Suspense and Surprise^{*}

Jeffrey Ely Northwestern University Alexander Frankel University of Chicago Emir Kamenica University of Chicago

October 2012

Abstract

We model demand for non-instrumental information, drawing on the idea that people derive entertainment utility from suspense and surprise. A period has more suspense if the variance of the next period's beliefs is greater. A period has more surprise if the current belief is further from the last period's belief. Under these definitions, we analyze the optimal way to reveal information over time so as to maximize expected suspense or surprise experienced by a Bayesian audience. We apply our results to the design of mystery novels, political primaries, casinos, game shows, auctions, and sports.

^{*}We would like to thank J.P. Benoit, Eric Budish, Sylvain Chassang, Eddie Dekel, Matthew Gentzkow, Ben Golub, Navin Kartik, Scott Kominers, Michael Ostrovsky, Devin Pope, Jesse Shapiro, Joel Sobel, and Satoru Takahashi for helpful comments. Jennifer Kwok, Ahmad Peivandi, and David Toniatti provided excellent research assistance. We thank the William Ladany Faculty Research Fund and the Initiative on Global Markets at the University of Chicago Booth School of Business for financial support.

1 Introduction

People frequently seek non-instrumental information. They follow international news and sports even when no contingent actions await. We postulate that a component of this demand for non-instrumental information is its entertainment value. It is exciting to refresh the *New York Times* webpage to find out whether Gaddafi has been captured and to watch a baseball pitch with full count and bases loaded.

In this paper, we formalize the idea that information provides entertainment and we analyze the optimal way to reveal information over time so as to maximize expected suspense or surprise experienced by a rational Bayesian audience. Our analysis informs two distinct sets of issues.

First, in a number of industries, provision of entertainment is a crucial service. Mystery novels, soap operas, sports events, and casinos all create value by revealing information over time in a manner that makes the experience more exciting. Of course, in each of these cases information revelation is bundled with other valuable features of the good – elegant prose, skilled acting, impressive athleticism, attractive waitstaff – but information revelation is the common component of these seemingly disparate industries. Moreover, entertainment is an important part of modern life. The American Time Use Survey reveals that adults in the United States spend roughly one fifth of their waking hours consuming entertainment (Aguiar, Hurst, and Karabarbounis (2011)).

Second, even if obtained for non-instrumental reasons, information can have substantial social consequences. Consider politics. As Downs (1957) has emphasized, the efficacy of democratic political systems is limited by voters' ignorance. This is particularly problematic because individual voters, who are unlikely to be pivotal, have little instrumental reason to obtain information about the candidates. Despite this lack of a direct incentive, many voters do in fact follow political news and watch political debates, thus becoming an informed electorate. A potential explanation is that the political process unfolds in a way that generates enjoyable suspense and surprise. Developing models of entertainment-based demand for non-instrumental information will thus inform the analysis of social issues, such as voting, that seem unrelated to entertainment.

In our model, there is a finite state space and a finite number of periods. The principal (the designer) reveals information about the state to the agent (the audience) over time. Specifically, the principal chooses the information policy: signals about the state, contingent on the current period and the current belief. The agent observes the realization of each

signal and forms beliefs by Bayes' rule. The agent has preferences over the stochastic path of his beliefs. A period generates more suspense if the variance of next period's beliefs is greater. A period generates more surprise if the current belief is further from last period's belief. The agent's utility in each period is an increasing function of suspense or surprise experienced in that period, and the principal seeks to maximize the expected undiscounted sum of per-period utilities.

To fix ideas, consider Figure 1 below. We plot the path of estimated beliefs about the winners of the 2011 US Open Semifinals over the course of these two tennis matches. Panel (a) shows Djoković versus Federer, and Panel (b) shows Murray versus Nadal.¹ The match between Djoković and Federer was exciting, with dramatic lead changes and key missed opportunities; Federer had multiple match points but went on to lose. In contrast, there was much less drama between Murray and Nadal, as Nadal dominated from the outset. Our model formalizes the notion that Federer-Djoković generated more suspense and more surprise than Murray-Nadal.





(a) Likelihood that Djoković beats Federer (b) Likelihood that Murray beats Nadal

We consider the problem of designing an optimal information policy, subject to a given prior and number of periods. We show that the suspense-optimal policy leads to decreasing residual uncertainty over time. Suspense is constant across periods and there is no variability in *ex post* suspense. This constant suspense is generated by asymmetric belief-swings – "plot twists" – which become both larger and less likely as time passes. The state is not fully revealed until the final period. One implication of our results is that most existing

¹We estimate the likelihood of victory given the current score using a simple model which assumes that the serving player wins the point with probability 0.64 (the overall fraction of points won by serving players in the tournament). The data and methodology are drawn from jeffsackmann.com.

sports cannot be suspense-optimal; we offer specific guidance on rules that would make sports more suspenseful.

Assuming additional structure, we also study the information policy that maximizes expected surprise. Under this policy, residual uncertainty may go up or down over time. Surprise in each period is variable, as is the *ex post* total surprise. In contrast to the suspense-optimal policy, when there are many periods the beliefs shift only a small amount in each period. There is a positive probability that the state is fully revealed before the final period. These features imply that a surprise-optimizing principal faces substantial commitment problems.

The previous results apply to the setting where the principal has no constraints on the technology by which she reveals information to the agent. We also briefly consider some specific constrained problems such as seeding teams in a tournament, determining the number of games in a finals series, and ordering sequential contests such as political primaries.

Our paper connects with four lines of research. It primarily contributes to the nascent literature on the design of informational environments. Kamenica and Gentzkow (2011) consider a static version of our model² where a principal chooses an arbitrary signal to reveal to an agent, who then takes an action that affects the welfare of both players. In that case, the principal has a value function over the agent's single-period posterior;³ our principal has a value function over the agent's multi-period belief martingale. In this sense our model shares features with Hörner and Skrzypacz (2011) who study a privately informed agent who is selling information to an investor. A scheme in which the agent sells information gradually maximizes the agent's ex ante incentives to acquire information.

In a broader sense, our paper also contributes to the literature on micro-foundations of preferences. Stigler and Becker (1977) advocate a research agenda of decomposing seemingly fundamental preferences into their constituent parts. This approach has been applied in a variety of settings, e.g., to demand for advertised (Becker and Murphy (1993)) and addictive (Becker and Murphy (1988)) goods. We apply it to drama-based entertainment. At first glance it may seem that the question of why one mystery novel is more enjoyable than another or the question of what makes a sports game exciting is outside the

 $^{^{2}}$ Brocas and Carrillo (2007), Rayo and Segal (2009), and Tamura (2012) examine special cases of this static model.

 $^{^{3}}$ A separate literature posits that agents have a direct preference for particular beliefs (e.g., Akerlof and Dickens (1982)); a number of papers analyze whether such preferences lead to demand for non-instrumental information (e.g., Eliaz and Spiegler (2006), Eliaz and Schotter (2010)).

purview of economics; such judgments may seem based on tastes that are inscrutable, like the preference for vanilla over chocolate ice cream. As our analysis reveals, however, reconceptualizing these judgments as being (partly) based on a taste for suspense and surprise reveals new insights about entertainment and demand for non-instrumental information more generally.

Third, our paper contributes to the literature on preferences over the timing of resolution of uncertainty. The original treatment of this subject goes back to Kreps and Porteus (1978). They axiomatize a representation of preferences for early or late resolution. Caplin and Leahy (2001) apply this framework to a setting where individuals have preferences over anticipatory emotions. They point out that agents might bet on their favorite team so as to increase the amount of suspense they will experience while watching a sports game. Kőszegi and Rabin (2009) and Dillenberger (2010) model agents who prefer one-shot rather than gradual revelation of information.

Finally, there is a small formal literature on suspense and surprise *per se.* Chan, Courty, and Li (2009) define suspense as valuing contestants' efforts more when the game is close and demonstrate that preference for suspense increases the appeal of rank-order incentive schemes over linear ones. The definition of surprise of Geanakoplos (1996) is similar to ours. He considers a psychological game (Geanakoplos, Pearce, and Stacchetti (1989)) where a principal wants to surprise an agent. Specifically, he examines the Hangman's Paradox (Gardner (1969)), the problem of choosing a date on which to hang a prisoner so that the prisoner is surprised. In his setting, the principal has no commitment power, so surprise is not possible in equilibrium. Borwein, Borwein, and Maréchal (2000) give the principal commitment power, and derive the surprise-optimal probabilities for hanging at each date.⁴

The remainder of the paper is structured as follows. Section 2 presents the model. Section 3 discusses the interpretation of the model. Section 4 and Section 5 derive the suspense- and surprise-optimal information policies. Section 6 compares these policies. Section 7 considers constrained problems. Section 8 concludes.

⁴Let p_t denote the conditional probability of being hanged on day t, given that the prisoner is still alive. They define surprise in period t as the entropy $(-\log p_t)$ if the prisoner is hanged, and 0 otherwise.

2 A model of suspense and surprise

We develop a model in which a principal reveals information over time about the state of the world to an agent.

2.1 Preferences, beliefs, and technology

There is a finite state space Ω . A typical state is denoted ω . A typical belief is denoted by $\mu \in \Delta(\Omega)$; μ^{ω} designates the probability of ω . The prior is μ_0 . Let $t \in \{0, 1, ..., T\}$ denote the period.

A signal π maps the state to a distribution over a finite signal realization space S, i.e., $\pi : \Omega \to \Delta(S)$. An information policy $\tilde{\pi}$ is a function that maps the current period and the current belief into a signal. Let $\tilde{\Pi}$ denote the set of all information policies.⁵

Any information policy generates a stochastic path of beliefs about the state. A belief martingale $\tilde{\mu}$ is a sequence $(\tilde{\mu}_t)_{t=0}^T$ s.t. (i) $\tilde{\mu}_t \in \Delta(\Delta(\Omega))$ for all t, (ii) $\tilde{\mu}_0$ is degenerate at μ_0 , and (iii) $E[\tilde{\mu}_t \mid \mu_0, ..., \mu_{t-1}] = \mu_{t-1}$ for all $t \in \{1, ..., T\}$. A realization of a belief martingale is a belief path $\eta = (\mu_t)_{t=0}^T$. Given the current belief μ_t , a signal induces a distribution of posteriors $\tilde{\mu}_{t+1}$ s.t. $E[\tilde{\mu}_{t+1}] = \mu_t$. Hence, an information policy induces a belief martingale. We denote the belief martingale induced by information policy $\tilde{\pi}$ (given the prior μ_0) by $\langle \tilde{\pi} \mid \mu_0 \rangle$.

There are two players: the agent and the principal. The agent has preferences over his belief path and the belief martingale. The agent has a *preference for suspense* if his utility function is

$$U_{\text{susp}}(\eta,\tilde{\mu}) = \sum_{t=0}^{T-1} u \left(E_t \sum_{\omega} \left(\tilde{\mu}_{t+1}^{\omega} - \mu_t^{\omega} \right)^2 \right)$$

for some increasing function $u(\cdot)$ with u(0) = 0. An agent has a *preference for surprise* if her utility function is

$$U_{\text{surp}}\left(\eta\right) = \sum_{t=1}^{T} u\left(\sum_{\omega} \left(\mu_{t}^{\omega} - \mu_{t-1}^{\omega}\right)^{2}\right)$$

for some increasing function $u(\cdot)$ with u(0) = 0. So suspense is induced by variance over the next period's beliefs, and surprise by change from the previous belief to the current

⁵The assumption that S is finite is innocuous since Ω is finite. The assumption that $\tilde{\pi}$ depends only on the current belief and period, rather than the full history, is without loss of generality in the sense that memoryless policies do not restrict the set of feasible outcomes.

one.

We will frequently focus on the *baseline specification* where suspense is the standard deviation of $\tilde{\mu}_{t+1}$ (aggregated across states) and surprise is the Euclidean distance between μ_t and μ_{t-1} . This corresponds to setting $u(x) = \sqrt{x}$.

