Fear of Missing Out

Pragmatic Beliefs and the Reception of New Technologies

Niels Boissonnet* Francesco Rocciolo†

November 2026

Abstract

When a new technology becomes available, people who use an older one often change their view of the future without having learned anything. We model this, with what people choose as the only data. A decision maker entertains several scenarios about the world and holds a baseline confidence in each. Facing the technologies on the table, one of which she uses, she gives more weight to the scenarios under which she would miss the most by standing pat: she fears missing out. She then maximizes expected utility. Three axioms on her choices, added to expected utility, characterize the model. The weak axiom of revealed preference holds across menus whose best payoffs, scenario by scenario, differ by a constant; translating a whole situation changes nothing; and a newcomer can move her choice among the old options only toward those that do well where it shines. Choices reveal her weights exactly, and an option that nobody chooses can reverse a choice. Evidence and stakes then enter her beliefs as two likelihoods that do not interact. A new technology is received through how far it pushes the frontier of what is already available, and radical technologies are favored. In a market, others' adoption raises what each firm would miss by staying behind, and fear of missing out turns this payoff externality into a belief externality: adoption can tip, and rest on a confidence that no information supports; when the value of the technology lies in not falling behind, more adoption is worse for everyone, and a useless technology can be adopted by all. After the launch, holding the new technology turns the fear around: adopters fear having missed the scenario under which they should have stayed. Markets then churn and reverse their waves, unless switching costs make adopters loyal.

1 Introduction

Consider a firm that runs its operations on an established system. A new system has been announced, and its specifications have been public for months. Nobody knows whether it is a modest improvement or the beginning of a transformation of the industry; the firm entertains both models of the future, which we call the incremental and the transformative *scenarios*. On the day of the release, nothing is learned about which scenario is true. What changes is that the new system, until then on paper, is now on the table. Yet on that day the firm's view of the future often changes. The question its managers ask is not "how likely is it that this technology is transformative?" but "can we afford to miss it?". The second question makes the transformative scenario loom larger, for no other reason than that the firm's current system would be a costly mistake under it. In everyday language, this is fear of missing out.

*Department of Economics, IMW, and JICE, Bielefeld University, Germany. Email: niels.boissonnet@uni-bielefeld.de.

†Graduate School of Business, Nazarbayev University, Astana, Kazakhstan. Email: francesco.rocciolo@nu.edu.kz.

Example 1 puts numbers on it. The firm's baseline confidence in the transformative scenario is 10%, and at that confidence the new system is worth less to it than its own: a Bayesian firm would keep what it has. On the day of the release, a firm moved by fear of missing out gives that scenario a weight of 45%, and adopts. It has learned nothing.

Nothing here is specific to technology. A buyer offered an extended warranty at the checkout faces the same situation: nothing about the product has changed between the shelf and the till, but the warranty is now on the table, and the scenario under which the product breaks looms larger. So does a citizen asked, for the first time, to vote on leaving a union.[1]

Standard tools have little to say about these cases. A Bayesian changes her beliefs only when she learns something, and putting an option on the table is not news about which scenario is true. Theories of salience and focusing (Bordalo et al., 2012, 2013; Kőszegi and Szeidl, 2013) let the options at hand affect how a decision maker weighs states or attributes, but they measure the contrast among the options, whatever she holds and in both directions: a scenario under which the new system would be a disaster is as salient as one under which it would be a triumph. Theories of motivated beliefs (Brunnermeier and Parker, 2005) let beliefs respond to payoffs, but toward what the decision maker would like to be true. That is not what the firm does. Its current system is excellent under the incremental scenario, and fear of missing out makes it think less of that scenario, not more.

This paper proposes a pragmatic theory of how the weights on scenarios are formed, pragmatic in the sense of the philosophical tradition for which a belief is to be understood through its consequences for action.[2] The decision maker (DM) entertains finitely many scenarios, probabilistic models of the world, and holds a prior Q_{base} over them, her baseline confidence: her assessment, on the evidence alone, of how well each is specified. She faces a menu $\mathcal{F}$ of acts, the options on the table, from a *status quo* $h \in \mathcal{F}$, the act she holds. For each scenario p, the *stakes* $\rho(p)$ measure, in utility units, by how much the best option on the table improves on the status quo if p is true: what she would miss by standing pat. Her weights tilt the baseline toward the scenarios with high stakes,

$$Q^*(p) \;\propto\; Q_{\text{base}}(p)\, e^{\alpha\, \rho(p)}, \qquad \alpha \geq 0,$$

and she maximizes expected utility under the prior that these weights put on the states. The parameter α measures how much she lets what she could miss move her. When $\alpha = 0$ she is a Bayesian whom the situation does not touch. Options she knows only on paper play no role, however well she knows them.

The contribution is twofold. First, we characterize the model from choices alone. The primitive is a choice correspondence: what the DM chooses from each menu when she holds each status quo. A theory for which a belief is to be understood through action should need no other data. The difficulty is well known: to offer a bet is to put it on the table, and so to change the situation it is meant to measure. This is true only of bets that push the *frontier*, the best that the table offers under each scenario. A bet below the frontier leaves the stakes where they are, and choices among such bets reveal the weights exactly (Proposition 1). A first axiom makes the DM an expected

[1] As usual in such applications, we read the voter as judging the outcome as if her vote decided it.

[2] The tradition starts with Peirce's maxim: "Consider what effects, that might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of these effects is the whole of our conception of the object" (Peirce, 1878, p. 132). Decision theory inherited the idea through Ramsey (1926), who measured a degree of belief by the extent to which one is prepared to act on it. Our theory goes one step further, toward James (1896): what is at stake for action also affects how much weight a model receives. James restricted this to options that are living, forced and momentous, and each of his conditions has a counterpart here. A choice is living when its options are on the table, forced because standing pat is itself one of them, and momentous when much is at stake. And the stakes shift weight only among the models that the evidence leaves standing (Corollary 1).

utility maximizer as long as what she holds and the frontier stay put; it says all that we need, and standard conditions on choices imply it (Appendix B). Three further axioms say how her choices change when they move. *Frontier WARP*: the weak axiom of revealed preference holds across menus whose frontiers differ by a constant. *Translation invariance*: only the gaps between options matter, not their levels. *Arrival monotonicity*: a newcomer that shines under one scenario can move the choice among the old options only toward those that do better there. With three or more scenarios, these axioms characterize the model (Theorem 1). The exponential form comes from the first of them, much as constant absolute risk aversion comes from asking that a sure change in wealth leave the attitude to risk alone. With two scenarios, the case of our market, one function stays free, and the exponential needs one more axiom, *Common reception*: a new option moves the odds of two users who face the same menu by the same factor, whatever each of them holds. The parameters are then identified. Evidence moves the baseline and the stakes move the weights away from it, and the two act as separate likelihoods: they do not interact, and α is an exchange rate that says how much evidence a unit of stakes is worth (Proposition 3). Stakes of bounded size can overturn only a bounded amount of evidence, so that pragmatic weights matter where the evidence leaves several scenarios standing (Corollary 1). This is the case of a new technology, whose scenarios agree on the past and differ on the future. When the menu contains a sure thing that is at least as good as every option under every scenario, the weights take the functional form of a known rule of payoff-dependent beliefs (Mayraz, 2011a), with the sign of pessimism: the mirror image of wishful thinking (Kovach, 2020). What the menu adds is that two menus move the DM's beliefs differently exactly when their frontiers differ by more than a constant (Proposition 4). Her weights are also those that an adversary would choose to maximize the regret of her status quo at an entropic cost (Proposition 5). Choice data also carry predictions of their own. An option that nobody chooses can reverse the choice between two others; on menus of two options, the model cannot be told apart from a model of tastes (Proposition 6); and a menu has at most one rest point, whatever its size (Proposition 11).

Second, we study the reception of a new technology, first by a single user and then in a market. For a single user, an entrant is received through how far it pushes the frontier of the installed base, the technologies already on the table: the shift in weights it causes is a Bayesian update with likelihood $e^{\alpha\Delta(p)}$, where $\Delta(p)$ measures this extension (Proposition 7). An entrant that beats the best incumbent under no scenario is inert, however bad it is under some. An entrant that is better than the status quo by D under one scenario and worse by L under another is adopted if and only if $L < (q_2/q_1)\, D\, e^{\alpha D}$: fear of missing out inflates the upside exponentially, and radical technologies are favored over timid ones with the same ratio of upside to downside (Proposition 8). Users who hold different technologies disagree, each overweighting the scenarios under which her own technology does badly (Proposition 9), and a user who judges from the technology she last chose may switch back and forth forever (Proposition 10). Foresight in valuation does not help her: when switching is free, a DM who values the future, but whose attention is set by the current menu, follows the cycle all the same (Proposition 12).

Part of what a firm fears to miss, however, is what its competitors will do with the new technology. Section 6 studies a market in which the share of adopters raises what each firm would lose, under the transformative scenario, by staying with the old technology. This is a payoff externality of the usual kind, but fear of missing out turns it into a belief externality: the more firms adopt, the more each firm attends to the transformative scenario, although others' adoption says nothing about whether the technology is transformative. The belief channel multiplies the response of the marginal adopter to others' adoption by a factor that grows exponentially with the upside (Proposition 14). Where a Bayesian market has a unique equilibrium, a FOMO market can have three, a low one and full adoption separated by a tipping point (Proposition 15); as the technology

improves, adoption creeps up and then jumps to everyone (Proposition 16). In equilibrium, firms give the transformative scenario a weight that rises with adoption and that no evidence supports: in our example, 0.86 at full adoption against a true probability of 0.10 (Proposition 17). This is herding without information. When the value of the technology lies in not falling behind, equilibria with more adoption are Pareto-worse, and a technology with no intrinsic value can be adopted by every firm (Proposition 18 and Corollary 2). When its value lies instead in joining others, firms that adopt in both of two equilibria, or in neither, prefer more adoption; the cost of fear of missing out then falls on the firms it pushes into adopting, which can be most of the market (Proposition 19).

These are properties of the market at the launch, when every firm still judges from the old technology. Section 7 follows the market afterwards, when each firm judges from what it holds. Holding the new technology turns the fear around: an adopter fears having missed the scenario under which it should have stayed. A positive mass of firms then switches back and forth forever, the long run averages two responses, and a launch wave is followed by a decline; in our example, full adoption collapses to 5% (Theorem 2 and Proposition 20). Fear of missing out makes firms eager and disloyal, so that in the long run it helps poor technologies and can hurt excellent ones (Proposition 21). Switching costs are what make adopters loyal, and they buy loyalty at the price of eagerness (Proposition 22).

Section 9 argues that the model formalizes an idea worth taking seriously: that attention, unlike evidence, may follow the stakes. It explains why fear of missing out is not wishful thinking, in what sense the DM remains a Bayesian, and why the same vigilance, once it is contagious, turns into collective credulity and then into collective remorse.

Related literature. The diffusion of new technologies has been studied since Griliches (1957) and Rogers (1962) through learning, network effects and heterogeneity in returns. Our DM learns nothing, and her reception of an entrant depends on what is already on the table. In informational cascades (Banerjee, 1992; Bikhchandani et al., 1992), others' adoption moves beliefs because it reveals information about the state; in our market it reveals nothing and moves beliefs anyway, because it raises the stakes. Network effects (Katz and Shapiro, 1985; Farrell and Saloner, 1985) and threshold models (Granovetter, 1978) produce tipping through complementarities in payoffs, and increasing returns can lock a market into a technology chosen by accident (Arthur, 1989). In our market, when downsides have a concave distribution, the payoff complementarity alone yields a unique equilibrium; multiplicity comes from beliefs amplifying it. Guimaraes et al. (2025) model fear of missing out as a cost borne by the non-users of a popular network, as our competitive pressure is, and find product market traps, which Bursztyn et al. (2025) measure for social media; we add the belief channel. Rational herding can generate booms and busts of technology adoption (Schaal and Taschereau-Dumouchel, 2023); ours arise without any information. The proofs use the monotone methods of games with strategic complementarities (Milgrom and Roberts, 1990, 1994), and our market with memory is a population game with revision opportunities (Sandholm, 2010) in which the best response depends on the strategy a firm currently plays. Salience theory (Bordalo et al., 2012, 2013) distorts the decision weights of states according to how much the options' payoffs contrast in them; focusing (Kőszegi and Szeidl, 2013) and relative thinking (Bushong et al., 2021) weigh attributes by the range of the options in the choice set. Our stakes measure contrast from the status quo, and only upward: what the DM could gain by moving, not what she could lose. This is what separates fear of missing out from salience, and what generates churn. Our stakes are the regret of the status quo in the sense of Savage (1951), computed scenario by scenario; regret theory (Bell, 1982; Loomes and Sugden, 1982) puts regret in the utility function, whereas our DM is not averse to regret but attentive to it, and in Hayashi (2008) anticipated regret makes the

evaluation of an act depend on the opportunities available, whereas in our model the menu moves the likelihoods and the evaluation stays expected utility. Our weights are those an adversary would choose to maximize the expected regret of the status quo at an entropic cost: the worst case of the multiplier preferences of Hansen and Sargent (2008) and Strzalecki (2011), with the status quo's regret in place of utility. As α grows, they concentrate on the scenarios under which this regret is largest; models of minimax regret derive beliefs in a loosely related way (Suryanarayanan, 2006). Status quo bias (Samuelson and Zeckhauser, 1988; Masatlioglu and Ok, 2005) makes the status quo sticky; fear of missing out makes it restless. Motivated beliefs (Brunnermeier and Parker, 2005) favor pleasant scenarios; ours favor pivotal ones. Closest to us are models in which beliefs depend on payoffs. In Mayraz (2011a), the payoff of an event under the act the DM chooses acts as evidence about the event, with a parameter whose sign separates optimism from pessimism; Kovach (2020) axiomatizes wishful thinking, in which an endowment tilts beliefs toward the states where it pays well. In both, beliefs depend on a single act. In ours they also depend on what else is on the table: when the menu contains a sure thing that is at least as good as every option under every scenario, our weights have Mayraz's functional form, with the sign of pessimism and the status quo as the reference act; otherwise the menu matters through its frontier, and only through it (Proposition 4). In Ortoleva (2010) beliefs do not move: the status quo changes the decision rule, which compares options under a fixed set of priors. All these models take as primitive a family of preferences: one for each reference act in Mayraz (2011a), one for each endowment in Kovach (2020), one for each menu in Piermont (2017), one for each status quo in Ortoleva (2010). Choice correspondences with a status quo are the primitive of Masatlioglu and Ok (2005, 2014) and, over acts, of Riella and Teper (2014); all of them assume the weak axiom for each status quo, which our DM violates across menus whose frontiers are not parallel. Choice correspondences over menus of acts, without a status quo, are the primitive of the theory of minimax regret (Hayashi, 2008; Stoye, 2011). Our Frontier WARP, for menus with the same frontier, is its "independence of never strictly optimal alternatives" (Stoye, 2011). Two things differ. There the menu enters the evaluation of each act; here it enters the weights, and the evaluation stays expected utility. And there, mixing a whole menu with an act leaves the choice unchanged; here, this shrinks the stakes and moves the weights, and Translation invariance is what survives of that axiom. To our knowledge, menu-dependent beliefs have not been characterized from a choice correspondence. Novák et al. (2024) obtain polarization by status quo through rational information acquisition; our users disagree without acquiring any information. The setting with given scenarios is that of decisions with objective models (Cerreia-Vioglio et al., 2013; Marinacci, 2015), from which we borrow the consistency axiom; models of ambiguity (Gilboa and Schmeidler, 1989; Klibanoff et al., 2005; Hansen and Sargent, 2008) describe attitudes toward a fixed model uncertainty, while our DM selects within it. Lanzani (2025) lets an agent's confidence in her models respond to how well they fit the data; our baseline may be formed as in that paper, and what we add is that the weights also respond to what is at stake. With misspecified models, Bayesian learning concentrates on the models closest to the truth (Berk, 1966; Esponda and Pouzo, 2016), and Corollary 1 says that bounded stakes do not prevent this. The exponential form is the one produced by an entropic attention cost, as in rational inattention (Sims, 2003; Matějka and McKay, 2015) and sparsity (Gabaix, 2014). Finally, psychologists describe fear of missing out as the apprehension that others are having rewarding experiences from which one is absent (Przybylski et al., 2013); our notion is its decision-theoretic counterpart, without the social comparison.

