dominance too, but with the recovery after 1982, the probability of VHS dominance rose again, never turning back after this year. Figure 6 shows the trade-off between observed market share and feedback strength, for a level of dominance defined as a limiting share of 95% or greater with 10% probability.

It is worth noting that even when VHS was the smaller player (39% market share) in 1976, the very high growth rate of the VCR market (168%) secured for it 14% probability of a limiting market share of 95% or greater (in comparison to 34% for Betamax).

9 Conclusions

Forecasting winners in dynamic competition is a useful art. In this paper we have presented a method based on a simple model which is applicable when the market is characterized by positive feedback, for example, due to increasing returns. Feedback, which is the source of history dependence, is recognized to be an general self-organizational feature in a large class of markets and systems.

We model the dynamics using a linear urn process and determine the trade-off between initial asymmetry in market shares and the feedback strength in such systems. In characterizing the trade-off, we introduce a probabilistic definition of the notion of dominance. This analysis is of relevance in finding the optimal strategy in dynamic contexts where externalities generate positive feedback. It is also of relevance as a diagnostic tool for tipping.

A Appendix: Proofs

A.1 Proof of Theorem 3.2

Proof. The result concerns the zeroes of δ. Notice that δ is smooth in $(0,1)$, continuous in $[0,1]$, and that $\delta(0) = \delta(1) = 0$. If we suppose that δ has n zeroes in the interval $(0,1)$, then it follows by Rolle's theorem that δ', the derivative of δ, has at least $n+1$ zeroes in this interval. Equivalently, if δ' has m zeroes, then δ must have at most $m-1$ zeroes.

We have that

$$\begin{aligned} \delta'(x) &= \frac{x^{\alpha_1 - 1}(1-x)^{\beta_1 - 1}}{B(\alpha_1, \beta_1)} - \frac{x^{\alpha_2 - 1}(1-x)^{\beta_2 - 1}}{B(\alpha_2, \beta_2)} \\ &= \frac{x^{\alpha_2 - 1}(1-x)^{\beta_2 - 1}}{B(\alpha_1, \beta_1)} g(x) \end{aligned} \tag{A.1}$$

where

$$g(x) = x^{\alpha_1 - \alpha_2}(1-x)^{\beta_1 - \beta_2} - \frac{B(\alpha_1, \beta_1)}{B(\alpha_2, \beta_2)}. \tag{A.2}$$

For $x \in (0,1)$, we have $\delta'(x) = 0$ if and only if $g(x) = 0$. We hence already know by Rolle's theorem that g has at least one zero in (0,1). We intend to use Rolle's theorem again, so differentiating g:

$$\begin{aligned} g'(x) &= (\alpha_1 - \alpha_2)x^{\alpha_1 - \alpha_2 - 1}(1-x)^{\beta_1 - \beta_2} - (\beta_1 - \beta_2)x^{\alpha_1 - \alpha_2}(1-x)^{\beta_1 - \beta_2 - 1} \\ &= x^{\alpha_1 - \alpha_2}(1-x)^{\beta_1 - \beta_2}\left(\frac{\alpha_1 - \alpha_2}{x} - \frac{\beta_1 - \beta_2}{1-x}\right). \end{aligned} \tag{A.3}$$

Now we consider the conditions laid out in the statement of the Proposition. Suppose that $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, so that $g(0) = g(1) = -\frac{B(\alpha_1,\beta_1)}{B(\alpha_2,\beta_2)} < 0$. By looking at the expression for g', we see that it has exactly one zero in $(0,1)$, specifically at

$$x = \frac{\alpha_1 - \alpha_2}{\alpha_1 - \alpha_2 + \beta_1 - \beta_2}. \tag{A.4}$$

By Rolle, g therefore has either one or two zeroes in $(0,1)$. The values of $g(0)$ and $g(1)$ are both negative so if g were to have exactly one zero in $(0,1)$, this zero must be at a stationary point (a maximum) of g. Thus δ' would be negative almost everywhere in $(0,1)$ and this would contradict $\delta(0) = \delta(1)$.

Thus g has exactly two zeroes in $(0,1)$. δ' must share these same two zeroes. By Rolle, δ therefore must have at most one zero in $(0,1)$. We can pinpoint the number of zeroes by observing that $\delta'(x)$ is negative when x is sufficiently close to either 0 or 1. Thus $\delta(x)$ must be negative when x is sufficiently close to 0 and postive when x is sufficiently close to 1. It finally follows from the intermediate value theorem that δ must have exactly one zero in $(0,1)$. We get the same result if $\alpha_1 < \alpha_2$ and $\beta_1 < \beta_2$.