The principal chooses the information policy to maximize the agent's expected utility. If the agent has a preference for suspense, the principal solves

$$\max_{\tilde{\pi}\in\tilde{\Pi}} E_{\langle\tilde{\pi}\mid\mu_0\rangle} U_{\text{susp}}\left(\eta,\langle\tilde{\pi}\mid\mu_0\rangle\right).$$

If the agent has a preference for surprise, the principal solves

$$\max_{\tilde{\pi}\in\tilde{\Pi}}E_{\langle\tilde{\pi}|\mu_{0}\rangle}U_{\mathrm{surp}}\left(\eta\right)$$

Note that the choice of the information policy affects the value of the objective function only through the belief martingale it induces. The additional details of the information policy are irrelevant for payoffs. Thus, it is convenient to recast the optimization problem as a direct choice of the belief martingale. An extension of the argument in Proposition 1 of Kamenica and Gentzkow (2011) shows that any belief martingale can be induced by some information policy.⁶

Lemma 1. Given any belief martingale $\tilde{\mu}$, there exists an information policy $\tilde{\pi}$ such that $\tilde{\mu} = \langle \tilde{\pi} \mid \mu_0 \rangle$.

In some settings the set of feasible information policies might be limited by tradition, complexity, or other institutional constraints. For example, the organizer of a tournament may have settled on an elimination format and is choosing between seeding procedures. Or a political party respects the rights of states to choose their own delegates, but may have control over the order in which states hold elections. Accordingly, let $P \subset \tilde{\Pi} \times (\Delta \Omega) \times \mathbb{N}$ be a subset of the Cartesian product of information policies, priors, and durations. In Section 7 we consider a constrained model in which the principal chooses $(\tilde{\pi}, \mu_0, T) \in P$ so as to maximize expected suspense or surprise.

⁶Kamenica and Gentzkow show in a static model that, when current belief is μ_t , any distribution of posteriors $\tilde{\mu}_{t+1}$ with mean μ_t can be induced by some signal. Applying this result period-by-period yields our Lemma 1.

2.2 Extensions

There are some natural extensions to our specification of the agent's utility function.

First and most obvious, the audience might value both suspense and surprise. While we cannot fully characterize the optimal information policy for such preferences, we discuss the tradeoff between suspense and surprise in Section 6.

Second, the audience would presumably have a distinct preference for learning the outcome by the end, i.e., from having μ_T degenerate. Including this term in the utility function would not have any effect on our results since any suspense- or surprise-maximizing policy reveals the state by the last period.

Third, the audience may experience additional utility from the realization of a particular state. For example, a sports fan enjoys games that are exciting, but is particularly happy should her team win. We can incorporate this in our model by supposing that the overall utility is a sum of the utility from suspense/surprise and a utility from the final belief being degenerate at some state. Such a modification does not affect the optimal policy; it only changes the payoff-maximizing prior.

Fourth, one may consider models with *state-dependent significance* where the audience cares more about changes in the likelihood of particular states. For instance, the reader of a mystery novel may be in great suspense about whether the protagonist committed the murder. But in the event of the protagonist's innocence, she cares less about whether the murderer was Stooge A or Stooge B. Or, if the New York Yankees have five times the audience of the Milwaukee Brewers, the league may value suspense/surprise about the Yankees' championship prospects five times as much as suspense about the Brewers. We can accordingly modify suspense utility to be

$$U_{\text{susp}}(\eta,\tilde{\mu}) = \sum_{t=0}^{T-1} u\left(E_t \sum_{\omega} \alpha^{\omega} \cdot \left(\tilde{\mu}_{t+1}^{\omega} - \mu_t^{\omega}\right)^2\right),$$

and likewise for surprise. More important characters or larger market sports teams have a larger state-dependent weight $\alpha^{\omega} \geq 0$.

Fifth, the audience might weigh suspense and surprise differently across periods. For example, later periods might be weighted more heavily if the reader becomes more invested in the characters as she advances through the novel. In models with *time-dependent* significance, we replace suspense utility with

$$U_{\text{susp}}(\eta,\tilde{\mu}) = \sum_{t=0}^{T-1} \beta_t u \left(E_t \sum_{\omega} \left(\tilde{\mu}_{t+1}^{\omega} - \mu_t^{\omega} \right)^2 \right),$$

and again likewise for surprise. A period in which the audience is more involved has a higher value of $\beta_t \ge 0$. In Section 4, we discuss the suspense-optimal information policies in cases of state- or time-dependent significance.

Finally, sometimes the agent might be invested only in some aspect of her belief such as the expectation. Consider a gambler who plays a sequence of fair gambles and experiences suspense (or surprise) when her expected earnings from the visit are about to change (or just did). Formally, ω is a bounded random variable and the agent has preferences over the path of its expectation; e.g., in case of surprise, agent's utility in period t is $u \left(E_{\mu_t} [\omega] - E_{\mu_{t-1}} [\omega]\right)^2$. While we will not discuss this extension at length, all of our results apply in this case as well.

3 The interpretation of the model

Before we proceed with the derivation of optimal information policies, we discuss some of the potential interpretations of the model above.

3.1 Interpretation of the technology

Suppose that the principal is a publishing house and the agent is a reader of mystery novels. In this case, a writer is associated with a belief martingale, and a particular book by writer $\tilde{\mu}$ is a belief path η drawn from $\tilde{\mu}$. For a concrete example, say that Mrs. X is a writer and all of her books follow a similar premise. In the opening pages of the novel, a dead body is found at a remote country house where n guests and staff are present. In every novel by Mrs. X, one of these n individuals single-handedly committed the murder.⁷ The opening pages establish a prior μ_0 over the likelihood that each individual $\omega \in \Omega$ is the culprit. There are then T chapters each revealing some information about the identity

⁷In Mrs. X's novels, it is never the case that the murderer is someone the reader has not been introduced to at the outset. This allows us to model uncertainty in a classical way without addressing issues of unawareness. We could easily allow, however, for the possibility that the murder was really a suicide (change n to n + 1) or, as in *Murder on the Orient Express*, that (spoiler alert) several of the suspects jointly committed the murder (change n to $2^n - 1$).

of the murderer. A literal (though perhaps not very literary) interpretation is that Mrs. X explicitly randomizes the plot of each chapter based on her information policy and her current belief (based on the content she has written thus far) and learns whodunnit only when she completes the novel.⁸

Alternatively, the principal is the rule-setting body of a sports association and the agent is a sports fan. In this case, we associate the feasible set of rules with some constraint set P, the chosen rule with an information policy $\tilde{\pi}$, a match-up with a prior and a belief martingale $\langle \tilde{\pi} | \mu_0 \rangle$, and a match with a belief path η . To see how modifying the rules changes the information policy, note that if it becomes more difficult to score as players get tired, rules that permit fewer substitutions increase the amount of information that is revealed in the earlier periods. Or, if it is easier to score when fewer players are on the field, rules that lower the threshold for issuing red cards may reveal more information later in the game. The rules of a sport also affect priors and belief martingales through the players' strategic responses to such rules. For instance, actions with low expected value but high variance may be played at the end of the game but will never be played at the beginning, e.g., 2-point conversions in football. Allowing such actions can alter the relative amount of information revealed early versus late. Different rules might also induce different priors. For instance, a worse tennis player will have a higher chance of winning under the tie-break rule for deciding sets.

Our model captures settings both where the state is realized *ex ante* and those where it is realized *ex post*. An example where ω is realized *ex ante* is a game show where a contestant receives either an empty or a prize-filled suitcase, and then information is slowly revealed about the suitcase's contents. In presidential primaries, on the other hand, ω is realized *ex post*. Whichever candidate gets more than 50% of the overall delegates wins the nomination. When a candidate wins a state's delegates, this outcome provides some information about whether she will win the nomination. In this case the state ω is not an aspect of the world that is fixed at the outset; it is determined by the signal realizations themselves.

Implicit in our formal structure is the assumption that the principal and the agent share

⁸This interpretation brings to mind the notion of "willful characters." For example, novelist Jodi Picoult writes, "Often, about 2/3 of the way through, the characters will take over and move the book in a different direction. I can fight them, but usually when I do that the book isn't as good as it could be. It sounds crazy, but the book really starts writing itself after a while. I often feel like I'm just transcribing a film that's being spooled in my head, and I have nothing to do with creating it. Certain scenes surprise me even after I have written them - I just stare at the computer screen, wondering how that happened."

a common language for conveying the informational content of a signal. For example, when the butler is found with a bloody glove in chapter 2 of Mrs. X's mystery novel, the reader knows to update his beliefs on the butler's guilt from (say) 27% to 51%. This assumption goes hand in hand with the requirement that beliefs are a martingale. If the reader believes that there is a 90% chance that the butler did it, then the final chapter must reveal the butler's guilt 90% of the time. The principal is constrained by the agent's rationality.

Like all writers of murder mysteries, Mrs. X faces a commitment problem. After giving a strong indication that the butler was the murderer, in the last chapter she may want to reveal that it was the maid. This would be very surprising. If rational readers expect Mrs. X to play tricks of this sort, though, they will not believe any early clues and thus their beliefs will not budge from the prior until the very last paragraph.

In our model Mrs. X can in fact send meaningful signals to the audience. This is because Mrs. X can commit to follow her information policy even when doing so results in lower *ex post* utility for the reader. She can commit to a small probability of plot twists on the final page, even if every novel with such a twist is more exciting. If her publisher refused to publish those boring books without plot twists, readers would find her remaining books less surprising. Note that this commitment problem is less of an issue in sports. The players involved want to win, so a team with a dominating lead won't slack off just to make the game more exciting. As we shall see, the commitment to allow *ex post* boring realizations is necessary for maximizing surprise but not suspense.

Finally, we consider a single principal who provides entertainment. In many settings, however, there are multiple entities who vie for the attention of the audience.⁹ Full analysis of such settings is beyond the scope of this paper, but we suspect that competition is likely to exacerbate commitment problems: the pressure to produce works that are more exciting than those of a competitor might induce the elimination of boring belief paths, which would in turn reduce the overall entertainment supplied in equilibrium.

3.2 Interpretation of the preferences

Under our definition, a moment is laden with suspense if some crucial uncertainty is about to be resolved. Suppose a college applicant is about to open an envelope that informs her whether she was accepted to her top-choice school, a soccer player steps up to take a free

⁹Ostrovsky and Schwarz (2010) and Gentzkow and Kamenica (2011) consider (static) environments where multiple senders independently choose signals in an attempt to influence a decision maker.

kick, or a pitcher faces a full-count with bases loaded. These situations seem suspenseful, and the key feature is that the belief about the state of the world (did she get in, which team will win) is about to change.¹⁰

For the purpose of aggregating suspense over multiple periods, it seems most plausible to assume that $u(\cdot)$ is strictly concave. Consider two mystery novels, both of which open with the same prior μ_0 and reveal the murderer by period T. Novel A slowly reveals clues over time, generating a suspense of say x in each period. The total suspense payoff from novel A is Tu(x). In novel B (for boring), nothing happens at all in the first third of the book, then the murderer is announced in a single chapter, and after that nothing at all happens again until the end of the book. This generates a suspense payoff of u(Tx). Taking novel A to be a better, more suspenseful novel amounts to assuming that $u(\cdot)$ is strictly concave.

We say that a moment generated a lot of surprise if the agent's belief just moved by a large amount.¹¹ Suppose our college applicant unexpectedly receives a letter rescinding the previous acceptance letter which had been mailed by mistake, or a soccer player scores a winning goal from 60 yards away in the final moments of the game. These events seem to generate surprise, and the key feature is that the belief about the state of the world changed dramatically.