2 The model

2.1 Primitives

Let S be a finite state space with $|S| \geq 2$, and $\Delta(S)$ the set of probability distributions on S, viewed as vectors of $\mathbb{R}^S$. Let $\mathcal{P} \subseteq \Delta(S)$ be a finite set of *scenarios*, with $|\mathcal{P}| \geq 2$, assumed linearly independent as vectors of $\mathbb{R}^S$. A scenario is a probabilistic model of the world. Linear independence holds whenever the scenarios are Dirac measures, or are two distinct distributions, and it is what allows the weights on scenarios to be read off choices. For $k \in \mathbb{R}^{\mathcal{P}}$ we write $k(p)$ for its coordinate at p, $\mathbf{1}_q$ for the unit vectors and $\mathbf{1} = \sum_q \mathbf{1}_q$; $\mathbb{R}^{\mathcal{P}}_+$ is the nonnegative orthant.

Let X be a set of consequences and $\Delta(X)$ the set of simple lotteries on X. An *act* is a map $f : S \to \Delta(X)$ (Anscombe and Aumann, 1963); acts are mixed pointwise, and a lottery is identified with the constant act that yields it. For an act f and a scenario p, the *p-reduction* of f is the lottery $f^p = \sum_s p(s)f(s)$: what the DM faces with f if p is the true scenario. In the applications, acts are technologies.

A *situation* is a pair $\sigma = (\mathcal{F}, h)$ where $\mathcal{F}$ is a nonempty finite set of acts, the *menu* or the options on the table, and $h \in \mathcal{F}$ is the *status quo*. The primitive is a *choice correspondence* C, which assigns to every situation a nonempty set $C(\mathcal{F}, h) \subseteq \mathcal{F}$: what the DM may choose from the menu $\mathcal{F}$ when she holds h. How she ranks acts that are not on the menu is not part of the data. A situation $(\{h\}, h)$ with a singleton menu is a *bare status quo*.

2.2 Stakes and the representation

Fix an affine utility index u on lotteries; Axiom 1 below provides it. The *profile* of an act f is the vector $u(f) = (u(f^p))_{p\in\mathcal{P}}$; when u is onto, every vector of $\mathbb{R}^{\mathcal{P}}$ is the profile of some act (Lemma A.1). Under a scenario p, the menu's *frontier* is $b_{\mathcal{F}}(p) = \max_{f\in\mathcal{F}} u(f^p)$, and the *stakes* of the situation $(\mathcal{F}, h)$ under p are

$$\rho_{\mathcal{F},h}(p) = b_{\mathcal{F}}(p) - u(h^p) \ \geq 0 :$$

what the DM would miss, under p, by standing pat while a better option is on the table. It is the regret of the status quo under p, in the sense of Savage (1951). We call $\rho_{\mathcal{F},h} \in \mathbb{R}^{\mathcal{P}}_+$ the *stakes profile* of the situation.

Definition 1. A *pragmatic representation* of C is a triple $(u, Q_{\text{base}}, \alpha)$, where $u : \Delta(X) \to \mathbb{R}$ is affine and onto, $Q_{\text{base}} \in \Delta(\mathcal{P})$ has full support and $\alpha \geq 0$, such that for every situation $(\mathcal{F}, h)$, $C(\mathcal{F}, h) = \arg\max_{g\in\mathcal{F}} V_{\mathcal{F},h}(g)$, where

$$V_{\mathcal{F},h}(g) = \sum_{p\in\mathcal{P}} Q^*_{\mathcal{F},h}(p)\, u(g^p), \qquad Q^*_{\mathcal{F},h}(p) = \frac{Q_{\text{base}}(p)\, e^{\alpha\rho_{\mathcal{F},h}(p)}}{\sum_{q\in\mathcal{P}} Q_{\text{base}}(q)\, e^{\alpha\rho_{\mathcal{F},h}(q)}}.$$

Since u is affine, $u(g^p)$ is the expected utility of g under p, and averaging commutes with expectation:

$$V_{\mathcal{F},h}(g) = \sum_{s\in S} \bar{p}_{\mathcal{F},h}(s)\, u(g(s)), \qquad \bar{p}_{\mathcal{F},h} = \sum_{p\in\mathcal{P}} Q^*_{\mathcal{F},h}(p)\, p.$$

The DM is a subjective expected utility maximizer in every situation. What the theory adds is a rule for her prior: a mixture of the scenarios, with weights that start from Q_{base} and tilt toward the scenarios with high stakes. The weights are Q_{base} whenever the stakes are the same under every scenario, whatever their size: Q_{base} is her confidence when she would miss the same amount

whatever the truth. The weights solve

$$Q^*_{\mathcal{F},h} \in \arg\max_{Q\in\Delta(\mathcal{P})} \; \alpha \sum_p Q(p)\,\rho_{\mathcal{F},h}(p) - \mathrm{KL}(Q \,\|\, Q_{\text{base}}) :$$

the DM shifts attention toward the scenarios where the choice matters, at a cost that grows with the distance from her baseline. The axioms below do not presuppose this reading; Section 4.3 gives it a second one.

Example 1 (The firm)**.** Two scenarios, "incremental" (p_1) and "transformative" (p_2); to keep the numbers simple, each pins down a state, so that reductions are state-wise consequences. The established system h yields utilities $(10, 4)$: it does well if the world does not change, and badly if competitors transform their operations. The new system e yields $(7, 12)$: switching costs something if the technology is incremental, and pays off if it is transformative. Baseline confidence is $Q_{\text{base}} = (0.9, 0.1)$. A Bayesian firm, which has these weights in every situation, keeps h when e arrives (9.4 against 7.5). For a firm that holds h with e on the table, the stakes profile of $(\{h, e\}, h)$ is $\rho = (0, 8)$: the firm would miss nothing by standing pat if the change is incremental, and 8 if it is transformative. With $\alpha = 1/4$, the weight on the transformative scenario rises from 0.10 to 0.45, and the firm adopts e (9.25 against 7.29). Nothing has been learned. Note which scenario the firm attends to: not the pleasant one, since its current system does well under the incremental scenario, but the one under which standing pat would be a mistake.

2.3 Choices reveal the weights

The DM's *odds* of p against p' are $o^*_{\mathcal{F},h}(p, p') = Q^*_{\mathcal{F},h}(p)/Q^*_{\mathcal{F},h}(p')$, and her baseline odds are $o_{\text{base}}(p, p') = Q_{\text{base}}(p)/Q_{\text{base}}(p')$. One may object that choices cannot reveal them: to offer a bet is to put it on the table, and so to change the situation it is meant to measure. The objection is right for bets that push the frontier. It fails for the others. An act e is *below the frontier* of $\mathcal{F}$ if $u(e) \le b_{\mathcal{F}}$: it beats the best option on the table under no scenario. Such an act leaves the stakes where they are, and since utility is unbounded, there is room for it in every situation in which something is at stake.

Proposition 1 (Exact revelation)**.** *Let C have a pragmatic representation $(u, Q_{\text{base}}, \alpha)$. Fix a situation $(\mathcal{F}, h)$, write $Q = Q^*_{\mathcal{F},h}$ and $V = V_{\mathcal{F},h}$, and let*

$$D = \sum_p Q(p)\, b_{\mathcal{F}}(p) - \max_{f\in\mathcal{F}} V(f) \;\ge 0.$$

(i) $D = 0$ if and only if some option of $\mathcal{F}$ is best under every scenario.

(ii) Let $D > 0$ and $\delta \in (0, D)$. Let g be an act with profile $b_{\mathcal{F}} - \delta\mathbf{1}$ and, for distinct scenarios p, p', $x \in (0, \delta]$ and $y > 0$, let $g_{x,y}$ be an act with profile $u(g) + x\mathbf{1}_p - y\mathbf{1}_{p'}$. Then

$$g_{x,y} \in C\big(\mathcal{F} \cup \{g, g_{x,y}\}, h\big) \quad\Longleftrightarrow\quad y \le o^*_{\mathcal{F},h}(p, p')\, x.$$

(iii) Let $D = 0$ and $\rho_{\mathcal{F},h} \neq 0$. Replace every option other than h by options e_p, one for each scenario with $\rho_{\mathcal{F},h}(p) > 0$, where e_p reaches the frontier under p and is strictly worse than h under every other scenario. The new menu has the same status quo, the same frontier and the same weights, no option of it is best under every scenario, and (ii) applies to it.

(iv) *If $\rho_{\mathcal{F},h} = 0$, the DM chooses h, or an act with the same profile, from every menu that contains h and has the frontier of $\mathcal{F}$. The baseline Q_{base} is the weight vector of every situation whose stakes are the same positive number under all scenarios.*

(v) (Bare status quo.) *Let h be an act, $p \neq p'$, $\varepsilon, t > 0$, and let $b_{\varepsilon,t}$ be an act with profile $u(h) + \varepsilon \mathbf{1}_p - t\varepsilon \mathbf{1}_{p'}$. Then $b_{\varepsilon,t} \in C(\{h, b_{\varepsilon,t}\}, h)$ if and only if $\log t \leq \log o_{\text{base}}(p, p') + \alpha\varepsilon$.*

The analyst need not know D: she lowers δ until g alone is chosen from $\mathcal{F} \cup \{g\}$. Nor must the menu have a unique best element. Part (iv) says what choices cannot do. With nothing at stake, the DM keeps what she holds, and her choices among the options on the table reveal nothing. The baseline can be read in two other places: in the situations in which she would miss the same amount whatever the truth, without extrapolation, and in part (v), where a bet against a bare status quo is the only alternative. There the bet does move the weights, and the odds at which she takes it rise exponentially with its upside: the intercept of the log-odds is her baseline and the slope is α. All parameters are thus identified from choices, together with choices among lotteries, which identify u.

What the baseline measures. Since Q_{base} is what the DM reveals when what she would miss does not depend on the scenario, it records her judgment of the scenarios on the evidence alone. We read it as her assessment of how well each scenario is specified relative to the others. If she is confident that one of her scenarios is true, $Q_{\text{base}}(p)$ is her probability that p is that scenario, as in decisions with objective models (Cerreia-Vioglio et al., 2013; Marinacci, 2015). If she suspects that every scenario is only an approximation, $Q_{\text{base}}(p)$ is her trust in p relative to the others, and Bayesian learning gives this a definite content: when all models are misspecified, the posterior concentrates on those closest to the truth (Berk, 1966; Esponda and Pouzo, 2016). What the baseline cannot express is the doubt that all scenarios may be wrong together; Lanzani (2025) represents that doubt by a separate parameter, and our model has no counterpart of it. The reading separates two forces that a Bayesian merges: evidence moves the baseline, and the situation moves the weights away from it. Section 4.1 makes the separation precise.

3 Axioms

3.1 Expected utility below the frontier

Two menus have *the same frontier* if $b_{\mathcal{F}} = b_{\mathcal{G}}$, and *parallel frontiers* if $b_{\mathcal{F}} - b_{\mathcal{G}}$ is constant across scenarios. The first axiom says what we need about choices for all that follows.

Axiom 1 (Expected utility below the frontier)**.** There is an affine and onto $u : \Delta(X) \to \mathbb{R}$ such that, with frontiers computed from u, every situation $(\mathcal{F}, h)$ has a full-support $Q_{\mathcal{F},h} \in \Delta(\mathcal{P})$ with

$$C(\mathcal{G}, h) = \arg\max_{g \in \mathcal{G}} \sum_p Q_{\mathcal{F},h}(p)\, u(g^p) \qquad \text{for every menu } \mathcal{G} \ni h \text{ with the frontier of } \mathcal{F}.$$

As long as what she holds and the frontier do not change, the DM is an expected utility maximizer. Two things are assumed. Within a situation she maximizes expected utility, and she dismisses no scenario. Across menus, her weights depend on the menu only through its frontier: an option below the frontier may be chosen, but it does not change the choice among the others. A DM who also attended to the downside of the menu would violate this. When nothing is at stake, the axiom says only that she keeps h or an act with the same profile, since every act of such a menu is weakly below h under every scenario.