Now we suppose "otherwise", i.e. either $\alpha_1 = \alpha_2$ or $\beta_1 = \beta_2$ or the two expressions $\alpha_1 - \alpha_2$ and $\beta_1 - \beta_2$ have different signs. In any of these cases, looking at the expression for g' we see that it has no zeroes in $(0,1)$. By Rolle, g (and therefore δ') must have exactly one zero in $(0,1)$. By Rolle again, δ must have no zeroes in $(0,1)$. □

A.2 Proof of Proposition 3.3

Proof. Recall the definition of g from the proof of Theorem 3.2, and recall that $(\alpha_1, \beta_1) \neq (\alpha_2, \beta_2)$ so at least one of the inequalities in the Proposition must be strict. Looking at the expression for g' we observe that it is positive in $(0,1)$, so if $r \in (0,1)$ is the unique zero of g (and thus also of δ'; we know this zero exists by Theorem 3.2) in $(0,1)$, then g (and also δ') is strictly increasing at r. In particular δ' is negative in $(0,r)$ and positive in $(r,1)$. Therefore δ attains a strict global minimum at r, so $\delta(r) < 0$. From Theorem 3.2 we know that the CDFs of the two distributions do not intersect in $(0,1)$, so $\delta(x) < 0$ for all $x \in (0,1)$ and we're done. □

A.3 Proof of Proposition 3.4

Proof. Let $r, s \in (0,1)$ be the two stationary points of δ (i.e. zeroes of δ', see Theorem 3.2), such that

$$0 < r < x_* < s < 1. \tag{A.5}$$

The point x_* is the unique zero of δ in $(0,1)$, and so must lie in between r and s by Rolle's theorem. Observe from the expressions for δ' and g given in the proof of Theorem 3.2 that $\delta'(x) < 0$ when $x \in (0,r) \cup (s,1)$. What is the sign of $\delta'(x)$ when $x \in (r,s)$? If it were negative, then δ' would be negative almost everywhere in $(0,1)$ and this would contradict $\delta(0) = \delta(1)$. So it's positive. So $\delta'(x_*) > 0$ and we're done. □

A.4 Proof of Proposition 3.6

Proof of Proposition 3.6. Define the function Δ by

$$\begin{aligned}\Delta(x) &= \int_0^x \delta(u)du \\ &= \int_0^x I(u;\alpha_1,\beta_1)du - \int_0^x I(u;\alpha_2,\beta_2)du.\end{aligned} \tag{A.6}$$

Obviously $\Delta' = \delta$, by the fundamental theorem of calculus. By the properties of δ and δ' that we derived in Theorem 3.2 and Lemma 3.4, we know that Δ has a single stationary point in $(0,1)$, and that this stationary point is a strict minimum. Let's call this point x_*, as it is the same point that we were concerned with in Lemma 3.4. Obviously $\Delta(0) = 0$, so $\Delta(x_*) < 0$. This is enough to show that $\text{Beta}(\alpha_2,\beta_2)$ does not have second-order stochastic dominance over $\text{Beta}(\alpha_1,\beta_1)$.

Does Δ have any zeroes in $(0,1)$? It cannot have any in the interval $(0,x_*)$, as this would imply the existence of a second stationary point of Δ by Rolle's theorem. We therefore have $\Delta(x) < 0$ for all $x \in (0,x_*)$. As for the interval $(x_*,1)$, we note that

$$\begin{aligned}\Delta(1) &= \frac{\beta_1}{\alpha_1+\beta_1} - \frac{\beta_2}{\alpha_2+\beta_2} \\ &= \frac{\alpha_2}{\alpha_2+\beta_2} - \frac{\alpha_1}{\alpha_1+\beta_1}.\end{aligned} \tag{A.7}$$

Since $\Delta' = \delta$ has a single zero $x_* \in (0,1)$ and is strictly increasing at x_*, it must be positive in $(x_*,1)$. This implies that Δ is strictly increasing in $(x_*,1)$, and in particular for all $x \in (x_*,1)$,

$$\Delta(x_*) < \Delta(x) < \frac{\alpha_2}{\alpha_2+\beta_2} - \frac{\alpha_1}{\alpha_1+\beta_1}. \tag{A.8}$$

Suppose that $\frac{\alpha_1}{\alpha_1+\beta_1} \geq \frac{\alpha_2}{\alpha_2+\beta_2}$. Then $\Delta(x) < 0$ for all $x \in (0,1)$, and additionally $\Delta(0) \leq 0$, $\Delta(1) \leq 0$. Thus $\text{Beta}(\alpha_1,\beta_1)$ has second-order stochastic dominance over $\text{Beta}(\alpha_2,\beta_2)$.

If on the other hand $\frac{\alpha_1}{\alpha_1+\beta_1} < \frac{\alpha_2}{\alpha_2+\beta_2}$ then immediately from $\Delta(1) > 0$ we have that $\text{Beta}(\alpha_1,\beta_1)$ does not have second-order stochastic dominance over $\text{Beta}(\alpha_2,\beta_2)$. □

A.5 Proof of Lemma 5.1

Proof of Lemma 5.1. Observe that

$$
\begin{aligned}
I(x;\alpha,\beta) &= \frac{B_x(\alpha,\beta)}{B(\alpha,\beta)} \\
&= \frac{\int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt}{\int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt + \int_x^1 t^{\alpha-1}(1-t)^{\beta-1}dt} \\
&= \frac{1}{1+Q_x(\alpha,\beta)}
\end{aligned}
\tag{A.9}
$$

where

$$
Q_x(\alpha,\beta) := \frac{\int_x^1 t^{\alpha-1}(1-t)^{\beta-1}dt}{\int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt}. \tag{A.10}
$$