Note that suspense is experienced *ex ante* whereas surprise is experienced *ex post*. There is another crucial distinction between the two concepts. The overall surprise depends solely on the belief path realized. In contrast, suspense depends on the belief martingale as well as the belief path. Recall Figure 1. The realized belief paths fully determine the surprise, but not the suspense, generated by each match. Suspense at a given point depends on the entire distribution of next period's beliefs. In Figure 2, we illustrate this distribution by adding gray tendrils that indicate what the belief would have been had the point turned out differently. Figure 2 makes it apparent that Djoković-Federer was a more suspenseful match than Murray-Nadal.

As we mentioned in the introduction, we believe suspense and surprise are important in many contexts.¹² Sports fans enjoy the drama of the shifting fortunes between players.

¹⁰Neurobiological evidence suggests that, in monkeys, suspense about whether a reward will be delivered induces sustained activation of dopamine neurons (Fiorillo, Tobler, and Schultz (2003)).

¹¹Kahneman and Miller (1986) offer a different conceptualization of surprise based on the notion of endogenously generated counterfactual alternatives.

¹²The applicability of our model probably depends on the stakes. On the one hand, when the stakes are too low, there is little scope for suspense and surprise since the agent is not invested in the outcome.



Figure 2: 2011 US Open Semifinals, Suspense and Surprise The gray lines indicate what the belief would have been, if the point had gone the other way.

(a) Likelihood that Djoković beats Federer

(b) Likelihood that Murray beats Nadal

Playing blackjack at the casino, a gambler knows the odds are against her but derives pleasure from the ups and downs of the game itself. Politicos and potential voters enjoy following the news when there is an exciting race for political office such as the 2008 Clinton-Obama primary. Figure 3 presents data on the belief movements in each of these settings.

For *Tennis*, we collected point-level data for every match played in Grand Slam tournaments in 2011. We focus exclusively on women's tennis in order to have a fixed bestout-of-three-sets structure. Each belief path is associated with a particular match. We estimate the likelihood that a given player will win the match by assuming that the serving player wins any given point with a fixed tournament-specific probability. This gives us, at each point, the likelihood of a win (thick blue line) and what this likelihood would have been had the point played out differently (thin gray tendrils). Hence, we can compute the suspense and surprise realized in each match.

For *Soccer*, we collected data on over 24,000 matches played between August 2011 and July 2012 across 67 professional leagues. We exclude knockout competitions where matches can end in overtime or penalty shootouts, so each match lasts approximately 94 minutes (inclusive of stoppage time). To parallel the binary state space in the rest of Figure 3, we focus on the likelihood that the home team will win (vs. tie/lose). We estimate this likelihood minute-by-minute, based on a league-specific hazard rate of goal scoring, computed separately for the home and away teams. For each minute we also

On the other hand, if the stakes are too high, the agent will not be in the mood for entertainment. These considerations are particularly important when stakes are endogenous (e.g., in a gambling context).



Figure 3: Suspense and surprise from a variety of processes

estimate the beliefs that would have realized if the home team, away team, or neither team had scored. Thus, as for tennis, we can compute the suspense and surprise generated by each soccer match.

For *Blackjack*, we simulate 20,000 visits to Las Vegas. On every visit, our artificial gambler begins with a budget of \$100 and plays \$10-hands of blackjack until he either increases his stack to \$200 or loses all his money. Each belief path is associated with a particular visit. The dealer's behavior in blackjack is regulated and our gambler strictly follows optimal (non-card-counting) play. Hence, following each individual card that is dealt, we can explicitly compute the updated probability that the gambler will walk away with the winnings rather than empty-handed. Also, we can determine what this probability would have been for every other card that could have been dealt. This allows us to determine the suspense and surprise realized in each visit. Moreover, we can use these simulations to examine how a given change in casino rules would influence the distribution of suspense and surprise. For example, we confirm that allowing doubling down and splitting, as all Vegas casinos do, indeed increases expected suspense and surprise.¹³

For *Clinton-Obama Primary*, we depict the daily average price of a security that pays out if Obama wins the 2008 Democratic National Convention. There is a single belief path for this particular primary. Unlike for *Tennis* and *Blackjack*, these data do not provide a way to estimate the distribution of next day's beliefs. Hence, we are able to compute realized surprise but not suspense.

For comparison we also draw 1,000 belief paths from the suspense- and surprise-optimal martingales.

In the right panel of Figure 3, we display the scatterplot of realized suspense and surprise for each setting.¹⁴ We identify the observation that is closest to the median level of suspense and median level of surprise. We mark this observation with a black circle in the right panel and depict its belief path in the left panel. Additional details about the construction of Figure 3 are in the Online Appendix.

These examples illustrate some of the ways that belief martingales can be estimated.

¹³Doubling down means the gambler is allowed to increase the initial bet by \$10 after his first two cards in exchange for committing to receive exactly one more card. Splitting means that, if the first two cards have the same value, the player can split them into two hands and wager an additional \$10 for the new hand.

¹⁴We utilize the baseline specification for $u(\cdot)$ and normalize realized suspense and surprise across settings by dividing by the square root of the number of periods. (As we show later, maximal suspense and surprise are proportional to \sqrt{T} .)

For *Blackjack*, it is possible to simulate the distribution of belief paths based solely on the structure of the rules; the data generating process is known. For the *Primary*, we use data from a prediction market. For *Tennis* and *Soccer*, we estimate the likelihood of the relevant events (server will win the point, home team will score, etc.) in each period and derive the belief path implied.¹⁵ Additionally, belief martingales might be elicited through laboratory experiments. For instance, a researcher could pay subjects to read a mystery novel and incentivize them to guess who the murderer is after each chapter. We hope that in future research this range of methods will allow for construction of datasets on suspense and surprise in a number of other contexts.

3.3 Illustrative examples

To develop basic intuition about the model, we consider a few examples where the principal's problem can be analyzed without any mathematics.

Suppose a principal wishes to auction a single object to bidders who have independent private values but also enjoy suspense or surprise about whether they will win the auction. The principal must choose between the English auction and the second-price sealed-bid auction. Conditional on standard bidding behavior, the English auction is preferable; it reveals information about the winner slowly rather than all at once, so it gives bidders a higher entertainment payoff.

Or, consider elimination announcements on a reality TV show. In each episode, one of two contestants, say Scottie or Haley, gets eliminated. The host of the show calls out one of the names, e.g., "Scottie, please step forward." Then, she either says "You are eliminated" or "You go through" to the person who was called. If the host always called the person who was to be eliminated, or the person who would go through, then the outcome would be revealed immediately when she asked Scottie to step forward. If the host chose participants without regard to elimination, then calling Scottie to step forward would convey no information at all. The second comment would reveal everything. To increase suspense or surprise, the host should make the initial call partially informative. Having called Scottie forward, the audience should believe that Scottie is now either more likely to go through or more

¹⁵Our estimation procedure for both *Tennis* and *Soccer* is admittedly crude, but it serves to illustrate a method for deriving belief paths. For *Tennis*, if we had more data we could estimate the likelihood a given player wins a given point conditional on the surface, the current score, the recent change in score, the players' rankings, etc. Or, one could directly estimate the likelihood a given player wins the entire match given these factors. Similar considerations apply to *Soccer*.

likely to be eliminated. Either policy works, as long as the audience understands how to interpret the signal. Anecdotally, many reality TV shows seem to follow this formula.

Economists and psychologists have extensively studied why people gamble – why they spend money in casinos. Our model gives a rationale for why people spend *time* in casinos. The weekend's monetary bets (a compound lottery) could be reduced to a single bet (a simple lottery). But this would deprive the gambler of an important element of the casino experience. Part of the fun of gambling is the suspense and surprise as the gambler anticipates and then observes each flip of a card, spin of the wheel, or roll of the dice.¹⁶

These three examples are specific instances of a more general feature of suspense and surprise: spoilers are bad. In other words, revealing all the information at once (as a spoiler does) yields the absolute minimum suspense and surprise (given that all information is revealed by the end). This feature of the preferences seems in accordance with real-world intuitions about suspense and surprise.¹⁷

Finally, when watching a sports match between two unfamiliar teams, spectators commonly root for the underdog or the trailing team. Our model gives an explanation for this kind of behavior. There is more surprise when the underdog wins. And when the trailing team scores, we have a closer match which generates more continuation suspense and surprise.

4 Suspense-Optimal Information Policies

In this section, we take the prior and the number of periods as given and derive the information policy that maximizes expected suspense. As we discussed above, any plausible utility function $u(\cdot)$ over suspense is strictly concave. Accordingly, this is the case we focus on in this section.¹⁸

¹⁶Barberis (2012) considers a model where gamblers behave according to prospect theory. Casinos then offer dynamic gambles so as to exploit the time inconsistency induced by non-linear probability weighting.

¹⁷Christenfeld and Leavitt (2011) claim that readers in fact prefer spoilers. However, they obtain this result only when spoilers are announced rather than embedded within the text, which raises concerns about experimenter demand effects.

¹⁸If $u(\cdot)$ is linear, then any information policy that is fully revealing by the end yields the same overall suspense. If $u(\cdot)$ is strictly convex, revealing the state in any single period is uniquely optimal.

4.1 Solving the Principal's Problem

Recall that given any belief martingale, there exists an information policy which induces it (Lemma 1). So we can think of the principal's choice of an information policy as equivalent to the choice of a martingale. To simplify notation, let the aggregated variance of period t + 1 beliefs, given information at time t, be denoted by

$$\sigma_t^2 \equiv E \sum_{\omega} \left(\tilde{\mu}_{t+1}^{\omega} - \mu_t^{\omega} \right)^2.$$

We can then write the principal's problem as maximizing $E_{\tilde{\mu}} \sum_{t=0}^{T-1} u(\sigma_t^2)$.

We begin by making two observations. First, any suspense-optimal martingale will be fully revealing by the end, i.e., the final belief μ_T is degenerate. To see this, take some policy that does not always fully reveal and modify the last signal to be fully informative. This increases σ_{T-1}^2 and leaves σ_t^2 unchanged at t < T - 1, raising suspense.

Second, all martingales that are fully revealing by the end yield the same expected sum of variances, $E_{\langle \tilde{\pi} | \mu_0 \rangle} \sum_{t=0}^{T-1} \sigma_t^2$. This follows from the fact that martingale differences are uncorrelated. For any collection of uncorrelated random variables, the sum of variances is equal to the variance of the sum. Hence, any fully revealing policy $\tilde{\pi}$ yields the same $E_{\langle \tilde{\pi} | \mu_0 \rangle} \sum_{t=0}^{T-1} \sigma_t^2$ as the policy that reveals the state in a single period.¹⁹ The same logic holds, going forward, from any current belief μ at any period. We denote this *residual variance* from full revelation by $\Psi(\mu) \equiv \sum_{\omega} \mu^{\omega} (1 - \mu^{\omega})$.

We summarize these two points in the following lemma.

Lemma 2. Any suspense-maximizing information policy is fully revealing by the end. Under any information policy $\tilde{\pi}$ that is fully revealing by the end, $E_{\langle \tilde{\pi} | \mu_0 \rangle} \sum_{t=0}^{T-1} \sigma_t^2 = \Psi(\mu_0)$.

Starting from the prior, the principal can thus be thought of as having a "budget of variance" equal to $\Psi(\mu_0)$. The principal then decides how to allocate this variance across periods so as to maximize $E_{\tilde{\mu}} \sum_t u(\sigma_t^2)$ subject to $E_{\tilde{\mu}} \sum_t \sigma_t^2 = \Psi(\mu_0)$. By concavity of $u(\cdot)$, it would be ideal to dole out variance evenly over time. Is it possible to construct an information policy so that σ_t^2 is equal to $\Psi(\mu_0)/T$ in each period t, on every path? If so,

¹⁹Augenblick and Rabin (2012) also point out this feature of belief martingales; they use it to construct a test of Bayesian rationality. Residual variance also plays a role in insider trading models (Kyle (1985), Ostrovsky (2012)) where it captures how much of the insider's private information has not yet been revealed to the market. This bounds the total variation in future prices which in turn places an upper bound on the insider's profits.

such a policy would be optimal. In fact, we can construct such a policy. Let

$$M_t \equiv \left\{ \mu \mid \Psi(\mu) = \frac{T-t}{T} \Psi(\mu_0) \right\}.$$

Proposition 1. Fix any strictly concave $u(\cdot)$. A belief martingale maximizes expected suspense if and only if $\mu_t \in M_t$ for all t. The agent's expected suspense from such a policy is $Tu\left(\frac{\Psi(\mu_0)}{T}\right)$.