The axiom has a representation in it, and this is its price. Appendix B shows that nothing is lost: it holds if and only if C satisfies five conditions that mention no weights (Proposition 23). On lotteries the weights play no role, and C chooses what maximizes u. Write $x \succsim y$ if $u(x) \geq u(y)$. Then $x \succsim y$ if and only if $x \in C(\{x, y\}, y)$, so that $\succsim$ is revealed by choices, and u is unique up to positive affine transformations by the von Neumann–Morgenstern theorem.[3] When something is at stake the weights are unique (Appendix A). The DM's *revealed odds* of p against p' in a situation $\sigma = (\mathcal{F}, h)$ with $\rho_{\mathcal{F},h} \neq 0$ are the worst terms at which she takes a bet on p against p': the largest $t > 0$ such that some menu $\mathcal{G} \ni h$ with the frontier of $\mathcal{F}$ has $f \in C(\mathcal{G}, h)$ and $g \in \mathcal{G}$ with $u(f) - u(g) = x(\mathbf{1}_p - t\mathbf{1}_{p'})$ for some $x > 0$. By Axiom 1, $o_\sigma(p, p') = Q_\sigma(p)/Q_\sigma(p')$, with no limit to take. The odds are thus an object of the data, as the stakes are.

3.2 How the situation moves the choice

Axiom 2 (Frontier WARP)**.** Let $\mathcal{F}$ and $\mathcal{G}$ contain h and have parallel frontiers. If $f, g \in \mathcal{F} \cap \mathcal{G}$, $f \in C(\mathcal{F}, h)$ and $g \in C(\mathcal{G}, h)$, then $f \in C(\mathcal{G}, h)$.

For a given status quo, the DM violates the weak axiom of revealed preference only across menus whose frontiers are not parallel. For menus with the same frontier, Axiom 1 already says so, and this half is known from minimax regret. Take now a frontier that lies above another by the same amount under every scenario. The axiom says that such a uniform improvement is irrelevant too. A situation in which the DM could miss a great deal, but the same amount whatever the truth, gives her no reason to favor one scenario: she attends to relative stakes. A single option cannot carry the improvement, since an option that beats the frontier by the same amount under every scenario is better than everything else on the table and is chosen whatever the weights. This is why the axiom compares menus. We have not found this half in the literature.

Axiom 3 (Translation invariance)**.** Let $\tau : \mathcal{F} \to \mathcal{G}$ be a bijection such that, for all $f, g \in \mathcal{F}$ and every scenario q,

$$\tfrac{1}{2}\,\tau(f)^q + \tfrac{1}{2}\,g^q \;\sim\; \tfrac{1}{2}\,\tau(g)^q + \tfrac{1}{2}\,f^q.$$

Then $C(\mathcal{G}, \tau(h)) = \tau\big(C(\mathcal{F}, h)\big)$ for every $h \in \mathcal{F}$.

The map τ moves every option by the same amount under each scenario: by the affinity of u, $u(\tau(f)) - u(f)$ is the same vector for every $f \in \mathcal{F}$. The axiom says that only the gaps between the options matter, not their levels. A wishful thinker, whose beliefs follow how well she fares with what she holds, violates the axiom: raise every option by the same amount under one scenario, and she comes to believe in that scenario. So does a DM whose attention follows monetary stakes while her utility is not linear in money, since equal gaps in utility are not equal gaps in money.

Axiom 4 (Arrival monotonicity)**.** Let e be an act that beats the frontier of $\mathcal{F}$ under the scenario p and under no other. If $f, g \in \mathcal{F}$, $f \in C(\mathcal{F}, h)$, $g \in C(\mathcal{F} \cup \{e\}, h)$ and $f^p \succsim g^p$, then $f \in C(\mathcal{F} \cup \{e\}, h)$.

A newcomer that shines under p can take the choice away from f only to give it to itself or to an old option that does strictly better than f under p. In Example 1, the arrival of e can push the firm only toward what does well if the technology is transformative. What the axiom rules out is a DM who turns away from a scenario because she would miss more under it.

With two scenarios, a last axiom is needed. It is stated on the odds that choices reveal.

[3] Appendix B.1 removes u from the wording of the axioms.

Axiom 5 (Common reception)**.** For every menu $\mathcal{F}$, all $h, h' \in \mathcal{F}$ with $\rho_{\mathcal{F},h} \neq 0 \neq \rho_{\mathcal{F},h'}$, every act e and all $p \neq p'$,

$$\frac{o_{\mathcal{F}\cup\{e\},h}(p,p')}{o_{\mathcal{F},h}(p,p')} = \frac{o_{\mathcal{F}\cup\{e\},h'}(p,p')}{o_{\mathcal{F},h'}(p,p')}.$$

This axiom says that a new option moves the odds of two users who face the same menu by the same factor, whatever each of them holds. The two users may hold different beliefs before the arrival, because they hold different technologies; the arrival revises these beliefs by a common likelihood. What the axiom rules out is a DM whose reaction to a newcomer depends on how far she already lags behind the frontier.

4 Representation

Definition 2. A *separable representation* of C is a triple $\big(u, Q_{\text{base}}, (\varphi_p)_{p\in\mathcal{P}}\big)$, where $u : \Delta(X) \to \mathbb{R}$ is affine and onto, $Q_{\text{base}} \in \Delta(\mathcal{P})$ has full support, and each $\varphi_p : \mathbb{R}_+ \to \mathbb{R}$ satisfies $\varphi_p(0) = 0$, such that for every situation $C(\mathcal{F}, h) = \arg\max_{g\in\mathcal{F}} \sum_p Q^{\varphi}_{\mathcal{F},h}(p)\, u(g^p)$, where

$$Q^{\varphi}_{\mathcal{F},h}(p) = \frac{Q_{\text{base}}(p)\, e^{\varphi_p(\rho_{\mathcal{F},h}(p))}}{\sum_{q\in\mathcal{P}} Q_{\text{base}}(q)\, e^{\varphi_q(\rho_{\mathcal{F},h}(q))}}.$$

Theorem 1. *Let $\mathcal{P}$ be finite and linearly independent. If $|\mathcal{P}| \geq 3$, a choice correspondence C satisfies Axioms 1–4 if and only if it has a pragmatic representation $(u, Q_{\text{base}}, \alpha)$. If $|\mathcal{P}| = 2$, the same holds with Axioms 1–5. In either case, (u', Q', α') is another pragmatic representation if and only if $u' = \gamma u + \delta$ for some $\gamma > 0$ and $\delta \in \mathbb{R}$, $Q' = Q_{\text{base}}$ and $\alpha' = \alpha/\gamma$; and $\alpha = 0$ if and only if $C(\mathcal{F}, h)$ does not depend on h.*

If $|\mathcal{P}| \geq 3$, moreover, under Axiom 1, Axioms 3 and 4 hold if and only if C has a separable representation with every φ_p nondecreasing. Then u is unique up to positive affine transformations, Q_{base} is unique, and, given u, each φ_p is unique.

Read $Q^{\varphi}_{\mathcal{F},h}$ as a posterior over scenarios, with prior Q_{base} and likelihood $e^{\varphi_p(\rho(p))}$. Translation invariance and Arrival monotonicity make stakes a likelihood: each scenario is judged on its own stakes alone, as each hypothesis is judged on its own likelihood in Bayesian updating, and higher stakes count in its favor. The likelihood is otherwise free. Frontier WARP, across parallel frontiers, then does all the rest. It makes all scenarios speak with the same voice, and it makes the log-likelihood linear. The exponential form thus has the same source as in the theory of risk. Constant absolute risk aversion is the only attitude to risk that a sure change in wealth leaves alone; the exponential is the only likelihood of stakes that a sure change in what is at stake leaves alone. The baseline Q_{base} is not a primitive: it is the confidence the DM reveals under uniform stakes, and u and α are likewise constructed from choices. The sensitivity is measured in inverse utility units, because stakes are measured in utility.

The proof is in Appendix A. Axiom 1 gives every situation with nonzero stakes a unique vector of weights, read off choices among options below the frontier. The weights then depend on the stakes alone (Translation invariance), move in the right direction and satisfy Luce's independence when one stake rises (Arrival monotonicity), and do not move when all stakes rise together (Frontier WARP). With three scenarios or more, the second property separates the log-odds into a term for each scenario, because odds are ratios of weights, and the third makes each term additive, hence linear.

Remark 1 (Two scenarios). Let $\mathcal{P} = \{p, p'\}$. Axioms 1–4 hold if and only if C maximizes expected utility in every situation with full-support weights whose odds are

$$\frac{Q_{\mathcal{F},h}(p)}{Q_{\mathcal{F},h}(p')} = \frac{Q_{\text{base}}(p)}{Q_{\text{base}}(p')} \exp \ell\big(\rho_{\mathcal{F},h}(p) - \rho_{\mathcal{F},h}(p')\big)$$

for a nondecreasing $\ell : \mathbb{R} \to \mathbb{R}$ with $\ell(0) = 0$. With no third scenario, odds being ratios of weights has no bite, and any such ℓ will do. Common reception makes ℓ additive, hence linear.

Revealed odds also give a second route to the model, for any number of scenarios. Three properties of the odds characterize it.

Proposition 2 (Properties of the odds)**.** *Let $|\mathcal{P}| \geq 2$ and let C satisfy Axiom 1. Then C has a pragmatic representation if and only if, among the situations with nonzero stakes:*

(i) *(*Pragmatism.*) For $p \neq p'$, if $\rho_\sigma(p) \geq \rho_{\sigma'}(p)$ and $\rho_\sigma(p') \leq \rho_{\sigma'}(p')$, then $o_\sigma(p, p') \geq o_{\sigma'}(p, p')$;*

(ii) *(*Uniform stakes.*) all situations with a stakes profile of the form $c\mathbf{1}$, $c > 0$, have the same odds;*

(iii) *(*Common reception.*) Axiom 5 holds.*

Pragmatism says that the DM's odds between two scenarios depend only on what she would miss under each, and rise with what she would miss under the first. Applied in both directions, it is an independence of irrelevant hypotheses in the spirit of Luce (1959), and it is the part of the theory that Bayesian updating shares, since there too the odds of two hypotheses depend only on how the evidence bears on them. In Example 1, the arrival of e raises the stakes under the transformative scenario and nowhere else, and the firm then bets more readily on that scenario. Part (ii) says that a situation in which the DM could miss the same amount whatever the truth leaves her at her baseline. It speaks of uniform stakes and not of menus of lotteries, because a menu of lotteries has a best option under every scenario, so that choices from it reveal no odds.

Example 2 (Thresholds of concern)**.** Let $\varphi_p(s) = \alpha \max(0, s - z_p)$, with $\alpha > 0$ and thresholds $z_p \geq 0$. A scenario gains weight only once what the DM would miss under it exceeds z_p. A firm that reacts to a new technology only when ignoring it would cost more than some materiality threshold is of this kind. This DM satisfies all the axioms but Frontier WARP, across parallel frontiers, if some $z_p > 0$: a uniform improvement can lift her stakes over a threshold, and so move her weights. If the thresholds differ, even a common stake that lies between two of them tilts her toward the scenario with the lower one.

Remark 2 (Inferior options are irrelevant). Let C have a pragmatic representation. If g is an act with $h^q \gtrsim g^q$ for every q, then $C(\mathcal{F} \cup \{g\}, h) \setminus \{g\} = C(\mathcal{F}, h)$ for every menu $\mathcal{F} \ni h$: the stakes profile is unchanged, and g is worth no more than h. An option that is worse than what the DM holds under every scenario, however dreadful, gives her nothing to attend to. This separates fear of missing out from attention to the downside of a menu: a DM who also attended to what she could lose by moving would be swayed by such options.

Remark 3 (Independence of the axioms). None of Axioms 2, 3 and 4 is implied by the others. In each of the following examples, C maximizes expected utility with the indicated weights and full support. Weights $Q(p) \propto Q_{\text{base}}(p)\, e^{-\alpha \rho_{\mathcal{F},h}(p)}$ with $\alpha > 0$, which turn attention away from the scenarios where the choice matters, violate Arrival monotonicity only. Weights $Q(p) \propto Q_{\text{base}}(p)\, e^{\alpha_p \rho_{\mathcal{F},h}(p)}$ with $\alpha_p \geq 0$ not all equal, and weights $Q(p) \propto Q_{\text{base}}(p)\, e^{\varphi(\rho_{\mathcal{F},h}(p))}$ with φ nondecreasing and not affine, violate Frontier WARP only, and only for frontiers that are parallel and distinct. Weights $Q(p) \propto Q_{\text{base}}(p)\, e^{\gamma u(h^p)}\, e^{\alpha \rho_{\mathcal{F},h}(p)}$ with $\gamma \neq 0$, those of a DM who is both wishful and afraid of missing out, violate Translation invariance only. The proof is in Appendix A.

4.1 Evidence and stakes

On the reading of Section 2, evidence should move the baseline, and nothing else. Suppose that, before she faces a situation, the DM observes a datum that has probability $\ell(p) > 0$ under each scenario p; we call ℓ its *likelihood*. She revises her assessment of the scenarios by Bayes' rule,

$$Q^{\ell}_{\text{base}}(p) = \frac{Q_{\text{base}}(p)\,\ell(p)}{\sum_{q\in\mathcal{P}} Q_{\text{base}}(q)\,\ell(q)},$$

and the datum leaves her tastes and her sensitivity to the stakes unchanged: after it, her choices have the pragmatic representation $(u, Q^{\ell}_{\text{base}}, \alpha)$. Write $Q^{*\ell}_{\mathcal{F},h}$ for her weights after the datum and $o^{*\ell}_{\mathcal{F},h}(p,p')$ for her odds. Without a datum, $\ell \equiv 1$, and we drop the superscript.

Proposition 3 (Evidence and stakes)**.** *Let $(u, Q_{\text{base}}, \alpha)$ be a pragmatic representation, let ℓ be a likelihood, and let p, p' be distinct scenarios. For every situation $(\mathcal{F}, h)$:*

(i) *(*Two likelihoods.*)* $Q^{*\ell}_{\mathcal{F},h}(p) \propto Q_{\text{base}}(p)\,\ell(p)\,e^{\alpha\rho_{\mathcal{F},h}(p)}$.

(ii) *(*No interaction.*) The log-odds decompose as*

$$\log o^{*\ell}_{\mathcal{F},h}(p,p') = \log o_{\text{base}}(p,p') + \log\frac{\ell(p)}{\ell(p')} + \alpha\big[\rho_{\mathcal{F},h}(p) - \rho_{\mathcal{F},h}(p')\big]:$$

the datum shifts them by the same amount in every situation, and the situation shifts them by the same amount whatever the datum.

