

THE VIRTUES OF HESITATION[‡]

URMEE KHAN AND MAXWELL B. STINCHCOMBE

ABSTRACT. In many economic, political and social situations, circumstances change at random points in time, reacting is costly, and reactions appropriate to present circumstances may become inappropriate upon future changes, requiring further costly reaction. Waiting is informative if arrival of the next change has non-constant hazard rate. We identify two classes of situations: in the first, delayed reaction is optimal only when the hazard rate of further changes is decreasing; in the second, it is optimal only when the hazard rate of further changes is increasing. These results in semi-Markovian decision theory provide motivations for building delay into decision systems.

Key Words and Phrases: Nonconstant hazard rates; renewal processes; non-stationary semi-Markovian decision theory; optimal timing and delay.

A handful of patience is worth more than a bushel of brains. - Dutch Proverb

Patience has its limits. Take it too far, and it's cowardice. - George Jackson

The essential ingredient of politics is timing. - Pierre Trudeau

Timing is everything - Attributed to various authors

1. INTRODUCTION

In many social, economic and political situations, there is a stochastic environment that changes at random points in time and responding to these changes entails significant costs. Given that a new state may give way to yet another, potentially making today's optimal action again obsolete, and given that actions

Date: September, 2014

[‡]Many thanks to Martin Dumav and Zahid Arefin Chowdhury for their timely and much appreciated help, to seminar participants at UC Riverside, Delhi School of Economics and University of Bielefeld and conference participants at the IHS Summer Research Conference (2010), Midwest Political Science Association Conference (2010), Midwest Theory Conference (2011), GAMES (2012) and SSCW (2012) for many helpful comments and discussions. We are very grateful to Larry Samuelson and three anonymous referees for very detailed and insightful feedback that has helped us improve the paper. All remaining errors are ours. The first author gratefully acknowledges the financial support of a Summer Research Fellowship from the Institute of Humane Studies.

are costly, the question is whether to respond to a change in the environment or to delay.

Variants of this problem have been extensively analyzed in economics (for example Boyarchenko and Levendorskiĭ [2007], Stokey [2009] and the references therein). However, a crucial aspect of most existing analyses is that the passage of time by itself does not reveal any information. By contrast, we study problems in which the passage of time without a change contains information about the arrival time of the next change, problems in which there may be value to delaying decisions beyond the usual option value of waiting.¹ We begin with examples where the length of time which a change has survived may be of crucial importance to its future longevity.

1.1. Political Change. Political processes in a democratic system are driven by ‘political issues’ and the configuration of opinions and attitudes of the polity on these issues. Such configurations are hardly, if ever, static. There are slow and gradual changes that take place side by side with rapid and explosive changes. Some changes are long-lasting, some short-lived. As Carmines and Stimson [1990] say:

... we shall see that issues, like species, can evolve to fit new niches as old ones disappear. But, unless they evolve to new forms, all issues are temporary. Most vanish at their birth. Some have the same duration as the wars, recessions, and scandals that created them. Some become highly associated with other similar issues or with the party system and thereby lose their independent impact. And some last so long as to reconstruct the political system that produced them

The Vietnam War and the Watergate scandal seem to have very little traces left today either in public attitude or legislative response to the issues of war and executive power respectively, but they were the biggest issues of their day. On the other hand, the Civil Rights Movement and its aftermath marked a fundamental realignment in US politics. In general, some ideas and opinions “wear out their welcome” after a time, perhaps through changes in the conditions that gave rise to them, perhaps by the accumulation of counterarguments to their veracity. Hence,

¹See McGuire and Kable [2012, 2013] for experimental evidence on how hazard rate patterns affect delay of gratification behavior.

the likelihood that such an idea would become irrelevant increases with time. By contrast, some types of issues or opinions tend to get more entrenched the longer they live. Political actors in various capacities try to cope and make decisions in the face of such ‘issue evolution’ [Carmines and Stimson, 1990]. Legislatures choose whether or not to change a law, Supreme Courts decides whether or not to re-interpret or overturn past precedents, political parties decide whether or not to realign politically and redefine the agenda. Oftentimes, the most crucial ingredient in such decisions is the aspect of timing.

Each of these decisions entails some fixed cost either to the society at large or the actor herself.² One would have to trade off the immediate gains with substantial future losses if the initial change that triggered the costly action turns to be relatively short-lived.

Another class of political decisions have to do with effective but costly response to a possibly impending crisis. Again, timing is of the essence, but in a very different way. When it appears that a costly prevention or crisis-averting action might need to be taken, the crux of an executive decision then becomes how to delay the cost as much as possible, without being too late to actually prevent, or at least lower the likelihood or the cost of, the crisis. Impending wars, international conflicts, painful and unpopular reforms or a system overhaul, all of these are examples.

The following two subsections discuss examples drawn from arenas of political decision making where the need for deliberation and caution interplay with the need for quick and decisive action, and one of the challenges of institution design turns out to be the demarcation of the different spheres where one should be emphasized at the expense of the other.

1.1.1. *Constitutional Amendments.* Constitutions establish the fundamental legal structures of a society. They are meta-institutions through which institutions are introduced, reformed and interpreted [Ostrom, 1990]. A constitution and the legal order it creates must have the support of, or at least tacit approval of, the governed to have legitimacy. Maintaining the legitimacy and relevance of a constitution requires a certain degree of adaptability or flexibility to change

²In case of legislative changes and court decisions, the citizens have to re-adjust and re-optimize with respect to the new rules. In case of a political party, realignment may mean losing a traditional support base and adjusting to a new one.

because technology, environment and public opinion are forever changing. On the other hand, the basic value of a constitution lies in its stability because it coordinates the actions and expectations of people and reduces the uncertainty in the environment [Hardin, 2003]. Hence the basic tradeoff between ‘commitment’ and ‘flexibility’ lies at the heart of the constitution design problem, as encapsulated in the famous exchange between Thomas Jefferson and James Madison in Smith [1995].³ It is also costly to change the constitution because it acts as a coordination device for peoples’ behavior, and changes are likely to impose large adjustment costs on significant parts of the population and disrupt ancillary institutions that grow around the constitution [Hardin, 2003].

From these perspectives, it is reasonable to presume that an optimal rule for constitutional change should be more sensitive to long-lasting changes than to transitory changes. It is clear that waiting longer will help answer whether a change will have a longer or shorter total life, but what matters for decisions is the longer or shorter *future* life of the change. One tradeoff is between costly unneeded or ultimately unwanted changes (e.g. Prohibition) and undermining the legitimacy of the constitutional regime by ignoring new realities. It is from this perspective that we study the general question of why some changes in laws should be more difficult to implement, and what this should depend on.

Under study is a class of explanations that we regard as complementary to the many previously offered ones, a class of explanations based on the observation that the persistence of changes in sentiment have predictive power for the future length of time the changes will last. For us the question becomes “How much longer should one wait before acting?” The dynamically consistent answer depends both on the costliness of the action and the costliness of its reversal.

The US constitution has had four different amendments that have extended voting rights to different parts of the population: Amendment XV (1870), which was passed at the end of the Civil War, extended suffrage to men independent of race or previous condition of servitude; XIX (1920) extended suffrage to women; XXIV (1964) made poll taxes illegal; and XXVI (1971) extended suffrage to those eighteen years of age or older. These formalized long-lived widely-shared changes

³For interesting empirical evidence to the effect that flexibility actually helps the sustainability of constitutions, see Elkins et al. [2009], for an analysis of the ways in which flexibility must be structured for the functioning and survival of formal systems that must adjust to changing circumstances, see Stinchcombe [2001].

in sentiment, but Amendment XVIII (1919), Prohibition, was an expensive and short-lived failure, being repealed fourteen years later by Amendment XXI (1933) (Amar [2006] is a detailed history of the amendments to the U.S. Constitution).

If one dates the beginning of the women's suffrage movement to the 1848 Seneca Falls Convention, it took 72 years for Amendment XIX to pass.⁴ At various points in the political process, there was evidence that the recognition of women's rights to vote would be long-lasting: the passage of suffrage at the state level in western states by the early 20'th century;⁵ the nation's westward expansion and the Civil War leading to an expanded need for women both in industrial settings and as teachers; the gradual increase in the numbers of college educated and professional women; unionization movements among female professions in the late 1800's and early 1900's. Even after one could perhaps clearly see that general sentiment had shifted in favor of the Nineteenth Amendment, there was (much) further delay in implementing what turns out to have been a long-lasting change in sentiment. This is perhaps consistent with unwillingness to believe that so drastic a change could be long-lasting, but our analysis suggests that the relatively high cost of depriving a class of people of liberties, once granted, exaggerates the effects of such doubts.

By contrast, Prohibition proved to be very costly to society, and was short-lived, repealed fourteen years later.⁶ The Temperance Movement had as long a history as the women's suffrage movement, and was used by some women's suffrage organizers as an occasion to teach women the necessity of having a voice in politics in order to achieve changes [Flexner, 1996]. The implementation of Prohibition very quickly gave rise to a host of social problems accompanied by a marked increase in crimes associated with illegal selling and consumption of alcohol. Public opinion quickly turned against it. From our point of view, this was a change of action that led to a change in the distribution of the time until general sentiment was reversed.

⁴Flexner [1996] make an argument for an earlier date for the beginning of the movement, emphasizing the experience of female abolitionists and agitators for women's education in the early 1800's as the roots of the suffrage movement.

⁵By 1915, Arizona, California, Colorado, Idaho, Illinois, Kansas, Montana, Nevada, Oregon, Utah, Washington, and Wyoming had granted full women's suffrage, and several other states or municipalities had granted suffrage in primary elections.

⁶Prohibition was repealed by Amendment XXI, the only amendment passed by state ratifying conventions rather than by the votes of state legislatures, sparing the legislators from re-election campaigns against a dedicated cadre of single-issue voters.

1.1.2. *Executive Power.* The US constitution has separation of powers between the executive and the legislative branches. While an extensive discussion of the virtues of such an approach, as opposed to one in which the two branches coincide, as in the United Kingdom, is outside our purview, one aspect of this separation is relevant to our motivating question here. Congress is the deliberative body, and congressional procedures are markedly more inclined towards caution and restraint against hasty decisions, in comparison with the structure of executive power, which enables and empowers the president to act in a quick and decisive fashion if need arises. Our observation is that it is one way to institutionally demarcate between decision problems where built-in delay could be desirable (making legislative actions more attuned to longer lasting changes), and those where not being too-late-to-be-of-any-good is the real concern (executive decisions, where delay, if at all, should come only out of prudence, not structural limitations).

We now turn to other settings in which our general methodology could be fruitfully applied.

1.2. **Marketing Strategy.** Research in consumer behavior has shown that when and how consumers switch brands depend on the last purchased brand and time since the last purchase. The inter-purchase time may exhibit increasing or decreasing hazard rates depending on consumer characteristics like “inertia” or “variety seeking tastes” [Chintagunta, 1998]. It has been suggested that optimal timing of targeting consumers for marketing should depend on such considerations instead of the traditional demographic variables [Gonul and Ter Hofstede, 2006, gives an empirical approach to optimal timing for catalog mailing]. The class of optimization models under study here are directly applicable to such situations.

1.3. **Labor Search.** One of the issues relating to long-term unemployment is depletion of human capital, which might make a candidate less and less attractive to the potential employers as the duration of unemployment gets longer and longer. This factor would have important implications for standard labor search models, as the value of the future discounted wage and thus the reservation wage for a job seeker would be affected. Ortego-Marti [2012] uses this insight to explain observed wage dispersion in the labor market. Also, the possibility that a long period searching without finding a job is taken, by potential employers,

as an indication that there is something wrong with the person, would mean a decreasing rate of arrival of a job of any given quality. This is independent of decay of human capital, and pushes the acceptance threshold downwards. One of the key ingredients of the standard Diamond, Mortensen, Pissarides (DMP) model is a stochastic description of labor turnover, along with a model of labor-market tightness, and a bargaining model of wage determination (Diamond [2011], Mortensen [2011], Pissarides [2011], Hall [2012] and the references therein). Our model results would be relevant for an attempt to capture non-stationarities in the first component. Recent empirical work based on the DMP model, (e.g. Shimer [2005]), highlights various aspects of the wage determination phenomenon that are not well explained by the stationary DMP models so far. It would be interesting to introduce non-stationarities in the job-arrival process and work out the implications. Our framework provides a minimal way of attempting that.