Both of the results we seek to prove are thus equivalent to results about Q_x. We prove the first item. We have

$$
\begin{aligned}
\int_x^1 t^{\alpha-1}(1-t)^{\beta-1}dt &\geq \min\{x^{\alpha-1},1\}\int_x^1 (1-t)^{\beta-1}dt \\
&= \min\{x^{\alpha-1},1\}\frac{(1-x)^\beta}{\beta}
\end{aligned}
\tag{A.11}
$$

and, assuming $\beta < 1$ without loss of generality,

$$
\begin{aligned}
\int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt &\leq (1-x)^{\beta-1}\int_0^x t^{\alpha-1}dt \\
&= \frac{x^\alpha(1-x)^{\beta-1}}{\alpha}.
\end{aligned}
\tag{A.12}
$$

Thus

$$
\begin{aligned}
Q_x(\alpha,\beta) &\geq \min\{x^{\alpha-1},1\}\frac{\alpha(1-x)}{\beta x^\alpha} \\
&\to \infty \qquad \text{as } \beta \to 0,
\end{aligned}
\tag{A.13}
$$

and so

$$
\lim_{\beta\to 0} I(x;\alpha,\beta) = 0. \tag{A.14}
$$

The proof of the second item follows almost identically, with the inequalities reversed. We have

$$
\begin{aligned}
\int_x^1 t^{\alpha-1}(1-t)^{\beta-1}dt &\leq \max\{x^{\alpha-1},1\}\int_x^1 (1-t)^{\beta-1}dt \\
&= \max\{x^{\alpha-1},1\}\frac{(1-x)^\beta}{\beta}
\end{aligned}
\tag{A.15}
$$

and, assuming $\beta > 1$ without loss of generality,

$$\begin{aligned}\int_0^x t^{\alpha-1}(1-t)^{\beta-1}dt &\geq (1-x)^{\beta-1}\int_0^x t^{\alpha-1}dt \\ &= \frac{x^\alpha(1-x)^{\beta-1}}{\alpha}.\end{aligned} \tag{A.16}$$

Thus

$$\begin{aligned}Q_x(\alpha,\beta) &\leq \max\{x^{\alpha-1},1\}\frac{\alpha(1-x)}{\beta x^\alpha} \\ &\to 0 \qquad \text{as } \beta\to\infty,\end{aligned} \tag{A.17}$$

and so

$$\lim_{\beta\to\infty} I(x;\alpha,\beta) = 1. \tag{A.18}$$

□

A.6 Proof of Theorem 5.3

Proof. Existence: Consider $I(x;\alpha,1)$. If $I(x;\alpha,1) = p$, then we're done. If not, suppose that $I(x;\alpha,1) > p$. We know that $\lim_{\beta\to 0} I(x;\alpha,\beta) = 0$ by Lemma 5.1, and in addition we know that I is continuous in (α,β). Thus by the intermediate value theorem, there exists $\beta > 0$ with $I(x;\alpha,\beta) = p$. A similar argument using the other item of Lemma 5.1 shows that the same is true when $I(x;\alpha,1) < p$.

Uniqueness: Suppose $\beta_1,\beta_2 \in (0,\infty)$ with $I(x;\alpha,\beta_1) = p = I(x;\alpha,\beta_2)$. If $\beta_1 \neq \beta_2$ then by Theorem 3.2, $I(x';\alpha,\beta_1)$ and $I(x';\alpha,\beta_2)$ do not intersect at any $x' \in (0,1)$. This is a contradiction, so $\beta_1 = \beta_2$. □

A.7 Proof of Proposition 5.5

Proof. Strictly increasing: Let $\alpha_1,\alpha_2 \in (0,\infty)$ with $\alpha_1 < \alpha_2$. The CDFs $I(\cdot;\alpha_1,\beta_{x,p}(\alpha_1))$ and $I(\cdot;\alpha_2,\beta_{x,p}(\alpha_2))$ satisfy $I(x;\alpha_1,\beta_{x,p}(\alpha_1)) = p = I(x;\alpha_2,\beta_{x,p}(\alpha_2))$, by definition. Thus by Theorem 3.2, it follows that $\beta_{x,p}(\alpha_1) < \beta_{x,p}(\alpha_2)$.