(iii) *(*Exchange rate.*) If two situations have the same stakes profile except that the stakes under p are higher by one unit of utility in the second, the weights in the second are those in the first after a datum with likelihood e^{α} under p and 1 under every other scenario.*

(iv) *(*What stakes can overturn.*) The situation moves the log-odds by at most α times the larger of the two stakes:*

$$\left|\log o^{*\ell}_{\mathcal{F},h}(p,p') - \log o_{\text{base}}(p,p') - \log\frac{\ell(p)}{\ell(p')}\right| \;\le\; \alpha\cdot\max\{\rho_{\mathcal{F},h}(p), \rho_{\mathcal{F},h}(p')\}.$$

Evidence and stakes act on the weights as two likelihoods on the same footing: the baseline is updated by what the DM observes, and the result is tilted by what is at stake. A datum moves the odds by the same factor whatever the situation, and a situation moves them by the same factor whatever the DM has observed. Part (iii) gives the sensitivity a unit: α is the log-likelihood ratio that one unit of stakes is worth. In Example 1, the release of the new system moves the firm exactly as a datum $e^2 \approx 7.4$ times more likely under the transformative scenario than under the incremental one would, and evidence of the same strength the other way would bring it back to its baseline. Part (iv) bounds what stakes can do: a situation whose stakes are at most K under each scenario can overturn evidence worth at most a likelihood ratio of $e^{\alpha K}$.

Corollary 1 (Evidence crowds out stakes)**.** *Let (ℓ_n) be a sequence of likelihoods, let p, p' be distinct scenarios, and let $K \ge 0$.*

(a) *If $\ell_n(p)/\ell_n(p') \to 0$, then $o^{*\ell_n}_{\mathcal{F},h}(p,p') \to 0$, uniformly over the situations whose stakes under p and p' are at most K.*

(b) *If $\ell_n(p) = \ell_n(p')$ for every n, then $o^{*\ell_n}_{\mathcal{F},h}(p,p') = o^{*}_{\mathcal{F},h}(p,p')$ for every n and every situation.*

When the data separate two scenarios, the stakes eventually lose their grip on them. With independent observations, the scenarios that Bayesian learning separates are those at different distances from the truth (Berk, 1966; Esponda and Pouzo, 2016), so that in the long run the weights, like the baseline, concentrate on the best-specified scenarios. The stakes keep their full force only among scenarios that the data cannot separate: scenarios that are equally well specified, or that agree on the past and differ on the future, as the incremental and the transformative scenarios do before a new technology has played out. Pragmatic weights are a phenomenon of the evidence leaving several scenarios standing. That is the condition under which James let the passions decide.

4.2 What the menu adds

The DM's weights depend on what she holds and on what is on the table. Models of payoff-dependent beliefs let them depend on a single act. The next result says where our model meets them and where it parts from them.

Proposition 4 (What the menu adds)**.** *Let $(u, Q_{\text{base}}, \alpha)$ be a pragmatic representation and h an act.*

(i) (Riskless alternatives.) *If a menu $\mathcal{F} \ni h$ contains a lottery x with $u(x) \geq u(f^p)$ for all $f \in \mathcal{F}$ and $p \in \mathcal{P}$, then $Q^*_{\mathcal{F},h}(p) \propto Q_{\text{base}}(p)\, e^{-\alpha u(h^p)}$.*

(ii) (Frontier dependence.) *Let $\alpha > 0$. For menus $\mathcal{F}, \mathcal{G} \ni h$, $Q^*_{\mathcal{F},h} = Q^*_{\mathcal{G},h}$ if and only if $b_{\mathcal{F}} - b_{\mathcal{G}}$ is constant across scenarios.*

Part (i) says that when the table holds a sure thing that is at least as good as every option under every scenario, the DM's weights depend only on what she holds: she tilts her baseline away from the scenarios under which her status quo does well. (She then chooses the sure thing, so the tilt shows in her choices from other menus with the same status quo and frontier, as in Proposition 1(iii).) The weights have the functional form of the belief map of Mayraz (2011a), in which the payoff of an event acts as evidence about its likelihood, with the negative sign that he calls pessimism and with the status quo in the role of his reference act. When scenarios are Dirac measures, they have the form of the "consequential distortion" of Kovach (2020), but with a decreasing distortion function, which his representation excludes. Kovach's wishful thinker overweights the states in which her endowment pays well; in these situations our DM does the opposite.

Part (ii) says what the menu adds. In models where beliefs depend on a single act, two menus that contain the same status quo induce the same weights. Here they do so exactly when their frontiers differ by a constant, that is, when neither menu makes one scenario more pivotal, relative to another, than the other menu does. An option that does not beat the frontier under any scenario leaves it unchanged, which is Proposition 7(iii), and part (ii) contains that result as the case $\mathcal{G} = \mathcal{F} \cup \{e\}$; options dominated by the status quo are the special case of Remark 2. Menu dependence through regret is not new: in Hayashi (2008), anticipated regret makes the evaluation of an act depend on the opportunities available. What is new is that the menu moves the likelihoods while the evaluation stays expected utility. The axioms differ accordingly. Pragmatism (Proposition 2) shares with the "independence of irrelevant payoffs" of Kovach (2020) the idea that the odds between two scenarios depend only on what happens under them; but what happens is measured by the stakes, which involve the whole menu, and not by the payoffs of a single act. Frontier WARP, across parallel frontiers, plays the role that a shift-invariance condition plays in Mayraz (2011a): in both, a Cauchy equation delivers the exponential form.

4.3 Robustness to regret

Section 2 noted that the weights solve a problem of attention. The problem has a second reading, familiar from robust decision making. Fix a full-support Q_{base} and a stakes profile $\rho \in \mathbb{R}^{\mathcal{P}}_{+}$, let $M = \arg\max_p \rho(p)$, and for $\alpha > 0$ let $Q^\alpha(p) \propto Q_{\text{base}}(p)\, e^{\alpha\rho(p)}$.

Proposition 5 (Robustness to regret). *(i) Q^α is the unique solution of*

$$\max_{Q \in \Delta(\mathcal{P})} \sum_p Q(p)\rho(p) - \tfrac{1}{\alpha}\,\mathrm{KL}(Q \,\|\, Q_{\text{base}}),$$

and the value of this problem is $R_\alpha = \frac{1}{\alpha}\log\sum_p Q_{\text{base}}(p)\, e^{\alpha\rho(p)}$.

(ii) R_α and $\sum_p Q^\alpha(p)\rho(p)$ are nondecreasing in α, the second is at least the first, and both tend to $\sum_p Q_{\text{base}}(p)\rho(p)$ as $\alpha \to 0$ and to $\max_p \rho(p)$ as $\alpha \to \infty$.

(iii) As $\alpha \to \infty$, Q^α tends to $Q_{\text{base}}(\cdot \mid M)$.

The DM acts as if an adversary chose her prior to maximize the expected regret of her status quo, at a cost $1/\alpha$ per unit of relative entropy. This is the worst case of the multiplier preferences of Hansen and Sargent (2008), axiomatized by Strzalecki (2011), with the regret of the status quo in place of the utility of the act being evaluated. Formally, $R_\alpha = -W_{1/\alpha}(-\rho)$, where $W_\theta(g) = \min_Q \left[\sum_p Q(p)g(p) + \theta\,\mathrm{KL}(Q \,\|\, Q_{\text{base}})\right]$ is the multiplier value of a profile g of utilities across scenarios. Two differences matter. In multiplier preferences, each act is evaluated under its own worst case, and the DM is averse to ambiguity; here a single prior, fixed by the status quo and the menu, evaluates every act, and the DM is an expected utility maximizer in each situation. And the adversary is not hostile: it points at the scenarios under which standing pat would be a mistake. The expected regret of the status quo under her own weights, $\sum_p Q^\alpha(p)\rho(p)$, and its entropic certainty equivalent R_α both run from the Bayesian expected regret, when α is small, to the maximal regret of Savage (1951), when α is large. In the limit, her prior sits on the scenarios under which the status quo is most regretted, weighted by her baseline, a selection among the priors that maximize this regret. Fear of missing out is, in this sense, a soft maximum of the regret of what one holds, loosely related to models that derive beliefs from minimax regret (Suryanarayanan, 2006).

4.4 What choices add

Choice data carry two predictions that we have not yet stated.

An unchosen option reverses a choice. Take Example 1: $Q_{\text{base}} = (0.9, 0.1)$, $\alpha = 1/4$, and the firm holds $h = (10, 4)$. Let $g = (9, 6)$ be a cautious upgrade. From $\{h, g\}$ the firm keeps h: the stakes are $(0, 2)$, the weight on the transformative scenario is 0.15, and h is worth 9.07 against 8.54. Add $e = (0, 12)$, a moonshot. The stakes become $(0, 8)$ and the weight 0.45; now g is worth 7.65, h is worth 7.29 and e is worth 5.41. The firm takes g. Nobody chooses e, and e decides between h and g. This is the prediction that availability moves beliefs, with no belief to elicit: the analyst compares choices from two menus. The reversal goes one way. If e shines under p only, the choice among the old options can move only to an option that does strictly better under p (Axiom 4). An option below the frontier, however dreadful, causes no reversal at all (Axiom 1).

Two options do not tell beliefs from tastes.

Proposition 6 (Binary menus). *Let C have a pragmatic representation $(u, Q_{\text{base}}, \alpha)$. For all acts h, f,*

$$f \in C(\{h, f\}, h) \quad \Longleftrightarrow \quad \sum_{p \in \mathcal{P}} Q_{\text{base}}(p)\, \phi_\alpha\big(u(f^p) - u(h^p)\big) \geq 0, \qquad \phi_\alpha(x) = x\, e^{\alpha \max(x,0)}.$$

Proposition 8(i) below is the case of two scenarios. When she compares what she holds with one alternative, the DM thus behaves as someone with fixed beliefs Q_{base} and reference-dependent tastes, applied scenario by scenario to the expected utility of the alternative minus that of what she holds: losses count as they are, and gains are inflated, the more so the larger they are. It is the mirror image of loss aversion. Why then speak of beliefs at all? On menus of two options, the question has no answer in choice data. On larger menus it has one. For a given status quo, a DM with reference-dependent tastes satisfies the weak axiom of revealed preference, and our DM can violate it across menus whose frontiers are not parallel. The reversal above is the test.

5 The reception of a new technology

Throughout this section, C has a pragmatic representation $(u, Q_{\text{base}}, \alpha)$. A situation $(\mathcal{F}, h)$ now describes the technologies on the table: $\mathcal{F}$ is the *installed base* of available technologies and h the one the DM uses. An *entrant* is an act $e \notin \mathcal{F}$. Write $Q^0 = Q^*_{\mathcal{F},h}$ and $Q^1 = Q^*_{\mathcal{F}\cup\{e\},h}$ for her weights before and after the arrival, V^0 and V^1 for the values under these weights, and

$$\Delta(p) = \big(u(e^p) - b_{\mathcal{F}}(p)\big)^+$$

for how far the entrant pushes the frontier up under p. The entrant is *adopted* in a situation if it alone is chosen: $C(\mathcal{F} \cup \{e\}, h) = \{e\}$.

5.1 The reception formula

Proposition 7 (Reception). *(i) The arrival of e updates the weights by the likelihood $e^{\alpha\Delta}$:*

$$Q^1(p) \;\propto\; Q^0(p)\, e^{\alpha\Delta(p)}.$$

(ii) The likelihood depends on the installed base $\mathcal{F}$ and on e, not on the status quo h: all users who face the same installed base revise their odds by the same factor.

(iii) The arrival leaves the weights unchanged if and only if $\alpha\Delta$ is constant across scenarios. In particular, an entrant that beats the frontier under no scenario is inert, however bad it is under some scenario.

(iv) (Availability never hurts the entrant.) *For every $f \in \mathcal{F}$, $V^0(e) \geq V^0(f)$ implies $V^1(e) \geq V^1(f)$, and $V^0(e) > V^0(f)$ implies $V^1(e) > V^1(f)$.*

The reception of a technology is an update, but not on information. The installed base plays the role of a prior and the entrant the role of a datum, and the datum is how far the entrant pushes the frontier of the available, scenario by scenario; part (ii) is Common reception. Figure 1 illustrates. What matters is the frontier, not the status quo and not the incumbents one by one: an entrant that does not beat the best incumbent under any scenario produces no shift at all, however extreme

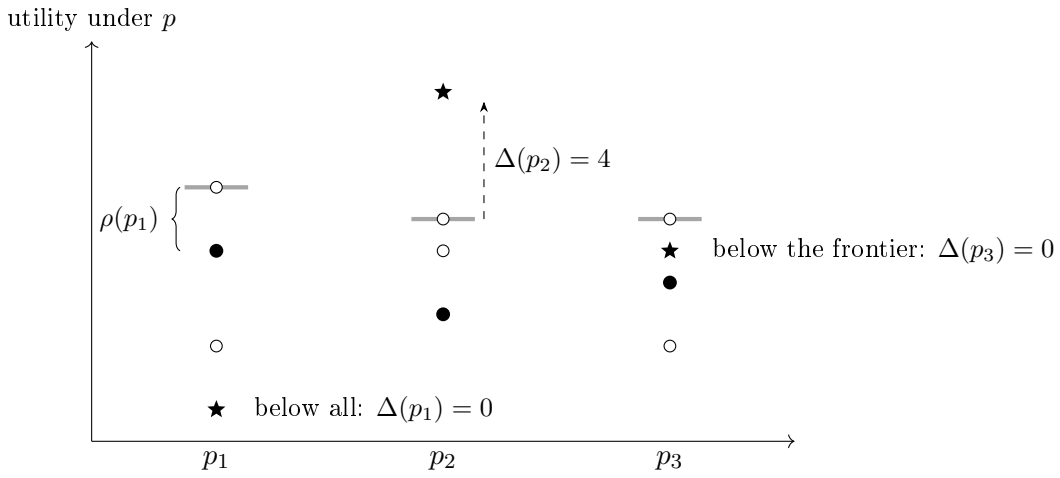


Figure 1: The frontier of an installed base under three scenarios. Gray bars mark the frontier, the best incumbent under each scenario; black dots mark the status quo, open circles the other incumbents, stars the entrant. The brace shows the stakes of the status quo under p_1. The entrant pushes the frontier up under p_2 and nowhere else; that it is worse than every incumbent under p_1 does not matter. Its arrival multiplies the odds of p_2 against each other scenario by $e^{4\alpha}$, for every user who faces this installed base.

it is relative to the DM's own technology. This is how a mature market immunizes its users against hype. In Example 1, if the installed base already contains an incumbent $g = (6, 12)$, the new system $e = (7, 12)$ beats the frontier nowhere and is inert; if it contains $g = (6, 10)$ instead, e extends the frontier by 2 under the transformative scenario, not by 8. An incumbent that occupies the scenario an entrant promises captures the belief shift for itself. Part (iv) gives the direction: the arrival tilts beliefs toward the scenarios under which the entrant beats every incumbent, so it can turn a technology that is worth less than an incumbent at the weights before the arrival into one she adopts, an *adoption out of FOMO*, but never the reverse. Before the arrival, of course, the entrant is not on the table, and $V^0(e)$ is not itself a choice. When something is at stake in $(\mathcal{F}, h)$, the analyst sees the comparison through a scaled copy: two acts below the frontier of $\mathcal{F}$ whose difference in profile is a positive multiple of $u(e) - u(f)$ (Proposition 1).