1.4. Currency Unions, International Treaties. The benefits of joining currency unions or international treaties are variable over time and leaving a currency union or abrogating a treaty is a costly decision. These two factors combine to provide the context for countries trying to devise strategies for entry, exit and crisis management. The recent crisis in the European Monetary Union (EMU) has rekindled the discussion around the impacts of countries exiting currency unions. Opinion is sharply divided, although the majority opinion seems to be that breaking up of the Euro will be disastrous [Eichengreen, 2010]. On the other hand, historically, there have been regular instances of countries leaving currency unions [Rose, 2012]. From the design perspective, one of the interesting features of the EMU is that there are no exit clauses. Omission of a well-defined exit clause, at the least, significantly raises the procedural cost of exiting the union, and thus can be seen as a mechanism to avoid hasty reaction by individual countries that could potentially threaten all the other members as a group. Overall, the European currency crisis and the responses to it by various parties highlight the importance of the main facets of our model: the trade-off between reacting quickly and flexibly on one hand, and the need to avoid precipitating a fast-moving crisis by unnecessarily hasty actions on the other.

1.5. Outline of the Paper. The next section contains three simple examples that give a sense of what is involved in the more general analyses that follow. The essential aspects of the model include: a starting state and action, i_0 and

a_0 ; random times Y_{k+1} at which the state changes from I_k to I_{k+1} according to a partly controlled, imbedded Markov process; and the option to engage in costly changes in actions during the inter-arrival stochastic intervals, that is, between Y_k and Y_{k+1} .

The within-interval maximization problems will be central to the dynamic analysis. The first two examples in the next section concern problems with only one interval, how long to search and how long to delay precautionary measures. With the intuitions about how the optimality conditions work for a single interval in place, the last example demonstrates how the value functions for an entire dynamic problem interact with the maximization problems within intervals.

The general model, existence of optima, and their recursive characterization through the value function are in the following section. The penultimate section develops the corresponding first order conditions (Euler equations) for a broad range of problems and discusses the implications of the results in the context of our examples. The last section concludes.

2. THREE EXAMPLES

Three non-stationary problems demonstrate the essential features of the optimization problems under study and of their solutions: the optimal delay before abandoning search involves searching until the hazard rate for search success is low enough and falling; the optimal hesitation before implementing expensive precautionary measures involves waiting until the hazard rate is high enough and increasing; and optimal adaptations to changing circumstances involves solving a sequence of problems much like the first two. We use the first two examples in analogy with the two classes of decision problems that we described in our discussion of political change.

We assume throughout that waiting times, W , have continuous densities, f , on $(0, \infty)$, and, possibly, an atom at ∞ . If W has an atom at ∞ , it is called an **incomplete** distribution, corresponding e.g. in search problems, to the object of search not existing or not being findable. Associated with W is the cdf, $F(t) = \int_0^t f(x) dx$, with $\lim_{t \uparrow \infty} F(t) < 1$ for incomplete distributions, and the hazard rate, $h(t) := \frac{f(t)}{1-F(t)}$.

The first and second order conditions (FOCs and SOCs) have common features: the net benefits to delaying from t to $t+dt$ take the form $(a+b \cdot h(t))dt$; solving the

FOCs requires $h(t_1^*) = -a/b$; this in turn requires that a and b have opposite signs. When a is negative, the costs of delay are certain and the benefits stochastic, when a is positive, this is reversed, there is a sure benefit to delay and a stochastic cost. Further, as the net benefits must be positive before t_1^* and negative afterwards, the SOC's require $h'(t_1^*) < 0$ if $a < 0$ and $b > 0$, and $h'(t_1^*) > 0$ if $a > 0$ and $b < 0$.

2.1. Hesitating to Cut Off Search. At a flow cost of $c \geq 0$, one can keep searching for a source of higher profits (a low cost source of a crucial input, a process breakthrough, a new product). If found, expected net flow profits of $\bar{\pi}$ result. If one abandons the search, the decision is, by assumption, irreversible, and the known alternative yields expected net flow profits of $\underline{\pi}$, $\bar{\pi} > \underline{\pi} > 0$.

The non-stationary choice problem is “At what time, t_1 , does one stop searching and accept the lower $\underline{\pi}$?” The results we give here are special cases of Theorem 2 (below), but we give an intuitive development here.

FOCs. The net benefits of waiting an extra instant dt at a pre- W time t_1 are $(a + b \cdot h(t_1))dt$ where: $a = -(c + \underline{\pi})$ incorporates the flow cost, c , and $\underline{\pi}$ is forgone flow value of the known alternative; and delay gives a probability $h(t_1)dt$ of moving the flow payoffs upward, which is worth, in present value terms, $b = \frac{1}{r}(\bar{\pi} - \underline{\pi})$. At an interior optimum, $0 < t_1^* < \infty$, the FOC's are therefore $h(t_1^*) = r(c + \underline{\pi})/(\bar{\pi} - \underline{\pi})$.

SOC's. In order for the solution just given to be a local maximum rather than a local minimum, the net benefits of waiting, $\frac{1}{r}(\bar{\pi} - \underline{\pi}) \cdot h(t_1) - (c + \underline{\pi})$, must be positive for $t_1 < t_1^*$ and negative for $t > t_1^*$. As $(\bar{\pi} - \underline{\pi}) > 0$, for this to be true, the hazard rate must be decreasing at t_1^* , $h'(t_1^*) < 0$.

Comparative statics. Because the hazard rate is decreasing, the optimal t_1^* is: higher for higher $\bar{\pi}$, it is worth searching longer when the reward is larger; lower for higher c , one searches less if searching is more costly; lower for higher r , one searches less if one is more impatient; lower for higher $\underline{\pi}$, one searches less when the fallback option is better; and higher for uniform upward shifts in $h(\cdot)$, one searches more if search is more productive.⁷

2.2. Hesitating to Inoculate. Present utility flows are $\bar{u} > 0$. At issue is the optimal timing of a costly inoculation or precautionary measure: an evacuation before a hurricane landfall; or a politically painful reform of a banking system

⁷One can also arrive at the FOC's and SOC's by explicitly writing the integrals giving the expected payoffs and differentiating. Details of the calculations for all three models in this section can be found in the online supplement.

before the next financial crisis. If the measure is *not* in place before W , then \bar{u} will be reduced to a utility flow normalized to 0. If the measure is in place, W will reduce utility flows to \underline{u} , with $\bar{u} > \underline{u} > 0$. As well as reducing the damages in case of crisis, inoculation at time t_1 reduces the probability of a crisis by changing the arrival density from $f(w)$ to

$$f_{\theta}(w; t_1) = \begin{cases} f(w) & \text{if } w < t_1 \\ (1 - \theta)f(w) & \text{if } w \geq t_1. \end{cases} \quad (1)$$

After the inoculation, $\theta \int_{t_1}^{\infty} f(w) dw$ of the probability shifts to the mass point at infinity, that is, to the disaster *not* arriving so that $\theta \in [0, 1]$ measures the efficacy of the inoculation.⁸ The non-stationary timing question is, “How long should one hesitate before incurring the cost of inoculation?”

FOCs. The net benefits of waiting an extra instant dt at a pre- W time t are $rCdt - (\frac{1}{r}[\theta\bar{u} + (1 - \theta)\underline{u}] - C)h(t)dt$ because rC is the annuitized flow value of the fixed cost C , and if W arrives in the instant $(t, t + dt)$, which happens with probability $h(t)dt$, one loses the expected net benefit of inoculation, $(\frac{1}{r}[\theta\bar{u} + (1 - \theta)\underline{u}] - C)$. This gives the FOCs $h(t_1^*) = rC / (\frac{1}{r}[\theta\bar{u} + (1 - \theta)\underline{u}] - C)$.

SOCs. Since the decision problem involves postponing the cost C as long as this outweighs the hazard of losing the benefits of inoculation, the SOC requires $h'(t_1^*) > 0$. One should wait until W is both imminent enough and becoming more so.

Comparative statics. Because the hazard rate is increasing at an optimum, increases in the cost, C , or in the rate at which the future is discounted, r , make a longer delay optimal, and one sees the same effect if the inoculation is less effective, that is, if $(\frac{1}{r}[\theta\bar{u} - (1 - \theta)\underline{u}] - C)$ is smaller.

2.3. Hazard Rate Patterns. It is useful to examine the interactions of the first and second order conditions with several non-constant hazard rate patterns.

2.3.1. Decreasing and Increasing Hazard Rates. For search problems, a decreasing hazard rate corresponds to a problem in which one searches in places in decreasing order of likelihood and/or the likelihood of the object never being found is

⁸If $\theta = 1$, then the precautionary measure precludes the arrival of the next crisis, e.g. an inoculation with a 100% success rate. If $\theta = 0$, then the precautionary measure has no effect on the likelihood of the disaster, e.g. if W is the landfall of a hurricane and the precautionary measure is an evacuation.

increasing over time, for inoculation timing problems, an everywhere decreasing hazard rate corresponds to being past the worst.⁹ For search problems with an everywhere decreasing hazard rate, satisfying the FOCs at t_1^* gives the global payoff maximum, and conditional on having already searched until any $T \geq t_1^*$, it is optimal to stop immediately. By contrast, for inoculation problems, satisfying the FOCs gives the global payoff *minimum*, the optimum is either at $t_1^* = 0$, immediate inoculation, or $t_1^* = \infty$, never inoculating. Further, if immediate inoculation is optimal, then conditional on not having inoculated yet, immediate inoculation is optimal until the time $T^\circ > 0$ at which one is indifferent between inoculating now and never inoculating.

These patterns are reversed for everywhere increasing hazard rates (e.g. Weibulls with parameters $\lambda > 0$ and $0 < \gamma < 1$). For inoculation problems, satisfying the FOCs at t_1^* gives the global payoff maximum and if inoculation has been delayed to any $T \geq t_1^*$, it is optimal to inoculate immediately. For search problems, satisfying the FOCs at t_1^* gives the global payoff minimum, the optimum is either at $t_1^* = 0$ or at $t_1^* = \infty$, if at $t_1^* = 0$, then immediate abandonment of search is optimal until the time $T^\circ > 0$ at which one is indifferent between abandoning search and continuing to search forever.

2.3.2. Some Mixed Cases. In quality control, one often sees a “bathtub shaped” hazard rate pattern: early failures are due to manufacturing flaws and tail off; late failures are due to systems wearing out, and increase over time.¹⁰ For search problems, a bathtub hazard rate pattern corresponds to a search process that first checks the likely and easy to check locations, and if the object is not found, then the remaining locations, however unlikely or difficult to check, must be where the object is located. If the first order conditions are first met at t'_1 at which the hazard rate is decreasing, then t'_1 is a local maximum. If the FOCs are not satisfied at a later time, then t'_1 must also be the global maximum. If the FOC are satisfied at a (necessarily single) later time, this is a local minimum, and t'_1 must be compared to $t_1 = \infty$ — early abandonment of search may be optimal, or unbounded search may be optimal.

⁹A tractable class of distributions with an everywhere decreasing hazard rate are the Weibulls with parameters $\lambda > 0, \gamma > 1$, i.e. $W = X^\gamma$ where X is a negative exponential with parameter λ . The hazard rate is $h(t) = \lambda t^{\frac{1-\gamma}{\gamma}} / \gamma$.

