Quarterly Journal of Political Science, 2006, 1: 87–115

Essay

Path Dependence

Scott E. Page*

Center for the Study of Complex Systems, University of Michigan, Ann Arbor 48104, E-mail: spage@umich.edu

> *I eat my peas with honey. I've done it all my life. It makes 'em taste quite funny, but it keeps them on the knife.*
>
> – an old Bostonian jump roping rhyme.

The concept of path dependence originated as an idea that a small initial advantage or a few minor random shocks along the way could alter the course of history (David 1985). Like many ideas, it has grown in scope and now encompasses almost any process in which someone can find or claim evidence of increasing returns, which are thought to be the causes of path dependence (Arthur 1994, Pierson 2000). This wider application of path dependence has dulled its value. In becoming a trendy way to say that history matters, path dependence no longer provides any analytic leverage.

Attempts to extend what is meant by path dependence reflect a need for a finer unpacking of historical causality. We need to differentiate between types of path dependence. The way to do that is with a formal framework. An obvious advantage of having such a framework is that we can conduct empirical analyses and discern whether the evidence supported or refuted a claim of the extent and scope of the sway of the past. That said, empirical testing of a framework of causality is far from the only reason for constructing a framework for modeling historical forces. Formal models discipline thicker, descriptive accounts (Gaddis 2002). By boiling down causes and effects to their spare fundamentals, they enable us to understand the hows and whys; they tell us where to look and where not to look for evidence. They also help us to identify conditions that are necessary and/or sufficient for past choices and outcomes to influence the present.

In light of the many benefits that formal theoretical frameworks would contribute, we cannot but be surprised by the lack of formal models that describe history-dependent

* The author thanks the MacArthur Foundation and the James S. McDonnell Foundation for financial support. Aaron Bramson, Jenna Bednar, Anna Grzymala-Busse, Jacob Hacker, John Jackson Ken Kollman, Skip Lupia, Robert Mickey, John Miller, Burt Monroe, Jim Morrow, Paul Pierson, Troy Tassier, Tarah Wheeler, Ken Zick and participants at the 2003 EITM summer workshops and the Santa Fe Institute Economics Workshop provided valuable comments. Tom Romer strongly influenced many steps on the path that resulted in this essay, and the author is certain this affected the final outcome for the better.

MS submitted 10 February 2005; final version received 16 November 2005
ISSN 1554-0626; DOI 10.1561/100.00000006

processes. One need only compare the concept of path dependence to that of equilibrium to see how far we have to go. Game theorists have refined the concept of equilibrium in myriad ways. In trying to understand or predict an outcome, a game theorist can choose among the quantile response equilibria, universally divine equilibria, evolutionary stable equilibria, and intuitive equilibria. Some may interpret these visions and revisions as splitting the hairs on the heads of so many angels dancing on pins. I disagree. These refinements clarify our thinking about what it means to be rational and provide alternative understandings of how equilibria are chosen and attained.

By comparision to equilibrium, the concept of path dependence seems almost metaphorical (David 1985). In common interpretations, path dependence means that current and future states, actions, or decisions depend on the path of previous states, actions, or decisions. Of late, path dependence has become a popular conveyor of the looser idea that history matters (Crouch and Farrell 2004, Pierson 2004). Theoretical, historical, and empirical studies of path dependence run the gamut, covering topics ranging from the selection of institutions (North 1991), to the formation of government policies (Hacker 2002), to the choice of technologies (Arthur 1994, David 1985), to the location of cities (Arthur 1994, Page 1998), to pest control strategies (Cowan and Gunby 1996), to the formation of languages and law (Hathaway 2001). But are all of these cases describing the same phenomena? Does history operate in each in the same way? Assuredly not.

Surely the micro-level processes that produce path-dependent government policies, pest control strategies, and laws differ markedly. The historical connections between laws differ in kind from those of pest control strategies. Historical forces constrain laws to be similar to past laws, whereas the opposite would seem to be true for pest control strategies. The optimistically named exterminator aims to kill precisely those pests that previous strategies spared. Thus, models should not treat these two cases as if they are the same. When examining histories, we should have the flexibility to choose from a box of lenses.

A survey of the literature on path dependence reveals four related causes: increasing returns, self-reinforcement, positive feedbacks, and lock-in. Though related, these causes differ. Increasing returns means that the more a choice is made or an action is taken, the greater its benefits. Self-reinforcement means that making a choice or taking an action puts in place a set of forces or complementary institutions that encourage that choice to be sustained. With positive feedbacks, an action or choice creates positive externalities when that same choice is made by other people. Positive feedbacks create something like increasing returns, but mathematically, they differ. Increasing returns can be thought of as benefits that rise smoothly as more people make a particular choice and positive feedbacks as little bonuses given to people who already made that choice or who will make that choice in the future. Finally, lock-in means that one choice or action becomes better than any other one because a sufficient number of people have already made that choice.

In this essay, I am less concerned with the causes of path dependence than with providing formal definitions that characterize different types of historical dependence and thinking through their consequences. The stakes here may be large. Path dependence may help explain why some countries succeed and others do not (Easterly 2001). Standard economic growth models predict that less developed countries should catch up with their richer counterparts, but that has not happened. Douglas North has suggested that country-level success depends on the proper build-up of institutions, behaviors, and law. Metaphorically speaking, some countries start out eating their peas with knives and honey, and never move on to the fork or spoon. Others move to new technologies gradually, while others take big jumps (Gerschenkron 1952). Understanding those paths can only be aided by a theory of what those paths might be.

To put some logical structure on the possible sway of history, I describe here a preliminary general framework within which I formally define several forms of history dependence. I differentiate between *path dependence*, where the path of previous outcomes matters, *state dependence* where the paths can be partitioned into a finite number of states which contain all relevant information, and what I call *phat dependence* where the events in the path matter, but not their order. I also distinguish between early and recent path dependence, and perhaps most importantly, between processes in which *outcomes* are history-dependent and those in which the *equilibria* depend on history. By equilibria here I mean limiting distributions over outcomes.

I also present examples of these various forms of path dependence within a dynamic systems framework using two broad classes of models. I base the first on dynamic systems and the second on decision theory. These models help to reveal the causes of path dependence. The proximate cause of history mattering differs in the two classes of models. In the first class of models, the past exerts sway over the present indirectly. For example, prohibiting women from voting affects how women see themselves in relation to men. When given the vote, women's actions are influenced by this past denial of rights. I model that influence as a change in the probabilities over outcomes. In the second class of models, I explicitly model externalities between choices. A train system that connects a central city to an airport creates a positive externality for the airport. When written in the dark lead of mathematics, the line between forces and externalities appears crisp, but that is not necessarily so. Externalities can change incentives, and incentives can be thought of as forces. Similarly, historical events that change how people respond in the future can be thought of as creating externalities with later decisions (Bednar and Page 2006, Medin and Atran 2004).

Within both classes of models, I show that increasing returns is neither necessary nor sufficient for historical dependence (Arthur 1994, David 1985, North 1991, Pierson 2000). This may strike some as counterintuitive given the conflation of the two concepts. Moreover, I show that, even when increasing returns do contribute to historical dependence, it may often only be phat dependence. Path dependence requires a build-up of behavioral routines, social connections, or cognitive structures around an institution. One could argue that the necessary micro-level stickiness and accumulation is implicit in historical accounts like those of Hacker (2002), North (1990), and Pierson (2004), but one could argue the other side – that these accounts present increasing returns alone as the sole cause.[1]

---

1 See Bednar and Page (2006) for a full account of the stickiness argument.

COMMON MISUNDERSTANDINGS

Before constructing the frameworks, I highlight five common misunderstandings that run through in the literature. The first misunderstanding is the aforementioned conflation of path dependence and increasing returns. They are logically distinct concepts. Increasing returns are neither necessary nor sufficient for path dependence. The conflation of increasing returns with path dependence rests on the following logic. If a process generates two possible paths, then some outcome must be more prevalent in one path than in the other. That is true. However, it need not be increasing returns that causes one outcome to be selected more often. Almost any externalities can alter the outcomes.

This first misunderstanding unnecessarily narrows our focus. John Von Neumann once referred to the term nonlinear functions as equivalent to the term non-elephant animals. The same sentiment might be expressed here. Increasing returns can create path dependence but so can almost any type of negative externalities. These could be caused by any kind of constraint: spatial, budgetary, or even cognitive. The set of all externality classes is so large that it is difficult to characterize. So far, we've only been considering the elephant of increasing returns. We need to investigate the other inhabitants of the ark as well.

Consider Hacker's (2002) analysis of health-care policy. As part of the New Deal, Roosevelt pushed for and got pension insurance, social security, but not national health care. Later, wage controls imposed during World War II prevented firms from increasing wages, though they did not prevent them from increasing benefits. Employer-provided health benefits became an obvious byproduct of increased demand for labor. Once some employers began to provide health care, the labor movement and worker expectations led to ever more employers providing health benefits. This history can be interpreted as increasing returns with respect to employer-provided health benefits, but that is an oversimplification. Hacker's full story is one of externalities accumulating on top of externalities. The wage and price controls imposed a negative externality on businesses by making the cost of some actions infinite. At the same time, the growth of a health-care system that catered to firms and not individuals created positive externalities with employer-provided benefit plans and negative externalities with individual health-care plans.