5.2 Hype

Consider the simplest contest: two scenarios with baseline weights q_1, q_2, a status quo h, and an entrant e that is better than h by $D > 0$ under p_2 and worse by $L > 0$ under p_1, in utility. The menu after the arrival is $\{h, e\}$.

Proposition 8 (Hype). *(i) (*The FOMO premium.*) The entrant is adopted if and only if*

$$L < \frac{q_2}{q_1} D\, e^{\alpha D}.$$

A Bayesian adopts if and only if $L < (q_2/q_1)D$: relative to a Bayesian, the threshold for the downside is multiplied by $e^{\alpha D}$.

*(ii) (*Radical technologies.*) For $s > 0$, let e_s be the entrant with upside sD and downside sL. If $\alpha > 0$, adoption of e_s is monotone in s, and there is $\bar{s}$ such that e_s is adopted for all $s > \bar{s}$. A Bayesian's decision does not depend on s.*

An entrant that promises a large gain under one scenario can carry a downside that a Bayesian would never accept, and the protection the DM enjoys against such a mirage shrinks exponentially as the promise grows. Part (ii) says that she is not neutral to the scale of a bet. A Bayesian evaluates a radical technology, with a large upside and a large downside, exactly as she evaluates a timid one with the same ratio of upside to downside. A DM who fears missing out favors the radical one, and adopts any technology of this shape once it is radical enough. A taste for moonshots is a prediction of the model.

5.3 Users who face the same technologies

Proposition 9 (Disagreement). *Let $(\mathcal{F}, h)$ and $(\mathcal{F}, h')$ be two situations with the same installed base. For all scenarios p, p',*

$$\log \frac{Q^*_{\mathcal{F},h'}(p)}{Q^*_{\mathcal{F},h'}(p')} - \log \frac{Q^*_{\mathcal{F},h}(p)}{Q^*_{\mathcal{F},h}(p')} = -\alpha\Big[\big(u(h'^p) - u(h^p)\big) - \big(u(h'^{p'}) - u(h^{p'})\big)\Big].$$

Two users who face the same technologies, with the same baseline and the same sensitivity, disagree only because they hold different ones. Each overweights the scenarios under which her own technology does relatively badly: the grass is greener elsewhere. With Proposition 7(ii), the result says that at the arrival of a new technology, users start from beliefs that differ according to what they hold and revise them by a common factor. What matters is the profile of what a user holds across scenarios, not how far it sits behind the frontier: the user of a technology that is uniformly far behind is not moved toward any scenario in particular.

5.4 Churn

Suppose now that the menu is fixed and that the status quo is the technology chosen in the previous period: the DM holds h_t, chooses from $C(\mathcal{F}, h_t)$, staying with h_t when it is in that set, and holds what she chose at $t+1$. A technology $h \in \mathcal{F}$ is a *rest point* if the DM who holds it keeps it: $h \in C(\mathcal{F}, h)$.

Proposition 10 (Churn). *Let $\mathcal{P} = \{p_1, p_2\}$ with baseline weights q_1, q_2, and let $\mathcal{F} = \{h, e\}$ with $u(h^{p_1}) - u(e^{p_1}) = L > 0$ and $u(e^{p_2}) - u(h^{p_2}) = D > 0$. Let $\kappa = \log\big(q_1 L/(q_2 D)\big)$. Then h is a rest point if and only if $\kappa \geq \alpha D$, and e is a rest point if and only if $\kappa \leq -\alpha L$. Consequently:*

(i) If $\alpha > 0$, there is at most one rest point, and if $-\alpha L < \kappa < \alpha D$ the DM switches technology in every period.

(ii) If $\alpha = 0$ and $\kappa \neq 0$, there is exactly one rest point, and the DM reaches it within one period from either starting point.

The mechanism is that of Proposition 9, run through time: holding a technology draws the DM's attention to the scenarios under which it is beaten. In Example 1, $\kappa = \log 3.375 \approx 1.22$, which lies between $-\alpha L = -0.75$ and $\alpha D = 2$. The firm adopts the new system, then attends to the scenario under which it should have kept the old one, reverts, attends again to the transformative scenario, and so on: adoption out of fear of missing out, followed by remorse. A Bayesian firm keeps the old system.

The bound on rest points does not depend on the example, nor on the size of the menu.

Proposition 11 (Rest points). *Let $\alpha > 0$. If h and h' are rest points of the same menu, then $u(h^p) = u(h'^p)$ for every p.*

A DM with status quo bias in the sense of Masatlioglu and Ok (2005) is the opposite. Their axiom makes an option chosen from a menu, whatever the status quo, a rest point of that menu, and a menu can have many. Here a menu has one rest point at most, up to acts with the same profile, and often none: in our simulations, about three random menus in ten had none (Appendix D).

The cycle is deterministic. Would a DM who sees it coming break it? Suppose that, at each date, the self who holds $x \in \mathcal{F}$ chooses $y \in \mathcal{F}$, receives $u(y^p)$ in that period if p is true, and holds y at the next date; changing technology costs $c \geq 0$ in the period of the change. She evaluates a stream of choices with her own weights Q_x and a discount factor $\delta \in [0,1)$: a stream $(y_0, y_1, \dots)$ is worth $\sum_k \delta^k \big[\sum_p Q_x(p)\, u(y_k^p) - c\,\mathbf{1}\{y_k \neq y_{k-1}\}\big]$, with $y_{-1} = x$. Her attention is set as before, by the stakes of the menu she faces in the current period; only her valuation looks ahead.[4] Thus Q_x is $Q^*_{\mathcal{F},x}$ when $c = 0$; with $c > 0$, it is computed from the stakes net of the cost, $\big(u(x'^p) - c - u(x^p)\big)^+$, where x' is the other technology, as in Section 7.5. Each self takes the behavior of her successors as given, and they follow a *Markov strategy* $\sigma : \mathcal{F} \to \mathcal{F}$, which says what a self who holds x chooses. The strategy σ is an *equilibrium* if, for each x, $\sigma(x)$ maximizes the value of the stream that starts with $\sigma(x)$ and continues according to σ, with $\sigma(x) = x$ when x is optimal. Write

$$G_x = \sum_p Q_x(p)\big[u(x'^p) - u(x^p)\big]$$

for the gain that the self holding x expects from one period with the other technology.

Proposition 12 (Foresight). *In the setting of Proposition 10, let $\delta \in [0,1)$.*

(i) If $c = 0$, then σ is an equilibrium if and only if each self makes the myopic choice: $\sigma(x) = x'$ if and only if $G_x > 0$. The Markov equilibrium is unique, and if $\alpha > 0$ and $-\alpha L < \kappa < \alpha D$, the DM switches technology in every period, whatever δ.

(ii) If $c > 0$: keeping both technologies is an equilibrium if and only if $G_h \leq (1-\delta)c$ and $G_e \leq (1-\delta)c$; always choosing x is an equilibrium if and only if $G_x \leq (1+\delta)c$ and $G_{x'} > (1-\delta)c$; and the cycle $\sigma(h) = e$, $\sigma(e) = h$ is an equilibrium if and only if $G_h > (1+\delta)c$ and $G_e > (1+\delta)c$.

Foresight in valuation cannot stop the cycle when switching is free. A self's choice decides only which phase of the cycle she lands in: her successor, who will hold what she chose, will fear what she now ignores. Since she cannot bind him, her best reply is to do what she would do if tomorrow did not exist. A switching cost changes this. A self who switches pays the cost now, but shares the gain with a successor who will switch back, so the cycle needs a gain above $(1+\delta)c$ rather than c: patience buys loyalty only when changing is costly. In Example 1 with $c = 1$, $G_h \approx 1.29$ and $G_e \approx 2.31$: the myopic firm cycles, while for $\delta > 0.29$ the unique Markov equilibrium has the firm keep its old system. With a larger cost, both "always h" and "always e" can be equilibria, and the DM's technology then depends on where she starts. The churn of Proposition 10 is therefore a prediction where switching is cheap, not an artifact of myopic valuation; whether it survives forward-looking attention is a different question (footnote 4).

6 FOMO in a market

Part of what a firm fears to miss is what its competitors will do. When more of them adopt the new technology, staying with the old one is costlier if the technology turns out to be transformative.

[4] If the stakes were instead computed from the continuation values of the two plans, keeping x or switching, each followed by the strategy below, the cycle would scale them by $1/(1+\delta)$. It would then be an equilibrium if and only if $-\alpha L/(1+\delta) < \kappa < \alpha D/(1+\delta)$. In Example 1 this fails for $\delta > 2/\kappa - 1 \approx 0.64$, and then no Markov strategy is an equilibrium, so that a patient DM of this kind has no stable pure plan at all.

The stakes under that scenario rise, the firm gives it more weight, and it is more willing to adopt. Others' adoption thus raises the weight each firm gives the transformative scenario, although it carries no information about it. This section studies the fixed points of this loop at the launch, when every firm still holds the old technology; Section 7 follows the market afterwards.

6.1 A market of pragmatic adopters

There are two scenarios, *incremental* (I) and *transformative* (T), with true probabilities $\pi_I, \pi_T > 0$; write $r = \pi_T/\pi_I$. A continuum of firms of mass one use an old technology, and a new technology e becomes available. Let $m \in [0,1]$ be the share of firms that adopt it. Firms differ in their *downside* $L \geq 0$, what adopting costs them if the technology is incremental. Downsides are distributed according to a continuous cdf G on $[0,\infty)$ with $G(0) = 0$. Each firm is a DM in the sense of the previous sections, with its own utility index.

Assumption 1 (Payoffs)**.** When a share m of firms adopt, the old technology is an act h_m, and for a firm with downside L,

$$u(h_m^I) = u_0, \quad u(e^I) = u_0 - L, \qquad u(h_m^T) = u_0 - \lambda m, \quad u(e^T) = u_0 + \theta,$$

with $\theta, \lambda \geq 0$.

Under the incremental scenario the new technology only costs. Under the transformative one it brings an intrinsic gain θ, and the old technology loses value in proportion to the share of competitors who have transformed their operations; λ measures this competitive pressure. Acts with these reductions exist (Lemma A.1). The gain from adopting is $-L$ under I and the *upside* $D(m) = \theta + \lambda m$ under T. Others' adoption affects a firm only through the cost of staying behind: in the baseline there are no network benefits, which Section 6.6 adds. Every firm's utility index is normalized as in Assumption 1, with common u_0, θ and λ, so that utility differences, the sensitivity α and the welfare comparisons below are comparable across firms. The payoffs of Example 1 are those of a firm with downside 3 when all its competitors have adopted, with $u_0 = 10$, $\theta = 2$ and $\lambda = 6$.

Every firm has a pragmatic representation with the true probabilities as baseline, $Q_{\text{base}} = \pi$, and a common sensitivity $\alpha \geq 0$. The baseline is correct, and the only departure from a Bayesian market, $\alpha = 0$, is the pragmatic tilt. No data can improve on this baseline before the technology has played out, since the two scenarios agree on what has been observed; by Corollary 1(b), the stakes then keep their full force. Each firm is negligible and takes m as given. Facing the menu $\{h_m, e\}$ from the status quo h_m, a firm has stakes 0 under I and $D(m)$ under T, whatever its downside, so its odds of T against I are $r\, e^{\alpha D(m)}$, the same for every firm.

Lemma 6.1 (Adoption rule)**.** *A firm with downside L adopts if and only if $L < L^*(m)$, where*

$$L^*(m) = r\, D(m)\, e^{\alpha D(m)}.$$

This is Proposition 8(i) with an upside that depends on m. A Bayesian firm adopts if and only if $L < rD(m)$. Given that a share m adopts, the share of firms that want to adopt is $\Phi(m) = G(L^*(m))$.

Definition 3. A *pragmatic equilibrium* is a share $m^* \in [0,1]$ such that $m^* = \Phi(m^*)$.

As in a large-market rational expectations equilibrium, every firm knows the share of adopters. What it does not know is the scenario, and the share carries no information about it. Since every

firm still judges from the old technology, a pragmatic equilibrium describes the market at the launch, and we also call it a *launch equilibrium*. We also consider *myopic adjustment*, $m_{t+1} = \Phi(m_t)$: a tâtonnement in which firms, all still on the old technology, respond to the last share they observed. Propositions 13, 14 and 17 use only the adoption rule and hold for any continuous nondecreasing upside $D(m) \geq 0$, differentiable for Proposition 14, except for their statements about θ and λ.

6.2 Existence and the FOMO multiplier

Proposition 13 (Existence and comparative statics). *(i) The map Φ is continuous and nondecreasing, and there are a least equilibrium $\underline{m}$ and a greatest equilibrium $\overline{m}$.*

(ii) Both are nondecreasing in α, θ, λ and r. In particular, the extremal equilibria of a FOMO market lie above those of a Bayesian market with the same primitives.

(iii) Myopic adjustment from $m_0 = 0$ increases to $\underline{m}$; from $m_0 = 1$ it decreases to $\overline{m}$.

That fear of missing out raises adoption is no surprise. The next result says by how much it amplifies the influence of others.

Proposition 14 (The FOMO multiplier). *At every m with $D(m) > 0$,*

$$\frac{d \log L^*}{d \log D} = 1 + \alpha D, \qquad \text{so that} \qquad \frac{dL^*}{dm} = r\, D'(m)\, e^{\alpha D}(1 + \alpha D).$$

For a Bayesian the elasticity is 1 and $dL^*/dm = r\, D'(m)$. Relative to a Bayesian facing the same upside, the response of the marginal adopter to others' adoption is therefore multiplied by

$$M = e^{\alpha D}(1 + \alpha D).$$

When others' adoption raises the upside, it moves the marginal adopter by more than it would move a Bayesian, for two reasons. Fear of missing out raises the threshold itself, by the factor $e^{\alpha D}$, so that a given proportional rise in the upside moves more firms. And it raises the elasticity of the threshold, by the factor $1 + \alpha D$. The multiplier grows exponentially with the upside. The amplification is a matter of beliefs only: the payoffs are those of the Bayesian market.