Continuous: We argue by contradiction. Let $(\alpha_n)_n$ be a sequence in $(0,\infty)$ that converges to some $\alpha_\infty \in (0,\infty)$, and suppose that the sequence $(\beta_{x,p}(\alpha_n))_n$ does *not* converge to $\beta_{x,p}(\alpha_\infty)$, i.e. suppose that there exists an $\varepsilon > 0$ such that for all $N \in \mathbb{N}$, there is an

$n > N$ with $|\beta_{x,p}(\alpha_n) - \beta_{x,p}(\alpha_\infty)| \geq \varepsilon$. Observe that, since $\alpha_n \to \alpha_\infty \in (0, \infty)$, we have that $\inf_n \alpha_n, \sup_n \alpha_n \in (0, \infty)$. Moreover, since $\beta_{x,p}$ is increasing, $(\beta_{x,p}(\alpha_n))_n$ is a sequence contained between $\beta_{x,p}(\inf_n \alpha_n)$ and $\beta_{x,p}(\sup_n \alpha_n)$ so we can use the Bolzano-Weierstrass theorem. Thus by taking a particular subsequence we can assume without loss of generality that $|\beta_{x,p}(\alpha_n) - \beta_{x,p}(\alpha_\infty)| \geq \varepsilon$ for all n, and also that $(\beta_{x,p}(\alpha_n))_n$ converges to some

$$\beta_\infty \in [\beta_{x,p}(\inf_n \alpha_n), \beta_{x,p}(\sup_n \alpha_n)] \subseteq (0, \infty). \tag{A.19}$$

In particular it follows that $|\beta_\infty - \beta_{x,p}(\alpha_\infty)| \geq \varepsilon$.

Now recall that, by definition, $I(x; \alpha_n, \beta_{x,p}(\alpha_n)) = p$ for all n. Thus by continuity of $I(x; \alpha, \beta)$ with respect to (α, β),

$$p = \lim_{n\to\infty} I(x; \alpha_n, \beta_{x,p}(\alpha_n)) = I(x; \alpha_\infty, \beta_\infty). \tag{A.20}$$

By uniqueness, this implies that $\beta_\infty = \beta_{x,p}(\alpha_\infty)$, a contradiction! Thus we must have

$$\lim_{n\to\infty} \beta_{x,p}(\alpha_n) = \beta_{x,p}(\alpha_\infty). \tag{A.21}$$

□

A.8 Proof of Proposition 5.7

Proof. We proceed in a manner similar to the proof of Lemma 5.1. The result in the case $x' = x$ is true, by definition. Suppose $x' > x$. Then

$$\begin{aligned} I(x'; \alpha, \beta_{x,p}(\alpha)) &= \mathbb{P}[X_\alpha \leq x'] \\ &= p + \mathbb{P}[x < X_\alpha \leq x'], \end{aligned} \tag{A.22}$$

where we have

$$\begin{aligned} \mathbb{P}[x < X_\alpha \leq x'] &= \frac{\int_x^{x'} t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1} dt}{B(\alpha, \beta_{x,p}(\alpha))} \\ &\leq \frac{\int_x^{x'} t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1} dt}{\int_0^{x_0} t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1} dt}. \end{aligned} \tag{A.23}$$

Assuming $\alpha < 1$ without loss of generality, the numerator in the above expression satisfies

$$\int_x^{x'} t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1} dt \leq x^{\alpha-1} \max\{(1-x)^{\beta_{x,p}(\alpha)-1}, (1-x')^{\beta_{x,p}(\alpha)-1}\} \tag{A.24}$$

and the denominator satisfies

$$\int_0^x t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1}dt \geq \frac{x^\alpha}{\alpha}\min\{1,(1-x)^{\beta_{x,p}(\alpha)-1}\}. \tag{A.25}$$

Now the function $\beta_{x,p}$ is positive and strictly increasing, so as $\alpha \to 0$, $\beta_{x,p}(\alpha)$ eventually becomes bounded between 0 and $\beta_{x,p}(1)$, say. This is enough to conclude that, for small enough α,

$$\begin{aligned}\mathbb{P}[x < X_\alpha \leq x'] &\leq \frac{\int_x^{x'} t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1}dt}{\int_0^x t^{\alpha-1}(1-t)^{\beta_{x,p}(\alpha)-1}dt} \\ &\leq \frac{\alpha}{x}\frac{\max\{(1-x)^{\beta_{x,p}(\alpha)-1},(1-x')^{\beta_{x,p}(\alpha)-1}\}}{\min\{1,(1-x)^{\beta_{x,p}(\alpha)-1}\}} \\ &\to 0 \qquad \text{as } \alpha \to 0,\end{aligned} \tag{A.26}$$

and hence

$$\lim_{\alpha\to 0} I(x';\alpha,\beta_{x,p}(\alpha)) = p. \tag{A.27}$$

If $x' < x$, we instead have

$$I(x';\alpha,\beta_{x,p}(\alpha)) = p - \mathbb{P}[x' < X_\alpha \leq x] \tag{A.28}$$

and the argument proceeds more or less identically. □

A.9 Proof of Proposition 5.8

Proof. We will take a slightly different approach with this proof. Recall that the variance of X_α is given by

$$\begin{aligned}\mathbb{V}\mathrm{ar}X_\alpha &= \frac{\alpha\beta_{x,p}(\alpha)}{(\alpha+\beta_{x,p}(\alpha))^2(\alpha+\beta_{x,p}(\alpha)+1)} \\ &\leq \frac{\alpha^2+2\alpha\beta_{x,p}(\alpha)+\beta_{x,p}(\alpha)^2}{2(\alpha+\beta_{x,p}(\alpha))^2(\alpha+\beta_{x,p}(\alpha)+1)} \\ &\leq \frac{1}{2(\alpha+1)} \\ &\to 0 \qquad \text{as } \alpha\to\infty.\end{aligned} \tag{A.29}$$