¹⁰This is often used as a quality control screen — overload elements for long enough that the first kinds of failures are weeded out [e.g. Rausand and Høyland, 2004, Ch. 2].

For search problems, a “hump-shaped” hazard rate pattern corresponds to a search process that needs time to begin to function well, but once the likely and/or easy-to-search locations have been exhausted, it becomes clearer and clearer that the object was not there in the first place. If the first order conditions are met while the hazard rate is increasing, this is a local minimum, so either $t_1^* = 0$ is optimal or some later time will be optimal. If a later time, t'_1 , corresponds to satisfying the FOCs a second (and necessarily last) time, then t'_1 is a local maximum, and if any search is optimal, it has a limited duration.

For inoculation problems, the hump-shaped hazard rate pattern corresponds to facing a gathering storm that may, or may not, yield a disaster, but if one is lucky enough to get through the most hazardous times without a disaster, the storm passes and the likelihood diminishes. If the first order conditions are met at a t'_1 where the hazard rate is increasing, then t'_1 gives the *locally* optimal tradeoff between delay of expensive preparations and imminent danger. Further, the next time the FOCs are satisfied will be a local minimum so that either $t_1^* = t'_1$ or it is $t_1^* = \infty$ is optimal — evacuate if the danger is imminent enough, otherwise try to ride it out and hope for the best.

2.4. Optimal Adaptation to Circumstances. We now study a simple model of the optimal timing of adaptations to a stochastic dynamic state. Solving the entire problem involves solving a sequence of single interval problems.

2.4.1. States and Actions. The first ingredients in the model are a state space, $S = \{i, j\}$, an action space, $A = \{a, b\}$, and a set of flow utilities to be discounted and integrated: for the ‘matched’ state action pairs (i, a) and (j, b) , we set $u(i, a) = 1$ and $u(j, b) = 2$; for the ‘mis-matched’ pairs, we set $u(i, b) = u(j, a) = 0$.

2.4.2. Change Costs. In our general model, change-of-action costs depend on the present action, the action one changes to in the set of available actions, and the present state. For the present model, there are four such costs. For simplicity here, we assume one can switch from a to b or back in either state, and all four switching costs are equal to $C > 0$.

2.4.3. Change Times. The states will change at random times. It is the costs of switching actions that may make instantaneous adjustment of the action to the state suboptimal — if another change in the state is expected soon, it may not

be worth incurring the cost to enjoy the extra flow. The non-stationary timing question is, “How long to hesitate before reacting?”

The state change times are random variables Y_k , $k = 0, 1, \dots$ with $Y_0 \equiv 0$, $Y_k \in (0, \infty]$ for $k \geq 1$, and $P(Y_k < Y_{k+1}) = 1$ if $Y_k < \infty$. With I_k denoting the state of the system for the k 'th state change time, starting state is i_0 , and the future states I_k , $k \geq 1$. The waiting time for the k 'th change is defined as $W_k := Y_k - Y_{k-1}$, and $P(W_k = \infty) > 0$ corresponds to the k 'th waiting time having an incomplete distribution and there being a positive probability that there is never a transition away from the state I_{k-1} . A policy is semi-stationary if it depends only on the initial state-action pair and the within-interval time, and we will see that such policies are optimal. For any such policy, the random sequence $(i_0, W_1), (I_1, W_2), \dots$ is a Markov process taking values in $S \times (0, \infty]$ that ‘ends’ at the first k , if any, at which $W_k = \infty$.

In the general model, the distribution of each W_k will depend on the state-action pair that is in effect at the start of each interval and the time path of action changes that occur within the interval. In this simpler version of the model, the distribution of the W_k depends only on the state-action pair, (i, a) , that prevails at Y_{k-1} , and the time, $Y_{k-1} + t_1$, $t_1 \in [0, \infty]$, at which the action is *first* changed. Because at most one change of action can be optimal and there is only one possible action to change to, a plan for an interval, denoted \mathbf{p} , can be specified by a single change time, $t_1 \in [0, \infty]$. Further, the stochastic structure can be summarized in a class of possibly incomplete probability density functions (pdfs) on $(0, \infty)$, $f_{i_0, a_0}(\cdot; t_1)_{t_1 \in [0, \infty]}$, $(i_0, a_0) \in S \times A$. Because changes of actions at t_1 can affect the future but not the past, we assume that for all $s \leq t_1 < t'_1 \leq \infty$ and all (i_0, a_0) , $f_{i_0, a_0}(s; t_1) = f_{i_0, a_0}(s; t'_1)$. To put it another way, the single pdf $f_{i_0, a_0}(\cdot; \infty)$ governs arrival the time of the state change up until a change of action, but the distribution of what happens after that change of action can depend on the time at which it occurred.

2.4.4. Expected Utility. All random variables are defined on a probability space (Ω, \mathcal{F}, P) . A **stochastic interval** is a subset of $\Omega \times [0, \infty)$ of the form $\llbracket Y_k, Y_{k+1} \llbracket = \llbracket Y_k, Y_k + W_k \llbracket := \{(\omega, t) : Y_k(\omega) \leq t < Y_{k+1}(\omega)\}$. If the interval starts with the action being a_k , we specify a plan to change from a_k to a'_k at a *within-interval time* t_1 using a *left continuous* $t \mapsto \mathbf{p}(t)$ with $\mathbf{p}(t) = a_k$ for all $t \leq t_1$ and $\mathbf{p}(t) = a'_k$

for all $t > t_1$. This means that $\mathbf{p}(Y_{k+1} - Y_k)$ is the action being used when the new interval, $\llbracket Y_{k+1}, Y_{k+2} \rrbracket$ starts.

For an interval $\llbracket Y_k, Y_{k+1} \rrbracket$ starting in the state-action pair (i_k, a_k) and a plan \mathbf{p}_k , the expected payoffs during the interval are defined as

$$\begin{aligned} & \int_0^\infty \left[\int_{Y_k}^{Y_k+w} u(i_k, \mathbf{p}(t)) e^{-rt} dt - C \sum_{\tau \in T_{\mathbf{p}}} 1_{(\tau, \infty)}(w) e^{-r\tau} \right] f_{i_k, a_k}(w; t_1) dw \\ &= \mathbb{E} \left[\int_{Y_k}^{Y_{k+1}} u(i_k, \mathbf{p}(t)) e^{-rt} dt - C \sum_{\tau \in T_{\mathbf{p}}} 1_{\llbracket Y_k, Y_{k+1} \rrbracket}(\tau) e^{-r\tau} \right] \end{aligned} \quad (2)$$

where $T_{\mathbf{p}}$ is the set of planned action change times and \mathbb{E} is the expectation operator given by the distribution of $Y_{k+1} = Y_k + W_k$, which in turn depends on \mathbf{p}_k . If we condition on $Y_k = t$, subtract t from Y_k , we denote the corresponding sum in (2) as $\mathcal{R}(\mathbf{p}_k; (i_k, a_k))$. This means that the total reward to starting in (i_0, a_0) and using a sequence of plans, \mathbf{p}_k , $k = 0, \dots$ is $\sum_{k \geq 0} \mathbb{E} e^{-rY_k} \mathcal{R}(\mathbf{p}_k; (I_k, \mathbf{p}_k(Y_{k+1} - Y_k)))$. In principle, \mathbf{p}_k can depend on any aspect of the history up to time Y_k , but we shall see that an optimum exists among the *semi-stationary* policies, i.e. those with \mathbf{p}_k depending only on the initial state-action pair at each Y_k .

2.4.5. The Value Function. The **value function**, $V^*(i_0, a_0)$, gives the maximal expected discounted utility to starting at $Y_0 = 0$ in state $i_0 \in S$ with the present action being a_0 , that is,

$$V^*(i, a) = \sup_{\mathbf{p}_0, \mathbf{p}_1, \dots} \mathbb{E} \sum_{k \geq 0} \mathcal{R}(\mathbf{p}_k; (I_k, \mathbf{p}_k(Y_{k+1} - Y_k))) e^{-rY_k}. \quad (3)$$

From Theorem 1 (below), $V^*(\cdot, \cdot)$ is well-defined, can be found as the fixed point to a contraction mapping, and given the value function, the optimal semi-stationary policy can be found by sequentially solving the within-interval optimization problems on the presumption that leaving the interval into the state-action pair (I_{k+1}, a_k) at time Y_{k+1} yields a payoff of $V^*(I_{k+1}, a_k) e^{-rY_{k+1}}$.

In this specific problem, for the matched state-action pairs (i, a) or (j, b) , we have $t_{i,a}^* = t_{j,b}^* = \infty$ being optimal because any change not only loses flow payoff, it unnecessarily incurs the change cost, C . Therefore,

$$V^*(i, a) = r_i + s_i V^*(j, a) \text{ and } V^*(j, b) = r_j + s_j V^*(i, b) \quad (4)$$

where r_i and r_j are the expected discounted flow payoffs until the next change state, and s_i and s_j are the corresponding expected discount factors. For the

mis-matched pair (i, b) , the value function is

$$V^*(i, b) = \max_{t_i \in [0, \infty]} V^*(j, b) \cdot \int_0^{t_i} e^{-rw} f_{i,b}(w; t_i) dw + \int_{t_i}^{\infty} \left[\left(\int_{t_i}^w 1e^{-rx} dx \right) + V^*(j, b)e^{-rw} \right] f_{i,b}(w; t_1) dw - Ce^{-rt_i}(1 - F_{i,b}(t)), \quad (5)$$

with a similar expression for the value of the other mis-matched pair, $V^*(j, a)$. The value functions reduce to two equations in two unknowns, rendering the numerical analysis and comparative statics moderately direct.

2.4.6. Optimality Conditions. The FOCs and SOC for the optimal amount of hesitation in the mismatched state-action pairs (j, a) and (i, b) are very similar to the search problem in §2.1 above. To most simply see what is involved, suppose that the hazard rates depend only on the action-state pair at the beginning of an interval. At time t in an interval that starts at (j, a) , the FOCs must capture indifference between switching and not switching between t_1^* and $t_1^* + dt$: the net flow benefit from switching is $[u(j, b) - u(j, a)]$, while it costs rC , the perpetual annuity flow value of the cost C ; $h(t_1^*)dt$ gives the probability that the state switches from j back to i in the next instant, and if this happens, then the decision maker has saved C plus the value difference $(V^*(i, a) - V^*(i, b))$. Setting these equal delivers a special case of the FOCs (Euler equations) in Theorem 2 below,

$$[u(j, b) - u(j, a)] - rC = h(t_1^*(j, a)) [C + (V^*(i, a) - V^*(i, b))]. \quad (6)$$

Note that the benefit of delay must be larger than the benefit of a change in action before t_1^* and smaller after. As in the search problem, this requires that the hazard rate be decreasing at the optimum. We will return to, and extend, this model in §4.2.

3. THE MODEL

We begin with a brief review of incomplete waiting times, hazard rates and their interpretations, then turn to the class of non-stationary stochastic optimization problems under consideration. The section ends with the two main results: Theorem 1 gives the Bellman equation, shows that the value function is the unique fixed point of a contraction mapping, and that one can solve the infinite horizon problem by solving the within-interval problems; Theorem 2 gives the Euler equations, that is, the necessary conditions, for an optimal amount of hesitation.

There are some delicate details in the description of the optimization problem, but the essential idea is simple: one starts a new stochastic interval in the state i with the inherited action, a ; a strategy specifies the times within the interval at which the action will be changed and what the new action will be; all of these changes are carried out so long as there is no change in state; a change in state opens a new interval, and the process begins again. Payoffs are given by the expected value of a discounted integrated flow of utility minus the expected value of change costs along the path. We call this a *semi*-Markovian structure because the arrival rate of the transition times can be non-stationary.¹¹

3.1. Summary. A **semi-Markovian decision problem (sMdp)** is specified by: $S \times A$, the product of a compact set of possible state and a compact set of actions; a continuous flow utility function, $u : S \times A \rightarrow \mathbb{R}$; a continuous, strictly positive cost-of-change function, $c : (A \times A) \times S \rightarrow \mathbb{R}$; the continuous feasible plan correspondence, $(i, a) \mapsto \mathbb{P}(i, a)$ as given in Definition 1; a continuous semi-Markovian stochastic structure, as given in Definition 2; and the mapping from plans to induced distributions as given in Definition 3.