The second misunderstanding stems from a credit assignment problem. In many of the examples of path dependence, while increasing returns do exist, negative externalities are the true cause. This is not merely a reframing of positive relative returns as negative relative returns. It requires a fundamental rethinking of the causes of path dependence. As just mentioned, any constraint, be it a budget constraint, a spatial constraint, or a time constraint, imposes negative externalities and can create path dependence. The logic of constraints applies to competing technologies, legal doctrines, and city locations. In each case, the exclusion of other options drives the path dependence. This occurs even in cases with extremely powerful increasing returns such as in the build-up of offensive forces and weaponry by nation states (Van Evera 1998). Later, I walk through this logic in some detail using Paul David's QWERTY example.

The third misunderstanding stems from the use of the Polya Process (which I formally define later in this essay) as the canonical example of path dependence.[2] As I show, outcomes in the the Polya Process do not depend on the order of past events. They only depend on the distribution over those events. Put in the formal language of this paper: the Polya Process is *phat*-dependent but not *path*-dependent. In a phat-dependent process, the order of events does not matter. Imagine if we were to read the pages of Washington's biography in random order, would the explanation for why he freed his slaves at the end of his life be just as coherent? I doubt it. My limited reading of history suggests that, while historians do not believe that all history is relevant, they often mean much more than phat dependence. They seem to care a great deal about the sequencing of events.

The fourth misunderstanding results from a failure to distinguish between outcomes that are path-dependent and path-dependent equilibria. If this period's outcome depends on the past, that does not imply that the long-run equilibrium does. When scholars refer to history mattering, they typically do not mean that it matters only for singular events. They mean that the course of the future has changed. And yet, evidence that what happened in one period depended on what happened earlier is not sufficient to make such claims. A system can exhibit outcome path dependence yet still have a unique equilibrium, as I show in the Balancing Process.

The final misunderstanding conflates *early path dependence* and *sensitivity to initial conditions*. Sensitivity to initial conditions, a term borrowed from chaos theory, refers to deterministic dynamic systems in which the trajectory or the equilibrium depends sensitively on the initial point of the system. Extreme sensitivity to initial conditions implies minor initial changes have enormous implications. Extreme sensitivity to initial conditions can be defined in deterministic systems. In the simplest examples, a single function f is recursively applied to an outcome, and small differences get magnified with each iteration. In contrast, early path dependence describes processes in which early random outcomes shape the probability distribution over future histories. They do not determine it. They shape it. For example, in the evolution of common law, the doctrine of *stare decisis et non quieta movere*, states that past decisions should stand, but this is not a hard-and-fast rule (Hathaway 2001). Decisions can, and do, get overturned as can the logic that underpins them. The doctrine of *stare decisis* implies that weight be placed on the early history of outcomes but it does not imply sensitivity to initial conditions. The future is not deterministic, but stochastic and biased toward early decisions. The process exhibits early path dependence and not extreme sensitivity to initial conditions.

DYNAMIC PROCESSES

A goal of this essay is to construct precise definitions of the various types of path dependence. With that in mind, I now introduce formal language and definitions. I begin with a general description of a dynamic process that produces *outcomes* at discrete time

[2] See Pierson (2004) and Arthur (1994).

intervals indexed by the integers, $t = 1, 2, \ldots$.[3] I denote the outcome at time t as x_t. In a more general model, in addition to the outcome there may also be other information, opportunities, or events that arise in a given time period. These can be described as the *environment at time t*. This contains exogenous factors that influence outcomes. A *history at time T*, h_T is the combination of all outcomes x_t up through time $T - 1$ and all other factors, the y_t, up through time T. In the ball and urn models that I describe in the next section, there will be no y_t terms, so the history consists only of past outcomes.

A dynamic process also has an *outcome function G* that maps the current history into the next outcome. The outcome generated by a dynamic process can then be written as follows:

$$x_{t+1} = G_t(h_t)$$

The outcome function can change over time so it is indexed by t. The function G_t is not necessarily deterministic. It can also generate a probability distribution over outcomes. This will be the case in many of the examples that I consider. History dependence need not imply deterministic dependence. It need only imply a shift in the probabilities of outcomes as a function of the past.

This stark framework enables me to distinguish between two important ways that history can matter. History can matter in determining the outcome at time t, x_t. I call this *outcome dependence*.

A process is **outcome-dependent** *if the outcome in a period depends on past outcomes or upon the time period.*

History can also matter for the limiting distribution over outcomes. I call this *equilibrium dependence*.[4]

A process is **equilibrium-dependent** *if the long-run distribution over outcomes depends on past outcomes.*

Throughout this paper, I adopt an expansive notion of equilibrium, that of convergence of the long-run distribution of outcomes. Alternatively, I might have required that the outcome functions G_t converge in each period to a common distribution over outcomes. Were I to do that, I would rule out processes that converge to equilibrium cycles or patterns.

Note that equilibrium dependence implies outcome dependence. If the equilibrium distribution over outcomes depends on the past, then so must the outcomes in individual time periods. If history determines which technology will be used in the long run, then it must also determine which technology is selected in some of the periods that make up the long run. As I mentioned earlier when discussing common misunderstandings,

---

3 Notice that this is a discrete time process. I could also write the process as occurring in continuous time. Discrete time processes are far easier to analyze.

4 Of course, a process need not attain an equilibrium distribution over outcomes, but for the purposes of this essay I restrict attention to that case with one exception.

the causality does not go in the other direction. History can matter for outcomes but not for equilibria.

Some examples help clarify this distinction between outcome and equilibrium dependence. Ferejohn (1991) describes historical forces and their impact on electoral competition in Early Stuart England. At that time, people, okay men, alternated holding elected office. No meaningful competition existed. Outcomes were history-dependent: the person elected in any one period depended on who had previously held the seat. However, the equilibrium distribution of who held the seat did not depend on history. Each eligible person held the seat roughly an equal number of times.

To give a more weighty example, the term "manifest destiny" introduced by US politicians in the 1840s implied an unstoppable process of continental expansion. Let us suppose that there was such a thing as manifest destiny. If we think through the counterfactual, at any moment in time, we might have expected the direction and extent of Western expansion to depend on the path taken. Whether farmers, ranchers, gold rushers, or trappers led the way would have determined whether the Northern or Southern part of the continent developed first. Thus, the process of Western expansion generated history-dependent outcomes: which regions gained population in 1852 depended on where people moved in 1849, 1850, and 1851. However, the eventual outcome – Western expansion – may not have depended on history. With or without the gold strike near Sutter's Mill, people would have eventually moved from Ohio to California.

That is not to say that equilibrium dependence cannot exist, and if we take Pierson's (2004) examples at face value, we cannot help but believe that it does. For example, his analysis of the creeping accumulation of judicial power suggests that what happens now *and* what happens in the long run are dependent on what happened in the past. But we have to be careful in defining what we mean by the long run. In general, temporary shifts in the balance of power between the branches of a government need not imply path-dependent equilibria. It may be that institutional safeguards maintain a balance of power between the branches of government. If so, specific outcomes and short-term trends may depend on the past but the long-run equilibrium need not. We see this explicitly in the Balancing Process that I describe in the next section.

THE BALL AND URN MODELS

To describe the specific ways that history can matter either for outcomes or for equilibria, I construct examples using simple ball and urn models. These models generalize the familiar Polya Process and prove incredibly flexible. It is relatively easy to construct a ball and urn example for each of the many definitions I present.

These models consist of a collection of various colored balls placed in an urn. In each period, a ball is selected from the urn and, depending on the color of the ball selected, other balls may be added or removed from the urn. The selection of the ball plays the role of the *outcome function*. Because the ball is selected randomly, the probability of an outcome depends on the composition of the urn: how many balls of each color it contains.

In almost all of the examples that follow, I assume two colors of balls: maroon, which I denote by M, and brown, which I denote by B. To provide some real-world context, rather than think of balls of various colors being drawn from an urn, you can think of a society adding new institutions. These institutions can be either market-based (M) or bureaucratic (B). Colors and institutional choices can be used interchangeably. An outcome M can be thought of as a maroon outcome or as a market outcome. Given this setup, a history of outcomes can be written as a sequence of M's and B's. I first describe processes in which outcomes do not depend on time or on these histories of outcomes. Such processes are called *independent*.

A process is **independent** *if the outcome in any period does not depend upon past outcomes or upon the time period. An independent process can be written as follows:*

$$x_{t+1} = G(\cdot)$$

My first example is called a Bernoulli Process. In a Bernoulli Process, no new balls are ever added to the urn, so the probabilities of selecting a maroon or brown ball never change.

Example 1 A Bernoulli Process The urn contains M maroon balls and B brown balls. Each period a ball is chosen randomly and then put back in the urn. The probability of drawing a maroon ball equals $\frac{M}{(M+B)}$ and the probability of drawing a brown ball equals $\frac{B}{(M+B)}$ in every period.