Proof. A martingale $\tilde{\mu}$ that is fully revealing by the end has σ_t^2 constant across t if and only if $\Psi(\mu_t) = \frac{T-t}{T}\Psi(\mu_0)$, or in other words $\mu_t \in M_t$, for all t. We are going to show that in fact there exists a martingale $\tilde{\mu}$ with $\mu_t \in M_t$ for all t (which therefore has constant σ_t^2). Then by Lemma 1, we know that there exists a policy $\tilde{\pi}$ s.t. $\langle \tilde{\pi} | \mu_0 \rangle = \tilde{\mu}$. It follows that $\tilde{\mu}$ is optimal, proving the sufficiency part of the Proposition. Any martingale with non-constant σ_t^2 gives a lower payoff, establishing necessity.

In general, given any sequence of sets (X_t) , there exists a martingale $\tilde{\mu}$ s.t. Supp $(\tilde{\mu}_t) \subset X_t \ \forall t \ \text{if} \ X_t \subset \operatorname{conv}(X_{t+1}) \ \forall t$. Therefore, it remains to show that $M_t \subset \operatorname{conv}(M_{t+1}) \ \forall t$. By definition of $\Psi(\cdot)$, for any $k \ge 0$ we have $\operatorname{conv}(\{\mu \mid \Psi(\mu) = k\}) = \{\mu \mid \Psi(\mu) \ge k\}$. Hence, $M_t = \{\mu \mid \Psi(\mu) = \frac{T-t}{T}\Psi(\mu_0)\} \subset \{\mu \mid \Psi(\mu) \ge \frac{T-(t+1)}{T}\Psi(\mu_0)\} = \operatorname{conv}(M_{t+1})$.

This proposition provides a recipe for constructing a suspense-optimal information policy. At the outset, the belief μ_0 is contained in M_0 . In each period t, given $\mu_t \in M_t$, the principal chooses a signal that induces a distribution of beliefs whose support is contained in the set M_{t+1} . Any distribution of beliefs $\tilde{\mu}_{t+1}$ can be induced by some signal so long as $E_{\tilde{\mu}_{t+1}}[\mu_{t+1}] = \mu_t$; this includes distributions with support in M_{t+1} , as shown in the proof of Proposition 1. In particular, given $\tilde{\mu}_{t+1}$, let $\pi(s \mid \omega) = \frac{\mu_s^{\omega} \tilde{\mu}_{t+1}(\mu_s)}{\mu_t^{\omega}}$. By Bayes' Rule, this signal induces $\tilde{\mu}_{t+1}$.

There is a natural geometric interpretation of the set M_t which sheds further light on the structure of suspense-optimal information policies. The set M_t is defined as those beliefs with residual variance $\frac{T-t}{T}\Psi(\mu_0)$. Following some simple algebra, we can equivalently characterize each M_t as a "circle" (a hypersphere) centered at the uniform belief. Denoting the uniform belief by $\mu_* \equiv \left(\frac{1}{|\Omega|}, ..., \frac{1}{|\Omega|}\right)$, we can write M_t as

$$M_t = \left\{ \mu \mid |\mu - \mu_*|^2 = |\mu_0 - \mu_*|^2 + \frac{t}{T} \Psi(\mu_0) \right\}$$

where $|\mu - \mu_*|^2 = \sum_{\omega} (\mu^{\omega} - \mu_*^{\omega})^2$ denotes the square of the Euclidean distance between μ

and μ_* . The uniform belief μ_* has the highest residual variance of all beliefs, and residual variance falls off with the square of the distance from μ_* . The residual variance of beliefs in M_t falls linearly in t, and hence beliefs lie on circles whose radius-squared increases linearly over time. At time T, the circle has the maximum radius and intersects $\Delta(\Omega)$ only at degenerate beliefs. Figure 4 illustrates this geometric characterization of suspense-optimal information policies.

Figure 4: The path of beliefs with $\Omega = \{0, 1, 2\}$



The triangle represents $\Delta(\Omega)$, the two-dimensional space of possible beliefs. The M_t sets are circles centered on the uniform belief μ_* , intersected with the triangle $\Delta(\Omega)$. The belief begins at μ_0 ; in this picture μ_0 is at μ_* . The belief μ_t at time t will be in M_t . Given current belief $\mu_t \in M_t$, any distribution $\tilde{\mu}_{t+1}$ over next-period beliefs with mean μ_t and support contained in M_{t+1} is consistent with a suspense-maximizing policy. At time T the uncertainty is resolved, so μ_T will be on a corner of the triangle.

Below we summarize some of the key qualitative features of suspense-optimal information revelation.

The state is revealed in the last period, and not before. As long as the prior is not degenerate, the residual variance is positive at any time t < T. In particular, the

residual variance at time t is $\frac{T-t}{T}\Psi(\mu_0)$.

- Uncertainty declines over time. Uncertainty, as measured by the residual variance, declines linearly over time from $\Psi(\mu_0)$ to 0.
- **Realized suspense is deterministic.** There is no *ex post* variation in suspense. Although the path of beliefs is random, the agent's suspense is the same across every realization. It is always exactly $Tu(\Psi(\mu_0)/T)$.
- Suspense is constant over time. The variance σ_t^2 in each period t is $\Psi(\mu_0)/T$. So the agent's experienced suspense is $u(\Psi(\mu_0)/T)$ in each period.
- The prior that maximizes suspense is the uniform belief. Total suspense is $Tu(\Psi(\mu_0)/T)$. The budget of variance $\Psi(\mu_0)$ is increasing in the proximity of μ_0 to μ_* ; $\Psi(\mu_0)$ is maximized at $\mu_0 = \mu_*$.
- The level of suspense increases in the number of periods T. It is immediate from the outset that suspense must be weakly increasing in T – any signals that can be sent over the course of T periods can also be sent in the first T periods of a longer game. In fact, the suspense $Tu(\Psi(\mu_0)/T)$ is strictly increasing in T. For $u(x) = x^{\gamma}$ with $0 < \gamma < 1$, for instance, total suspense is proportional to $T^{1-\gamma}$.
- Suspense-optimal information policies are independent of the stage utility function. Under any concave $u(\cdot)$, any optimal policy induces beliefs in M_t . The expression for M_t is independent of $u(\cdot)$.

4.2 Illustration of Suspense-Optimal Policies

4.2.1 Two states

We first illustrate these policies in the case of a binary state space $\Omega = \{A, B\}$. In a sporting event, will team A or team B win? In a mystery novel, is the main character guilty or not? In this case, each set M_t consists of just two points. This leads to a unique suspense maximizing belief martingale, depicted in Figure 5.

The suspense-optimal policy gives rise to the following dynamics. At period t the belief $\mu_t \equiv \Pr(A)$ is either a high value $H_t > \frac{1}{2}$ or a low value $L_t = 1 - H_t$.²⁰ In each period, one of two things happens. With high probability the agent observes additional confirmation

²⁰In this binary case we abuse notation by associating the belief with the probability of one of the states.



This picture depicts the case where $\mu_0 = \frac{1}{2}$. The belief at time t will be either $H_t > \frac{1}{2}$ or $L_t < \frac{1}{2}$. The probability of a plot twist (which takes beliefs from H_t to L_{t+1} or from L_t to H_{t+1}) declines over time.

– the high belief H_t moves to a slightly higher belief H_{t+1} , or the low belief L_t moves to L_{t+1} . With a smaller probability, there is a *plot twist*. In the event of a plot twist, beliefs jump from the high path to the low one, or vice versa. As time passes, plot twists become larger but less likely.²¹ If we consider a limit as T goes to infinity, the arrival of plot twists approaches a Poisson process with an arrival rate that decreases over time.²²

In the context of a mystery novel, these dynamics imply the following familiar plot structure. At each point in the book, the reader thinks that the weight of evidence either suggests that the protagonist accused of murder is guilty or is innocent. But in any given chapter, there is a chance of a plot twist that reverses the reader's beliefs. As the book

$$H_t = \frac{1}{2} + \sqrt{\left(\mu_0 - \frac{1}{2}\right)^2 + \frac{t}{T}\mu_0\left(1 - \mu_0\right)}, \ L_t = \frac{1}{2} - \sqrt{\left(\mu_0 - \frac{1}{2}\right)^2 + \frac{t}{T}\mu_0\left(1 - \mu_0\right)},$$

The probability of a plot twist is $\frac{1}{2} - \frac{1}{2} \frac{\sqrt{(\mu_0 - \frac{1}{2})^2 + \frac{t-1}{T}\mu_0(1-\mu_0)}}{\sqrt{(\mu_0 - \frac{1}{2})^2 + \frac{t}{T}\mu_0(1-\mu_0)}}.$

²²To derive the limit, take T to infinity and rescale time to s = t/T. This yields a Poisson process with plot twist arrival intensity of $\frac{\mu_0(1-\mu_0)}{1-4(1-s)\mu_0(1-\mu_0)}$.

²¹ We can solve for H_t and L_t explicitly as

continues along, plot twists become less likely but more dramatic.

In the context of sports, optimal dynamics could be induced by the following set of rules. We declare the winner to be the *last* team to score. Moreover, scoring becomes more difficult as the game progresses (e.g., the goal shrinks over time). The former ensures that uncertainty declines over time while the latter generates a decreasing arrival rate of plot twists. (In this context, plot twists are lead changes).

Note that existing rules of most sports cannot be suspense optimal. In soccer, for example, the probability that the leading team will win depends not only on the period of the game, but also on whether it is a tight game or a blowout. Moreover, the team that is behind can come back to tie up the game, in which case uncertainty will have increased rather than decreased over time.

To conclude the discussion of binary states, we note the following qualitative points that apply in this case.

- Beliefs can jump by a large amount in a single period. In each period, either beliefs are confirmed or there is a plot twist. A plot twist takes beliefs from μ_t to something further away than $1 - \mu_t$.
- Belief paths are smooth with rare discrete jumps when there are many periods. Beliefs move along the increasing H_t or decreasing L_t curves with occasional plot twists when beliefs jump from one curve to the other. In the limit as T gets large, the expected number of total plot twists stays small. Expected absolute variation $\sum_{t=0}^{T-1} |\mu_{t+1} - \mu_t|$ converges to a finite value as T goes to infinity.²³

4.2.2 Three or more states

With more than two possible outcomes, there is additional flexibility in the design of a suspense-maximizing martingale. Say that there is a mystery novel with three suspects, and we currently believe A to be the most likely murderer. In the next chapter we will see a clue (a signal) which alters our beliefs, either providing further evidence of A's guilt or pointing to B or C as suspects. It may be the case that there are three possible clues we can see, or five, or fifty. The only restriction is that after observing the clue, our belief has the right amount of uncertainty as measured by residual variance.

²³In the limit as $T \to \infty$, the expected absolute variation converges to $2\min\{\mu_0, 1-\mu_0\}$.