3.2. Hazard Rates. A random variable, $W \geq 0$, is **incomplete** if it has a mass point at ∞ . For a possibly incomplete W with density on $[0, \infty)$ and $0 \leq t < \infty$, we have the following relations between the density, $f(t)$, the cumulative distribution function (cdf), $F(t)$, the reverse cdf, $G(t)$, the hazard rate, $h(t)$, the cumulative hazard, $H(t)$, and the mass at infinity, q :

$$F(t) = \int_0^t f(x) dx; \quad G(t) = 1 - F(t); \quad h(t) = \frac{f(t)}{G(t)} \text{ or } f(t) = h(t)G(t);$$

$$H(t) = \int_0^t h(x) dx; \quad G(t) = e^{-H(t)}; \quad EW = \int_0^\infty G(t) dt; \text{ and } q = e^{-H(\infty)}. \quad (7)$$

If the incompleteness parameter, q , is strictly positive, then $EW = \infty$, and, as time goes on, and one becomes surer and surer that the event will never happen, $P(W = \infty | W > t) \uparrow 1$ as $t \uparrow \infty$. For complete distributions, if the

¹¹Semi-Markovian dynamic programming often imposes the restriction that any change in action occurs at the beginning of an interval e.g. in the textbooks Ross [1970] and Hu and Yue [2008] separated by almost four decades in time. As has been recognized in the operations research literature, and as the examples in §2 demonstrate, this is often sub-optimal. We will give a more detailed comparison with the literature below.

hazard rate is decreasing (resp. increasing), then the expected *future* waiting time, $E(W|W > t) - t$, is increasing (resp. decreasing).

3.2.1. *Examples.* The following are well-studied examples.

- (1) An incomplete negative exponential with parameters $\lambda > 0$ and $q \geq 0$ has cdf $F(t) = (1 - q)(1 - e^{-\lambda t})$. For $q > 0$, the hazard rate $h(t) = \lambda \cdot ((1 - q)e^{-\lambda t}) / (q + (1 - q)e^{-\lambda t})$ is everywhere strictly decreasing. If $q = 0$, then the hazard rate is constant at λ , that is, the waiting time is *memoryless*.
- (2) An incomplete Weibull distribution with parameters $\gamma, \lambda > 0$ and $q \geq 0$ is of the form $W = X^\gamma$, where X is an incomplete negative exponential with parameters λ, q . The cdf is $F(t) = (1 - q)(1 - e^{-\lambda t^{1/\gamma}})$, and the hazard rate is $h(t) = \frac{\lambda}{\gamma} t^{\frac{1-\gamma}{\gamma}} \left[(q/(1 - q))e^{\lambda t^{1/\gamma}} + 1 \right]^{-1}$.
 - (a) If $\gamma > 1$, the Weibull is a convex transformation of the negative exponential, $h(0+) = \infty$, and the hazard rate strictly decreases to 0 whether or not $q > 0$.
 - (b) If $\gamma < 1$, the Weibull is a concave transformation of the negative exponential and $h(0) = 0$. If $q > 0$, the hazard rate is hump-shaped, first increasing then decreasing back to 0. If $q = 0$, the hazard rate is strictly increasing, and $\lim_{t \uparrow \infty} h(t) = \infty$.
- (3) In increasing order of generality, we have the following distributions of waiting for events to happen.
 - (a) An Erlang(λ, M) distribution of the time one waits until M iid negative exponential events with parameter λ occur. It has cdf $F(t) = 1 - \sum_{m=0}^{M-1} \frac{1}{m!} e^{-\lambda t} (\lambda t)^m$, $h(0) = 0$, and the hazard rate is increasing, concave, and bounded above with $\lim_{t \uparrow \infty} h(t) = \lambda$.
 - (b) The hypoexponential distribution corresponds to waiting for M independent negative exponentials with parameters $\lambda_1, \dots, \lambda_M$ to happen.
 - (c) The Coxian distribution generalizes the hypoexponential by allowing, at random, the m 'th event to be the last one, where $m \leq M$ is possible.
 - (d) The phase-type distributions generalize the previous three types of distributions. Suppose that a process begins state $m \in \{1, \dots, M\}$ with probability $\alpha = (\alpha_1, \dots, \alpha_M)$, and transitions from the m 'th state either to an absorbing state, $m = 0$, or to one of the other non-absorbing states with constant hazard rates $(\lambda_{m,i})_{i \neq m}$. A phase-type

distribution is parametrized by α and the hazard rate matrix, and is distribution of the waiting time until the transition to the absorbing state, $m = 0$.

3.2.2. Dense Class Interpretations. Up to any degree of precision, arbitrary hazard rate patterns can be interpreted as the result of waiting for unobserved random sequences of events to be realized. For example, a phase-type distribution with $M = 2$, α putting most of its mass on $m = 2$, a low transition rate, $\lambda_{2,1}$ from 2 to 1, and a higher $\lambda_{1,0}$ has a low initial hazard rate that increases — initially, the unobserved state is most likely to be 2, but as time passes, the likelihood of being in 1, conditional on transition to $m = 0$ not having happened, increases. With more elaborate state spaces and transition patterns, the hazard rate can decline again, then increase. More generally, the set of phase-type waiting distributions is L_1 -dense in the set of waiting time distributions — for every (possibly incomplete) density f and $\epsilon > 0$, there is a phase-type density g with $\|f - g\|_1 := \int_0^\infty |f(s) - g(s)| ds < \epsilon$.

3.2.3. Updating Interpretations. There are also a wide variety of mixture model or updating interpretations of hazard rates.

If β is a prior distribution on the set of possible distributions of the next arrival time in $(0, \infty]$, denoted $\beta \in \Delta(\Delta(0, \infty])$, then the hazard rate at t is

$$h_\beta(t) = \lim_{\delta \downarrow 0} \left(\frac{1}{\delta} \int_{\Delta(0, \infty]} p((t, t + \delta)) d\beta(p) \right) / \left(\int_{\Delta(0, \infty]} p((t, \infty]) d\beta(p) \right). \quad (8)$$

If β and β' are two priors with $h_\beta(\cdot) = h_{\beta'}(\cdot)$, then the optimal course of action is the same so that hazard rate analyses can cover a very broad array of models.

Known hazard rates. With δ_t denoting point mass on t , if $\beta(\{\delta_t : t \in A\}) = \int_A f(x) dx$, then $h_\beta(t)$ is the hazard rate associated with the density $f(\cdot)$, yielding an interpretation of a known hazard rate.

Bi-modal hazard rates. In our models of Constitutional amendments, bi-modal hazard rates for reversals of opinions are of interest. They can be understood as a mixture of three effects: with probability $\alpha_1 > 0$, the previous change was a fad, and the density for the reversal time of fads is given by f_1 with a low upper bound to its support; with probability $\alpha_2 > 0$, it is a longer lasting but still temporary change, and the density, f_2 , of such reversals puts mass on a much longer set of times; and with probability $q > 0$, there is never a reversal. The hazard

rate is initially high, tails off, increases, then decreases again as the probability that $W = \infty$ given $W > t$ converges to 1 as $t \uparrow$.

Smooth mixtures. Mixtures of smooth classes of distributions are often dense. For example, if W is an incomplete Weibull distribution with parameters λ, γ, q where (λ, γ, q) is itself random with mixture distribution β , then the class of densities that W may have is again L_1 -dense.¹²

3.2.4. *Interpretations with Expanded State Spaces.* The passage of time, changes of the state, and changes of actions contain information about the distribution of future developments. Adding more states allows one to expand the set of dynamic factors that are modeled: if the state i_k contains information about how many localities and/or states have approved womens' suffrage, then every time this number increases, we expect that the future hazard rate path for opinion reversal should be lower; if the state contains information about how many times a unit has been repaired or upgraded rather than replaced, then every time this number increases we expect the new hazard rate path for a breakdown should be higher than the previous hazard rate path in repair/replace models; if the state contains information about the number of unsuccessful job interviews, then every time this number increases, we might expect the hazard rate for and the distribution of a good job offer to decline.

3.3. **Controlled Semi-Markov Processes.** There is a compact *state space*, S , with metric D_S , generic elements denoted i, j, i_0, i_1, \dots , and generic random elements of S will be denoted $I_k, k = 1, 2, \dots$. There is also a compact *action space*, A , with generic elements denoted a, b, a_0, a_1, \dots . As above, $Y_k, k = 0, 1, \dots$ are the random variables giving the state-change times, and a **stochastic interval** is a subset of $\Omega \times [0, \infty)$ of the form $\llbracket Y_k, Y_{k+1} \llbracket := \{(\omega, t) : Y_k(\omega) \leq t < Y_{k+1}(\omega)\}$.

We first define the class of plans/policies, their associated outcomes and utilities, and then the stochastic structure, that is, the determinants of the joint distributions of the Y_k and the states, I_k .

3.3.1. *Semi-Markovian Plans.* We use functions $\mathbf{p} : [0, \infty) \rightarrow A$, to specify the chosen plan of action *within* a stochastic interval. We measure distance between plans using the Hausdorff distance between (the closures of) their graphs in a

¹²Proofs of such can be given using the approximation techniques in [Hornik et al., 1994]. See Cho and White [2010] for references to mixture models of exponentials and Weibulls in both financial econometrics and in duration models.

fashion less sensitive to differences in the right tails of the plans. Specifically, we introduce a distance $D(\cdot, \cdot)$ between points (t, a) and (t', a') in $[0, \infty) \times A$ that takes the form¹³ $D((t, a), (t', a')) = (|e^{-\rho t} - e^{-\rho t'}| + \|e^{-\rho t} a - e^{-\rho t'} a'\|)$ in the case that A is a compact subset of \mathbb{R}^ℓ . The distance between $\mathbf{p} \in \mathbb{P}(i, a)$ and $\mathbf{p}' \in \mathbb{P}(i', a')$ is defined as

$$d_H(\mathbf{p}, \mathbf{p}') = \inf\{\epsilon > 0 : gr(\mathbf{p}) \subset gr(\mathbf{p}')^\epsilon \text{ and } gr(\mathbf{p}') \subset gr(\mathbf{p})^\epsilon\} \quad (9)$$

where $gr(\mathbf{p})$ and $gr(\mathbf{p}')$ are the graphs of the plans and $gr(\mathbf{p})^\epsilon$ and $gr(\mathbf{p}')^\epsilon$ are the sets of points at D -distance less than ϵ from the graphs.

This metric has the property that differences in the right tails are down-weighted. For example, if r is the diameter of A and $e^{-\rho T} < \epsilon / \max(r, 1)$, then two plans that disagree only on $[T, \infty)$ can be at d_H -distance no more than ϵ from each other. This means that, up to ϵ , only changes in plan before T matter. We will see that expected payoffs are d_H -continuous and that the search for optima can be, without loss, restricted to d_H -compact sets of plans.

Definition 1. A **plan starting from** (i_0, a_0) is a function $\mathbf{p} : [0, \infty) \rightarrow A$ that satisfies $\mathbf{p}(0) = a_0$, is piecewise-constant, continuous from the left, and has at most finitely many discontinuities over any finite interval. The set of **feasible plans** is a non-empty set of plans denoted $\mathbb{P}(i_0, a_0)$ where the correspondence from (i_0, a_0) to $\mathbb{P}(i_0, a_0)$ is assumed closed-valued and continuous. A **complete semi-stationary plan** is a measurable selection from the set of feasible plans, that is, a measurable function from state-action pairs to plans, $(i_0, a_0) \mapsto \mathbf{p}_{(i_0, a_0)} \in \mathbb{P}(i_0, a_0)$.