Many canonical random processes are assumed to be independent: the outcome of a coin flip or the roll of a die, or the sex of a child.[5] Examples of independent processes from the political and economic world are harder to come by, as we typically think that the past matters in some way.

A process that is not independent can be *history-dependent*: the current outcome or both the current outcome and the equilibrium distribution over outcomes could depend on the past history of outcomes. In either case, there remains the question of how much and to what extent history matters for outcomes and equilibria. I distinguish among three types of history dependence: *state dependence*, *phat dependence*, and *path dependence*. These types can be thought of as levels of history dependence with state-dependent processes being the least and path-dependent processes being the most history-dependent.

In some cases, it is possible to partition the space of all histories into a finite number of sets: $\{s_1, \ldots, s_N\}$ such that the outcome function at each moment in time depends only on the set to which the current history belongs. These sets are then called *states*. In the ball and urn models, the number of maroon and brown balls contained in the urn can represent the state of the process, but only if the number of possible combinations of balls is finite. If there can be infinitely many balls in the urn, then there is no state, per se.

Suppose, for example, that two political parties alternate top position on the ballot from election to election and that the outcome in period t is the party named on the

[5] Given the human tendency to predict and recognize patterns in random sequences, people often see path dependency when none exists (Gilovich, Vallone, and Tversky 1985).

top of the ballot. To know the outcome at time $t + 1$, we need only know the outcome at time t. All histories can be partitioned into two states: those in which the first party was last at the top of the ballot and those in which the second party was at the top of the ballot. To give another example, suppose that the Federal Reserve Board used a fixed rule for setting the prime rate (as some fear and some hope) and that this rule depended only on the current inflation rate. If so, the prime rate (the outcome) would depend on history only in so far as the current inflation rate depended on history. History matters, but any two histories that result in the same inflation rate (the same state) are equivalent with respect to outcomes.

Given that there exist only a finite number of states, it is possible to write a mapping from each history into one of these N states. All that remains is to describe how the states change over time. This is determined by a *state transition rule*, T, that maps the current state s_t and (possibly) the current outcome x_t, into the next period's state. This can be written as $s_{t+1} = T(s_t, x_t)$. The state transition rule can be random or deterministic, but it cannot depend on the entire history. It can only depend on the finite states.

A process is **state-dependent**, *if the outcome in any period period depends only upon the state of the process at that time. A state-dependent process can be written as follows:*

$$x_{t+1} = G(s_t) \quad \text{where} \quad s_{t+1} = T(s_t, x_t)$$

Since the outcome only depends on the state, this implies that $G_t = G$ for all time periods t. Such processes are commonly called *Markov Processes.* I refer to them here as state-dependent processes to highlight their differences from phat-dependent and path-dependent processes. For obvious reasons, these processes generate history-dependent outcomes. The history determines the state and the state in turn determines the distribution over outcomes.

I now add two further restrictions. A state-dependent process is said to be *stationary* if the state transition rule T is the same in every time period. A state-dependent process is said to be *ergodic* if through some series of states it is possible to get from any one state to any other. I now state one of the most important, and I dare say neglected, theorems in historical analysis, the Ergodic Theorem.

The Ergodic Theorem *A stationary, ergodic, state-dependent process generates a unique equilibrium distribution over outcomes.*

The Ergodic Theorem raises the bar quite high for anyone who wants to claim that history matters in the long run. The theorem says that (i) if it is possible to define a finite set of relevant states that determine the next outcome, (ii) if the mapping from states to states as a function of outcomes does not change, and (iii) if it is possible (no matter how unlikely) to get from any state to any other, then the process converges to a unique distribution over outcomes. That equilibrium, being the only one, is also stable.[6]

[6] The phrase "E.T. phone home" can be a helpful reminder here. The Ergodic Theorem (E.T.) states that the distribution is called to a single place (home).

The Ergodic Theorem does not deny outcome dependence. In fact, only the most trivial state-dependent processes do not exhibit outcome dependence. It says that if we were to run a process many times, and if each time we were to bin the outcomes and create a distribution, we would find that the distributions were all the same. In the long run, the history of outcomes would not have mattered. This leads to my second observation, which is a restatement, albeit an important one, of the Ergodic Theorem.

Observation 1 *Equilibrium dependence requires a changing process, a big space, or divergent paths.*

This observation can be restated more formally. To have multiple equilibria, one of the four core assumptions of the Ergodic Theorem must be violated: (i) the outcome function must change over time, or (ii) the transition function between states must change over time, or (iii) the number of states must not be finite, or (iv) getting from some states to some other states must be impossible via any sequence of states. To assume the first is to say that the world changes over time. If it changed as a function of an outcome, then this would, in effect, assume path dependence. That's perfectly admissible. One interpretation of the concept of *lever points* is that they are moments in time in which an outcome changes the course of history, and in doing so changes the outcome function. If that happens, then certainly history would matter.

A process might well violate stationarity. If the state represented a voter's political ideology: conservative, liberal, or independent, the transition probabilities between ideologies could be a function of the ideologies of other voters. If so, the transition probabilities would change over time and multiple equilibria could be sustained (Page 2006). It is also possible to violate the assumption of a finite number of states, but more relevantly, sometimes the set of states can be so large that the time to get to the equilibrium may be so long that it becomes a meaningless prediction. Finally, a system can violate ergodicity. Sometimes, two roads really do diverge in a yellow wood, and when we take one we cannot go back and take the other just as fair.[7] Making the case that a process is not ergodic is not as easy as it sounds. It has to be shown that some state can never be reached by *any path* from some other state.

I now turn to formal definitions of path and phat dependence. If the history of outcomes matters, but not the order in which they occurred, I define the process as *phat-dependent*. I chose the word phat for two reasons beyond the obvious desire to establish my hip-hop bona-fides. Phat is not only an acronym for Pretty Hot And Tempting (which assuredly the concept of phat dependence is) but also an anagram for path. As such, the word phat reminds us that the order does not matter, even though the outcomes do. Phat also sounds like "fat" which is a synonym for "thick." This serves to remind us that logically consistent historical narratives, unlike the casual, potted histories that I present here, require thick description.

---

[7] If Frost had been more mathematically inclined, rather than writing, "Yet knowing how way leads on to way, I doubted if I should ever come back," he may have written "Yet knowing the paths were not ergodic, I knew that I should never come back."

A process is **phat-dependent** *if the outcome in any period depends on the set of outcomes and opportunities that arose in a history but not upon their order. A phat-dependent process can be written as follows:*

$$x_{t+1} = G_t(\{h_t\})$$

Where $\{h_t\}$ denotes the set of outcomes up to time t.

If the order of the history of outcomes also matters, I define the process as *path dependence*. The definition of path dependence looks almost identical to that of phat-dependence. The lone difference is that for a path-dependent process the outcome function depends on the vector of history h_t. Changing the order of x_1 and x_2 could change the outcome produced by G_t.

A process is **path-dependent** *if the outcome in any period depends on history and can depend on their order. A path-dependent process can be written as follows:*

$$x_{t+1} = G_t(h_t)$$

To put this in the context of the ball and urn models, the *path* is the ordered set of all previous outcomes. The paths MBM, BMM, and MMB create the same set of outcomes, namely $\{M, M, B\}$. If only this set matters, and not the order in which the outcomes arose, the process is *phat dependence*. For example, jurors voting on a defendant's guilt or innocence often sequentially reveal their opinions. A juror's opinion might change in response to the expressed opinions of others. If each juror considered only the number of previous jurors voting guilty and innocent when making his or her own opinion known, the process would be phat-dependent. But if each juror also took into account the order in which the other opinions were voiced, the process would be path-dependent. Testing for phat dependence requires a different econometric model than testing for path dependence, a point I return to later.[8]

In contrast, the distinction between state dependence and phat and path dependence is blurrier. Any history of past outcomes could be written as a state, but in some cases the set of possible states would be enormous. This means that we can transform any process, be it path- or phat-dependent, into a state-dependent process.[9] This would seem to imply that all processes are state-dependent, and in some trivial sense they are. However, in many cases, the number of histories and sets of histories would get larger each period. Here, I assume that the set of states is finite and fixed. This assumption precludes me from reinterpreting paths and sets of paths as states.

My next example is the famous Polya Process. In the Polya Process, the urn initially contains one brown and one maroon ball. In each period, a ball is selected and returned to the urn, and another ball is added to the urn of the same color as the selected ball. At least metaphorically, this process is thought to capture the phenomenon of increasing returns that Arthur (1994), David (1985), Pierson (2004) and others describe. The more

---

[8] I thank John Jackson and Ken Kollman for this observation.

[9] This is trivially accomplished. Define the state to be the past history of outcomes.

often that market solutions (maroon balls) are chosen, the more likely that they will be chosen in the future.

Example 2 The Polya Process Initially, $M = B = 1$. In any period, if a brown (resp. a maroon) ball is selected then it is put back in the urn together with an additional ball of the same color.