We highlight two classes of optimal policies with natural interpretations. Figure 7(a) illustrates what we call an *alive till the end* policy. In each period there is news favoring one of the states. No piece of news ever eliminates any state entirely until the last period. This corresponds to a mystery novel where the reader becomes more confident of the murderer over time, but always anticipates the possibility of a plot twist pointing towards any other suspect; or to a race where no participant is entirely ruled out until the very end. Figure 7(b) shows a different kind of policy, *sequential elimination*. Here, one of the probabilities is quickly taken to 0. At that point the policy maximizes suspense over the remaining states. In a mystery novel, suspects are eliminated by being killed off one-by-one. Or in a sports tournament, each game eliminates one of the players.



Figure 6: Two optimal policies

Beliefs travel outwards over the M_t circles, with possible belief paths indicated by the arrows.

4.3 Extensions

When we introduced the model we discussed potential extensions to state-dependent or time-dependent significances. In the case of state-dependent significance (with weights α^{ω}), all of the results apply when we redefine $\Psi(\mu)$ as $\sum_{\omega} \alpha^{\omega} \mu^{\omega} (1 - \mu^{\omega})$. Geometrically, the M_t sets are ellipses rather than circles (ellipsoids rather than hyperspheres). Also, the optimal prior is no longer necessarily uniform when there are more than two states: more significant states are given priors closer to $\frac{1}{2}$. Sufficiently insignificant states may be given a prior of 0, and as the significance of one state begins to dominate all others the prior on that state goes to $\frac{1}{2}$.

An alternative extension is a setting where the principal chooses α^{ω} 's and μ_0 .²⁴ (For example, a sports league may be able to influence the market share across teams, or the novelist chooses how much reader empathy to generate for each character.) In this case, the principal's choice is optimal if and only if $\mu_0^{\omega} = \frac{1}{2}$ for each state with $\alpha^{\omega} > 0$.²⁵ This means that there are two basic ways to maximize suspense. One way is to have only two states of interest, with 50/50 odds between those two. Good versus evil, Democrat versus Republican, Barcelona vs. Real Madrid. Alternatively, there may be a single state of interest, realized with probability 50%. The reader cares only about whether the protagonist is found innocent or guilty. Conditional on the protagonist's innocence, any of the irrelevant characters may be the murderer with any probabilities.

In the case of time-dependent significances (with weights β_t), it is no longer optimal to divide suspense evenly over time. Instead, more important periods are made to be more suspenseful. For example, in the baseline specification for $u(\cdot)$, we set σ_t proportional to β_t .

5 Surprise-Optimal Information Policies

Solving for the surprise-optimal martingale is difficult in general and, in contrast to the case of suspense, sensitive to the choice of $u(\cdot)$. So in this section we restrict attention to binary states $\Omega = \{A, B\}$ and the baseline specification where surprise in period t equals $|\mu_t - \mu_{t-1}|$ with $\mu_t \equiv \Pr(A)$.²⁶ We derive an exact characterization of optimal belief martingales for very small T, and discuss properties of the solution for large T.²⁷

Let $W_T(\mu)$ be the value function of the surprise maximization problem, where T is the number of periods remaining and μ is the current belief. We can express the value function

²⁴Utility increases in each α^{ω} , so to make this problem well-posed we assume that the principal is constrained by a fixed sum of significances $\sum_{\omega} \alpha^{\omega}$.

²⁵Proof of this and other claims from this subsection are in the Online Appendix.

²⁶Recall that under the baseline specification, surprise in period t is the Euclidean distance between the belief vectors in periods t and t - 1. To avoid introducing a nuisance term, however, in this section we set surprise in period t to be the Euclidean distance between the scalars μ_t and μ_{t-1} . This amounts to scaling all surprise payoffs by $\frac{1}{\sqrt{2}}$.

²⁷In the Online Appendix, we discuss how some features of the surprise optimum change if we depart from the baseline specification.

recursively by setting $W_0(\mu) \equiv 0$, and

$$W_T(\mu) = \max_{\tilde{\mu}' \in \Delta(\Delta(\Omega))} E_{\tilde{\mu}'} \left[\left| \mu' - \mu \right| + W_{T-1}(\mu') \right] \text{ s.t. } E_{\tilde{\mu}'} \left[\mu' \right] = \mu.$$

The single-period problem above can always be solved by some $\tilde{\mu}'$ with binary support.²⁸ That is, for any current belief, there is a surprise-maximizing martingale such that next period's belief is either some μ_l or $\mu_h \geq \mu_l$.

The solution can be derived by working backwards from the last period. In the final period, it is optimal to fully reveal from any prior: $\mu_l = 0$ and $\mu_h = 1$. This yields a value function of $W_1(\mu) = 2\mu(1-\mu)$.

With two periods remaining, it is optimal to set $\mu_l = \mu - \frac{1}{4}$ and $\mu_h = \mu + \frac{1}{4}$, as long as $\mu \in [\frac{1}{4}, \frac{3}{4}]$. Therefore if $\mu_0 = \frac{1}{2}$ and T = 2, the surprise-optimal martingale induces beliefs $\mu_l = \frac{1}{4}$ or $\mu_h = \frac{3}{4}$ in period 1, and then fully reveals the state in the second period. The details of this and the next derivation are in Appendix A.1.

The solution for T = 3 with the prior of $\mu_0 = \frac{1}{2}$ is displayed below in Figure 7. In the first period the belief moves to either 1/4 or 3/4 with equal probability. Then, the belief either moves to the boundary or returns to 1/2. Any remaining uncertainty is resolved in the final period.



We see that, unlike the solution for optimal suspense, there are paths in which all

²⁸Fixing μ and W_{T-1} , this single-period problem of choosing $\tilde{\mu}'$ is a special case of the problem considered by Kamenica and Gentzkow (2011).

uncertainty is resolved before the final period T. These paths have lower overall surprise than the paths that resolve only at the end. But the optimal information policy accepts a positive probability of early resolution in return for a chance to move beliefs back to the interior and set the stage for later surprises. Also unlike the suspense solution, uncertainty can both increase or decrease over time. Beliefs may move towards an edge, or back towards $\frac{1}{2}$. In the suspense-optimal martingale, uncertainty only increases.

These features underscore the commitment problem facing a surprise-maximizing principal. The surprise-optimal martingale has paths that generate very little surprise. In order to implement the optimal policy, the principal requires the commitment power to follow such paths. Otherwise, he is tempted to prune such paths and choose a path with maximal surprise. The agent would expect this deviation and the chosen path would no longer be surprising.²⁹ In contrast, maximizing suspense does not involve this form of commitment because the suspense-optimal information policy yields equal suspense across all realized paths.

This importance of commitment sheds some light on the phenomenon of dedicated sports fans. It may seem tempting to record a game and let others tell you whether it was exciting before you decide whether to watch it.³⁰ However, such a strategy is self-defeating: the very knowledge that the game was exciting reduces its excitement. Similarly, if ESPN Classic shows only those games with comeback victories, the audience would never be surprised at the comeback. To make the comebacks surprising, ESPN Classic would have to show some games in which one team took a commanding early lead and never looked back.

For arbitrary T, it is difficult to analytically solve for surprise-optimal belief martingales. The properties of the value function $W_T(\mu)$ in the limit as T goes to infinity, however, have previously been studied by Mertens and Zamir (1977). Their interest in this limiting variation of a bounded martingale arose in the study of repeated games with asymmetric information.³¹

²⁹Without commitment, all paths (including ones where the belief moves monotonically to a boundary) must generate the same surprise in equilibrium. Hence the principal's payoff cannot be greater than if all information is revealed at once. This echoes Geanakoplos's (1996) result about the Hangman's Paradox.

³⁰The web site http://ShouldIWatch.com, for example, provides this information.

³¹De Meyer (1998) extends their results to the more general L^q variation, i.e., he considers the problem of maximizing $\mathbb{E}\left[\sum_{t=0}^{T-1} (\mathbb{E}|\mu_{k+1} - \mu_k|^q)^{\frac{1}{q}}\right]$. Under binary states and the baseline specification, this is our surprise problem if q = 1 and our suspense problem if q = 2. De Meyer searches for the limit of the L^q value function divided by \sqrt{T} , as T goes to infinity. He finds that for $q \in [1, 2)$, this limit is constant in q and is equal to $\phi(\mu)$, as given in Proposition 2 for q = 1. For q > 2, de Meyer finds that the limit approaches

Proposition 2 (Mertens and Zamir (1977), Equation (4.22)). For any μ ,

$$\lim_{T \to \infty} \frac{W_T(\mu)}{\sqrt{T}} = \phi(\mu),$$

where $\phi(\mu)$ is the pdf of the standard normal distribution evaluated at its μ^{th} -quantile:

$$\phi(\mu) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x_{\mu}^2}$$
, with x_{μ} defined by $\int_{-\infty}^{x_{\mu}} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx = \mu$

In particular, $\sqrt{T}\phi(\mu) - \alpha \leq W_T(\mu) \leq \sqrt{T}\phi(\mu) + \alpha$ for some constant $\alpha > 0$ independent of μ and T.

The fact that the *ex ante* surprise payoff – i.e., the expected absolute variation – goes to infinity tells us that paths are "spiky" rather than smooth as T gets large. Recall that in the suspense optimal martingale, expected absolute variation was bounded in T.

Another difference between optimal surprise and suspense is the range of possible belief changes in a given period. In each period of the suspense problem, there is a chance of a twist that leads to a large shift in beliefs. In the surprise problem, however, beliefs move up or down only a small amount in periods when there is a lot of time remaining.

Proposition 3. For all $\epsilon > 0$, if T - t is sufficiently large then for any belief path in the support of any surprise-optimal martingale, $|\mu_{t+1} - \mu_t| < \epsilon$.

The proof, in Appendix A.2, builds on Proposition 2. We can now summarize some qualitative features of surprise-optimal information revelation.

- The state is fully revealed, possibly before the final period. Any time the belief at T-2 is below $\frac{1}{4}$ or above $\frac{3}{4}$, for instance, there is a chance of full revelation at period T-1.
- Uncertainty may increase or decrease over time. Beliefs often move toward $\mu = \frac{1}{2}$. With sufficiently many periods remaining, residual uncertainty in the next period can always either increase or decrease (except in the special case of $\mu_t = \frac{1}{2}$).

infinity for any $\mu \in (0, 1)$. However, de Meyer incorrectly suggests that the methods of Mertens and Zamir can be used to show that the value function for q = 2 will be identical to that for q < 2. In fact, our solution to the suspense problem with binary states and the baseline specification shows this to be false. The limiting value function is $\sqrt{\mu(1-\mu)}$ rather than $\phi(\mu)$.

- **Realized surprise is stochastic.** In an optimal martingale there are low surprise paths (e.g., ones in which beliefs move monotonically to an edge) or high surprise paths (with a lot of movement up and down).
- Surprise varies over time. Both realized and expected surprise can vary over time. Consider T = 2 where $\mu_0 = \frac{1}{2}$, $\mu_1 \in {\{\frac{1}{4}, \frac{3}{4}\}}$, and $\mu_2 \in {\{0, 1\}}$. On a particular belief path, say $(\frac{1}{2}, \frac{1}{4}, 1)$, realized surprise in period 1 is different from realized surprise in period 2. Moreover, from the *ex ante* perspective, the expected surprise in period 1 is $\frac{1}{4}$ while the expected surprise in period 2 is $\frac{3}{8}$.
- The prior that maximizes surprise is the uniform belief. We show that for $T \leq 3$, surprise is maximized at the uniform prior. We conjecture that this holds for all T. Proposition 2 shows that surprise is maximized at the uniform prior in the limit as $T \to \infty$.
- The level of surprise increases in the number of periods T. While we do not have a general expression for surprise as a function of T, we know that in the limit surprise increases proportionally with \sqrt{T} . It is obvious that surprise is weakly increasing in T.
- Beliefs change little when there are many periods remaining. By Proposition 3, $|\mu_{t+1} \mu_t|$ is small when T t is large.
- Belief paths are spiky when there are many periods. Expected absolute variation, which is equal to the surprise value function, goes to infinity as T gets large.
- Surprise-optimal information policies depend on the stage utility function. In the Online Appendix, we consider alternative $u(\cdot)$ functions. For very concave $u(\cdot)$, the surprise-optimal policy can be non-fully revealing by the end. For convex $u(\cdot)$, the surprise-optimal prior can be non-uniform.