We will first prove that $\lim_{\alpha\to\infty}\mathbb{E}X_\alpha = x$. Let $\varepsilon > 0$. By Chebyshev's inequality,

$$\mathbb{P}\left[|X_\alpha - \mathbb{E}X_\alpha| \geq \varepsilon\right] \leq \frac{\mathbb{V}\mathrm{ar}X_\alpha}{\varepsilon^2}. \tag{A.30}$$

Since $\lim_{\alpha\to\infty} \mathbb{V}\text{ar}X_\alpha = 0$, there exists an $A > 0$ such that, for all $\alpha > A$,

$$\mathbb{P}\left[|X_\alpha - \mathbb{E}X_\alpha| \geq \varepsilon\right] < \min\{p, 1-p\}. \tag{A.31}$$

It follows[1] that $|\mathbb{E}X_\alpha - x| \leq \varepsilon$ for all such α. Thus $\lim_{\alpha\to\infty} \mathbb{E}X_\alpha = x$.

Now if $x' \in (0, x)$, by conditioning on the event $\left\{|\mathbb{E}X_\alpha - x| \leq \frac{1}{2}(x - x')\right\}$ we have

$$\begin{aligned} I(x'; \alpha, \beta_{x,p}(\alpha)) &= \mathbb{P}[X_\alpha \leq x'] \\ &\leq \mathbb{P}\left[|X_\alpha - x| \geq x - x'\right] \\ &\leq \mathbb{P}\left[|X_\alpha - \mathbb{E}X_\alpha| \geq \frac{1}{2}(x - x')\right] + 1_{\left\{|\mathbb{E}X_\alpha - x| > \frac{1}{2}(x-x')\right\}} \\ &\to 0 \qquad \text{as } \alpha \to \infty. \end{aligned} \tag{A.32}$$

The case $x' \in (x, 1)$ is similar:

$$\begin{aligned} I(x'; \alpha, \beta_{x,p}(\alpha)) &= \mathbb{P}[X_\alpha \leq x'] \\ &\geq \mathbb{P}\left[|X_\alpha - x| \leq x' - x\right] \\ &= 1 - \mathbb{P}\left[|X_\alpha - x| > x' - x\right] \\ &\to 1 \qquad \text{as } \alpha \to \infty. \end{aligned} \tag{A.33}$$

□

A.10 Proof of Corollary 5.9

Proof. The proof relies on the definition of convergence in distribution, which for real-valued random variables $(Z_n)_n$ and Z is equivalent to the following: $Z_n \to Z$ in distribution as $n \to \infty$ if and only if for all bounded and continuous functions f,

$$\mathbb{E}f(Z_n) \to \mathbb{E}f(Z) \qquad \text{as } n \to \infty. \tag{A.34}$$

We will take advantage of the fact that the Beta distribution can only take values in $[0, 1]$. Define the function $f : \mathbb{R} \to \mathbb{R}$ by

$$f(x) = \begin{cases} 0 & \text{if } x < 0, \\ x & \text{if } 0 \leq x \leq 1, \\ 1 & \text{if } x > 1. \end{cases} \tag{A.35}$$

[1] e.g. if $x < \mathbb{E}X_\alpha - \varepsilon$ then $p = \mathbb{P}[X_\alpha \leq x] \leq \mathbb{P}\left[|X_\alpha - \mathbb{E}X_\alpha| \geq \varepsilon\right] < \min\{p, 1-p\}$, a contradiction.

Observe that f is bounded and continuous, and if a random variable Z only takes values in $[0,1]$ then $\mathbb{E}f(Z) \equiv \mathbb{E}Z$. Thus Propositions 5.7 and 5.8 give us

$$\lim_{\alpha \to 0} \mathbb{E}X_\alpha = \mathbb{E}X_0 = 1 - p \tag{A.36}$$

and

$$\lim_{\alpha \to \infty} \mathbb{E}X_\alpha = \mathbb{E}X_\infty = x \tag{A.37}$$

respectively. The final step comes from the observation that for $\alpha \in (0, \infty)$,

$$\mathbb{E}X_\alpha = \frac{\alpha}{\alpha + \beta_{x,p}(\alpha)} = \frac{1}{1 + \frac{\beta_{x,p}(\alpha)}{\alpha}}. \tag{A.38}$$

The statements of the Corollary follow by continuity of the function $t \mapsto \frac{1}{t} - 1$ for $t \in (0, 1)$. □

A.11 Proof of Proposition 5.11

Proof. Let's get the "otherwise" out of the way first. The CDF of the Beta distribution is always strictly increasing, so in this case there's no way that a Beta CDF could pass through both points. So we're done.

Now assume without loss of generality that $x_1 > x_2$ and $p_1 > p_2$.