The Polya Process is equilibrium-dependent. Not only can the process converge to more than one ratio of maroon and brown balls, it can converge to *any* ratio of maroon and brown balls. Depending upon the history of outcomes the urn could eventually contain 80% maroon balls and 20% brown balls, or it could contain 63% maroon balls and 37% brown balls. At some point, the urn contains enough balls that the ratio converges, and balls continue to be selected in those proportions.[10] Thus, the Polya Process is equilibrium-dependent and also outcome-dependent.

The Polya Process is not, however, path-dependent, but it is phat-dependent. The outcome at time t only depends on the set of past outcomes, not on their order. In period five, if there are three additional maroon balls in the urn and one additional brown ball, the outcome is the same regardless of the order in which those balls were selected. It does not matter whether the three maroon balls were chosen first and then one brown ball or whether the one brown ball was chosen and then three maroon balls.[11]

Observation 2 *The Polya Process is phat-dependent.*

The fact that the Polya Process is only phat-dependent does not imply that the real-world situations it has been used to describe are only phat-dependent as well. Nor is phat dependence only a peculiarity of the Polya Process. Social scientists often assume phat dependence in analyzing outcomes. In calculating the ideologies of congresspeople, the order of the votes within a congress is not considered, just the votes themselves.[12] A model that uses ideology to predict an outcome – in this instance, a current vote – implicitly assumes only phat dependence. In predicting the number of votes that a congressperson, let's say my congressperson, John Dingell, will receive in his next congressional race, we might regress an equation based on a host of variables: his ideology, his reputation, his positions on the issues, his opponent, his supporters and his opponent's supporters, the amount of money he and his opponent have, and characteristics of his district. If this regression equation does not include any time lags it captures only phat dependence.

---

[10] The Polya Process has many interesting features that lie beyond the scope of this analysis. See Arthur (1994).

[11] In addition, both of these paths are equally likely. This is easily shown. The probability of the first path is $(1/2)*(2/3)*(3/4)*(1/5)$. The probability of the second path is $(1/2)*(1/3)*(2/4)*(3/5)$. This holds in general. For any distribution of balls, any history consistent with that history is equally likely. This means that the set does not tell us much about the past.

[12] Note: If the number of ideologies were finite, then this would be a state-dependent process.

An obvious question to ask is whether every phat-dependent process is equilibrium phat-dependent or whether it might be only outcome phat-dependent. My next example, the Balancing Process, exhibits only outcome phat-dependence: it is has a unique equilibrium distribution.

In the Balancing Process, the probability of selecting a ball of a given color again depends on the colors of the previous balls chosen, but in this case the feedbacks are negative. Instead of inserting a ball that is the same color as the ball that was drawn, as in the Polya Process, now a ball of the *opposite* color is added to the urn. As more maroon balls are chosen brown balls become more prevalent in the urn and become more likely to be chosen.

Example 3 The Balancing Process Initially $M = B = 1$. In any period, if a brown (resp. a maroon) ball is selected then it is put back in the urn together with an additional ball of the opposite color.

To see why the Balancing Process cannot generate multiple equilibria, suppose that the process converged to something other than an equal number of maroon and brown balls. Imagine an urn with a large number of balls, 60% of which are maroon and 40% of which are brown. From that point onward, maroon balls would be more likely to be selected. Selecting these maroon balls would add brown balls to the urn, increasing the proportion of brown balls above 40%.

I previously discussed the balance of powers in a government, but balancing occurs in other contexts as well. If one political constituency prefers market solutions and another prefers bureaucratic solutions, then successes by one constituency may result in the mustering of greater political forces by the other: The addition of another market-based institution may create more pressure for future bureaucratic solutions. A concern for fairness might also create balancing forces. Rotation schemes (Kollman 2003) in the European Union are an extreme example of balancing, but more subtle balancing may occur in the selection of locations for such things as political conventions, the World Cup, or the summer and winter Olympics.[13]

As I discussed earlier, much of the literature on path dependence emphasizes increasing returns. The informal definition of increasing returns goes as follows: the more an outcome occurs, the higher the relative return to that outcome, and therefore, the more likely it occurs in the future. Increasing returns within the class of urn processes is a simpler concept. It leaves out the middle step in the causal chain. I present here a precise definition of increasing returns in the class of urn models.

A dynamic process generates **increasing returns** *if an outcome of any type in period t increases the probability of generating that outcome in the next period.*

Given this definition, the Polya Process satisfies increasing returns. It is thus possible for an urn process to exhibit increasing returns and to generate multiple equilibria. That does

13 I thank participants at the NSF-sponsored EITM at the University of Michigan for these examples.

not mean that all processes with increasing returns generate multiple equilibria, nor does it imply that all process that generate multiple equilibria satisfy increasing returns. In fact, no logical implication exists in either direction, as I show in the next two examples. The first example relies on red (R) and green (G) balls as well as maroon and brown balls. I call it the Balancing Polya Process. It is a combination of the Balancing Process and the Polya Process.

Example 4 The Balancing Polya Process Initially, $M = B = R = G = 1$. In each period, the ball selected is returned to the urn. In addition, if a red ball is selected, a maroon ball is added to the urn. If a maroon ball is selected, a red ball is added to the urn. If a green ball is selected, a brown ball is added to the urn. And if a brown ball is selected, a green ball is added to the urn.

The rules for adding balls in the Balancing Polya Process can be written as follows:

$$\begin{array}{rcl} \text{Pick} & R \rightarrow \text{Add} & M \\ \text{Pick} & M \rightarrow \text{Add} & R \\ \text{Pick} & G \rightarrow \text{Add} & B \\ \text{Pick} & B \rightarrow \text{Add} & G \end{array}$$

To show that this process exhibits equilibrium phat-dependence, paint the red balls maroon and the green balls brown. Doing so creates the Polya Process which, as we know, is equilibrium phat-dependent. It is also easy to show that the Balancing Polya Process does not satisfy increasing returns. In any given period, choosing any color ball decreases the probability of picking that ball in the next period. Thus, increasing returns are not necessary for equilibrium dependence.[14]

The Balancing Polya Process also provides a hint as to how complementarities between outcomes – red outcomes creating an environment favorable for maroon outcomes in the future – can generate equilibrium dependence. Many scholars, e.g. Ikenberry (2001), North (1990), and Pierson (2004), argue that there exist strong complementarities between institutions. The Balancing Polya Process provides some intuition for how those complementarities, and not increasing returns per se, can generate equilibrium dependence.

My next example proves a lack of sufficiency of increasing returns for path dependence. It exhibits increasing returns for both maroon and brown balls but does not generate path- or even phat-dependent equilibria. I call this the Biased Polya Process, as the brown balls have an advantage.

[14] The Balancing Polya Process is not the only example that demonstrates a lack of necessity. Another example is the *Locked-out Process*, which is defined as follows. Initially, $M = B = 10$. In the first 19 periods, a ball is selected and removed from the urn. In subsequent periods, the one remaining ball is repeatedly selected and returned to the urn. The Locked-out Process exhibits equilibrium dependence. It converges to one of two equilibria. Eventually, either all brown or all maroon balls are selected.

Example 5 The Biased Polya Process Initially, $M = 1$ and $B = 2$. In each period a ball is selected. If a maroon ball is selected, it is put back in the urn together with another maroon ball and another brown ball. If a brown ball is selected in period t, it is put back in the urn together with $2t$ additional brown balls.

With probability one, eventually a brown ball is selected. Once that brown ball is selected, the probability that the next ball is brown exceeds 75%. Eventually, the proportion of brown balls in the urn converges to 100%, so this process generates a unique equilibrium. Selecting a brown ball clearly satisfies increasing returns. Select a brown ball in one period, and a brown ball is more likely to be selected in the next period. Surprisingly, though, maroon balls also satisfy increasing returns. Select a maroon ball, and the probability of selecting a maroon ball in the next period also increases.[15]

Observation 3 *Increasing returns are neither necessary nor sufficient for equilibrium dependence.*

In light of this observation, why the common conflation between increasing returns and path dependence? The logic goes as follows: for a process to exhibit equilibrium dependence it must converge to at least two distinct equilibrium probability distributions. For this to happen in an urn model, in one of these equilibria the proportion of maroon balls must strictly exceed the proportion in the other equilibrium. Consider two paths, one that leads to the first equilibrium and one that leads to the second. Along the first path, an outcome in a given period may or may not have some effect on the equilibrium. But in some periods these outcomes have to shift the probabilities of future outcomes. Theoretically, weight could be assigned to each outcome in terms of how much influence it has on future outcomes. These influences would have to exaggerate one outcome at the expense of another. This is why people think path dependence requires increasing returns. The flaw in this logic is that the reinforcement or exclusion need not be in the form of increasing returns. They could occur through complementarities as was true in the Balancing Polya Process. The logic that increasing returns implies path dependence seems obvious, but it too is incorrect. All of the outcomes could have increasing returns, but if one outcome has much stronger increasing returns than the others, it will always win. That is true of the brown balls in the Biased Polya Process.