6 Comparing Suspense and Surprise

As the two preceding sections reveal, suspense-optimal and surprise-optimal belief martingales are qualitatively different. Another way to appreciate these differences is to consider sample belief paths drawn from a suspense-optimal and a surprise-optimal martingale, as shown in Figure 8. Here we show three representative belief paths for each of the two processes: a path at the 25th percentile of suspense and surprise from the simulations of Figure 3, at the median, and at the 75th percentile. The suspense paths reveal the distinctive plot-twist structure whereas the surprise paths show the spiky nature of surprise-optimal martingales.



While no existing sport would induce the exact distributions of belief paths we derive, we can think of *soccer* and *basketball* as representing extreme examples of sports with the qualitative features of optimum suspense and surprise. In any given minute of a soccer game it is very likely that nothing consequential happens. Whichever team is currently ahead becomes slightly more likely to win (since less time remains). There is a small chance that a team scores a goal, however, which would have a huge impact on beliefs. So (as Figure 3 illustrates), belief paths in soccer are smooth, with few rare jumps. This sustained small probability of large belief shifts makes soccer a very suspenseful game. In basketball, points are scored every minute. With every possession, a team becomes slightly more likely

to win if it scores and slightly less likely to win if it does not. But no single possession can have a very large impact on beliefs, at least not until the final minutes of the game. Belief paths are spiky, with a high frequency of small jumps up and down; basketball is a game with lots of surprise.

The distinction between suspense-optimal and surprise-optimal martingales somewhat clashes with an intuition that more suspenseful events also generate more surprise. This intuition is indeed valid in the following two senses. First, given a martingale, belief paths with high realized suspense tend to have high realized surprise; this can be seen in the right column of Figure 3. Moreover, the *ex ante* suspense and surprise are highly correlated across martingales generated by "random" information policies. Specifically, suppose T = 10, $\Omega = \{L, R\}$, and $\mu_0 = \frac{1}{2}$. In periods 1 through 9, a signal realization l or r is observed. When the true state is ω , the signal $\pi_{\omega,t}$ at period t is l with probability $\rho_{\omega,t}$ and r with probability $1 - \rho_{\omega,t}$. The values of $\rho_{\omega,t}$ are drawn iid uniformly from [0,1] for each ω and t. The state is revealed in period 10. Figure 9 depicts a scatterplot of ex ante suspense and surprise of 250 such random policies; it is clear that policies that generate more suspense also tend to generate more surprise. Note that these policies are history-independent in the sense that the signal sent at period t depends only on t and not on μ_t . The figure also shows the numerically derived production possibilities set for suspense and surprise over all fully revealing policies. As this set reveals, the suspenseoptimal martingale does not generate much surprise while the surprise-optimal martingale only reduces suspense a little below its maximum. This suggests that maximizing a convex combination of suspense and surprise is likely to lead to belief paths that resemble the surprise optimum.

7 Constrained Information Policies

In practice there are often institutional restrictions which impose constraints on the information the principal can release over time. Recall that we formalize these situations as the principal's choosing $(\tilde{\pi}, \mu_0, T) \in P$ so as to maximize expected suspense or surprise. In this section we will study the nature of the constraint set and the constrained-optimal policies in some specific examples. Throughout this section, we impose the baseline specification for $u(\cdot)$.



Figure 9: The Surprise-Suspense frontier

7.1 Tournament Seeding

Consider the problem of designing an elimination tournament to maximize spectator interest. Elimination tournaments begin by "seeding" teams into a bracket. The traditional seeding pits stronger teams against weaker teams in early rounds thereby amplifying their relative advantage. We analyze the effect of this choice on the suspense and surprise generated by the tournament. The tradeoff is clear: by further disadvantaging the weaker team, the traditional seeding reduces the chance of an upset but increases the drama when an upset does occur.

The simplest example of tournament seeding occurs when there are three teams. Two teams play in a first round and the winner plays the remaining team in the final. This remaining team is said to have the first-round "bye." Which team should have the bye?

Formally, the state of the world ω is identified with the team that wins the overall tournament, so $\omega \in \{1, 2, 3\}$. Assume that the probability that a team wins an individual contest is determined by the difference in the ranking of the two teams. Let $p > \frac{1}{2}$ denote the probability that a team defeats the team that is just below it in the ranking, e.g. that *i* beats team i + 1; and let q > p denote the probability that team 1 defeats team 3. The principal chooses which team will be awarded the bye. This determines the prior as well as the sequence of signals. For example, team 1's prior probability of winning the tournament is $p^2 + (1-p)q$ if team 1 has the bye and pq if team 2 has the bye. This choice of seeding implies that first it will be revealed whether team 2 or team 3 has lost, and then it will be

revealed which of the remaining teams has won. Figure 10 illustrates the belief paths for each of the possible tournament structures.



Figure 10: Beliefs paths for alternate seedings.

Notice that one of the shortcomings of the traditional seeding in which the strongest team has the bye is that it has low residual uncertainty. In fact, for any p and q, straightforward algebra shows that this traditional seeding generates the *least* surprise; it is optimal to give the third team the bye. In the case of suspense, the conclusions are less clear cut but for many reasonable parameters this same ordering holds.

There are of course many other reasons for the traditional seeding which favors the stronger teams. Incentives are an important consideration: teams that perform well from tournament to tournament improve their rankings and are rewarded with better seedings in subsequent tournaments. Our analysis suggests that optimizing the tournament seeding for its incentive properties can have a cost in short-run suspense and surprise.

7.2 Number of Games in a Playoff Series

Each round of the NBA playoffs consists of a best-of-seven series. Major League Baseball playoffs use best-of-five series for early rounds, and best-of-seven for later rounds. In the NFL playoffs, each elimination round consists of a single game. Of course the length of the series is partly determined by logistical considerations, but it also influences the suspense and surprise. On the one hand, having more games leads to slower information revelation, increasing suspense and surprise. On the other hand, in a long series the team which is better on average is more likely to eventually win, and this reduces both suspense and surprise. If a team wins 60% of the matches, there is much more uncertainty over the

outcome of a single match than over the outcome of a best-of-seven series, or a best-ofseventeen one. With less uncertainty there is less scope for suspense and surprise.

Formally, consider two teams playing a sequence of games against each other. The favored team has an independent probability p > .5 of winning any given game. The organizer chooses some odd number T and declares that the winner of the series is the first team to win $\frac{T+1}{2}$ out of T total games. The organizer chooses T to maximize suspense or surprise.

As we show in Appendix B, suspense and surprise are proportional to one another for any choice of T. Therefore, for this class of constrained problems, maximizing suspense is equivalent to maximizing surprise.

In Table 1 we display the suspense- and surprise-maximizing series length T^* as a function of p. As the table shows, the maximizing series length is increasing in the proximity of p to $\frac{1}{2}$. The intuition behind this is simple. When neither team is much better than the other, the cost of a large T becomes small (μ_0 does not move far from μ_*) while the benefit of releasing information slowly remains.

Table 1: Suspense- and Surprise-Maximizing Length of a Playoff Series



7.3 Sequential Contests

The order of sequential primary elections may affect which political candidate is ultimately chosen as a party's nominee for President (Knight and Schiff (2007)). If voters in late states converge around early winners, then it surely matters whether Iowa goes first, or New Hampshire, or Florida. As a number of researchers have analyzed, e.g. Hummel and Holden (2012), political parties may want to choose the order of primaries so as to maximize the expected quality of a nominee.

We now pose a different question: in what order should the states vote if the goal is to maximize the suspense or surprise of the race? A more exciting primary season may elicit more attention from citizens, yielding a more engaged and informed electorate. In order to highlight the relevant mechanisms, our analysis here will assume that the order of primaries does not affect the likelihood any given candidate wins.

We model a primary campaign as follows. Two candidates, A and B, compete against each other in a series of winner-take-all state primary elections. State i has n_i delegates and will be won by candidate A with probability p_i . The probability A wins each state is independent. Candidate A wins the nomination if and only if she gets at least $n^* \in [0, \sum_i n_i]$ delegates.

Early states are sure to have a small but positive impact on beliefs about the nominee, whereas late states have some chance of being extremely important and some chance of having zero impact. Perhaps smaller or more partial states should go first, so that information about the nominee is revealed as slowly as possible. Or should large swing states go first, to guarantee that these potentially exciting votes don't take place after a nominee has already been chosen?

Remarkably, for any distribution of delegates n_i , for any set of probabilities p_i , for any cutoff n^* , the order of the primaries has no effect on expected suspense or surprise.³² Small or large states, partian or swing states – they may go in any order.

This neutrality result applies not only to political primaries but to many other settings where a pair of players engage in sequential contests. For instance, in a televised game show, players compete in a variety of tasks with different amounts of points at stake. Family Feud doubles and then triples the points awarded in later rounds. Our results imply that this leads to no more total excitement than if the high stakes rounds were at the beginning. Likewise, in a typical best-of-seven sports series one advantaged team will be scheduled to play four home games and three away games. We find that there is an equal amount of suspense and surprise in the 2-3-2 format of the NBA Finals (two home, three away, then two home games) as in the 2-2-1-1-1 format of the earlier NBA playoff rounds.

8 Conclusion

Our analysis is predicated on the assumption that people value suspense (variance of next period's beliefs) or surprise (movements in beliefs). We derive implications of this assumption, focusing on the optimal way to release information over time.

One way to test our model would be to examine whether the information policies we identify indeed attract a greater audience than other policies. More generally, a dataset

³²Proof is in Appendix B.

that combines information about revealed preference with estimates of belief dynamics would allow us to directly examine what aspects of belief dynamics generate entertainment utility.

We have already discussed various ways one can estimate belief dynamics: (i) relying on the explicit structure of the data generating process, (ii) using prediction markets, (iii) estimating probabilities in each period based on outcomes of many matches, or (iv) eliciting beliefs through incentivized laboratory experiments. In principle, it should be possible to complement such data with measures of revealed preference. TV ratings would reveal whether a show becomes more popular when previous episodes generate more suspense and surprise. Or, detailed data collected by The Nielsen Company could be used to examine whether the preferences we postulate explain channel-switching behavior during a sports game. Whether suspense or surprise or some other aspects of belief dynamics drive demand for non-instrumental information is fundamentally an empirical question, one which we hope will be addressed by future research.

A Surprise-Optimal Policies

A.1 Surprise for $T \leq 3$

Let the function f_T be defined by

$$f_T(\mu';\mu) \equiv |\mu'-\mu| + W_{T-1}(\mu')$$

The notation $f'_T(\mu';\mu)$ will indicate the derivative with respect to the first component.

The recursive problem of maximizing surprise in a given period, starting from belief μ_t at time t, can be written as

$$W_{T-t}(\mu_t) = \max_{\mu_l \le \mu_h} p f_{T-t}(\mu_h; \mu_t) + (1-p) f_{T-t}(\mu_l; \mu_t)$$
(1)

where p is the probability given by $p\mu_h + (1-p)\mu_l = \mu_t$. Without loss of generality, $\mu_l \leq \mu \leq \mu_h$.