For $\mathbf{p} \in \mathbb{P}(i_0, a_0)$, $T_{\mathbf{p}} = \{t \in [0, \infty) : \mathbf{p}(t) \neq \mathbf{p}(t+)\}$ denotes the set of action change times (where $\mathbf{p}(t+) := \lim_{\epsilon \downarrow 0} \mathbf{p}(t + \epsilon)$). The “continuous from the left” aspect of plans accomodates instant reactions to changes of states — a plan that instantly changes to an action a_1 in response to a state change to (i_0, a_0) has $0 \in T_{\mathbf{p}}$, satisfies $\mathbf{p}(0) = a_0$ and $\mathbf{p}(0+) = a_1$. We interpret the times in $T_{\mathbf{p}}$ as specifying the chosen action as a function of the *within interval time*, that is, as a function of the time since the last state change. Formally, for $A_k := \{\omega : Y_k(\omega) < \infty\}$, $\omega \in A_k$ and $t \in [Y_k(\omega), Y_{k+1}(\omega))$, the **within interval time** s is given by $s_k(\omega, t) = t - Y_k(\omega)$. Using \mathbf{p} from an interval starting in (i_k, a_k) at time Y_k corresponds to changing actions at the within-interval times s_k with $\mathbf{p}(s_k) \neq \mathbf{p}(s_k+)$.

¹³That is, has the same uniformity as. See the online supplement for the general definition.

The feasibility constraint is a flexible modeling tool: irreversibility of the action change after, say, changing from (i_0, a_0) to a_1 is modeled by $\mathbb{P}(i_0, a_1)$ containing only the plan satisfying $\mathbf{p}(t) \equiv a_1$; a minimal reaction time of r times involve $\mathbf{p}(t) = a_1$ for all $t \in [0, r]$.

3.3.2. The Stochastic Structure. We now specify the density and the conditional change probabilities as a function of the past history that satisfying two kinds of continuity conditions, one to guarantee the validity of the Euler equations in Theorem 2 and one to guarantee the existence of optima in the Bellman equation in Theorem 1. In the following, the distance between densities, complete or not, is the L_1 distance, $\|f - g\|_1 = \int_0^\infty |f(s) - g(s)| ds$.

Definition 2. A **semi-Markovian stochastic structure** is a function $(s, \mathbf{p}, i) \mapsto (f(s; \mathbf{p}, i), \mu(s; \mathbf{p}, i))$ from $(0, \infty) \times \mathbb{P} \times S$ to $[0, \infty) \times \Delta(S)$ satisfying the following conditions:

- a. progressivity, for all state-action pairs (i, a) , all $\mathbf{p}, \mathbf{p}' \in \mathbb{P}(i, a)$ and all intervals $[0, t]$, if $\mathbf{p}|_{[0, t]} = \mathbf{p}'|_{[0, t]}$, then for all $s \in [0, t]$, $(f(s; \mathbf{p}, i), \mu(s; \mathbf{p}, i)) = (f(s; \mathbf{p}', i), \mu(s; \mathbf{p}', i))$;
- b. joint and L_1 -continuity, $(s, \mathbf{p}, i) \mapsto (f(s; \mathbf{p}, i), \mu(s; \mathbf{p}, i))$ is jointly continuous, and $(\mathbf{p}, i) \mapsto f(\cdot; \mathbf{p}, i)$ is L_1 -continuous;
- c. non-redundancy, for all s , $\mu(s; \mathbf{p}, i)(\{i\}) \equiv 0$; and
- d. waiting times are uniformly smooth, for every $\epsilon > 0$, there exists a $\tau \in (0, 1)$ such for all $(i, a) \in S \times A$, all $\mathbf{p} \in \mathbb{P}(i, a)$, and all $t \geq 0$, $\int_t^{t+\tau} f(s; \mathbf{p}, i) ds < \epsilon$.

Using a plan $\mathbf{p} \in \mathbb{P}(i_k, a_k)$ in an interval $\llbracket Y_k, Y_{k+1} \rrbracket$ that starts in (i_k, a_k) means that the waiting time, $W_{k+1} := Y_{k+1} - Y_k$, has the possibly incomplete density $f(s; \mathbf{p}, i_k)$ on $(0, \infty)$, and the distribution of the waiting time must be L_1 -continuous in the chosen plan and starting state. Unbounded densities, such as the Weibulls with $f(0+) = h(0+) = \infty$, are allowed because the continuity is only required on $(0, \infty)$, but the uniform smoothness condition requires that plans cannot push waiting times arbitrarily close to having mass points, e.g. at 0, of any size. While densities are required to be continuous on the within-interval time $(0, \infty)$ when there is no change in action, continuity in plans allows for the

density and the change probability to be discontinuously different before and after a change in action.¹⁴ Conditional on $W_{k+1} = w < \infty$, the distribution of the new state is given by $\mu(w; \mathbf{p}, i_k)$, and non-redundancy requires changes in states to actually be changes. The progressivity condition guarantees that only past changes in actions and states can affect the current density and conditional change probability.

The following specifies how plans interact with stochastic structures to give rise to distributions of waiting times and state changes. Recall the convention than $\min \emptyset = \infty$.

Definition 3. Using a plan $\mathbf{p}_0 \in \mathbb{P}(i_0, a_0)$ starting from (i_0, a_0) at $Y_0 = 0$ gives rise to a random waiting time-state pair, (W_1, I_1) with the **induced distribution**, possibly incomplete, defined by $\nu_1(E \times F) = \int_E f(s; \mathbf{p}_0, i_0) \cdot \mu(s; \mathbf{p}_0, i_0)(F) ds$ for $E \times F \subset (0, \infty) \times S$. For $k \geq 1$, if $W_k = w_k < \infty$, $a_k = \mathbf{p}_{k-1}(W_k)$ and $I_k = i_k$, using a plan $\mathbf{p}_k \in \mathbb{P}(i_k, a_k)$ gives rise to a random waiting time-state pair, (W_{k+1}, I_{k+1}) with the **induced distribution** $\nu_{k+1}(E \times F) = \int_E f(s; \mathbf{p}_k, i_k) \cdot \mu(s; \mathbf{p}_k, i_k)(F) ds$. Continuing inductively, a sequence of plans $\mathbf{p}_0, \mathbf{p}_1, \dots$ with $\mathbf{p}_k \in \mathbb{P}(i_k, \mathbf{p}_{k-1}(W_k))$ gives rise to a sequence of conditionally independent random waiting times-states pairs, (W_k, I_k) , $k = 1, \dots, K$, with $K := \min\{k : W_k = \infty\}$.

For any semi-stationary plan, i.e. any measurable function $(i, a) \mapsto \mathbf{p}_{(i,a)} \in \mathbb{P}(i, a)$, the induced distribution of (W_k, I_k) depends only on the value of I_{k-1} and $\mathbf{p}_{k-1}(W_k)$. Therefore, the induced distributions for a complete semi-stationary plan give a stochastic kernel [Feller, 1971, VI.11, Definition 1] so that the process (W_k, I_k) , $k = 1, \dots, K$ is a Markov chain with a random, possibly infinite, number of realizations. Theorem 1 will show that semi-stationary plans are optimal.

3.3.3. *Utilities.* Changing action from a to b in state i has a cost, $c(a, b; i)$. We assume that $c : (A \times A) \times S \rightarrow \mathbb{R}_{++}$ is continuous, has a strictly positive minimum value, and that direct changes are always cheaper than multi-step changes of action, $c(a, a'; i) + c(a', a''; i) > c(a, a''; i)$. Flow utilities are given by a continuous $u : S \times A \rightarrow \mathbb{R}$. For an interval starting in (i, a) , the utility associated with a plan $\mathbf{p} \in \mathbb{P}(i, a)$ has two components, the expected flow utility,

$$\int_0^\infty \left[\int_0^s u(i, \mathbf{p}(t)) e^{-rt} dt \right] f(s; \mathbf{p}, i) ds, \quad (10)$$

¹⁴An example of an allowed discontinuity in the density is in eqn. (1), specifying the stochastics for the inoculation model.

and, with $T_{\mathbf{p}}$ denoting the action change times, the expected costs,

$$\int_0^\infty \left[\sum_{t \in T_{\mathbf{p}}, t < s} c(\mathbf{p}(t), \mathbf{p}(t+); i) e^{-rt} \right] f(s; \mathbf{p}, i) ds. \quad (11)$$

We denote the sum of the terms in (10) and (11) by $\mathcal{R}(\mathbf{p}; (i, a))$ so that the utility of a sequence of plans $\mathbf{p}_0, \mathbf{p}_1, \dots$ starting in (i_0, a_0) at $Y_0 = 0$ is

$$\mathcal{R}(\mathbf{p}_0, (i_0, a_0)) + \mathbb{E} e^{-rY_1} \mathcal{R}(\mathbf{p}_1, (I_1, \mathbf{p}_0(Y_1 - Y_0))) + \dots \quad (12)$$

where the expectations of the I_k and the e^{-rY_k} are taken with respect to the distributions induced by the plans, and we assume that $\mathbf{p}_k \in \mathbb{P}(I_k, \mathbf{p}_{k-1}(Y_k - Y_{k-1}))$ because this is the state-action pair in effect when the state change at Y_k happens.

3.4. Existence and the Bellman Equation. The value function for an sMdp, $V^* : S \times A \rightarrow \mathbb{R}$, is defined as

$$V^*(i_0, a_0) = \sup_{\mathbf{p}_0, \mathbf{p}_1, \dots} \mathcal{R}(\mathbf{p}_0; (i_0, a_0)) + \sum_{k \geq 1} \mathbb{E} e^{-rY_k} \mathcal{R}(\mathbf{p}_k; (I_k, \mathbf{p}_{k-1}(Y_k - Y_{k-1}))) \quad (13)$$

where $\mathbf{p}_0 \in \mathbb{P}(i_0, a_0)$ and $\mathbf{p}_k \in \mathbb{P}(I_k, \mathbf{p}_{k-1}(Y_k - Y_{k-1}))$ for $k \geq 1$. The following parallels the existence results for discrete time dynamic programming in Blackwell [1965, §6].

Theorem 1. *For an sMdp, the mapping $T : C(S \times A) \rightarrow C(S \times A)$ defined by*

$$T(V)(i, a) = \max_{\mathbf{p} \in \mathbb{P}(i, a)} \left[\mathcal{R}(\mathbf{p}; (i, a)) + \int_{(0, \infty) \times S} V(i, \mathbf{p}(s)) e^{-rs} d\nu_{\mathbf{p}}(s, i) \right] \quad (14)$$

has contraction factor at most $\beta := \sup_{(i, a), \mathbf{p} \in \mathbb{P}(i, a)} \int_0^\infty e^{-rs} f(s; \mathbf{p}, i) ds < 1$, its unique fixed point is the value function, V^ , and following the complete semi-stationary plan given by the argmax policy for (14) in every stochastic interval delivers expected payoffs $V^*(i, a)$ for every $(i, a) \in S \times A$.*

Thus, the stochastic intervals deliver the recursive structure that allows us to use semi-stationary plans to solve the problem in (13): if one leaves a stochastic interval $\llbracket Y_k, Y_{k+1} \rrbracket$ to the state I_{k+1} while taking an action a_k , one starts the new stochastic interval in the state-action pair (I_{k+1}, a_k) ; if one behaves optimally thereafter, one receives $V^*(I_{k+1}, a_k) e^{-rY_{k+1}}$; and this is the maximization problem in (14). In the proof, the assumption that costs are bounded away from 0 is used only to guarantee that only a compact set of plans need be considered. As an immediate Corollary, if the feasible set of plans is already compact, the assumption that costs be bounded from 0 is not needed and the same result holds.