Does this mean that claims that link path dependence to increasing returns are incorrect? Not entirely. Much of this research claims that a particular institutional arrangement, e.g. rent seeking by oil-rich states (Karl 1997) or party patronage systems (Shefter 1977) comes into being and then because of increasing returns results in path-dependent outcomes. Karl's example of oil-rich states proves instructive here. That example can be more accurately described as including both increasing returns and negative externalities. Easy profits in the oil industry destroy the incentives for developing other industries. So while it is true that oil begets oil, it is equally true that oil

[15] The increase can be explained as follows: in each period there are fewer maroon balls than brown balls, so adding one maroon and one brown increases the relative proportion of maroon balls.

precludes cars and semiconductors, and to quote Frost (with a sigh), "that has made all the difference."

More generally, when someone finds evidence of increasing returns and positive feedbacks, whether it be with respect to technology or international institutions (Weber 1997) this evidence is not sufficient proof that multiple equilibria exist. It could well be that the system has a single equilibrium. All that must be true is that one configuration's positive feedbacks and increasing returns are so large as to swamp any others. An historical example of this would be the competition between gas-, steam-, and electric-powered cars that took place a century ago. All three technologies exhibited the classic properties of increasing returns to scale technologies: falling production costs, networked delivery systems, etc. Initially, electric cars outsold gasoline cars. However, gasoline-powered cars had much larger increasing returns. A person cannot carry electricity in a can, and what electric infrastructure existed was inside cities. Getting gas to the farm was much easier. And America's population was mostly rural. Rerun history ten times, a thousand times, even a million times, and the gasoline engine probably wins almost every time. Increasing returns? Yes. Path dependence? Probably not.

The Biased Polya Process is path-dependent. If in the first five periods only one brown ball is selected, the number of balls in the urn depends on the period in which it was selected. However, there are multiple paths that generate the same outcome probabilities. If, in the fourth period, there are ten brown balls and three maroon balls, it could either be that a brown ball was selected only in period 3, or that a brown ball was selected in periods 1 and 2 but not in period 3. It is therefore possible to construct an even stronger notion of path dependence, namely that *any two distinct paths lead to different outcome probabilities*, I refer to this as *strong path dependence*.

A process is **strong path-dependent** *if, for any two distinct histories, the outcome function differs. A strong path-dependent process can be written as follows:*

$$x_{t+1} = G_t(h_t)$$

Where $G_t(h_t) \neq G_t(\hat{h}_t)$ if $h_i \neq \hat{h}_i$ for some $i = 1$ to t.

Strong path dependence implies path dependence. Strong path dependence might also be called *order of the path* dependence. The next two examples show that it is possible to construct a strong path-dependent process using a simple urn model, and that it is also possible (up to a set of paths of measure zero) to construct a strong path-dependent process that is equilibrium-dependent.

Example 6 A Strong Path-dependent Process Initially $M = B = 1$. In period t, a ball is chosen and 2^{t-1} balls are added to the urn of the color of the chosen ball.

To see that this process is strong path-dependent, it helps to consider a specific path of outcomes, say $MBMMB$. After period 1, a maroon ball is added. After period 2, two brown balls are added. After periods 3 and 4, four and eight maroon balls are added, and after period 5, 16 brown balls are added. Therefore, before the ball is chosen in period 6 there are 14 maroon balls $(1 + 1 + 4 + 8)$ and there are 19 brown balls $(1 + 2 + 16)$. It can be shown that $MBMMMB$ is the unique history that generates 14 maroon balls and 19 brown balls.[16] Therefore, the process is strong path-dependent. Unfortunately, this process does not converge to any fixed probability of selecting a maroon ball. The next process, however, does converge up to a set of measure zero.

Example 7 The Burden of History Process Initially $M = B = 1$. In period t, a ball is chosen and put back in the urn together with a ball of the same color. In addition, for each period $s < t$, $2^{t-s} - 2^{t-s-1}$ balls are added to the urn of the color of the ball chosen in period s.

This process relies on the same basic construction as the previous example. After T periods, there will be 2^{T-1} balls placed in the urn that match the ball selected in the first period, 2^{T-2} that match the ball selected in the second period and so on. In this process, the first ball selected always determines the color of approximately one half of the balls added to the urn, the second ball selected determines approximately one fourth of the balls added to the urn and so on. Later periods matter exponentially less. The process can be shown to converge to a unique equilibrium distribution for any history up to a set of measure zero.[17]

This last process further clarifies the relationship between path dependence and increasing returns. The Polya Process notwithstanding, path dependence is loosely conceived as implying that the entire path matters. Scholars often make an implicit assumption that the weight of early history matters more, that early decisions, actions, and choices grow more and more important over time. In this last process, as time unfolds the past takes on more and more weight. In doing so it creates strong path-dependent equilibria. These two loose conceptualizations: that of strong path dependence and that of an increasing weight of past history, are linked. Even though increasing returns are neither necessary nor sufficient for equilibrium dependence, some form of increasing returns, or increasing complementarities, can create path dependence and not just phat dependence. I return to this insight later in this essay, when I discuss the possibility of path dependence in the externality models.

Early and Recent Path Dependence

Some processes do not depend on the entire path, but on the initial history or on the early path. In these cases the histories can be partitioned into a finite number of sets, so the processes can be thought of as state-dependent. However, if the early part of the

---

[16] Change M's to 1's and B's to 0's. This gives a binary sequence. The number of maroon balls in the urn equals the integer conversion of the binary sequence which is unique.

[17] It is possible in this process that two histories converge to the same probability distribution. For example, the history of one maroon ball and then all brown balls and the history of one brown ball and then all maroon balls both converge to equal probability of maroon and brown balls. But this is a zero probability event.

path affects later outcomes, then these processes cannot satisfy all of the assumptions of the Ergodic Theorem. Typically, they violate ergodicity. Once an outcome occurs several times, it becomes *locked in*.

A process is **initial outcome-dependent** *if all subsequent outcomes depend only on the first outcome. An initial outcome-dependent process can be written as follows:*

$$x_t = G(h_2)$$

The next example describes an initial outcome-dependent process in which a founder makes a random decision which subsequently charts the future course of events deterministically.

Example 8 The Founder Process $M = B = 1$. If the ball chosen in period 1 is maroon, the maroon ball is put back in the urn and the brown ball is removed. Similarly, if the ball chosen is brown, the brown ball is put back in the urn and the maroon ball removed.

This process has only two paths. All future outcomes must be the same as the first outcome. This process is an example of strictly initial outcome dependence.

A process is **early path-dependent** *if the outcome in any subsequent period depends only upon the history up to some period T. An early path-dependent process can be written as follows:*

$$x_{t+1} = G_t(h_t) \quad \text{for} \quad t \leq T, \quad \text{and} \quad x_{t+1} = G(h_T) \quad \text{for} \quad t > T$$

This phenomenon has been popularized in the information cascades literature (Bikhchandani, Hirshleifer, and Welch 1992, 1998; Lee 1993) When enough consecutive people vote yes or buy a stock, others follow like lemmings regardless of their information, so only the early part of the path matters. Example 9 provides a simplified version of a cascade.

Example 9 A Cascade Initially $M = B = 1$. Balls are selected and replaced in the urn until three consecutive balls of the same color are selected. When this occurs, the ball of the other color is removed from the urn.

In this process, market and bureaucratic outcomes are equally likely until three consecutive outcomes are identical, at which point the process locks into that outcome. This captures the phenomenon that the early history matters exclusively.

Processes can also depend not on the early path but on the recent path. Retrospective voting processes would be an example of recent path dependence. The notion of recent path dependence runs counter to common conceptions of path dependence which emphasize early decisions. Many empirical investigations suggest the prevalence of recent path dependence. If the recent path matters, then the depth of a process can be measured as the number of periods back that influence the next state and outcome.[18]

---

[18] This notion of depth is borrowed from the physical concept of thermodynamic depth.

In the interests of brevity, I construct examples of last outcome dependence and recent path dependence but only formally define recent path dependence.

A process is **recent path-dependent** *if the outcome in any subsequent period depends only upon the outcomes and opportunities in the recent past. A recent path-dependent process can be written as follows:*

$$x_{t+1} = G_t(h_t) \quad \text{for} \quad t \leq T, \quad x_{t+1} = G(h_t/h_{t-T}) \quad \text{for} \quad t > T$$

I call the example the Unstable Government Process. Imagine that there are two parties and that each party has a different preferred institution, either M or B. If a party's preferred institution is selected, then the party stays in power. If not, the party falls out of power.

Example 10 The Unstable Government Process Either of two parties, D or R, can be in power. If D is in power then $M = 1$ and $B = 2$. If R is power then $M = 2$, and $B = 1$. A randomly chosen party is in power in the first period.[19] Each period, a ball is chosen. In subsequent periods, the party in power equals D if a B was chosen and equals R if an M was chosen.

This process satisfies all of the conditions of the Ergodic Theorem, and in the unique equilibrium distribution, maroon and brown balls are equally likely to be selected. The second example I call a Forgetting Process. The process has finite memory; only the recent past matters.

Example 11 A Forgetting Process Initially $M = B = 1$. An additional ball of the same color as the selected ball is added for $K > 0$ periods and then removed from the urn.