6.3 Tipping and sudden waves

Write $\xi(m) = \Phi(m) - m$ for the excess response at the share m.

Proposition 15 (Tipping). *(i) If $\alpha = 0$, $\theta > 0$ and G is concave on $[0, \infty)$, for instance uniform, the equilibrium is unique.*[5]

(ii) Let $\alpha > 0$, $\theta > 0$, $\lambda > 0$, and let G be uniform on $[0, \bar{L}]$. If $L^*(1) > \bar{L}$ and $\xi(\hat{m}) < 0$ for some $\hat{m} \in (0, 1)$, there are exactly three equilibria, $0 < m_L < m_T < m_H = 1$. Myopic adjustment from any $m_0 < m_T$ converges to m_L, and from any $m_0 > m_T$ to full adoption. If $L^*(1) < \bar{L}$, or if $\xi > 0$ on $[0, 1)$, the equilibrium is unique.

The Bayesian aggregate response $G(rD(m))$ inherits the curvature of G; with a concave G it crosses the diagonal once. Fear of missing out composes G with the convex map $D \mapsto rDe^{\alpha D}$, and with a uniform G the response becomes convex: it can dip below the diagonal and come back

[5] With $\theta = 0$, $m = 0$ is always an equilibrium, and a Bayesian market with concave G can have others; with uniform G and $r\lambda = \bar{L}$, every $m \in [0, 1]$ is one. Corollary 2 treats $\theta = 0$ directly.

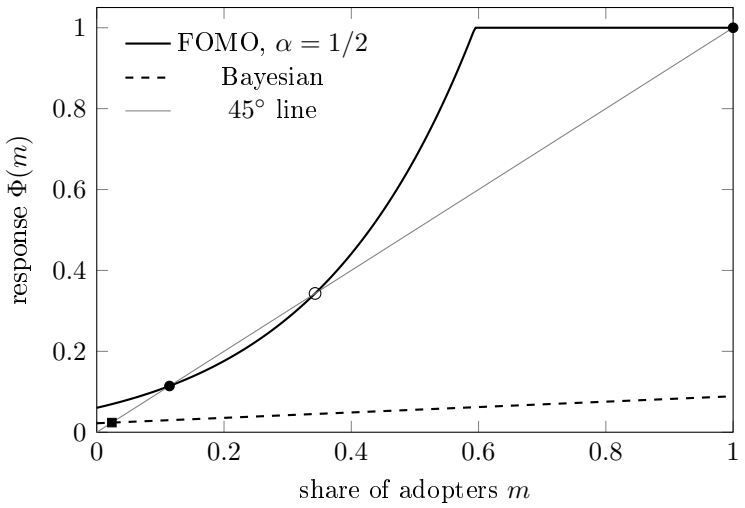


Figure 2: The aggregate response in Example 3. The Bayesian response is linear and crosses the diagonal once, at $m^B \approx 0.024$ (square). The FOMO response is convex until it reaches one; it crosses the diagonal at a stable low equilibrium $m_L \approx 0.114$ and at a tipping point $m_T \approx 0.343$ (open circle), and full adoption is a stable equilibrium.

above it. The market then has a low equilibrium, where few firms adopt and few are worried, and a high one, where everyone adopts because everyone is worried. Between the two lies a tipping point: a market pushed beyond it, by a shock or a campaign, goes on to full adoption. Figure 2 shows the two responses in the following example, which we use for the rest of the paper. Which distributions of downsides allow a market to tip, with and without fear of missing out, is settled in Appendix C: when downsides have a density, fear of missing out creates tipping exactly when it pushes the elasticity of adoption with respect to the stakes above one, and it multiplies that elasticity by $1 + \alpha D$ (Theorem 3).

Example 3 (A market)**.** Let $\pi_T = 0.1$ as in Example 1, so that $r = 1/9$, and let $\theta = 2$, $\lambda = 6$, downsides be uniform on $[0, 10]$ and $\alpha = 1/2$. The FOMO market has three equilibria, $m_L \approx 0.114$, $m_T \approx 0.343$ and $m_H = 1$. The Bayesian market has one, $m^B = 1/42 \approx 0.024$. At m_L the upside is $D \approx 2.69$, and the FOMO multiplier is $M = e^{\alpha D}(1 + \alpha D) \approx 3.83 \times 2.34 \approx 8.97$.

Suppose now that the technology improves, so that its intrinsic gain θ rises. A Bayesian market follows smoothly. A FOMO market does not.

Proposition 16 (Sudden waves and hysteresis)**.** *Let $\alpha > 0$, $\lambda > 0$, and let G be uniform on $[0, \bar{L}]$ with $\bar{L} > r\lambda$. Let $\hat{D} > 0$ solve $(1 + \alpha\hat{D})\, e^{\alpha\hat{D}} = \bar{L}/(r\lambda)$, and set*

$$\mu^* = \frac{\hat{D}}{\lambda(1 + \alpha\hat{D})}, \qquad \theta^* = \frac{\alpha\hat{D}^2}{1 + \alpha\hat{D}}.$$

Assume $\mu^* < 1$ and $r(\theta^* + \lambda)e^{\alpha(\theta^* + \lambda)} > \bar{L}$, and write $\underline{m}(\theta)$ for the least equilibrium as a function of the intrinsic gain.

(i) For $\theta \in (\max\{0, \hat{D} - \lambda\}, \theta^*)$, $\underline{m}(\theta) < (\hat{D} - \theta)/\lambda$. Hence $\limsup_{\theta \uparrow \theta^*} \underline{m}(\theta) \leq \mu^*$.

(ii) For $\theta > \theta^*$, full adoption is the unique equilibrium.

(iii) For every $\theta \geq 0$ with $r(\theta+\lambda)e^{\alpha(\theta+\lambda)} \geq \bar{L}$, full adoption is an equilibrium.

(iv) (Bayesian benchmark.) *If instead $\alpha = 0$ and G is concave, the unique equilibrium is continuous in $\theta > 0$; if G is uniform on $[0, \bar{L}]$ with $\bar{L} > r\lambda$, it equals* $\min\{1, r\theta/(\bar{L} - r\lambda)\}$.

The least equilibrium jumps, at θ^*, from at most μ^* to one: a technology that improves slowly is adopted slowly, then all at once. Full adoption remains an equilibrium after quality falls back below θ^*, as long as $r(\theta+\lambda)e^{\alpha(\theta+\lambda)} \geq \bar{L}$; when $r\lambda e^{\alpha\lambda} \geq \bar{L}$, this holds at every quality, and the wave is not reversed by conditions that would not have started it. In Example 3, $\hat{D} \approx 3.42$, $\mu^* \approx 0.21$ and $\theta^* \approx 2.16$, and $r\lambda e^{\alpha\lambda} \approx 13.4 \geq 10$. Let θ take the successive values 1, 1.5, 2, 2.1, 2.2, 2.3, and then 2, 1.5, 1, 0.5, and let the market settle by myopic adjustment at each value. Adoption goes from 0.02 to 0.05, 0.11 and 0.15, jumps to 1 at $\theta = 2.2$, and stays at 1 all the way down. The Bayesian market moves between 0.006 and 0.027 and retraces its path. This hysteresis is a property of the market at the launch, in which every firm judges the new technology from the old one; it does not survive once adopters judge from what they have adopted (Corollary 3).

6.4 Hype without information

Proposition 17 (Hype without information)**.** *In an equilibrium m^*:*

(i) All firms act on the same probability of the transformative scenario,

$$Q^*(T) = \frac{r\, e^{\alpha D(m^*)}}{1 + r\, e^{\alpha D(m^*)}},$$

which exceeds π_T whenever $\alpha D(m^*) > 0$ and is strictly increasing in m^* when $\alpha\lambda > 0$. Across equilibria, more adoption goes with more weight on the transformative scenario, although the true probability is the same.

(ii) The firms with $rD(m^*) < L < L^*(m^*)$ adopt although, at the true probabilities, adopting lowers their expected payoff. Their mass is $G(L^*(m^*)) - G(rD(m^*))$.

In Example 3, firms act as if the transformative scenario had probability 0.30 at m_L, 0.46 at m_T and 0.86 at full adoption; the true probability is 0.10. At m_L, 8.4% of firms adopt although adopting lowers their expected payoff at the true probabilities. This is herding, but not of the informational kind: others' adoption reveals nothing, and still it moves beliefs.

6.5 Adoption traps

Proposition 18 (Adoption traps)**.** *Let $m_1 < m_2$ be two equilibria. At the true probabilities, every firm's expected payoff is at least as high at m_1 as at m_2. It is strictly higher for every firm that switches (adopts at m_2 but not at m_1) and has $L > rD(m_1)$, which is every switcher but a set of mass zero, and, if $\lambda > 0$, for every firm that adopts in neither.*

When adoption is driven by competitive pressure, the high-adoption equilibrium is not a coordination success but a rat race. Firms that adopt in both equilibria are indifferent between them, and firms that adopt in neither suffer from the extra pressure. Firms that switch are worse off than at m_1: the pressure at m_2 pushes them into adopting, whether or not adopting at m_2 is itself a mistake at the true probabilities. In Example 3, the average expected payoff relative to u_0 is -0.10 at m_L, -0.60 at m_T and -4.30 at full adoption, against -0.01 in the Bayesian equilibrium.

Market	α	Launch equilibria	Long-run share (Section 7)
Bayesian	0	0.024	0.024
Mild FOMO	0.25	0.044	0.033
FOMO	0.5	0.114, 0.343, 1	0.052
Strong FOMO	0.75	1	0.521
Very strong FOMO	1	1	0.519

Table 1: Example 3 ($r = 1/9$, $\theta = 2$, $\lambda = 6$, downsides uniform on $[0, 10]$) for five sensitivities, computed numerically (Appendix D). The last column is the unique stationary share of the market with memory, with equal revision probabilities.

Corollary 2 (Useless technologies)**.** *Let $\theta = 0$ and let G be uniform on $[0, \bar{L}]$. Then $m = 0$ is an equilibrium, and $m = 1$ is an equilibrium if and only if $\bar{L} \leq r\lambda\, e^{\alpha\lambda}$; for a Bayesian, if and only if $\bar{L} \leq r\lambda$. At $m = 1$, every firm with $L > 0$ is worse off, at the true probabilities, than at $m = 0$.*

A technology that brings nothing by itself, and whose only effect is to make laggards look bad if it catches on, can thus be adopted by every firm in a FOMO market. In Example 3 with $\theta = 0$, the Bayesian market has the unique equilibrium $m = 0$, while the FOMO market has three: $m = 0$, a tipping point at $m \approx 0.90$, and $m = 1$. The condition for full adoption is weaker than for a Bayesian by the factor $e^{\alpha\lambda}$: $r\lambda e^{\alpha\lambda} \approx 13.4 \geq 10$, while $r\lambda \approx 0.67 < 10$.

Table 1 shows how the launch depends on the strength of the fear. Mild fear raises adoption a little; at $\alpha = 1/2$ the market has three equilibria; stronger fear makes full adoption the unique launch equilibrium. The last column anticipates Section 7: once adopters judge from what they hold, the long run is far from the launch.

6.6 Network benefits

The welfare results so far rest on competitive pressure: the new technology is valuable because laggards fall behind. Suppose instead that adopters also gain from others' adoption, $u(e^T) = u_0 + \theta + \beta m$ with $\beta \geq 0$, the other payoffs being those of Assumption 1. The upside becomes $D(m) = \theta + (\lambda + \beta)m$, so that Lemma 6.1 and Propositions 13–17 hold with $\lambda + \beta$ in place of λ. Network benefits and competitive pressure have the same effect on adoption. They differ in welfare.

Proposition 19 (Network benefits)**.** *Let $m_1 < m_2$ be two equilibria. At the true probabilities:*

(i) a firm that adopts in both is better off at m_2 by $\pi_T\beta(m_2 - m_1)$, and a firm that adopts in neither is worse off at m_2 by $\pi_T\lambda(m_2 - m_1)$;

(ii) a firm that adopts at m_2 only is better off at m_2 if and only if $L < r(\theta + \beta m_2 + \lambda m_1)$;

(iii) if $\lambda = 0$, every firm weakly prefers m_2, except the firms with $L^*(m_1) \leq L < L^*(m_2)$ and $L > rD(m_2)$, which adopt at m_2 only. If $\alpha = 0$ there are none.

Moreover, let $\lambda = 0$, let m_B be an equilibrium of the Bayesian market and $m_F \geq m_B$ an equilibrium of the FOMO market with the same primitives. If $rD(m_F) \geq \sup \operatorname{supp} G$, every firm weakly prefers m_F to m_B.

With pure network benefits, firms that adopt in both equilibria, or in neither, prefer more adoption, and when $\alpha = 0$ so does every firm: the ranking of Proposition 18 is reversed. With fear of missing out, the exceptions are the firms that fear pushes into adopting at m_2 although,

at the true probabilities, adopting there lowers their payoff, and they can be most of the market. Take Example 3 with $\lambda = 0$ and $\beta = 6$: the upside is unchanged, and so are the equilibria. Comparing $m_L \approx 0.114$ with full adoption, 89% of the firms adopt only at full adoption and are worse off there; the average payoff relative to u_0 is -0.03 at m_L and -3.70 at full adoption, against $+0.003$ at the unique Bayesian equilibrium. Fear of missing out helps in the case of the last part of the proposition, which forces $m_F = 1$: if the Bayesian market can rest at a low equilibrium although full adoption, at which every firm gains from adopting, is also an equilibrium, raising the least equilibrium (Proposition 13) is a remedy for the excess inertia of Farrell and Saloner (1985). Whether the value of a technology lies in joining others or in not falling behind them thus matters, but it does not by itself make fear of missing out benign; adoption traps belong to the second kind of technology.

7 Markets with memory

In Section 6, every firm judges the new technology from the old one. This is right at the launch, when nobody has adopted yet. It is not right afterwards. A firm that has adopted holds the new technology, and Proposition 10 says that what a DM holds changes what she attends to. This section follows the market after the launch, when each firm judges from what it holds. The key observation is that *holding the new technology turns the fear around*. A firm on the old technology fears missing the transformative scenario, under which it would be left behind. A firm on the new technology fears missing the incremental scenario, under which it should have stayed. After the launch, the market is therefore governed by two responses: the launch response Φ, which says who adopts, and a *keeping response* $\tilde{\Phi}$, which says who stays.