3.5. Euler Equations. We now turn to the FOCs and SOC for finite but strictly positive delay to be optimal for those intervals in which one change of action is optimal. Here we restrict attention to models in which the optimality conditions for a $t_1^* \in (0, \infty)$ involve only the value function and properties of the stochastic structure at t_1^* . For plans with only one action change at t_1 , we can express the stochastic structure as a function $w \mapsto (f(w; t_1), \mu(w; t_1))$. Progressivity requires the existence a function $w \mapsto (f^-(w), \mu^-(w))$ such that for all $w < t_1 < t_1'$, $(f(w; t_1), \mu(w; t_1)) = (f(w; t_1'), \mu(w; t_1')) = (f^-(w), \mu^-(w))$. For $t_1 < t_1' < w$ we can have $(f(w; t_1), \mu(w; t_1)) \neq (f(w; t_1'), \mu(w; t_1'))$, which means that the general FOCs for an optimal t_1^* must involve the discounted future path of the *derivative* of $(f(\cdot; t_1^*), \mu(\cdot; t_1^*))$. These are non-local aspects of the future that are **not** captured by the value function. The following is a minimal assumption for FOCs and SOC to involve only the value function and properties of the stochastic structure local to t_1^* .

Definition 4. A stochastic structure is **localized** if for all starting points (i_0, a_0) and a_0 to a_1 changes, there exists a function $w \mapsto (f^+(w), \mu^+(w))$ such that for all $0 < t_1 < t_1' < w$, $(f(w; t_1), \mu(w; t_1)) = (f(w; t_1'), \mu(w; t_1')) = (f^+(w), \mu^+(w))$.

We use $h^+(w)$ and \mathbb{E}^+ to denote the hazard rate and the expectation operator for $\mu^+(w)$ at times $w > t_1$, and use $h^-(w)$ and $\mathbb{E}^-(w)$ to denote the corresponding quantities for $w < t_1$.

Theorem 2. *In an sMdp with a localized¹⁵ stochastic structure, if \mathbf{p}^* solves*

$$\max_{\mathbf{p} \in \mathbb{P}(i_0, a_0)} \left[\mathcal{R}(\mathbf{p}; (i_0, a_0)) + \int_{(0, \infty) \times S} V^*(i, \mathbf{p}(s)) e^{-rs} d\nu_{\mathbf{p}}(s, i) \right], \quad (15)$$

and hesitates until a time $t_1^ \in (0, \infty)$ for its single action change to a_1^* at cost c^* , then t_1^* must satisfy*

$$rc^* - [u(i_0, a_1^*) - u(i_0, a_0)] = \quad (16)$$

$$h^-(t_1^*) \left(\frac{1}{r} u(i_0, a_1^*) (1 - \gamma(t_1^*)) + [\gamma(t_1^*) \mathbb{E}^+ V^*(I_1, a_1^*) - \mathbb{E}^- V^*(I_1, a_0)] - c^* \right)$$

where $\gamma(t_1^) := \frac{h^+(t_1^*)}{h^-(t_1^*)}$. Further, if the term multiplying $h^-(t_1^*)$ in (16) is positive (resp. negative), then the SOC are strictly satisfied iff the hazard rate h^- is increasing (resp. decreasing) at t_1^* .*

The class of problems for which this is applicable models situations where: there is only one realistic opportunity for a change of action; or the cost of changes is high relative to the differences in the flow utility; or when uncertainty about the durability of the circumstances pushing one to change actions is initially high but then settles down for good. A broad class of institutional decision making problems fit these descriptions. Examples below demonstrate that the one-move necessary condition in (16) can also be useful when several moves are optimal within an interval and when the choice of possible action is continuous rather than discrete.

4. APPLICATIONS OF THE BELLMAN AND EULER EQUATIONS

The increasing and decreasing hazard rates cases in Theorem 2 have very different intuitions, one consonant with the searching till success seems too unlikely, the other consonant with delaying expensive preventive measures until danger presses. To see this in more detail, we now examine the Euler equations in various models, starting with the simplest, and proceeding through their use in models that include multiple action switches within a stochastic interval and continuous choices of action. We then turn to an analysis of optima based more directly on the Bellman equations.

4.1. Memoryless Processes. When $f(\cdot; \mathbf{p}, i)$, the density part of stochastic structure has a constant hazard rate between action changes and change probability, $\mu(\cdot; \mathbf{p}, i)$, is constant except, possibly, at action changes, the right-hand side of the Euler equation, (16), is constant for any choice of a_1^* . Within this class of problems, at most one change of action is ever optimal, and there are three possibilities for the optimal hesitation.

First, if flow benefits to the best change of action outweigh the benefits of delay, that is if the left-hand side of (16) is larger than the right-hand side, then $t_1^* = 0$. Further, because the hazard rates and transition probabilities are constant, if this is true at any t_1 , then it is true for all t_1 .

Second, if the benefits of delay outweigh the flow benefits of the best change of action, then the optimal $t_1^* = \infty$ because the right-hand side of (16) being larger than the left-hand side. Again, if this is true at any t_1 , then it is true for all t_1 .

Third, if we have equality in (16), then any $t_1^* \in [0, \infty]$ is optimal.

¹⁵The FOCs for non-localized stochastic structures can be found in the on-line supplement.

4.2. Two Types of Search. A standard search procedure often starts with a quick scan which will find anything obvious, and, if this is not successful, then moves on to a slower, more thorough search. We first review the simple search model above then extend it to encompass the quick-scan-then-grind search model. The model of §2.1 has two states, i_0 , no search results, and i_1 , successful search; and two actions, s_0 , searching, and b , abandoning search. The set of feasible plans is D -compact: it has $\mathbb{P}(i_k, s_0)$ being the set of all functions taking the value s over some interval $[0, t_1]$, $t_1 \in [0, \infty]$, and taking the value b thereafter, $k = 0, 1$; and, to capture irreversibility, has $\mathbb{P}(i_k, b)$ being the set of all functions constant at b . The density before t_1 is given by $f(\cdot)$, after the change to b , it is given by $f(\cdot; t_1) \equiv 0$ so that the post- to pre-change hazard ratio is $\gamma(t_1) = 0$. The flow utilities are $u(i_0, s_0) = -c$, $u(i_0, b) = \underline{\pi}$, $u(i_1, s_0) = \bar{\pi} - c$ and $u(i_1, b) = \bar{\pi}$. Change costs are set equal to 0. With these choices, the FOCs from Theorem 2 are

$$0 - [\underline{\pi} + c] = h^-(t_1^*) \left(\frac{1}{r} u(i_0, b)(1 - 0) + [0\mathbb{E}^+V^*(I_1, b) - \mathbb{E}^-V^*(I_1, s_0)] - 0 \right). \quad (17)$$

Since the only possible change of state is from i_0 to i_1 , the new state, I_1 is, with probability 1, equal to i_1 so that $\mathbb{E}^-V^*(I_1, s_0)$ is equal to $\frac{1}{r}\bar{\pi}$, the expected discounted value of the flow of $\bar{\pi}$. With these, the FOCs are $h(t_1^*) = r(c + \underline{\pi})/(\bar{\pi} - \underline{\pi})$ and the SOC requires $h'(t_1^*) < 0$.

Suppose now that there are two search techniques: s_1 with flow costs c_1 and hazard rate h_1 that starts high, quickly declines to 0 and has $\int_0^\infty h_1(x) dx < \infty$, that is, a strictly positive incompleteness parameter; and s_2 with corresponding c_2 and h_2 that declines much more slowly and has $\int_0^\infty h_2(x) dx = \infty$. Suppose further that switching from s_1 to s_2 at $t_{1,2}$ means that the hazard rate at $t_{1,2} + s$ is $h_2(s)$, that the value if only s_2 is available is $V_2 > \frac{1}{r}\underline{\pi}$ is the value if only s_2 , and that it is optimal to use s_1 first. The FOCs for $t_{1,2}^*$ are $h_1(t_{1,2}^*) = \frac{r(c_1 + rV_2)}{\bar{\pi} - rV_2}$ because rV_2 is the annuitized flow value of V_2 . Because the SOC is $h_1'(t_{1,2}^*) < 0$, the comparative statics are straightforward: increases in r , i.e. becoming less patient, means more time spent on scanning, that is, more time hoping for a quick resolution; increases in c_1 or V_2 as well as everywhere decreases in h_1 mean less time spent scanning.

The hazard rate at $t_{1,2} + s$ being $h_2(s)$ is an independence assumption that does not make sense if the distributions are incomplete and our interpretation of incomplete distributions for search times is that the object is not there to be

found.¹⁶ One can capture and analyze such a situation by using $V_2(q_t)$ in the FOCs where q_t is posterior probability that the object is not findable using s_2 given that using s_1 until t has not found the object and $V_2(q_t)$ is the value to using s_2 starting from that point.

4.3. Continuous Degrees of Innoculation. In the simple inoculation model of §2.2 there are: two states, i_0 , the crisis not having arrived, and i_1 , the crisis having arrived; and two actions, a_0 , not inoculating, and inoculating, a_1 . The cost of switching from a_0 to a_1 is C , the pre-disaster flow utilities are $u(i_0, a_0) = u(i_0, a_1) = \bar{u}$, while the post-disaster flows depend on whether or not one was inoculated, $u(i_1, a_0) = 0$ and $u(i_1, a_1) = \underline{u}$, $\bar{u} > \underline{u} > 0$. To capture inoculation being useless after the disaster has arrived, it is simplest to specify that $\mathbb{P}(i_1, a_0)$ contains only the time path constant at a_0 . The arrival density at times w prior to inoculation at t_1 is $f(w)$, at $w \geq t_1$, it is $(1 - \theta)f(w)$ so that the post- to pre-hazard rate ratio is $\gamma(t_1) = (1 - \theta)$. With these, the FOCs from Theorem 2 are

$$rC - [\bar{u} - \underline{u}] = h^-(t_1^*) \left(\frac{1}{r} \bar{u} \theta + [(1 - \theta) \mathbb{E}^+ V^*(I_1, a_1) - \mathbb{E}^- V^*(I_1, a_0)] - C \right). \quad (18)$$

Having inoculated when the state change $I_1 = i_1$ arrives yields expected discounted payoffs of $\mathbb{E}^+ V^*(I_1, a_1) = \frac{1}{r} \underline{u}$, having failed to inoculate by that time yields $\mathbb{E}^- V^*(I_1, a_0) = 0$. Combining, the FOCs are $h(t_1^*) = rC / (\frac{1}{r} [\theta \bar{u} + (1 - \theta) \underline{u}] - C)$, the SOCs are $h'(t_1^*) > 0$, and increases in either θ or \underline{u} make less hesitation optimal.

Let us now suppose that the benefits of inoculation, the efficacy, θ , and the post-disaster utility, \underline{u} , are continuously adjustable, and that their cost is $C(\theta, \underline{u})$. In this context, optimal design of the inoculation equates expected marginal costs and benefits. With $C_{\underline{u}} du$ denoting the cost of increasing \underline{u} by du , the optimal choice of \underline{u} equates the marginal cost, $C_{\underline{u}}$, with the expected value of the increased flow of utility if inoculation precedes disaster, $C_{\underline{u}} = (1 - \theta^*) \frac{1}{r} \mathbb{E}(e^{-r(W-t_1^*)} | W > t_1^*)$, and the optimal θ^* equates the marginal cost of prevention to the discounted expected value of avoiding the disaster, $C_{\theta} = (\bar{u} - \underline{u}^*) \frac{1}{r} \mathbb{E}(e^{-r(W-t_1^*)} | W > t_1^*)$. These FOCs yield the following conclusions: increases in one aspect of the quality of an inoculation lower the marginal value of the other aspect; irrespective of

¹⁶Different search techniques can be blind in different ways, so “not being findable using s_1 ” need not have implications for the likelihood of “not being findable using s_2 .”

the distribution of $(W|W > t_1)$, the (absolute value of the) slope of the iso-cost curve at the optimum is $C_\theta/C_u = (1 - \theta^*)/(\bar{u} - \underline{u}^*)$; and the signs of comparative statics results for θ^* and \underline{u}^* must be equal, θ^* and \underline{u}^* increase with more patience (lower r), with more at stake (higher \bar{u}), and with a more hazardous environment (an everywhere increase hazard rate).