The Forgetting Process depends on more than just the last outcome, so it differs from the Unstable Government Process. If a process discounts the past at some rate, then a process like the Forgetting Process may not be a bad approximation.

SUMMARY OF URN MODELS

The previous definitions do not create a classification. Multiple definitions often apply to a single case. For example, a recent outcome path-dependent process can also generate early path-dependent equilibria. Consider an urn that begins with five maroon and five brown balls that are numbered from one to ten. In each new period, when a ball is drawn it is replaced and another ball of the same color is added. If, in addition, the lowest numbered ball is removed from the urn, then this process generates both an early path-dependent equilibrium and a recent path-dependent equilibrium, as at some point the urn contains only one color of ball. Table 1 tells which properties each example exhibits.

[19] Parties are each chosen with probability one half.

Table 1. Ball and urn models summary

Process	Properties
1 Bernoulli	Independent
2 Polya	Phat-dependent: multiple equilibria
3 Balancing	Phat-dependent, unique equilibria
4 Balancing Polya	Phat-dependent equilibria, not increasing returns
5 Biased Polya	Increasing returns, unique equilibria
6 Strong Path-Dependent	Path-dependent, may not converge
7 Burden of History	Path-dependent, converges
8 Founder	Initial outcome-dependent
9 Cascade	Early path-dependent
10 Unstable Government	Last outcome-dependent
11 Forgetting	Recent path-dependent

A MODEL WITH EXTERNALITIES

I now describe a second class of models that rely on externalities between actions and choices. These models are decision theoretic. Though their notation differs from the ball and urn models, they have the same basic structure. In each period there is an outcome, and that outcome may or may not change the probabilities over future outcomes. In the ball and urn models, those changes in probabilities depend on a fixed rule. In the externality models, those probabilities shift owing to the calculations of a payoff-maximizing decision maker. This second class of models enables a deeper unpacking of the relationship between increasing returns, positive externalities and path and phat dependence. It also suggests other ways one might extend the ball and urn models to make them more useful to social scientists.

I assume that there exists a decision maker who takes sequential actions. I describe these actions as decisions over project proposals such as whether to build a highway or bridge or to create a new institution. Thus, an action can be to accept or reject a proposed project. I assume that the sequence is infinite so that the decision maker cannot wait until seeing all possible proposals prior to taking an action.

In each period the decision maker considers a single proposal. Proposals can then be identified by the period in which they are considered. An extended model in which multiple decisions arise in any period yields similar results.[20] I further assume that, once approved, a proposal cannot be reversed, but that a rejected proposal can be approved in any later period. This assumption captures that many public projects cannot be undone without some cost. I am, in effect, assuming that this cost is infinite.[21]

[20] Multiple periods in the current model can be condensed to a single period in a more general model.

[21] This is a strong form of irreversibility. Without some type of irreversibility, history dependence would rarely occur because anything done could be undone – at any time, a decision maker could

In the model, the value of the proposal in period t equals its *isolated value*, v_t, which may be either positive or negative plus the value of any externalities it generates with other approved proposals. I let e_{st} denote the externality between the proposal in period t and the proposal in period s, where $s < t$. These e_{st}'s can be positive, negative, or of zero value.[22] It is important that externalities are only realized among proposals that are approved. The decision maker knows this isolated value as well as the values of all the externalities. This is a strong assumption, but appropriate in light of the ambitions of this essay. Similar results could be derived in a model with uncertainty.

To see how the model unfolds, suppose that the proposals considered in period 1 and period 3 have been approved. Even though the proposal considered in period 3 generates externalities e_{13} and e_{23}, only the former is realized. The value of the set of approved proposals after period 3 equals $v_1 + v_3 + e_{13}$, and the marginal value of adding the proposal considered in period 4 equals its isolated value, v_4, plus the the sum of the values of the externalities it creates with the approved proposals $e_{14} + e_{34}$. Without externalities, this model is not interesting. History dependence cannot exist. Each proposal can be made in isolation. With externalities, earlier proposals can constrain or influence later proposals, and the decision rule can be history-dependent.

I assume there that the decision maker does not consider the future. The decision maker therefore approves a proposal (or a set of proposals) if and only if its expected value is positive. This makes path-dependent outcomes very likely yet, as I show, they still need not happen. I next introduce a model-specific definition of a history-dependent decision rule. For a given finite set of proposals, an *ordering* of those proposals is a sequence and a *permutation* of that ordering is a rearrangement of that sequence.

A decision rule is **history-dependent** *if there exists a reordering of some finite set of proposals that yields a different set of approved proposals.*

This is an analog of my previous definition of history dependence redefined within this model. The outcome x_t can be thought of as the set of proposals. The definition requires that the history of proposals influences the set of proposals approved. History-independent rules, such as a rule that accepts or rejects all proposals, trivially exist. But neither of those rules would necessarily be a good rule to follow. The relevant question is whether there exists a decision rule that makes optimal choices which is not history-dependent. As the sequence of proposals is infinite, this necessitates a clarification of what is optimal. I could define optimality with respect to the present discounted value of proposals. I could also define optimality with respect to each period. The main results of this section hold given either convention. I choose to adopt the latter convention of optimality because it makes the analysis more transparent. Define the *acceptance set in*

undo the past and implement the optimal set of choices. The assumption of infinite costs can be relaxed without changing the main result that history dependence is not implied by positive externalities or increasing returns but the mathematics would be more cumbersome.

[22] Note: I am assuming that all externalities are among pairs of proposals. The results still hold if I admit higher-order externalities.

period t as the set of proposals approved in the first t periods. I can then define optimality as follows:

A decision rule is **optimal** *in period t if the acceptance set in period t obtains the highest possible value for all subsets of the first t proposals.*

Suppose that these decisions were being made by a government. If an election were held after period t, the incumbent government would want to have approved the best set of proposals from among those possible.

Notice that this construction implies that a decision rule that is optimal up to period t need not be optimal up to period $t + 1$. If a proposal arises in period 7 that creates negative externalities with proposal 3, then the optimal set of proposals at the end of period 6 may include proposal 3 and not proposal 7. But if proposal 7 has a high isolated value, then the optimal set of proposals at the end of period 7 may include proposal 7 and exclude proposal 3. Given that acceptances are irreversible, once proposal 3 is approved, it cannot be undone. Therefore, any rule that is optimal in period 6 cannot be optimal in period 7.

To avoid the messiness that arises with multiple optimal acceptance sets, I assume that for any t there is a unique optimal acceptance set. Without this restriction, it is possible to construct examples in which the proposal rule is history-dependent, but the paths chosen have the same value and therefore history is irrelevant.[23] I can now state the following observation.

Observation 4 *A decision rule that is optimal in every period cannot be history-dependent.*

The proof is by contradiction. Suppose that a decision-making rule is history-dependent. It follows that it cannot be optimal for all t. Consider an ordering of proposals and any permutation of that ordering that generate distinct acceptance sets in some period $\hat{t}$. In period $\hat{t}$ only one of these acceptance sets can be optimal; a contradiction.

Showing that optimality and history dependence cannot coexist does not settle matters. It remains to show that there exists an optimal decision rule in nontrivial contexts (otherwise the observation is vacuous), but that is easily accomplished. Assuming only positive externalities, the optimal decision rule in period t is as follows: (i) accept the proposal in period t if it has a positive isolated value; (ii) accept the proposal in period t if its isolated value plus its externalities with all approved proposals is positive; and (iii) after the proposal in period t has been accepted or rejected using criteria (i) and (ii), accept any subset of previously rejected proposals if when considered collectively increase the value function.

To see how to apply this rule, suppose that each of the three proposals has an isolated value of negative 5 but that each generates a positive externality of 6 with every other proposal. Proposal 1 considered alone would be rejected using this rule. Proposal 2

[23] This assumption can be relaxed in several ways, but I want to make the intuition as clean as possible.

considered alone and as a pair with proposal 1 would also be rejected. Proposal 3 would be rejected when considered alone but, when considered with both proposals 1 and 2, the three of them would be approved because they create a total value of three ($6 + 6 + 6 - 5 - 5 - 5$). The intuition behind the optimality of this rule is straightforward. Since all externalities are positive, any proposal or subset of proposals that contributes positively to the total value in period t will continue to do so as more proposals are added. This verbal argument can be formalized.

Observation 5 *With only positive externalities between proposals, there exists a decision rule that is optimal up to period* t *for all* t.

For a proof of this observation see Page (1997). A corollary to this observation is that increasing returns need not imply a path-dependent decision rule.

Corollary 1 *With only positive externalities between proposals, there exists an optimal decision rule, which is therefore not history-dependent.*

I next show that this observation and its corollary apply to environments with increasing returns. I first define increasing returns using the externalities framework I have constructed. To accomplish this, I create types. These types represent the characteristics of proposals or institutions that create increasing returns. I assume that each proposal belongs to one of a finite number of types: $\{A, B, C, \ldots, L\}$. I then say that a collection of proposals exhibits *exclusively increasing returns* if proposals belonging to the same type create positive externalities with one another but create no other externalities. Formally, this means that e_{st} is positive if the proposals considered in periods s and t belong to the same type and is zero otherwise.