Existence: The pair (α_*, β_*) we are looking for must necessarily satisfy $\beta_* = \beta_{x_2,p_2}(\alpha_*)$ so the problem reduces to finding an $\alpha_* \in (0, \infty)$ such that

$$I(x_1; \alpha_*, \beta_{x_2,p_2}(\alpha_*)) = p_1. \tag{A.39}$$

We try $\alpha_* = 1$. Note that since $x_1 > x_2$, we must have $I(x_1; 1, \beta_{x_2,p_2}(1)) \in (p_2, 1)$. If $I(x_1; 1, \beta_{x_2,p_2}(1)) = p_1$, then we're done. If not, then consider that the function $\alpha \mapsto I(x_1; \alpha, \beta_{x_2,p_2}(\alpha))$ is continuous. If $I(x_1; 1, \beta_{x_2,p_2}(1)) > p_1$, then our desired α_* must exist by Proposition 5.7 and the intermediate value theorem. If $I(x_1; 1, \beta_{x_2,p_2}(1)) < p_1$ then we get the same result via Proposition 5.8.

Uniqueness: If two Beta CDFs pass though both of the points, then they obviously intersect at least twice. By Theorem 3.2 these CDFs must be equal. □

A.12 Proof of Proposition 6.1

Proof. We use induction on n. For $i \in \{1, \ldots, d\}$, let e_i denote the unit vector in the ith component. Now for $n = 1$, the probability of picking any colour from the urn is equal to that colour's initial proportion in the urn. The colour that is picked is then replaced along with one more unit of the same colour, so the only values that $X_1 - x_0$ can take are unit vectors. For $i \in \{1, \ldots, d\}$,

$$\begin{aligned}\mathbb{P}\left[X_1 - x_0 = e_i\right] &= \frac{x_0^i}{A} \\ &= \frac{1!}{\prod_{j=1}^d (e_i^j!)} \frac{\Gamma(A)}{\Gamma(1+A)} \prod_{j=1}^d \frac{\Gamma(e_i^j + x_0^j)}{\Gamma(x_0^j)},\end{aligned} \tag{A.40}$$

where $A = \sum_{j=1}^d x_0^j$. So the Proposition is true for $n = 1$.

For the inductive step, suppose the Proposition is true for some $n \geq 1$. We consider the distribution of $X_{n+1} - x_0$. Let $r \in \mathbb{Z}_+^d$ such that $\sum_{i=1}^d r^i = n + 1$. By the properties of the urn model,

$$\begin{aligned}\mathbb{P}\left[X_{n+1} - x_0 = r\right] &= \sum_{i=1}^d \frac{x_0^i + r^i - 1}{A + n} \mathbb{P}\left[X_n - x_0 = r - e_i\right] 1_{\{r^i > 0\}} \\ &= \sum_{i=1}^d \frac{x_0^i + r^i - 1}{n + A} \frac{n!}{\prod_{j=1}^d ((r^j - e_i^j)!)} \frac{\Gamma(A)}{\Gamma(n+A)} \prod_{j=1}^d \frac{\Gamma(r^j - e_i^j + x_0^j)}{\Gamma(x_0^j)} 1_{\{r^i > 0\}} \\ &= \sum_{i=1}^d \frac{x_0^i + r^i - 1}{n + A} \frac{r^i n!}{\prod_{j=1}^d (r^j!)} \frac{\Gamma(A)}{\Gamma(n+A)} \prod_{j=1}^d \frac{\Gamma(r^j - e_i^j + x_0^j)}{\Gamma(x_0^j)} \\ &= \sum_{i=1}^d \frac{r^i}{n+1} \frac{(n+1)!}{\prod_{j=1}^d (r^j!)} \frac{\Gamma(A)}{\Gamma(n+1+A)} \prod_{j=1}^d \frac{\Gamma(r^j + x_0^j)}{\Gamma(x_0^j)} \\ &= \frac{(n+1)!}{\prod_{j=1}^d (r^j!)} \frac{\Gamma(A)}{\Gamma(n+1+A)} \prod_{j=1}^d \frac{\Gamma(r^j + x_0^j)}{\Gamma(x_0^j)}.\end{aligned} \tag{A.41}$$

The numbers of balls of each colour added to the urn by time $n + 1$ must be integers and they must sum up to $n + 1$, so for any other value of r the above probability is 0. This completes the inductive step. □

A.13 Proof of Proposition 6.2

Proof. We first recall that all Pólya urn processes are Markov. Moreover, kX inherits the Markov property from X. It is therefore sufficient to check only the initial value of the process and its transition probabilities:

1. By definition, $kX_0 = kx_0$.

2. For $i \in \{1, \dots, d\}$, let e_i denote the unit vector in the ith component. Fix an integer $n \geq 0$. It is clear (since X is a Pólya urn process) that kX_n can only take values of the form $kx_0 + kSr$ where $r \in \mathbb{Z}_+^d$ with $\sum_{j=1}^d r^j = n$. Then for all such r and all $i \in \{1, \dots, d\}$,

$$\begin{aligned}\mathbb{P}\left[kX_{n+1} = kx_0 + kS(r + e_i) | kX_n = kx_0 + kSr\right] \\ &= \mathbb{P}\left[X_{n+1} = x_0 + S(r + e_i) | X_n = x_0 + Sr\right] \\ &= \frac{x_0^i + Sr^i}{\sum_{j=1}^d \left(x_0^j + Sr^j\right)} \\ &= \frac{kx_0^i + kSr^i}{\sum_{j=1}^d \left(kx_0^j + kSr^j\right)}.\end{aligned} \tag{A.42}$$