4.4. Revisiting Adaptation to Changing Circumstances. Taking voting rights away from a newly enfranchised group that can vote against being disenfranchised is difficult in a democracy.¹⁷ By contrast, restoring privileges that have been taken away, such as the legal consumption of alcohol, seems relatively easy. We first use the Euler equations for the model of optimal adaptation to ever changing circumstances of §2.4 to show that higher costs of either passing or repealing an Amendment lead to more hesitation (one can also use the Bellman equations for this). We then show that, in many but not all cases, if passage of a law reduces the hazard that opinion will switch against it, this decreases optimal hesitation.

4.4.1. The Costs of Repeal. Recall that the two matched state-action pairs, (i, a) and (j, b) give positive flow utilities, while the mis-matched pairs, (i, b) and (j, a) , give flow utilities of 0. In either matched state-action pair, there is no incentive to incur the cost of changing action so that the Euler equations are never satisfied. For the mis-matched pair (i, b) , provided that $\gamma(t) = 1$, the Euler equation is

$$[rc(b, a; i) - u(i, a)] = h_i(t_i) (0 + [V^*(j, a) - V^*(j, b)] - c(b, a; i)) \quad (19)$$

with a similar expression for the t_j^* associated with the other mis-matched pair, (j, a) .

Since $rc(b, a; i) < u(i, a)$ is necessary for it ever to be optimal to switch actions, if the Euler equations are satisfied, then both sides of (19) are negative and the SOC's require the hazard rate to be decreasing. Rewriting with the hazard as a ratio of two positive numbers we have

$$h_i(t_i^*) = \frac{u(i, a) - rc(b, a; i)}{[V^*(j, b) - V^*(j, a)] + c(b, a; i)}. \quad (20)$$

Since the hazard rates must be decreasing, the direct effect of an increase in $c(b, a; i)$ is to require more delay, that is, a higher t_i^* . There are also indirect

¹⁷Witness the decades of effort to disenfranchise freedmen in the southern United States after the passage of Amendment XV granting suffrage to former (male) slaves.

effects of increases in either $c(b, a; i)$, the cost of passing an Amendment, or of increases in $c(a, b; j)$ the cost of repealing an Amendment, but these indirect effects also push for more delay. Increases in the change costs decrease all of the values, but, e.g. in (20), the difference, $[V^*(j, b) - V^*(j, a)]$, increases as the action change costs increase — (j, b) is a matched pair so no change will occur in the stochastic interval beginning at (j, b) while the pair (j, a) is mis-matched. With declining hazard rates, the larger delays correspond to waiting until one is more sure that any future change of opinion will not happen, or is, at least, further in the future.

4.4.2. *The Power of Laws to Affect Attitudes.* Several contemporary sources suggest that both the opponents and proponents believed, mistakenly as it turns out, that in the event Prohibition was adopted, the laws would be widely followed and the social norm would also switch in the direction of temperance. In the model, this corresponds to lowering the post-change hazard rate for a switch back to the status quo ante. Intuitively, this pushes for less hesitation because a swifter passage of the amendment allows one to ‘lock in’ the benefits of having the action match the present state. In the model, this is true for a wide range of cases, but not in complete generality.

Recall that different matched state-action pairs have different flow utilities. If these flow utilities are approximately equal, as we think reasonable for the present application, or if being matched in the present state is more valuable than being matched in the other state(s), then acting so as to get the lock-in is optimal. However, if being matched to the present state is very much less desirable than being matched in the other state(s), locking in corresponds to choosing a long time at a very low level of utility, and this can be sub-optimal. In terms of the FOCs from Theorem 2, this corresponds to the negative term multiplying $h^-(t_1^*)$ being very large in absolute value, meaning that it is optimal to delay until the hazard rate for future change is very low indeed.

5. CONCLUSIONS AND EXTENSIONS

We began with the question as to what the optimal time is for changing a status quo policy in response to a change of the environment when policy change is costly and there is the possibility of subsequent changes in the environment. The tradeoff is between optimizing with respect to the current state and optimizing

with respect to the expected future state when the passage of time contains information about how soon the next change is likely to occur and what it is likely to entail. The Euler equations of Theorem 2 for the optimal amount of hesitation before a change of policy contain information about the optimal time, how that optimal time interacts with the new action, and how it interacts with different hazard rate patterns.

5.1. Decreasing Hazard Rate Problems. Every instant of hesitation in matching the current action to the current state has an opportunity cost, but with a decreasing hazard rate, hesitation increases the likelihood that any change of action will be enduring. When the costs of changing action, or of changing it back, are higher, this kind of ‘informative waiting’ is more valuable, an observation that contains some insights into the examples presented at the beginning, and suggests other areas of applicability.

Prohibition and women’s suffrage have an intimately tied history with large overlap in the activists promoting both Amendments [Flexner, 1996]. Our model suggests two cost-based explanations for the passage of Amendment XVIII (Prohibition) and XIX (Suffrage). Prohibition was passed three years after Amendment XVI, which allowed the federal government to collect income taxes. By providing an alternative to the large tax receipts from alcohol, this drastically lowered the cost of Prohibition, lowering the ‘optimal’ delay. Perhaps contributing to this effect was the widespread belief that, if passed, temperance would become the new norm. Women’s suffrage was delayed long past the accumulation of a great deal of evidence of majority approval, and part of the delay may be explained as ‘optimality’ in the face of the high costs of reversing any enfranchisement.

A fascinating institutional design question is how much delay to build in. James Madison [Hamilton et al., 1966, Federalist #62] argued for “The necessity of a senate” with “tenure of considerable duration” because of “the propensity of all single and numerous assemblies to yield to the impulse of sudden and violent passions, and to be seduced by factious leaders into intemperate and pernicious resolutions.” From the point of view of the model here, “intemperate and pernicious resolutions” correspond to actions that are either very costly or very costly to reverse, and “sudden and violent passions” correspond to changes of state that have high hazard rates for reversal, both of which lead to larger optimal delays.

5.2. Increasing Hazard Rate Analyses. While the previous models have a common intuition about trading off the informative value of waiting against the opportunity cost of immediate action, an entirely different class of intuitions is at work when hazard rates are increasing.

The costs of inoculation against financial system crises are, politically at least, very high. Our investigation of the optimal delay highlights the tradeoff between interventions that lower the arrival rate of the next crisis and those that lower the damage that it will cause: a policy change that reduces one of these makes the other less valuable at the margin. In such cases, hesitation saves incurring the cost of inoculation, but becomes sub-optimal as the hazard rate rises and the danger comes closer, an issue that seems relevant to policy responses to international conflicts and to climate change.

Issues of limitations on executive power, especially with respect to war, have been discussed, at least since the invention of the telegraph, in terms of the need for quick decisions. The Vietnam war, the Watergate scandal — neither led to any particularly extensive checks on, or built-in delays to the use of, executive power (though there is debate about the legality and effects of the War Powers Act). From the point of view of the model here, this is optimal if delays in changing policy in the face of increasing danger are very costly, if prompt action is a form of inoculation against future disasters in the context of spiking hazard rates. This seems a good description of many of the international relations problems the executive branch is called on to solve.

5.3. Our Methodology in Perspective. Our interest throughout has been focused on the optimal timing of discrete changes of action. There is a part of the semi-Markovian decision theory literature that studies continuously adjustable controls: Stone [1973] gives a differential inclusion condition for the value function when actions are continuously adjustable and control both transition times and distributions; Doshi [1979] generalized this work to models in which the states not only follow jump processes but may also have a diffusion-like component. The modeling tradeoffs are between discrete changes and their costs, which is our focus, versus continuous adjustability of a policy, which is theirs. This is reflected in the solutions methodologies: FOCs and more analytic tractability in our approach; the near necessity for numerical solutions characterized by differential inclusions in theirs.

An interesting perspective that sheds light on the approach taken here is in Ghosh and Saha [2012], who reduce sMdp's with continuously adjustable controls to imbedded discrete-time Markovian decision problems with compact action spaces representing the set of measurable time paths of the control variables. For stationary, discrete-time Markovian decision problems, in state z_k in the k 'th period, the value, $v(z_k)$, is the maximum, over available actions, $a_k \in A(z_k)$, of the expected one period reward, $r(a_k, z_k)$, plus the expected discounted (at rate δ) value of the state in one period, $\delta E(v(z_{k+1})|z_k, a_k)$. In our sMdp's: being in z_k in the k 'th period is replaced with being in the state-action pair (i_k, a_k) at the beginning of the k 'th stochastic interval, $\llbracket Y_k, Y_{k+1} \rrbracket$; the set of actions becomes the set of plans $\mathbf{p} \in \mathbb{P}(i_k, a_k)$; $r(a_k, z_k)$ becomes the expected payoffs over the interval $\llbracket Y_k, Y_{k+1} \rrbracket$; and the expected discounted value of the state in one period becomes the expected discounted value of $V^*(I_{k+1}, \mathbf{p}(Y_{k+1} - Y_k))e^{-rY_{k+1}}$ because $\mathbf{p}(Y_{k+1} - Y_k)$ is the action being used as Y_{k+1} arrives. Theorem 1 shows that one can reduce the problem to the sequence of problems encountered at the beginning of the intervals. This directly parallels Ghosh and Saha [2012, Theorem 3.1], who take the set of actions to be the set of progressively measurable, continuous time control paths as opposed to our $\mathbb{P}(i_k, a_k)$, a set of left-continuous piecewise constant functions.

5.4. Extensions and Further Research. In this paper we have assumed that the decision maker's knowledge of the stochastic structure can be expressed in terms of well-defined probabilities. One set of extensions that we are working on is to incorporate ambiguity into the model and analyze the implications. In real world problems like climate change, where lack of consensus about the exact stochastic structure is a key issue, such an approach seems more reasonable.

Economic models of climate change have traditionally treated the process as one of gradual change to new, stable state. Recent research in climate science has found evidence of both very rapid changes over a short period of time (around a decade), as well as significant fluctuations or 'environmental flickering' over periods as short as a year [Hall and Behl, 2006, Stern and Treasury, 2007]. These phases of rapid change and/or flickering seem to be triggered once a threshold point is reached in the ecological system. Quick changes in climate are more expensive to adapt to, and if we are in a situation where the arrival times of subsequent changes have a *range* of high arrival rates, the expense is further

increased, pushing policy recommendations in the direction of those that arise from some forms of the precautionary principle.¹⁸ While ‘gradual change’ models have usually prescribed ‘adaptation to climate change’ as opposed to ‘intervention to avert it’ [Nordhaus and Boyer, 2003], the decision problem takes a new shape when we incorporate the uncertainty over the expected arrival time of a possible catastrophic change and over the issue of whether or not we are moving towards such a critical threshold. In this context, the cost of controlling the intensity of the arrival rate process, and its interplay with ambiguity and ambiguity aversion need to be factored into the any discussion of optimal policy.

REFERENCES

- A. R. Amar. *America’s Constitution: A Biography*. Random House, 2006.
- David Blackwell. Discounted dynamic programming. *Ann. Math. Statist.*, 36: 226–235, 1965. ISSN 0003-4851.
- Svetlana I. Boyarchenko and Sergei Levendorskiĭ. *Irreversible Decisions Under Uncertainty: Optimal Stopping Made Easy*, volume 27. Springer Verlag, 2007.
- E. G. Carmines and J. A. Stimson. *Issue Evolution: Race and the Transformation of American Politics*. Princeton University Press, 1990.
- P. K. Chintagunta. Inertia and variety seeking in a model of brand-purchase timing. *Marketing Science*, 17(3):253–270, 1998.
- Jin Seo Cho and Halbert White. Testing for unobserved heterogeneity in exponential and Weibull duration models. *Journal of Econometrics*, 157(2):458–480, 2010.
- D. Corbae, M. Stinchcombe, and J. Zeman. *An Introduction to Mathematical Analysis for Economic Theory and Econometrics*. Princeton University Press, 2009.
- Peter Diamond. Unemployment, vacancies, wages. *American Economic Review*, 101(4):1045 – 1072, 2011.
- Bharat T. Doshi. Generalized semi-Markov decision processes. *J. Appl. Probab.*, 16(3):618–630, 1979. ISSN 0021-9002.