7.1 Holding turns the fear around

We keep the primitives of Section 6, with any upside $D(m)$ that is continuous and nondecreasing in the share m of adopters; in Assumption 1, $D(m) = \theta + \lambda m$. Throughout this section, $D(m) > 0$ for all m, as when $\theta > 0$.

Firms now carry their technology. A *holding profile* is a measurable map $x : [0, \infty) \to [0, 1]$, where $x(L)$ is the share of firms with downside L that hold the new technology e; the share of adopters is $m(x) = \int x \, dG$. Time is discrete. In each period, each firm independently has the opportunity to revise with probability $\nu \in (0, 1]$. A revising firm chooses from the menu, given the current share m, with the technology it holds as its status quo; as in Proposition 10, it keeps what it holds when indifferent. The other firms keep their technology. With $\nu = 1$ all firms revise at the same dates; with $\nu < 1$ revisions are staggered.

Lemma 7.1 (Holding turns the fear around)**.** *When the share of adopters is m:*

(i) a firm with downside L that holds the old technology adopts if and only if $L < L^*(m) = rD(m)\, e^{\alpha D(m)}$;

(ii) a firm with downside L that holds the new technology keeps it if and only if $L \leq \tilde{L}(m)$, where $\tilde{L}(m)$ solves $\tilde{L}\, e^{\alpha \tilde{L}} = rD(m)$.

If $\alpha > 0$, then $\tilde{L}(m) < rD(m) < L^*(m)$.

Part (ii) is the computation of Lemma 6.1 from the other status quo. For an adopter, the downside L is what it would miss under the incremental scenario by keeping the new technology, and it is now this scenario that fear of missing out inflates. A firm on the old technology adopts

more readily than a Bayesian, and a firm on the new technology abandons it more readily. The firms with $\tilde{L}(m) < L < L^*(m)$ switch whatever they hold; we call them *restless*, and the interval the *restless band*. The firms with $L \leq \tilde{L}(m)$, which keep the new technology once they hold it, are *loyal adopters*.

By the law of large numbers, the profile evolves deterministically, $x_{t+1} = \mathcal{T}_\nu x_t$, where

$$(\mathcal{T}_\nu x)(L) = x(L) + \nu\Big[\mathbf{1}\{L < L^*(m)\}\big(1 - x(L)\big) - \mathbf{1}\{L > \tilde{L}(m)\}\, x(L)\Big], \qquad m = m(x).$$

A profile is *stationary* if $\mathcal{T}_\nu x = x$ almost everywhere (with respect to G), and its share $m(x)$ is then a *stationary share*. Stationary profiles do not depend on ν, and at a stationary profile the share does not move, so the myopia of revising firms plays no role: a stationary profile is an equilibrium in which every revising firm responds to the share that actually prevails. The *launch profile* of a pragmatic equilibrium m^S of Section 6 is $x^S = \mathbf{1}\{L < L^*(m^S)\}$. Three modelling choices deserve mention. Switching is free and payoffs accrue per period: the downside L is a flow cost of using the new technology, not a sunk cost of acquiring it; Section 7.5 adds a cost of switching. Revising firms are myopic, as the DM of Proposition 10; with free switching this is without loss at stationary profiles (Remark 4). And the deterministic law of motion rests on an exact law of large numbers for independent revisions by a continuum of firms.

7.2 The long run

Write $\Phi(m) = G(L^*(m))$ for the launch response of Section 6, $\tilde{\Phi}(m) = G(\tilde{L}(m))$ for the keeping response, and $\Psi = \frac{1}{2}(\Phi + \tilde{\Phi})$.

Theorem 2 (The long run). *(i) If $\alpha = 0$, the stationary profiles are the launch profiles of the Bayesian equilibria. At each of them, no firm switches.*

(ii) If $\alpha > 0$, a profile is stationary if and only if, almost everywhere, $x = 1$ on $[0, \tilde{L}(m)]$, $x = \frac{1}{2}$ on the restless band $(\tilde{L}(m), L^*(m))$ and $x = 0$ on $[L^*(m), \infty)$, with $m = \Psi(m)$. Stationary profiles exist. At each of them, a mass $\nu[\Phi(m) - \tilde{\Phi}(m)]$ of firms switches in every period, half in each direction. The least (greatest) stationary share lies between the least (greatest) fixed points of $\tilde{\Phi}$ and of Φ. Every stationary share is at most $\frac{1}{2}[1 + \tilde{\Phi}(1)]$, and full adoption is stationary if and only if $\tilde{\Phi}(1) = 1$.

Without fear of missing out, the market at the launch is the market forever. With it, the market cannot rest. The restless firms switch at every revision, and the long-run share averages the two responses. Full adoption then requires that every adopter stay loyal despite its fear of having missed the incremental scenario, $\tilde{\Phi}(1) = 1$. This is stronger than the launch condition $\Phi(1) = 1$, and much stronger in the examples below. The factor $\frac{1}{2}$ comes from equal revision probabilities. If holders of the old and of the new technology revise with probabilities ν_h and ν_e, restless firms hold the new technology in proportion $\omega = \nu_h/(\nu_h + \nu_e)$, Ψ becomes $\omega\Phi + (1-\omega)\tilde{\Phi}$, and the rest of the theorem is unchanged with the corresponding switching mass and bound.[6]

Example 4 (Example 3 with memory). Take the market of Example 3: $r = 1/9$, $\theta = 2$, $\lambda = 6$, downsides uniform on $[0, 10]$ and $\alpha = 1/2$. The launch equilibria are 0.114, 0.343 and 1. With memory there is a unique stationary share, 0.052. The restless band is $(0.23, 0.82)$: these firms,

[6] The switching mass becomes $2\nu_h\nu_e(\nu_h + \nu_e)^{-1}[\Phi(m) - \tilde{\Phi}(m)]$ and the bound $\omega + (1-\omega)\tilde{\Phi}(1)$. Proposition 21(ii) below uses $\omega = \frac{1}{2}$ and can fail for small ω: in Example 3, the long-run share is 0.052 for $\omega = \frac{1}{2}$, 0.026 for $\omega = 0.1$ and 0.022 for $\omega = 0.01$, against 0.024 for a Bayesian market.

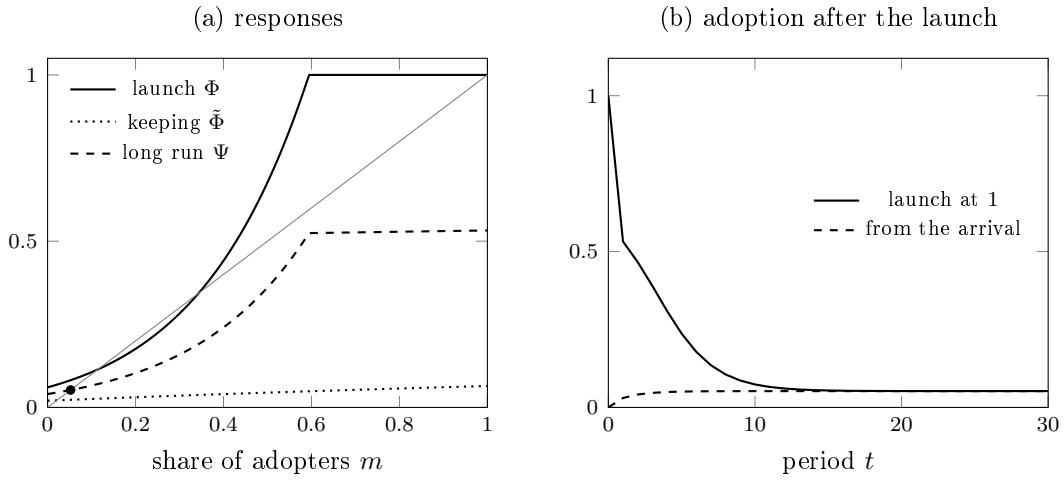


Figure 3: Example 3 with staggered revisions ($\nu = 1/2$). (a) The launch response Φ has three fixed points, but the long-run response Ψ, the average of Φ and of the keeping response $\tilde{\Phi}$, has one, at 0.052 (dot). (b) Starting from full adoption at the launch, the market collapses to 0.052; starting from no adoption it rises to the same share.

5.9% of the market, switch at each of their revisions. Since $\tilde{\Phi}(1) \approx 0.064$, full adoption is far from stationary. Figure 3(a) shows the three responses, and Table 1 gives the long-run share for other sensitivities: with strong fear, full adoption at the launch gives way to a long-run share near one half.

Memory does not remove tipping in general. The stationary shares are the fixed points of Ψ, and with uniform downsides and $r < 1$ the map Ψ can cross the diagonal three times for some θ and λ (Remark 6 in Appendix C); in Example 4 it happens to cross it once.

Remark 4 (Forward-looking firms). Let a revising firm that holds x evaluate the stream of its payoffs with its own weights, set by the current period's stakes as in Proposition 12, and a discount factor $\delta \in [0, 1)$, knowing that it keeps what it chooses until its next revision; let a revision, when it occurs, come before that period's payoff, and let the firm's future selves follow a Markov strategy. At a stationary profile the share does not move. With free switching, the best reply of a revising firm is then the myopic one, whatever Markov strategy its future selves follow; the stationary profiles of Theorem 2 are therefore also those of a market of forward-looking firms. With a switching cost $c > 0$, alternating is an equilibrium of a firm's problem, at a given share, if and only if $G_x > (1 - \delta + 2\delta\nu)\, c$ for both technologies, where G_x is the expected one-period gain from switching, computed with the weights of a firm that holds x; the myopic firm of Section 7.5 needs $G_x > c$. Under foresight, alternation thus survives a larger switching cost when revisions are rare, $\nu < 1/2$, and a smaller one when they are frequent, $\nu > 1/2$: a firm that revises rarely keeps what it chooses for a long time, which the myopic benchmark ignores. We do not characterize the stationary profiles of a market of forward-looking firms with switching costs.

7.3 After the launch

The next result follows the market from two starting points: the arrival of the technology, when nobody holds it, and the launch profile of a pragmatic equilibrium, as when a wave of Section 6 has just taken place.

Proposition 20 (Boom and bust). *Let $\nu \leq 1/2$, or let $\alpha = 0$ and $\nu \in (0, 1]$.*

(i) The map $\mathcal{T}_\nu$ is order-preserving: if $x \leq y$ almost everywhere, then $\mathcal{T}_\nu x \leq \mathcal{T}_\nu y$ almost everywhere.

(ii) From the arrival profile $x \equiv 0$, the share of adopters is nondecreasing and converges to the least stationary share.

(iii) Let m^S be a pragmatic equilibrium with launch profile x^S. If no adopter is restless at m^S, that is, if $\Phi(m^S) = \tilde{\Phi}(m^S)$, then x^S is stationary. Otherwise the share of adopters is nonincreasing, falls by $\nu[\Phi(m^S) - \tilde{\Phi}(m^S)]$ in the first period, and converges to the greatest fixed point of Ψ in $[0, m^S]$, which is therefore at most $m^S - \nu[\Phi(m^S) - \tilde{\Phi}(m^S)]$.

(iv) If $\alpha > 0$ and Ψ has a unique fixed point, the share of adopters converges to it from every initial profile.

Part (iii) is a cycle of hype and disillusionment without any information. At the launch, firms judge from the old technology and fear of missing out pushes them in; once they hold the new one, it turns against it. In Example 3 with $\nu = 1/2$, full adoption at the launch is followed by shares of 0.53, 0.47, 0.39, 0.31, and so on down to 0.052 (Figure 3(b)). With $\nu = 1/4$ the market also ends at 0.052, by part (iv), since the stationary share is unique; with $\nu = 3/4$, which part (iv) does not cover, it does so in our computations. Part (ii) says that a market that starts from nothing settles at its least stationary share, which is at most the least launch equilibrium: in Example 3 the market rises from 0 to 0.052, not to 0.114.

What makes a wave last is the loyalty of the adopters at the post-wave share, and fear of missing out works against it.

Corollary 3 (Waves with memory). *Let the assumptions of Proposition 16 hold, and write Φ_θ, $\tilde{\Phi}_\theta$ and Ψ_θ for the three responses at the intrinsic gain θ. Then $\tilde{\Phi}_\theta(1) < 1$ if and only if $\theta < \theta^{**} = \bar{L}\, e^{\alpha \bar{L}}/r - \lambda$, and $\theta^{**} \geq \theta^*$. Hence, at every $\theta \in (0, \theta^{**})$ at which full adoption is a launch equilibrium, as when a wave has just occurred or when quality has fallen back after one, the share of adopters declines from full adoption, and with $\nu \leq 1/2$ it converges to the greatest fixed point of Ψ_θ, which is below one.*

In the running example, $\theta^{**} \approx 13{,}351$, against $\theta^* \approx 2.16$: with memory, the wave of Proposition 16 does not last at any quality the example considers, and the hysteresis of Proposition 16(iii) does not survive. A wave that is sudden and lasting needs loyal adopters, and here only a cost of switching can provide them (Section 7.5). Without fear of missing out and with concave G, on the other hand, there is no sudden wave at all, since the equilibrium is unique and continuous in θ.

Remark 5 (Synchronized revisions). With $\nu = 1$ and $\alpha > 0$, every restless firm switches in every period, as the DM of Proposition 10, and $\mathcal{T}_\nu$ is not order-preserving. Without competitive pressure ($\lambda = 0$), the share then alternates forever between $\Phi(0)$ and $\tilde{\Phi}(0)$ from the arrival profile; with $\lambda > 0$ the cycles persist in our computations, between 0.048 and 0.060 in Example 3. A little staggering restores convergence: Appendix C gives conditions under which staggered revisions guarantee it (Proposition 25), and shows that no two-period cycle survives any staggering (Proposition 26). Market-level cycles of adoption and remorse thus require synchronized revisions, such as a common budget cycle.