These are exactly the properties of the $\mathrm{PU}(d; kx_0, kS)$ process. □

A.14 Proof of Proposition 6.4

Proof. By the construction of the DM distribution, we can interpret Z as a d-variate $\mathrm{Multin}(n, P)$ random variable with d-variate random parameter $P \sim \mathrm{Dir}(\alpha)$. The multinomial distribution satisfies the (easily verifiable) aggregation property

$$\left(\sum_{j \in A_1} Z^j, \dots, \sum_{j \in A_r} Z^j\right) \sim \mathrm{Multin}\left(n, \left(\sum_{j \in A_1} P^j, \dots, \sum_{j \in A_r} P^j\right)\right), \tag{A.43}$$

and the Dirichlet distribution satisfies the aggregation property

$$\left(\sum_{j \in A_1} P^j, \dots, \sum_{j \in A_r} P^j\right) \sim \mathrm{Dir}\left(\sum_{j \in A_1} \alpha^j, \dots, \sum_{j \in A_r} \alpha^j\right). \tag{A.44}$$

The Proposition follows. □

A.15 Proof of Proposition 6.5

Proof. Recall that $(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)$ simply means that $\alpha_1 \geq \alpha_2$, $\beta_1 \leq \beta_2$ and $(\alpha_1, \beta_1) \neq (\alpha_2, \beta_2)$.

Define random variables $X_1 \sim \mathrm{BB}(n, \alpha_1, \beta_1)$ and $X_2 \sim \mathrm{BB}(n, \alpha_2, \beta_2)$. We will use the compound interpretation of the Beta-binomial distribution: for $i = \{1, 2\}$, we have that $X_i \sim \mathrm{Bin}(n, P_i)$ where $P_i \sim \mathrm{Beta}(\alpha_i, \beta_i)$. Then for $k \in \{0, 1, \ldots, n-1\}$,

$$\begin{aligned} \mathbb{P}[X_i > k] &= \mathbb{E}\left[1 - I(1 - P_i; n-k, 1+k)\right] \\ &= \mathbb{E}\left[\int_0^1 1_{\{t \leq 1 - I(1-P_i; n-k, 1+k)\}} dt\right] \\ &= \int_0^1 \mathbb{P}\left[t \leq 1 - I(1 - P_i; n-k, 1+k)\right] dt \end{aligned} \tag{A.45}$$

where the last equality follows from Fubini's theorem. Now the expression $1 - I(1-p; n-k, 1+k)$ is a strictly increasing function of p, so it admits a strictly increasing inverse function which we will simply call I^{-1}. In particular, I^{-1} maps the interval $(0, 1)$ to itself. From the above we have that

$$\mathbb{P}[X_i > k] = \int_0^1 \mathbb{P}\left[I^{-1}(t) \leq P_i\right] dt. \tag{A.46}$$

We supposed that $(\alpha_1, \beta_1) >_1 (\alpha_2, \beta_2)$, so by our results in section 3 we know that

$$\mathbb{P}\left[P_1 \geq x\right] > \mathbb{P}\left[P_2 \geq x\right] \tag{A.47}$$

for all $x \in (0, 1)$. Thus

$$\begin{aligned} \mathbb{P}[X_1 > k] - \mathbb{P}[X_2 > k] &= \int_0^1 \mathbb{P}\left[I^{-1}(t) \leq P_1\right] - \mathbb{P}\left[I^{-1}(t) \leq P_2\right] dt \\ &> 0 \end{aligned} \tag{A.48}$$

and the result follows. □

A.16 Proof of Lemma 6.7

Proof. Define a family $(X_\beta)_{\beta \in (0,\infty)}$ of random variables such that $X_\beta \sim \mathrm{BB}(n, \alpha, \beta)$ for each β. We again use the compound interpretation of the Beta-binomial distribution: for each $\beta \in (0, \infty)$, we have that $X_\beta \sim \mathrm{Bin}(n, P_\beta)$ where $P_\beta \sim \mathrm{Beta}(\alpha, \beta)$. Lemma 5.1 implies that as $\beta \to 0$, P_β converges in distribution to a degenerate random variable that takes the value 1 almost surely. The expression $I(1-p; n-k, 1+k)$ is a bounded and continuous function of p. Thus

$$\begin{aligned} \lim_{\beta \to 0} I^n(k; \alpha, \beta) &= \lim_{\beta \to 0} \mathbb{E}\left[I(1-P_\beta; n-k, 1+k)\right] \\ &= \mathbb{E}\left[I(0; n-k, 1+k)\right] \\ &= 0. \end{aligned} \tag{A.49}$$

The second item follows similarly. □

A.17 Proof of Theorem 6.8

Proof. Existence: Just as in the proof of Theorem 5.3, except that we use Lemma 6.7 instead of Lemma 5.1.