¹⁸While the ‘Precautionary Principle’ has many variants, it has a common core: for actions or policies with significant risk of long term or irrevocable harm to the public or to the environment, regulation is appropriate even before all of burden of proof that the risk is sensible falls on those taking the action. Though it is the basis of major environmental policy frameworks like the Kyoto Protocol, some versions of it have been criticized for not having proper analytical basis e.g. Sunstein [2002].

- Barry Eichengreen. The breakup of the Euro area. In *Europe and the Euro*, pages 11–51. University of Chicago Press, 2010.
- Zachary Elkins, Tom Ginsburg, and James Melton. *The Endurance of National Constitutions*. Cambridge University Press, 2009.
- William Feller. *An introduction to probability theory and its applications. Vol. II*. Second edition. John Wiley & Sons Inc., New York, 1971.
- Eleanor Flexner. *Century of Struggle: The Woman’s Rights Movement in the United States*. Harvard University Press, 1996.
- Mrinal K. Ghosh and Subhamay Saha. Optimal control of Markov processes with age-dependent transition rates. *Appl. Math. Optim.*, 66(2):257–271, 2012. ISSN 0095-4616.
- F. F. Gonul and F. Ter Hofstede. How to compute optimal catalog mailing decisions. *Marketing Science*, 25(1):65–74, 2006.
- Darwin C. Hall and Richard J. Behl. Integrating economic analysis and the science of climate instability. *Ecological Economics*, 57(3):442 – 465, 2006.
- Robert E. Hall. The 2010 Nobel Prize in Economics: How the DMP model explains current high unemployment. *Working paper, Department of Economics, Stanford*, 2012.
- Alexander Hamilton, James Madison, John Jay, and Roy P Fairfield. *The Federalist Papers: A Collection of Essays Written in Support of the Constitution of the United States: from the Original Text of Alexander Hamilton, James Madison, John Jay*. Johns Hopkins University Press, 1966.
- R. Hardin. *Liberalism, Constitutionalism, and Democracy*. Oxford University Press, USA, 2003.
- Kurt Hornik, Maxwell Stinchcombe, Halbert White, and Peter Auer. Degree of approximation results for feedforward networks approximating unknown mappings and their derivatives. *Neural Computation*, 6(6):1262–1275, 1994.
- Qiyang Hu and Wuyi Yue. *Markov Decision Processes with their Applications*, volume 14 of *Advances in Mechanics and Mathematics*. Springer, New York, 2008. ISBN 978-0-387-36950-1.
- Joseph T. McGuire and Joseph W. Kable. Decision makers calibrate behavioral persistence on the basis of time-interval experience. *Cognition*, 124(2):216–226, 2012.

- Joseph T. McGuire and Joseph W. Kable. Rational temporal predictions can underlie apparent failures to delay gratification. *Psychological Review*, 120(2): 395–410, 2013.
- Dale T. Mortensen. Markets with search friction and the DMP model. *American Economic Review*, 101(4):1073, 2011.
- W. D. Nordhaus and J. Boyer. *Warming the World: Economic Models of Global Warming*. MIT Press, 2003.
- Victor Ortego-Marti. Unemployment history and frictional wage dispersion. Technical report, Mimeo, Department of Economics, University of California Riverside, 2012.
- E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- Christopher A. Pissarides. Equilibrium in the labor market with search frictions. *American Economic Review*, 101(4):1092 – 1105, 2011.
- Marvin Rausand and Arnljot Høyland. *System Reliability Theory: Models, Statistical Methods, and Applications*, volume 396 of *Wiley series in probability and statistics. Applied probability and statistics*. Wiley Interscience, 2004.
- Andrew Rose. Checking out: exits from currency unions. *Journal of Financial Transformation*, 19:121–128, 2012.
- Sheldon M. Ross. *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco, Calif.-London-Amsterdam, 1970.
- Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, pages 25–49, 2005.
- J. M. Smith. The Republic of Letters: The correspondence between Thomas Jefferson and James Madison 1776-1826, 3 vols. *New York: WW Norton*, 3: 1364, 1995.
- N. H. Stern and G. B. Treasury. *The Economics of Climate Change: The Stern Review*. Cambridge University Press, 2007.
- Arthur L. Stinchcombe. *When Formality Works: Authority and Abstraction in Law and Organizations*. University of Chicago Press, 2001.
- N. L. Stokey. *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press, 2009.

Lawrence D. Stone. Necessary and sufficient conditions for optimal control of semi-Markov jump processes. *SIAM J. Control*, 11:187–201, 1973. ISSN 0363-0129.

C. R. Sunstein. The paralyzing principle. *Regulation*, 25(4):32–37, 2002.

6. APPENDIX

The arguments here give the essential ideas but suppress many details. These can be found in the online supplement.

Proof of Theorem 1: We show that the mapping T is well-defined, has the given contraction factor, that its unique fixed point is the true value function, V^* , and that following the argmax policy in (14) delivers V^* .

To appeal to Berge’s theorem to show that T is well-defined, we must, first, replace the correspondence $(i, a) \mapsto \mathbb{P}(i, a)$ with a continuous, d_H -compact value correspondence, $(i, a) \mapsto K_\tau(i, a) \subset \mathbb{P}(i, a)$ without changing the value function, and second, show that the function being maximized is jointly continuous.

Because d_H pays vanishingly small attention to differences on $[T, \infty)$ as $T \uparrow \infty$, the set of plans starting at any (i, a) with action change times at least $\tau > 0$ apart has at most $1 + T/\tau$ action changes by any T , hence is d_H -compact. We denote this set of plans $\mathbb{K}_\tau(i, a)$ and note that the correspondence is continuous. To see that we do not change the value function by this restriction, it is sufficient to show that there exists a sufficiently small $\tau \in (0, 1)$ such that it will never be optimal to change actions at times satisfying $t_k < t_{k+1} < t_k + \tau$. To this end, consider two plans, equal outside of $[t_k, t_{k+1}]$ with the following properties: $\mathbf{p}(2)$ changes twice, from a_{k-1} to a_k at t_k and from a_k to a_{k+1} at t_{k+1} ; and $\mathbf{p}(1)$ that skips the intermediate change, changing only once, from a_{k-1} to a_{k+1} at some $t \in [t_k, t_{k+1}]$. Because indirect changes of action are more costly than direct changes, the cost savings in $\mathbf{p}(1)$ are bounded away from 0 uniformly. Define $u^\circ = \max_{(i,a)} u(i, a) - \min_{(j,b)} u(j, b)$, and to avoid trivialities, assume that $u^\circ > 0$. The benefit difference is bounded by $u^\circ \cdot (t_{k+1} - t_k)$ plus the expected change in V resulting from the extra action. By uniform smoothness, both additional benefit terms go to 0 uniformly as $(t_{k+1} - t_k) \downarrow 0$.

From Blackwell’s Lemma for contraction mappings (e.g. Corbae et al. [2009, Lemma 6.2.33]), to show that T is a contraction mapping with contraction factor β , it is sufficient to show that (a) T is monotonic, and (b) that for any constant

$\kappa \geq 0$ and $V \in C(S \times A)$, $T(V + \kappa) \leq T(V) + \beta\kappa$. Monotonicity is immediate. For (b), define $\Psi(V, \mathbf{p}) = \left[\mathcal{R}(\mathbf{p}; (i, a)) + \int_{(0, \infty) \times S} V(i, \mathbf{p}(s)) e^{-rs} d\nu_{\mathbf{p}}(s, i) \right]$, let \mathbf{p}^* and \mathbf{p}^κ solve

$$\max_{\mathbf{p} \in \mathbb{P}(i, a)} \Psi(V, \mathbf{p}) \text{ and } \max_{\mathbf{p} \in \mathbb{P}(i, a)} \Psi(V + \kappa, \mathbf{p}) \quad (21)$$

so that $\Psi(V, \mathbf{p}^*) = T(V)$ and $\Psi(V + \kappa, \mathbf{p}^\kappa) = T(V + \kappa)$. By the optimality of \mathbf{p}^* , $\Psi(V + \kappa, \mathbf{p}^*) \leq \Psi(V + \kappa, \mathbf{p}^\kappa)$. Further,

$$\Psi(V + \kappa, \mathbf{p}^*) = \Psi(V) + \kappa \cdot \int_{(0, \infty) \times S} e^{-rs} d\nu_{\mathbf{p}^*}(s, i) \leq \kappa\beta \quad (22)$$

where both the last inequality and $\beta < 1$ come from uniform smoothness at $t = 0$.

The arguments for the last two parts of the theorem parallel the standard discrete-time dynamic programming arguments: if V_\dagger is the unique fixed point for T , then following the argmax \mathbf{p}^* for the first stochastic interval and receiving $V_\dagger(I_1, a)$ when the interval ends must have value V_\dagger . Therefore following \mathbf{p}^* for the first two stochastic intervals has value V_\dagger . Since the strictly positive W_k have expectation uniformly bounded away from 0, following \mathbf{p}^* in each interval yields V_\dagger . By definition, $V^* \geq V_\dagger$, and if $V^*(i_0, a_0) > V_\dagger(i_0, a_0)$ for some (i_0, a_0) , then V_\dagger is not the fixed point of T . \square

Proof of Theorem 2. At an optimal $0 < t_1 < \infty$, the difference between the payoff to changing at t_1 and changing at $t_1 + dt$, $dt \simeq 0$ must be equal to 0. There are four components to the difference: the difference in the annuitized flow cost of the action change, $-rc^*dt$ where c^* denotes $c(a_0, a_1^*; i_0)$; the difference in flow utility, $[u(i_0, a_1^*) - u(i_0, a_0)] dt$; the difference in the expected value of a state change net of the change cost c^* ,

$$\left[\left(\int_S V(i, a_1^*) \mu^+(t_1)(di) - c^* \right) h^+(t_1) - \left(\int_S V(i, a_0) \mu^-(t_1)(di) \right) h^-(t_1) \right] dt; \quad (23)$$

and the difference in the expected utilities accrued from the change in the size in the atom at ∞ conditional on the change not having arrived by t_1 ,

$$\left[\frac{\int_{t_1}^{t_1+dt} (f^- - f^+)(s) ds}{1 - F(t_1)} \right] u(i_0, a_1^*) \frac{1}{r} = \left[u(i_0, a_1^*) \frac{1}{r} \right] (h^-(t_1) - h^+(t_1)) dt. \quad (24)$$

Adding the terms, setting them equal to 0 and solving yields the FOCs

$$h^-(t_1^*) = \frac{rc^* - [u(i_0, a_1^*) - u(i_0, a_0)]}{\frac{1}{r} u(i_0, a_1^*) (1 - \gamma(t_1^*)) + [\gamma(t_1^*) \mathbb{E}^+ V(I_1, a_1^*) - \mathbb{E}^- V(I_1, a_0)] - c^*} \quad (25)$$

where $\gamma(t)$ is defined as the ratio of post- and pre-change hazard rates, $\gamma(t) = \frac{h^+(t)}{h^-(t)}$, and \mathbb{E}^- and \mathbb{E}^+ are defined as expectations of the next state, I_1 , under the distributions $\mu^-(t_1)$ and $\mu^+(t_1)$ respectively. The arguments for the SOCs and the sign of the derivative of the hazard rate directly parallel the arguments given in §2. □