Observation 6 *If proposals exhibit exclusively increasing returns then there exists an optimal decision rule which is therefore not history-dependent.*

Thus, positive externalities do not imply path dependence, but the presence of negative externalities makes path dependence difficult to avoid. The reason for this is obvious. Negative externalities imply that, once a proposal is made, future proposals are constrained in a way that compromises optimality. This can be stated formally.

Observation 7 *If there does not exist an optimal decision rule for some subset of proposals for all periods up to period* t*, then there exists at least one negative externality between proposals.*

For a proof of this observation see the appendix.

With these observations in mind, I return to the example of the QWERTY typewriter keyboard. In the great sweep of history, keyboard configurations are significantly less than a minor footnote, nevertheless the QWERTY keyboard is by far the best known example of a path-dependent equilibrium. The use of the Polya Process to explain the QWERTY example contributes to why people confuse path dependence and increasing returns.

What follows is a canned version of history of QWERTY together with an argument that negative externalities are the cause of the path dependence.[24]

The crude history goes as follows: though there were many possible keyboard arrangements, for a variety of reasons, initially QWERTY typewriters dominated the market. Typing on a keyboard with a different key configuration requires learning to type anew. This meant that, as more people bought QWERTY keyboards, QWERTY keyboards became locked-in. QWERTY keyboards exhibit increasing returns and they created a path-dependent process. Eventuallly, people then could only buy typewriters with QWERTY keyboards.[25]

This simple version contains a subtle but important blurring of the causes of path dependence. While it is true that the QWERTY typewriter keyboard exhibits increasing returns, these are increasing *relative* returns, not increasing absolute returns. The QWERTY typewriter becomes relatively more valuable than other typewriters because it creates both positive and negative externalities. The positive externality is between QWERTY keyboards. The more people who could type on QWERTY keyboards, the more valuable a QWERTY keyboard becomes as other QWERTY typists can also use them. But the sale of more QWERTY typewriters also creates two negative externalities. Once a person has one typewriter, that person derives little benefit from having another one, especially one with a different keyboard configuration. This is an *intra-personal* negative externality. Second, as more people buy QWERTY keyboards, the value to a person of learning on another keyboard decreases because of the prevalence of QWERTY keyboards. This is an *interpersonal* negative externality.

Therefore, the increasing relative returns to QWERTY are due to both *positive externalities* for QWERTY typewriters and *negative externalities* imposed on other typewriters. The same logic applies in the case of VHS and BETA videotape technologies. As more people bought VHS machines, this created positive externalities for future VHS purchasers (more available titles) and negative externalities for future BETA purchasers (relatively fewer available titles).

This decomposition of increasing relative returns into positive and negative externalities enables a recasting of the QWERTY example in which the negative externalities can be seen as the primary cause of path dependence. To show this, I construct a simple model with two keyboard configurations: the QWERTY keyboard, Q, and the ZRJSOC keyboard, Z. Let N_Q and N_Z denote the number of people who have bought each type and let $x_i \in \{0, Q, Z, B\}$, denote whether a person i has purchased no typewriter, a Q typewriter, a Z typewriter, or both typewriters. In this simple model each person has a value for each keyboard. The values depend on the number of people who have bought each type as well as whether or not the person has bought a typewriter. I write person i's values for the Q and Z typewriters as $V_{iQ}(N_Q, N_Z, x_i)$ and $V_{iZ}(N_Q, N_Z, x_i)$, respectively, and I denote prices by p_Q and p_Z.

[24] See Liebowitz and Margolis (2002) for a more complete analysis of QWERTY.

[25] My favorite explanation for the QWERTY configuration is that it allowed people to type the word *typewriter* without moving their fingers from the top row of keys.

I assume that there are many periods and that, in each period, a salesperson shows up at a random person's door and offers for sale either a Q or a Z typewriter. This choice over keyboards is assumed to be random. I first assume only positive externalities. If the first person buys a Q typewriter, this increases the value of Q typewriters for other people. This first purchase may even cause an initial cascade of Q purchases. However, it has no effect on people's values for, or the price of, the Z typewriters. Therefore, at some point, people whose values for the Z typewriter exceed the price will purchase Z typewriters as well. This in turn induces others to buy them. In the long run, Z purchases are unaffected by the path of P purchases. So, although the outcome in any period may depend on the path, the long-run distribution over outcomes does not. The process is not path-dependent.

To see this in an example, assume that there are five people numbered from 1 to 5 and indexed by i. Let $V_{ix} = 2i + 2N_x$ for $x = Q, Z$. Assume that the price of any typewriter equals \$5. Initially, persons 3, 4, and 5 would buy either a Q or a Z typewriter, but the other two people would buy neither. Suppose that person 5 buys a Q typewriter. Now, person 2 would also be willing to buy a Q typewriter, even though person 2 would not be willing to buy a Z. If any two people buy Q's, then person 1 will also want to buy a Q, thus, eventually, everyone owns a Q. However, by the exact same logic, eventually everyone owns a Z as well.[26]

To see this another way, suppose that the two products are VHS machines and QWERTY typewriters. Each creates positive externalities within its own type but no externalities with the other type: buying a QWERTY typewriter creates no negative externality for a VHS machine, and buying a VHS machine creates no negative externality for a QWERTY typewriter. Therefore, the number of QWERTY typewriters and VHS machines sold does not depend at all on the path of purchases.

I next assume that there are only interpersonal negative externalities. To formalize this, I let $V_{iQ} = 2i - 2N_Z$ and $V_{iZ} = 2i - 2N_Q$. Now, when a person buys a Q typewriter, this lowers the value of Z typewriters for others. Consider two scenarios. In the first scenario, person 5 and person 4 both buy Q typewriters. Person 3's value for a Z typewriter falls from \$6 to \$2, so person 3 will now only buy a Q typewriter. Suppose person 3 buys a Q typewriter. Person 5's value for a Z typewriter now falls to \$4, and person 4's value falls to \$2, so no more sales of Z typewriters occur. In the second scenario, person 5 buys a Q typewriter, and person 4 buys a Z typewriter. Now, person 3 buys neither a Q or a Z typewriter owing to the negative externalities. She values each at \$4. Person 5, however, also purchases a Z typewriter. Person 4 does not purchase a Q typewriter, because two people own Z typewriters reducing person 4's value for a Q to \$4. Thus, these two different paths lead to different equilibria.

Thus, interpersonal positive externalities do not create path dependence, but interpersonal negative externalities do. None of this implies that positive externalities do not

---

[26] If each person is only offered one typewriter to buy, then path dependence would occur. If the first person buys a Q, then this raises the value of Q's relative to the Z's typewriters, and the result would be path-dependent. However, the assumption of a single purchasing opportunity is equivalent to a negative externality.

contribute to path dependence. They do. Positive interpersonal externalities reinforce the path dependence. If I include both types of externalities $V_{iQ} = 2i + 2N_Q - 2N_Z$ and $V_{iZ} = 2i + 2N_Z - 2N_Q$. Now, if persons 3 , 4, and 5 buy QWERTY typewriters, everyone buys a QWERTY and no-one buys a ZRJSOC. Thus, if both positive and negative externalities exist, then path dependence will be powerful and ubiquitous. The positive externalities create a bias in favor of some proposals and the negative externalities create a bias against others. Both forces operate in the QWERTY example, which is why there is such exaggerated path dependence. Yet, to refer to this as a case of increasing returns causing path dependence is misleading given that the negative externalities are the true cause and the positive externalities only exaggerate the phenomenon.

If negative externalities create history dependence, where do they come from? One answer is constraints. Temporal, financial, and spatial constraints all create negative externalities. Large public projects such as schools and road systems take up space, money, and time, creating negative externalities with future public projects. Projects that demand more of these limited resources may well have greater impact on the path of projects. Yet, although smaller projects may be less likely to influence the path of history on their own (their pecuniary externalities are much smaller) these less significant projects can accumulate over time and restrict history to certain paths. Limited power also creates constraints. A president has only so much political capital, and must decide where to spend it. Whatever choices are made may exclude other opportunities. Again, this can result in path or phat dependence. Finally, cognition imposes constraints. Our heads only have room for so much.

I should also point out that the externality framework as I have formulated it only admits phat dependence. The externalities created by any collection of approved proposals do not depend on the order in which the proposals were approved. For path dependence to exist, the order in which the proposals are approved must affect the existence or size of the externalities. To make this clear, I introduce some formal definitions.

A decision rule is **independent** *if it does not depend on the set of approved proposals. It is* **phat-dependent** *if it depends on the set of approved proposals, and it is* **path-dependent** *if it depends on the order in which proposals were approved.*

With these definitions, I can then make the following observation.