Uniqueness: Suppose $\beta_1, \beta_2 \in (0, \infty)$. If, for example, $\beta_1 > \beta_2$, then $(\alpha, \beta_1) >_1 (\alpha, \beta_2)$. Thus by Proposition 6.5, $I^n(k; \alpha, \beta_1) < I^n(k; \alpha, \beta_2)$ for all $k \in \{0, 1, \ldots, n-1\}$. This proves uniqueness. □

References

[1] Arthur, W.B., Ermoliev, Yur.M., and Kaniovski, Yu.M. (1983). 'A Generalized Urn Problem and its Applications', *Cybernetics*, 19, pp. 61-71.

[2] Arthur, W.B., Ermoliev, Yur.M., and Kaniovski, Yu.M. (1986). 'Strong Laws for a Class of Path-Dependent Urn Processes', *Proceedings of the International Conference on Stochastic Optimization, Kiev 1984*, Arkin, Shiryayev and Wets (eds.) New York: Springer (Springer Lecture Notes in Control and Information Sciences).

[3] Arthur, W.B., Ermoliev, Yur.M., and Kaniovski, Yu.M. (1987). 'Path Dependent Processes and the Emergence of Macro-Structure', *European Journal of Operational Research*, 30, pp. 294-303.

[4] Athreya, K.B. (1969). 'On a Characteristic Property of Pólya's Urn', *Studia Scientiarum Mathematicarum Hungarica*, 4, pp. 31-35.

[5] Bass, F.M. (1969) .'A New Product Growth Model for Consumer Durables', *Management Science*, 15(5), pp. 215-227.

[6] Cusumano, M.A., Mylonadis, Y., and Rosenbloom, R.S. (1992) 'Strategic Maneuvering and Mas-Market Dynamics: The Triumph of VHS over Beta', *Business History Review*, 66, pp. 51-94.

[7] David, P. (1985). 'Clio and the economics of QWERTY', *American Economic Review*, 75, pp. 332-337.

[8] Ellison and Fudenberg (2006). 'Knife Edge or Plateau: When do Markets Tip?', MIT Department of Economics Working Paper No. 03-07.

[9] Fisher, R.A. (1929). 'Tests of Significance in Harmonic Analysis', *Proc. of the Royal Society of London, Series A*, 125(796), pp. 54-59.

[10] Johnson, N.L. and Kotz, S. (1977). 'Urn Models and their Applications, an Approach to Modern Discrete Probability Theory', John Wiley and Sons.

[11] Freedman, D.A. (1965). 'Bernard Friedman's Urn', *The Annals of Maths. Stats.*, 36(3), pp. 956-970.

[12] Frigyik, B. A., Kapila, A. and Gupta, M. R. (2010). Introduction to the Dirichlet Distribution and Related Processes. Seattle, Washington: *UWEE Technical Report* Number UWEETR-2010-0006.

[13] Gladwell, M. (2000). 'The Tipping Point: How little things can make a big difference', Little, Brown and Company, Boston, New York, London.

[14] Grindley, P. (1992) 'Standards, Business Strategy and Policy: A Casebook', London: London Business School.

[15] Hill, B. M., Lane, D. and Sudderth, W. (1980). 'A Strong Law for some Generalized Urn Processes', *The Annals of Probability*, 8(2), pp. 214-226.

[16] Kattuman, P.A., (1998). 'The role of history in the transition to an industrial district', In P. Cadène & M. Holmström (eds.), 'Decentralized Production in India: Industrial Districts, Flexible Specialization, and Employment', Thousand Oaks, CA: Sage Publications, pp. 230–250.

[17] Kermack, W.O. and McKendrick, A.G., 'A Contribution to the Mathematical Theory of Epidemics,' *Proc. Roy. Soc. Lond. A* (1927) 115, pp. 700-721.

[18] Liebowitz, S.J., and Margolis, S.E. (1999). 'Winners, Losers & Microsoft: Competition and Antitrust in High Technology', The Independent Institute, Oakland, California.

[19] Mauldon, J.G. (1951) 'Random Division of an Interval', *Proc.Cambridge Philos. Soc.*, 47, pp. 331-336.

[20] Mosimann, J.E.,(1962) 'On the Compound Multinomial Distribution, the Multivariate β-Distribution, and Correlations Among Proportions', *Biometrika*, 49, pp. 65-82.

[21] Özçaḡ, E., Ege, İ., and Gürçay, H. (2008) 'An Extension of the Incomplete Beta Function for Negative Integers', *Journal of Mathematical Analysis and Applications* 338, pp. 984-992

[22] Page, S. (2006) 'Path Dependence', *The Quarterly Journal of Political Science*, 1(1), pp. 87-115.

[23] Pemantle, R. (2006) 'A survey of random processes with reinforcement', Mimeo, University of Pennsylvania, Department of Mathematics.

[24] Pólya, G. (1930) 'Sur Quelques Points de la Théorie des Probabilités', *Poincare Inst. B*, pp. 117-161.