We set W_0 to be identically 0, and so for any prior μ_0 in the one-period problem T = 1it is optimal to set $\mu_l = 0$ and $\mu_h = 1$ so that $p = \mu_0$ and the maximized value is

$$W_1(\mu_0) = 2\mu_0(1-\mu_0).$$

Consider now the problem for T > 1, starting at period 0 (which can be embedded as the last T periods of a longer problem). We can derive two first-order conditions for optimality of the posteriors (μ_l, μ_h) , holding the prior μ_0 constant.³³ First, consider moving the pair (μ_l, μ_h) in the direction $d(\mu_l, \mu_h) = (1, -\frac{1-p}{p})$, i.e., moving the posteriors toward one another along the line that keeps p, the probability of μ_h constant. The derivative of the objective function in this direction is the inner product

$$d(\mu_l,\mu_h) \cdot ((1-p)f'_T(\mu_l;\mu_0), pf'_T(\mu_h;\mu_0)).$$

At an optimum this derivative must be non-positive, and equal to zero in the case of an interior optimum, i.e.

$$f'_T(\mu_l;\mu_0) - f'_T(\mu_h;\mu_0) \le 0$$
(2)

³³The objective function is not differentiable when, say, $\mu_h = \mu_0$, but it is straightforward to show that releasing zero information is never optimal. The differentiability of W_{T-1} will be verified directly.

with equality when $0 < \mu_l < \mu_h < 1$. Expanding the derivatives of $f_T(\cdot)$, this translates to

$$\frac{d}{d\mu_l} W_{T-1}(\mu_l) - \frac{d}{d\mu_h} W_{T-1}(\mu_h) \le 2$$
(3)

with equality in the interior. This condition illustrates the tradeoff in determining whether to reveal more informative signals, increasing current surprise while decreasing the "stock" of future surprises.

To derive an additional first-order condition, consider moving μ_l closer to μ_0 , holding μ_h constant. The derivative of the objective function in this direction is

$$\frac{dp}{d\mu_l} \left[f_T(\mu_h;\mu_0) - f_T(\mu_l;\mu_0) \right] + (1-p) \frac{df_T}{d\mu_l}(\mu_l;\mu_0)$$

which must be non-positive at an optimum and equal to zero when $\mu_l > 0$. Substituting

$$(1-p) = \frac{\mu_h - \mu_t}{\mu_h - \mu_l}$$

and

$$\frac{dp}{d\mu_l} = \frac{\mu_t - \mu_h}{(\mu_h - \mu_l)^2}$$

and re-arranging, we obtain

$$\frac{df}{d\mu_l}(\mu_l;\mu_0) \le \frac{f_T(\mu_h;\mu_0) - f_T(\mu_l;\mu_0)}{\mu_h - \mu_l}$$

or, directly in terms of p and W_{T-1} ,

$$\frac{d}{d\mu_l}W_{T-1}(\mu_l) - \frac{W_{T-1}(\mu_h) - W_{T-1}(\mu_l)}{\mu_h - \mu_l} \le 2(1-p)$$
(4)

with equality when μ_l is interior. This condition connects the marginal gain in continuation value at μ_l with the "average" continuation value. According to the condition, the larger the difference between these two measures, the larger should be the probability (1 - p) of μ_l .

With these two conditions we can solve the surprise problems for T = 2 and T = 3. First, recall that $W_1(\mu) = 2\mu(1-\mu)$, so that $W'_1(\mu) = 2 - 4\mu$. Then, to solve to problem for T = 2 we use Equation 3 to derive $2 - 4\mu_l - (1 - 4\mu_h) = 2$ or $\mu_h - \mu_l = 1/2$ at an interior optimum. Substituting this equation into Equation 4 yields, after some algebra, $\mu_l = \mu_0 - 1/4$, implying that for any $\mu_0 \in [1/4, 3/4]$ the optimal signal is symmetric with $\mu_l = \mu_0 - 1/4$ and $\mu_h = \mu_0 + 1/4$. In particular, note that when the prior is $\mu_0 = 3/4$, there is full revelation of the state 1, i.e., $\mu_h = 1$ in the first period.

Indeed, when the prior $\mu_0 \geq 3/4$, there is no interior solution and Equation 3 holds with a strict inequality. In that case we have $\mu_h = 1$ and $W_1(\mu_h) = 0$, which by Equation 4 yields $\mu_l = p = 1 - \sqrt{1 - \mu_0}$. By symmetry, when $\mu_0 < 1/4$ we have $\mu_l = 0$ and $\mu_h = \sqrt{\mu_0}$. This gives the following value function:

$$W_2(\mu) = \begin{cases} 4\mu(1-\sqrt{\mu}) & \text{if } \mu \in [0,1/4] \\ 1/8 + 2\mu(1-\mu) & \text{if } \mu \in [1/4,3/4] \\ 4(1-\mu)(1-\sqrt{1-\mu}) & \text{if } \mu \in [3/4,1] \end{cases}$$

The surprise-maximizing prior is $\mu_0 = 1/2$, after which the optimal signals are $(\mu_l, \mu_h) =$ (1/4, 3/4) in the first period followed by full revelation in the last period.

For the T = 3 problem, fix a prior to of 1/2. It can be verified that $W_2(\cdot)$ is differentiable³⁴ and so the first order conditions apply. By symmetry, the optimal signal starting from a prior $\mu_0 = 1/2$ is also symmetric and thus Equation 4 reduces to $\frac{d}{d\mu_l} W_{T-1}(\mu_l) = 1$ so that $\mu_l = 1/4$ and by symmetry $\mu_h = 3/4$. This yields a surprise payoff at the prior $\mu_0 = \frac{1}{2}$ of $W_3(\frac{1}{2}) = \frac{3}{4}$. In fact, it can be shown that $W_3(\mu) \leq \frac{3}{4}$ for all μ . So in fact the prior of $\mu_0 = \frac{1}{2}$ maximizes $W_3(\mu_0)$ over all possible μ_0 .³⁵

A.2**Proof of Proposition 3**

Proof of Proposition 3. As shown by (Mertens and Zamir, 1977, Equation (3.5)) the function ϕ satisfies the differential equation $\phi''(\mu) = -1/\phi(\mu)$. We will make use of this fact below, e.g., in observing that ϕ is concave.

Optimal policies are history-independent conditional on the current belief and the number of periods remaining, so without loss of generality we will show that $|\mu_1 - \mu_0| < \epsilon$ for any μ_0 , if the number of periods T is large enough. In particular we will show that μ_l converges to μ_0 as T gets large, and that this convergence is uniform over μ_0 .³⁶ The argument

 $^{^{34}}$ In particular, the left and right derivatives at 1/4 and 3/4 are equal.

³⁵From the expression for $W_2(\mu)$, we see that $W_2(\mu) \leq 1/8 + 2\mu(1-\mu)$ for each μ ; it is equal on the interval $\left[\frac{1}{4}, \frac{3}{4}\right]$ and below at other points. In other words, $W_2(\mu) \leq \frac{1}{8} + W_1(\mu)$. By the recursive definition of W_T , then, $W_3(\mu) \leq \frac{1}{8} + W_2(\mu)$. This implies that $W_3(\mu) \leq \frac{1}{4} + 2\mu(1-\mu)$, which has a maximum value of $\frac{3}{4}$ at $\mu = \frac{1}{2}$. ³⁶There may exist optimal policies that are not binary. If we take μ_l to be the infimum of all points in

for μ_h would follow similarly.

Step 1: For each μ and μ' , for fixed T, the function $f_T(\mu';\mu)$ is in the interval $[W_{T-1}(\mu'), W_{T-1}(\mu') + 1]$. So by Proposition 2, there exists $\alpha > 0$ such that for all $\mu, \mu' \in [0, 1]$

$$\sqrt{T-1}\phi(\mu') - \alpha \le f_T(\mu';\mu) \le \sqrt{T-1}\phi(\mu') + \alpha$$

Consider the concavification of f_T with respect to μ' :

$$\hat{f}_T(\mu';\mu) \equiv \max\{z: (\mu',z) \in \operatorname{co}(f_T(\mu';\mu))\}$$

where $co(f_T(\mu'; \mu))$ is the convex hull of the graph of $f_T(\mu'; \mu)$ viewed as a function of μ' . The function $\hat{f}_T(\mu'; \mu)$ is concave in μ' and is in fact the pointwise minimum of all concave functions that are pointwise larger than $f_T(\mu'; \mu)$ (Rockafellar, 1997, Chapter 12). Since the function ϕ is concave so we also have

$$\sqrt{T-1}\phi(\mu') - \alpha \le \hat{f}_T(\mu';\mu) \le \sqrt{T-1}\phi(\mu') + \alpha \tag{5}$$

and all three terms are concave in μ' .

Step 2: It is immediate from the definitions and from the concavity of $f_T(\mu;\mu)$ that

$$W_T(\mu) = \hat{f}_T(\mu;\mu) = \max_{\mu_l \le \mu_h} p \hat{f}_T(\mu_h;\mu) + (1-p) \hat{f}_T(\mu_l;\mu),$$

i.e., we can replace f_T with \hat{f}_T in the optimization problem in Equation 1.

Therefore, any optimal μ_l and μ_h will satisfy the first-order condition given in Equation 2:

$$\hat{f}'_T(\mu_l;\mu_0) = \hat{f}'_T(\mu_h;\mu_0)$$

where, in case \hat{f}_T is not differentiable, $\hat{f}'_T(\mu';\mu_0)$ is some supergradient of $\hat{f}(\mu';\mu_0)$ with respect to μ' . By concavity, since these supergradients are equal, the function $\hat{f}(\mu';\mu_0)$ is linear between μ_l and μ_h and the supergradients are equal to the derivative at $\mu' = \mu$:

$$\hat{f}'_{T}(\mu_{l};\mu_{0}) = \hat{f}'_{T}(\mu_{0};\mu_{0}) = \hat{f}'(\mu_{h};\mu_{0}).$$
(6)

Step 3: Refer to Figure 11. Because $\hat{f}_T(\mu';\mu)$ is concave, its graph lies everywhere below its tangent line (i.e. the derivative) at $\mu' = \mu$. Because of the bound in Equation 5,

the support then an identical argument shows that $\mu_l \to \mu_0$.



Figure 11: The dotted line shows a possible $f(\mu';\mu)$ curve as a function of μ' and the dashed line shows its concavification $\hat{f}(\mu';\mu)$. For any f bounded between the blue curves, μ_l must be greater than μ_b .

its graph must also lie between the graphs of $\sqrt{T-1}\phi(\mu') - \alpha$ and $\sqrt{T-1}\phi(\mu') + \alpha$. Thus, an upper bound for $\hat{f}'_T(\mu_0;\mu_0)$ is the slope of the line through $(\mu_0,\sqrt{T-1}\phi(\mu_0)+\alpha)$ which is tangent to the function $\sqrt{T-1}\phi(\mu') - \alpha$. Let $\mu' = \mu_a < \mu$ be the point of tangency, and let $\mu_b < \mu_a$ be the point of intersection with the graph of $\sqrt{T}\phi(\mu') + \alpha$.

By Equation 6, μ_l satisfies $\hat{f}'(\mu_l;\mu_0) = \hat{f}'(\mu_0;\mu_0)$. Suppose μ_l were less than μ_b . Then, by the concavity of $\hat{f}_T(\mu';\mu)$ the value at μ_a , i.e. $\hat{f}_T(\mu_a;\mu)$ would be less than $\sqrt{T-1}\phi(\mu_a) - \alpha$. Hence μ_b is a lower bound for μ_l . We will show that $\mu_b \to \mu_0$, uniformly in μ_0 , as T approaches infinity.