Observation 8 *If the magnitude or number of externalities associated with a proposal does not change over time, then a decision rule can at most be phat-dependent. It cannot be path-dependent.*

This observation implies that, to make the case for path dependence in environments with externalities, we need those externalities to accumulate or abate over time. In the case of institutional choices, one way in which the externalities could grow is if smaller complementary institutions arise that increase attachment to early institutional choices. (Ikenberry 2001). If the longer an institution has been around, the more it creates incentives for complementary institutions, and therefore the larger the effect of the institution, then the externalities associated with the original institution could increase over time. This alone would not imply path dependence for, as we have just seen, if all of the externalities are positive, then there would be no path dependence. Therefore, the smaller institutions that are complementary for some larger institutions must create negative externalities with other institutions for path dependence to occur.

Relatedly, culture has been proposed as a possible mechanism for institutional path dependence. By culture, here I mean many things – cognitive and behavioral repertoires, social networks, trust relationships and the like.[27] One can easily find evidence that institutional choices depend in some way on a society's culture (broadly defined) and that culture is deeply rooted in the past. But to make that argument formally requires a formal definition of culture as well as a mechanism for choosing institutions. Only then can we begin to understand how culture and institutions co-evolve and the extent to which that co-evolution produces path-dependent (see Bednar and Page 2005).

CONCLUDING THOUGHTS

In this essay I have provided some basic definitions of various types of path dependence, as well as examples that clarify those distinctions. Most importantly, I have highlighted the difference between path-dependent outcomes and path-dependent equilibria and between path and phat dependence. I also proved a lack of equivalence between path dependence and increasing likelihood of outcomes. One implication that we can draw from this last result is that the focus on increasing returns and positive externalities is misplaced. Scholars searching for evidence of path dependence should instead be looking for evidence of negative externalities as well. Why that has not happened is not surprising. Negative externalities are neither a new idea nor sexy. Social scientists have long been aware that budget, time, spatial, and power constraints restrict choices and actions. And, as my analysis of the QWERTY example shows, positive externalities exaggerate the degree of path dependence: the most compelling and extreme examples of path dependence probably include positive externalities as well as negative externalities. But extreme cases may be rare, evidenced by the attention given to typewriter keyboards. Finally, the existence of negative externalities need not imply path dependence either. They provide only a necessary but not a sufficient condition.

In concluding, I would be remiss if I did not at least touch upon the relationship between path dependence and suboptimality. Historical paths can lead to coordination on not only keyboard configurations but on more important conventions like languages and systems of weights and measures which, if viewed with the benefit of hindsight, appear suboptimal. List (2003) shows that a lack of internal consistency, a type of irrationality, can lead to path dependence.[28] Examples of coordination failure and List's results and others

---

[27] This role of culture can be found in Chong (2000), Grief (1994), North (1991) and Pierson (2004), among others.

[28] List's results are important given that policy decisions are made by collectives and not by individuals. While an individual's choices may be consistent, no such claims can be made for groups, especially

like them do not imply that an optimal decision maker cannot exhibit path-dependent actions. If a decision maker discounts the future, or makes choices under uncertainty, early choices may determine or restrict later choices. What appears suboptimal ex post need not have been suboptimal ex ante.

I might add that I intend for this essay to be read as introductory, not definitive. I offer up these primitive ball and urn models not as necessarily realistic or testable models of political and economic phenomena, but to clarify our thinking. My hope is that they will help to move us in the direction of richer, more detailed models that we can then use to construct a science of the various ways that histories unfold.

APPENDIX

Proof of the observation that negative externalities are necessary for path-dependent decision rules.

By assumption there exists a period s such that the optimal acceptance set at time s, X_s^*, is not a subset of the optimal acceptance set at time $s+1$, X_{s+1}^*. Let $V(X)$ denote the value of the subset of proposals X. Given that there is a unique optimal acceptance set, it follows that

$$V(X_s^*) - V(X_s \cap X_{s+1}^*) > 0$$

Suppose that all externalities are positive. It follows that for any subset of proposals denoted by A

$$V(X_s^* \cup A) - V((X_s^* \cap X_{s+1}^*) \cup A) > 0$$

This must be true for A equals X_{s+1}^* Therefore,

$$V(X_s \cup X_{s+1}^*) - V(X_{s+1}^*) > 0$$

which contradicts the optimality of X_{s+1}^*. ■

REFERENCES

Arthur, Brian. 1994. *Increasing Returns and Path Dependence in the Economy.* Ann Arber: University of Michigan Press.

Bednar, Jenna, and Scott Page. 2005. "A Model of Institutional Path Dependence." Presented at the annual meetings of the American Political Science Association, Washington, D.C.

Bednar, Jenna, and Scott Page. 2006 "Can Game(s) Theory Explain Culture?" *Rationality and Society*. Forthcoming. Available as Santa Fe Institute Working Paper #04-12-039.

Bikhchandani, Sunil, David Hirshleifer, and Ian Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100(5): 992–1026.

Bikhchandani, Sunil, David Hirshleifer, and Ian Welch. 1998. "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives* 12(3): 151–70.

groups with changing compositions. Thus, even forward-looking political bodies can make path-dependent decisions in sequential contexts.

Chong, Dennis. 2000. *Rational Lives: Norms and Values in Politics and Society.* Chicago: University of Chicago Press.

Cowan, Robin, and Philip Gunby. 1996. "Sprayed to Death: Path Dependence, Lock-in and Pest Control Strategies." *Economic Journal* 106(436): 521–42.

Crouch, Colin, and Henry Farrell. 2004. "Breaking the Path of Institutional Development? Alternatives to the New Determinism." *Rationality and Society* 16(1): 5–43.

David, Paul. 1985. "Clio and the Economics of QWERTY." *American Economic Review, 75(2), Papers and Proceedings of the Ninety-Seventh Annual Meeting of the American Economic Association*, pp. 332–37.

Easterly, W. 2001. *The Elusive Quest for Growth: Economists' Adventures and Misadventures in the Tropics.* Cambridge, MA: MIT Press.

Ferejohn, John. 1991. "Rationality and Interpretation: Parliamentary Elections in Early Stuart England," in *The Economic Approach to Politics: A Critical Reassessment of the Theory of Rational Action*, ed. K.R. Monroe. New York: HarperCollins.

Gaddis, John L. 2002. *The Landscape of History: How Historians Map the Past*. Oxford University Press.

Gerschenkron, Alexander. 1952. "Economic Backwardness in Historical Perspective." in *The Sociology of Economic Life*, ed. Mark Granovetter and Richard Swedberg. Boulder, CO: Westview Press.

Gilovich, T, R. Vallone, and A. Tversky. 1985. "The Hot Hand in Basketball: On the Misperception of Random Sequences." *Cognitive Psychology* 17: 295–314.

Grief, Avner. 1994. "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies." *Journal of Political Economy* 102(5): 912–50.

Hacker, Jacob. 2002. *The Divided Welfare State: The Battle over Public and Private Social Benefits in the United States*. Cambridge: Cambridge University Press.

Hathaway, Oona A. 2001. "Path Dependence in the Law: The Course and Pattern of Legal Change in a Common Law System." *The lowa Law Review*, 86(2): January 2001.

Ikenberry, G. J. 2001. *After Victory: Institutions, Strategic Restraint, and the Rebuilding of Order After Major Wars*. Princeton, NJ: Princeton University Press.

Karl, Terry. 1997. *The Paradox of Plenty: Oil Booms and Petro-States*. Berkeley: University of California Press.

Kollman, Ken. 2003. "The Rotating Presidency of the European Council as a Search for Good Policies." *European Union Politics* 4 (Mar): 51-71.

Lee In Ho. 1993. "On the Convergence of Informational Cascades." *Journal of Economic Theory* 61(2): 395–411.

Liebowitz, Stan J., and Stephen E. Margolis. 1990. "The Fable of the Keys," *Journal of Law and Economics* 33: 1–27.

Liebowitz, Stan J., and Stephen E. Margolis. 2002. *The Economics of Qwerty: Papers by Stan Liebowitz and Stephen Margolis*, ed. Peter Lewin. MacMillan NYU Press.

List, Christian. 2003. "A Model of Path Dependence in Decisions Over Multiple Propositions." Working paper, London School of Economics.

North, Douglas C. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.

North, Douglas C. 1991. *Journal of Economic Perspective* 5(1): 97–112.

Page, Scott. 1997. "An Appending Efficient Algorithm for Allocating Public Projects with Complementarities," *Journal of Public Economics* 64(3): 291–322.

Page, Scott. 1998. "On the Emergence of Cities." *Journal of Urban Economics* 45: 184–208.

Page, Scott. 2006. "Type Interactions and the Rule of Six." *Economic Theory*, forthcoming.

Pierson, Paul. 2000. "Path Dependence, Increasing Returns, and the Study of Politics," *American Political Science Review* 94(2): 251–67.

Pierson, Paul. 2004. *Politics in Time: History, Institutions, and Social Analysis*. NJ: Princeton, Princeton University Press.

Shefter, Martin. 1977. "Party Patronage: Germany, England, and Italy." *Politics and Society* 7: 403–52.

Van Evera, S. 1998. "Offense, Defense, and the Causes of War." *International Security* 22: 5–43.

Weber, S. 1997. "Institutions and Change." In *New Thinking in International Relations Theory*, eds. M. W. Doyle and G. J. Ikenberry. Boulder, CO: Westview Press.