Step 4: By construction of μ_a ,

$$\sqrt{T-1}\phi(\mu_a) - \alpha + (\mu - \mu_a)\sqrt{T-1}\phi'(\mu_a) = \sqrt{T-1}\phi(\mu_0) + \alpha.$$

which we can rewrite as follows

$$\sqrt{T-1} \left(\phi(\mu_a) + (\mu_0 - \mu_a) \phi'(\mu_a) - \phi(\mu_0) \right) = 2\alpha.$$
(7)

Because $\phi''(\mu) = -1/\phi(\mu)$, it holds that $\phi''(\mu)$ is everywhere weakly below³⁷

$$\underline{\phi}'' = \max_{\mu} \left[-\frac{1}{\phi(\mu)} \right] = -\frac{1}{\phi\left(\frac{1}{2}\right)} < 0,$$

and we obtain by simple integration that

$$\phi(\mu_0) \le \phi(\mu_a) + (\mu_0 - \mu_a)\phi'(\mu_a) + \frac{1}{2}(\mu_0 - \mu_a)^2 \underline{\phi}''.$$

Substituting into Equation 7

$$2\alpha = \sqrt{T-1} \left(\phi(\mu_a) + (\mu_0 - \mu_a) \phi'(\mu_a) - \phi(\mu_0) \right)$$

$$\geq -\frac{1}{2} \sqrt{T-1} (\mu_0 - \mu_a)^2 \underline{\phi}''$$

and rearranging we obtain

$$\mu_0 - \mu_a \le \sqrt{\frac{2\alpha}{-\frac{1}{2}\sqrt{T - 1}\underline{\phi}''}}.$$
(8)

Analogous manipulations applied to μ_b yield

$$\mu_a - \mu_b \le \sqrt{\frac{2\alpha}{-\frac{1}{2}\sqrt{T - 1}\underline{\phi}''}}.$$
(9)

Combining Equation 8 and Equation 9

$$\mu_0 - \mu_l \le 2\sqrt{\frac{2\alpha}{-\frac{1}{2}\sqrt{T-1}\underline{\phi}''}}.$$

The right-hand side is independent of μ_0 , and goes to 0 as T goes to infinity. This completes the proof.

B Constrained Information Policies

Throughout this section, we set surprise in period t to be $|\mu_t - \mu_{t-1}|$ and suspense to be the standard deviation of $\tilde{\mu}_{t+1}$.

 $^{^{37}\}phi\left(\frac{1}{2}\right)$ evaluates to the height of the standard normal pdf at 0, i.e., $\frac{1}{\sqrt{2\pi}}$.

B.1 Number of Games in a Finals Series

Proposition 4. For each $X \in \{1, ..., N\}$, define S(X, N) by

$$S(X,N) = Np^{X-1}(1-p)^{N-X} \binom{N-1}{X-1}$$

Then if the favored team needs to win X out of the remaining N games, the remaining suspense is $\sqrt{p(1-p)}S(X,N)$. The surprise is 2p(1-p)S(X,N).

Proof. Given X and N, the belief that the favored team wins the series is $\mu(X, N)$. If the team wins the current game, the belief jumps to $\mu(X - 1, N - 1)$; if the team loses, it falls to $\mu(X, N - 1)$. Let $\Delta(X, N) = \mu(X - 1, N - 1) - \mu(X, N - 1)$ be the range of next-period beliefs. $\Delta(X, N)$ is equal to the probability that the current game is marginal, i.e., that the favored team wins exactly X - 1 out of the following N - 1 games: $\Delta(X, N) = p^{X-1}(1-p)^{N-X} {N-1 \choose X-1}$.

If the favored team wins, the belief jumps by

$$\mu(X - 1, N - 1) - \mu(X, N) = (1 - p)\Delta(X, N).$$

If the favored team loses, the belief falls by

$$\mu(X, N) - \mu(X, N - 1) = p\Delta(X, N).$$

The expected surprise in the current period can thus be calculated to be $2p(1-p)\Delta(X,N)$, and the suspense is $\sqrt{p(1-p)}\Delta(X,N)$.

Define S(X, N) by induction on N, with S(X, N) = 0 if X = 0 or X > N and

$$S(X,N) = \Delta(X,N) + pS(X-1,N-1) + (1-p)S(X,N-1)$$

= $p^{X-1}(1-p)^{N-X} {N-1 \choose X-1} + pS(X-1,N-1) + (1-p)S(X,N-1),$

The suspense and surprise payoffs are constructed by the same recursion, inserting the appropriate coefficient on $\Delta(X, N)$. Suspense is thus $\sqrt{p(1-p)}S(X, N)$ and surprise is 2p(1-p)S(x, N). It remains only to show that S(X, N) has the explicit formula of $Np^{X-1}(1-p)^{N-X}\binom{N-1}{X-1}$. This follows from a simple induction on N, applying the binomial identity $\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k}$.

B.2 Political Primaries

Consider two states, *i* and *j*, which are holding consecutive primaries. Without loss of generality, suppose that *i* has weakly more delegates than $j: n_i \ge n_j$. Candidate *A* has a probability p_i of winning state *i*, and a probability p_j of winning state *j*.

Let s_i and s_j indicate the winners of the respective state primaries, A or B. Given past states' primary outcomes and taking expectations over the future states' primary outcomes, let Candidate A have a probability $\mu_{s_is_j}$ of winning the nomination given vote realizations s_i and s_j .

We seek to show that suspense and surprise are identical whether i votes before j or after. If we cannot affect expected suspense or surprise by swapping any such pair of states, then the suspense and surprise must be independent of the order of votes.

By the structure of the game, the candidate who wins a primary has a weakly higher chance of winning the nomination. That implies a monotonicity condition $\mu_{BB} \leq \mu_{BA}, \mu_{AB} \leq \mu_{AA}$. Because $n_i \geq n_j$, it also holds that $\mu_{BA} \leq \mu_{AB}$.

Given these definitions, the belief prior to the two primaries is $(1 - p_i)(1 - p_j)\mu_{BB} + p_i(1 - p_j)\mu_{AB} + (1 - p_i)p_j\mu_{BA} + p_ip_j\mu_{AA}$. If state *i* has its primary first, then the belief conditional on outcome s_i is $p_j\mu_{s_iA} + (1 - p_j)\mu_{s_iB}$. If state *j* has its primary first, then the belief conditional on outcome s_j is $p_i\mu_{As_j} + (1 - p_i)\mu_{Bs_j}$.

After some algebra we can express the expected surprise associated with either ordering as follows.

$$2\Big((1-p_j)p_j(\mu_{BA}-\mu_{BB})-p_i^2\big((-1+p_j)\mu_{BB}+\mu_{AB}-p_j(\mu_{BA}+\mu_{AB}-\mu_{AA})\big) +p_i\big((-1+p_j)^2\mu_{BB}-\mu_{AB}-(-2+p_j)p_j(\mu_{BA}+\mu_{AB}-\mu_{AA})\big)\Big).$$

For either ordering, the expected suspense payoff is

$$\sqrt{p_i(1-p_i)} \left(-(1-p_i)\mu_{BB} + \mu_{BA} - p_i(\mu_{BA} + \mu_{AB} + \mu_{AA}) \right) + \sqrt{p_j(1-p_j)} \left(-(1-p_j)\mu_{BB} + \mu_{AB} - p_j(\mu_{BA} + \mu_{AB} + \mu_{AA}) \right)$$

References

- AGUIAR, M., E. HURST, AND L. KARABARBOUNIS (2011): "Time use during recessions," Discussion paper, National Bureau of Economic Research.
- AKERLOF, G., AND W. DICKENS (1982): "The economic consequences of cognitive dissonance," *The American Economic Review*, 72(3), 307–319.
- AUGENBLICK, N., AND M. RABIN (2012): "A Model-Free Test of Excess Volatility in Beliefs Streams," unpublished.
- BARBERIS, N. (2012): "A Model of Casino Gambling," Management Science, 58(1), 35.
- BECKER, G., AND K. MURPHY (1988): "A theory of rational addiction," *The Journal of Political Economy*, pp. 675–700.
- (1993): "A simple theory of advertising as a good or bad," *The Quarterly Journal* of *Economics*, 108(4), 941–964.
- BORWEIN, D., J. BORWEIN, AND P. MARÉCHAL (2000): "Surprise maximization," *The American Mathematical Monthly*, 107(6), 517–527.
- BROCAS, I., AND J. CARRILLO (2007): "Influence through ignorance," *The RAND Journal* of *Economics*, 38(4), 931–947.
- CAPLIN, A., AND J. LEAHY (2001): "Psychological expected utility theory and anticipatory feelings," *The Quarterly Journal of Economics*, 116(1), 55–79.
- CHAN, W., P. COURTY, AND H. LI (2009): "Suspense: Dynamic Incentives in Sports Contests"," *The Economic Journal*, 119(534), 24–46.
- CHRISTENFELD, N., AND J. LEAVITT (2011): "Story Spoilers Don't Spoil Stories," *Psychological Science*, 22(9).
- DE MEYER, B. (1998): "The maximal variation of a bounded martingale and the central limit theorem," in Annales de l'Institut Henri Poincare (B) Probability and Statistics, vol. 34, pp. 49–59. Elsevier.
- DILLENBERGER, D. (2010): "Preferences for One-Shot Resolution of Uncertainty and Allais-Type Behavior," *Econometrica*, 78(6), 1973–2004.

- DOWNS, A. (1957): "An economic theory of political action in a democracy," *The Journal* of *Political Economy*, 65(2), 135–150.
- ELIAZ, K., AND A. SCHOTTER (2010): "Paying for confidence: An experimental study of the demand for non-instrumental information," *Games and Economic Behavior*, 70(2), 304–324.
- ELIAZ, K., AND R. SPIEGLER (2006): "Can anticipatory feelings explain anomalous choices of information sources?," *Games and Economic Behavior*, 56(1), 87–104.
- FIORILLO, C., P. TOBLER, AND W. SCHULTZ (2003): "Discrete coding of reward probability and uncertainty by dopamine neurons," *Science's STKE*, 299(5614), 1898.
- GARDNER, M. (1969): The Unexpected Hanging And Other Mathematical Diversions. Simon and Schuster.
- GEANAKOPLOS, J. (1996): "The Hangman's Paradox and Newcomb's Paradox as psychological games," *Cowles Foundation Discussion Paper No. 1128*.
- GEANAKOPLOS, J., D. PEARCE, AND E. STACCHETTI (1989): "Psychological games and sequential rationality," *Games and Economic Behavior*, 1(1), 60–79.
- GENTZKOW, M., AND E. KAMENICA (2011): "Competition in persuasion," Discussion paper, National Bureau of Economic Research.
- HÖRNER, J., AND A. SKRZYPACZ (2011): "Selling information," Cowles Foundation Discussion Paper No. 1743R.
- HUMMEL, P., AND R. HOLDEN (2012): "Optimal primaries," http://patrickhummel. webs.com/wpo.pdf.
- KAHNEMAN, D., AND D. MILLER (1986): "Norm theory: Comparing reality to its alternatives.," *Psychological review*, 93(2), 136.
- KAMENICA, E., AND M. GENTZKOW (2011): "Bayesian Persuasion," The American Economic Review, 101(6), 2590–2615.
- KŐSZEGI, B., AND M. RABIN (2009): "Reference-dependent consumption plans," *The American Economic Review*, 99(3), 909–936.

- KNIGHT, B., AND N. SCHIFF (2007): "Momentum and social learning in presidential primaries," NBER Working Paper No. 13637.
- KREPS, D., AND E. PORTEUS (1978): "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, pp. 185–200.
- KYLE, A. (1985): "Continuous auctions and insider trading," Econometrica: Journal of the Econometric Society, pp. 1315–1335.
- MERTENS, J., AND S. ZAMIR (1977): "The maximal variation of a bounded martingale," Israel Journal of Mathematics, 27(3), 252–276.
- OSTROVSKY, M. (2012): "Information Aggregation in Dynamic Markets with Strategic Traders," *Econometrica*.
- OSTROVSKY, M., AND M. SCHWARZ (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2(2), 34–63.
- RAYO, L., AND I. SEGAL (2009): "Optimal Information Disclosure," *Journal of Political Economy*.
- ROCKAFELLAR, R. (1997): Convex analysis. Princeton Univ Press.
- STIGLER, G., AND G. BECKER (1977): "De gustibus non est disputandum," The American Economic Review, 67(2), 76–90.
- TAMURA, W. (2012): "A Theory of Multidimensional Information Disclosure," ISER Discussion Paper No. 828.