7.4 Eager and disloyal

Proposition 21 (Eager and disloyal). *Let $\alpha > 0$ and $D(m) = \theta + \lambda m$. FOMO firms are more eager to adopt than Bayesian firms, and less loyal once they have adopted: $\tilde{L}(m) < rD(m) < L^*(m)$ for every m. Moreover:*

(i) (Excellent technologies.) *If G has bounded support $[0, \bar{L}]$ and $r\theta \geq \bar{L}$, the Bayesian market has the unique stationary share 1. If moreover $\tilde{\Phi}(1) < 1$, every stationary share of the FOMO market is at most $\frac{1}{2}[1 + \tilde{\Phi}(1)] < 1$.*

(ii) (Poor technologies.) *If G is uniform on $[0, \bar{L}]$, $r \leq 1$ and $rD(1) \leq \bar{L}/2$, then $\Psi(m) > G(rD(m))$ for every m, and every stationary share of the FOMO market exceeds the Bayesian equilibrium.*

At the launch, fear of missing out helps any technology (Proposition 13). In the long run it helps poor technologies and can hurt excellent ones. A technology good enough to be adopted by every Bayesian firm is abandoned in part by FOMO firms whenever some adopters are restless at full adoption, because each of them keeps wondering whether it should have stayed. Example 3 satisfies the conditions of part (ii): the long-run share is 0.052 against 0.024 for a Bayesian market. For part (i), take $r = 1/9$, $\theta = 9$, $\lambda = 0$, downsides uniform on $[0, 1]$ and $\alpha = 1/2$. Every Bayesian firm adopts. In the FOMO market $\tilde{L} \approx 0.70$, the long-run share is 0.85, and 30% of the firms are restless. Churn has a welfare cost of its own, borne by the restless firms, and the rat race of Proposition 18 survives memory with one exception: the firms that would be better off with the new technology at the true probabilities, but that fear of missing out makes restless, gain from the extra pressure that makes them loyal (Proposition 24 in Appendix C).

7.5 Switching costs

Suppose that a firm pays a cost $c \geq 0$, in utility and under both scenarios, whenever it changes technology. A switching cost enters the stakes as well as the payoffs: it lowers what a firm could gain by switching in either direction, and so what it could miss by standing pat. For $D > 0$ write $L^*(D) = rDe^{\alpha D}$ and let $\tilde{L}(D)$ solve $\tilde{L}e^{\alpha \tilde{L}} = rD$, and set $L^*(D) = 0$ for $D \leq 0$. Lemma C.2 in Appendix C extends Theorem 2 to any pair of continuous nondecreasing thresholds $A(m)$ for adopting and $K(m)$ for keeping, including the case $A \leq K$, in which the firms in $[A(m), K(m)]$ form a *loyal band* that keeps whatever it holds and every share with $G(A(m)) \leq m \leq G(K(m))$ is stationary.

Proposition 22 (Switching costs). *(i) A firm that holds the old technology adopts if and only if $L < A_c(m) = \big[L^*(D(m) - c) - c\big]^+$, and one that holds the new technology keeps it if and only if $L \leq K_c(m) = c + \tilde{L}(D(m) + c)$. Lemma C.2 applies. The adoption threshold falls and the keeping threshold rises with c.*

(ii) Let $\alpha > 0$. At the share m, the restless band $(K_c(m), A_c(m))$ is empty if and only if $c \geq \bar{c}(m)$, where $\bar{c}(m) \in (0, D(m))$ is the unique solution of $2c = L^*(D(m) - c) - \tilde{L}(D(m) + c)$. Moreover $\bar{c}(m) \leq \frac{1}{2}[L^*(D(m)) - \tilde{L}(D(m))]$: half the width of the restless band without costs suffices to stop churn.

With a large enough cost the restless band closes, a loyal band opens, and a market moved by fear of missing out remembers instead of churning. In Example 3, $\bar{c}$ is 0.21 at the stationary share without costs, 0.052, where the upside is 2.3, and 3.0 at full adoption, where it is 8. With $c = 2$, the launch equilibria are 0, 0.965 and 1. With $\nu = 1/2$, after a full-adoption launch the market settles

at 0.25 instead of 0.052, whereas from the arrival it stays at 0: costs restore part of the hysteresis. They cannot restore all of it here. Every adopter is loyal at full adoption only if $c \geq 8.9$, while full adoption is a launch equilibrium only if $c \leq 2.15$. A switching cost buys loyalty at the price of eagerness.

8 Testable implications

The model's predictions can be tested on choices alone. Bets placed below the frontier reveal the weights exactly (Proposition 1), so that odds, below, are choice data. The sharpest predictions separate fear of missing out from accounts in which beliefs move only with information.

(1) *Availability, not information.* Odds on scenarios shift when a technology becomes available, compared with when it is only described, by the factor of Proposition 7(i). An account in which beliefs move only with information predicts no shift.

(2) *Frontier dependence.* An entrant that does not beat the best incumbent under any scenario leaves odds unchanged, however bad it is under some scenario, and adding an option worse than the status quo under every scenario leaves them unchanged too (Proposition 7(iii) and Remark 2). The shift caused by a given entrant falls as the frontier rises.

(3) *An unchosen option.* Adding an option that nobody chooses, and that shines under one scenario only, can reverse the choice between two others, and only toward the one that does better under that scenario (Section 4.4). This tests the belief channel with no bet at all. Menus of two options cannot do it, since a DM with fixed beliefs and gain-inflating tastes behaves in the same way on them (Proposition 6).

(4) *Common reception and disagreement.* Users who face the same technologies but hold different ones disagree, each overweighting the scenarios under which her own technology does badly, and revise their odds by the same factor at an arrival (Propositions 9 and 7(ii)). The size of the disagreement identifies α.

(5) *Scale.* Bayesians are indifferent to scaling upside and downside together; DMs who fear missing out favor the scaled-up, radical version (Proposition 8(ii)). In bets offered against a bare status quo, the log of the odds at which the DM bets on a scenario is affine in the bet's upside, with slope α (Proposition 1(v)).

(6) *Stakes, not cascades.* Exogenous variation in others' adoption that carries no information about the scenario, such as adoption decisions of computer players programmed independently of the state, shifts beliefs about the scenario provided it changes the cost of staying behind (Proposition 17); it does not under informational cascades. Conversely, information about others' private signals that leaves the stakes unchanged shifts beliefs under cascades and not here.

(7) *Evidence and stakes.* Data of known likelihood ratio and the arrival of a technology move the elicited log-odds additively, with no interaction; the datum that exactly cancels an arrival measures α in units of evidence; and as data that separate the scenarios accumulate, the effect of an arrival fades (Proposition 3 and Corollary 1).

(8) *Churn.* After a launch there is persistent two-way switching and, unless switching costs make adopters loyal, waves followed by abandonment; aggregate two-period cycles of adoption and

remorse appear only where firms revise at common dates (Theorem 2, Propositions 20, 22 and 26).

9 Interpretation

The theorems say what the DM does. This section says what she is doing.

Attention follows the stakes; evidence does not. The firm of the introduction is not revising its evidence. Nothing has been observed; the two scenarios have the support they had. What changes is which of them it holds before its mind while it decides. The model separates two things that a Bayesian keeps together: the DM's assessment Q_{base} of how well each scenario is specified, which answers to evidence, and the weight each scenario carries in the decision at hand, which answers to the stakes as well. Proposition 3 makes the separation exact. The idea that practical stakes bear on belief is known in epistemology as pragmatic encroachment (Fantl and McGrath, 2002; Stanley, 2005): whether one knows that the bank is open on Saturday may depend, on this view, on how much rides on it. The model is a decision-theoretic cousin of that view, with a difference. Here stakes do not raise the bar that evidence must clear; they redirect attention toward the scenarios that bear on the choice. A scenario may be well supported and ignored, because the decision is settled under it; another may be poorly supported and dwelt upon, because the decision turns on it.

Fear of missing out is not wishful thinking. William James held that when an option is live, forced and momentous, and the evidence does not decide, our "passional nature" may legitimately decide for us (James, 1896). William Clifford had held that it is wrong, always and for anyone, to believe anything upon insufficient evidence (Clifford, 1877). A DM who fears missing out looks like a Jamesian: she lets the momentousness of a choice shape which scenarios she credits. The danger in James's position is well known. If the stakes may decide, what stops them from deciding in favor of what we would like to be true? The stakes as defined here close this door. A wishful thinker tilts toward the scenarios under which she is well off. Our DM tilts toward those under which she is wrong to stand pat. A scenario under which her current technology is excellent is a pleasant one, and she gives it no extra weight, because nothing turns on it. What she attends to is neither pleasant nor unpleasant scenarios but pivotal ones, and pivotality is a property of the gap between what she holds and what she could hold, not of the level of either. Fear of missing out, so understood, is a vigilance against one's own complacency.

The two dispositions need not compete for the same evidence. Wishful thinking is documented where no alternative to what subjects hold is on the table: subjects randomly made to gain from a high price predict higher prices (Mayraz, 2011b), and subjects facing a possible shock or loss are less likely to recognize the signals that announce it (Engelmann et al., 2024). In such settings our DM's only stakes come from the payments for accuracy. When these are the same whatever the outcome, the stakes are roughly equal across scenarios, and the model predicts no tilt. Once an alternative is within reach, it predicts a tilt toward the scenarios under which the status quo is most regretted, which is the opposite of wishful thinking when the alternative is safe. The sharp test is to put a safe alternative on the table (Section 8).

The vigilance is not irrational, and it is not safe. A DM who fears missing out can be sold a mirage: an option dazzling under one scenario and ruinous under the others captures her attention, and her protection shrinks exponentially as the dazzle grows (Proposition 8). She cannot be sold a scare: an option worse than what she holds under every scenario, however dreadful, does not move

her (Remark 2). The model describes a vigilance against complacency, not against error in general. A physician should not overlook a treatable variant, and fear of missing out is the right vigilance for her; an investor shown a prospectus should look first at what could go wrong, and it is the wrong one for him.

Live options. What makes a technology "available"? James distinguished live options, which appeal as real possibilities, from dead ones. The menu of the model is best read as the set of live options: those the DM considers as things she could do. A product that is announced may or may not be live; one that is on the shelf of a competitor may be live for some users and not for others. The prediction of the model is that beliefs move when an option becomes live, not when information about it arrives, and whether an announcement makes an option live is an empirical question the model leaves open.

Two readings of the weights. The model licenses stakes-driven attention under one condition: the evidence has left several scenarios standing (Corollary 1). In that space of underdetermination something must decide which scenario to act on, and the model's answer is that the decision itself decides. There is a precise sense in which the DM remains a Bayesian while doing so. The arrival of an entrant updates her weights on scenarios by the likelihood $e^{\alpha\Delta}$ (Proposition 7): at the level of scenarios she conditions, but on what is at stake rather than on what is observed. At the level of states she does not condition at all, since her prior over states moves for reasons unrelated to them. Is $Q^*_{\mathcal{F},h}$ then what she believes? On one reading it is not: her belief is Q_{base} throughout, and $Q^*_{\mathcal{F},h}$ records how her limited deliberative capacity is allocated among scenarios when a decision is at hand, the entropic cost being the cost of that capacity. On this reading, which the expected utility form of the model favors, the DM commits no epistemic fault, and Clifford has nothing to object. On the other reading, $Q^*_{\mathcal{F},h}$ is what she believes while engaged, and Clifford's objection applies in full. The formalism does not choose; it shows that the same behavior supports both stories, and that whether fear of missing out is a lapse or a virtue is a question about attention, not about evidence.

Contagious vigilance. In a market, the attention reading meets a new difficulty. Each firm's vigilance is reasonable on its own terms: under the transformative scenario, the more competitors adopt, the more the firm would lose by standing pat, and that is a good reason to look at that scenario more closely. Yet the vigilance of all becomes the credulity of each. In the full-adoption equilibrium of Example 3, every firm acts as if the transformative scenario had probability 0.86, and its reason for attending to that scenario is that the others have adopted. No evidence has been exchanged. Such a market looks like a herd, but not like the herds of informational cascades: there, each agent defers to what the others know; here, each attends to what the others' choices put at stake. Attention that follows the stakes is harmless when the stakes are given. When the choices of others create them, attention produces the stakes that justify it, the market as a whole acts on a confidence that nothing supports, and Proposition 18 shows that it may pay for it. The same vigilance then turns against what it started. A firm that has adopted attends to the scenario under which it should have stayed, and a market of FOMO firms stages its own disillusionment (Section 7). Credulity and remorse are two faces of one disposition: the attention that follows what one could miss.

10 Final remarks

We have characterized, from her choices alone, a decision maker who, facing the technologies on the table, one of which she uses, gives more weight to the scenarios under which standing pat would make her miss the most, and who otherwise maximizes expected utility. Evidence moves her baseline, the stakes move her weights away from it, and the two act as separate likelihoods. In a market, others' adoption raises the stakes: at the launch, fear of missing out can make adoption tip and rest on weights that no evidence supports; afterwards, holding the technology turns the fear around, and markets churn and reverse their waves unless switching costs make adopters loyal.

Foresight in valuation does not remove the churn when switching is free (Proposition 12 and Remark 4). The main limitation lies elsewhere. Our DM looks ahead when she values a plan, but not when she decides where to look: her attention follows the stakes of the current period. When attention also looks ahead, a patient DM may have no stable plan at all (footnote 4), and what she then does is open; so is how firms behave when they anticipate the path of adoption after a launch. Whether the value of a technology lies in joining others or in not falling behind them changes who bears the cost of fear of missing out, and when it can cure inertia (Section 6.6); which of the two a given technology is seems to us the question that matters most for policy. One extension seems real. A DM who also attended to what she could lose by moving, what the computer industry calls fear, uncertainty and doubt, would add the downside of the menu to the stakes and be swayed by the options that Remark 2 declares irrelevant; such a fear would make adopters loyal rather than restless.

Whether attention should follow the stakes is a question as old as the quarrel between James and Clifford. New technologies put it to every user, every time one arrives, and markets put it to all of them at once.

A Proofs

In a pragmatic representation, for every situation σ and all $p \neq p'$,

$$\log o^*_\sigma(p,p') = \log o_{\text{base}}(p,p') + \alpha\big[\rho_\sigma(p) - \rho_\sigma(p')\big]. \tag{1}$$

Stakes profiles are elements k of $\mathbb{R}^{\mathcal{P}}_+$, and 0 is the zero profile.

A.1 Preliminary lemmas

Lemma A.1 (Realizability)**.** *Let $\mathcal{P}$ be linearly independent and u affine and onto. For every $\eta \in \mathbb{R}^{\mathcal{P}}$ there is an act f with $u(f^p) = \eta(p)$ for all $p \in \mathcal{P}$.*

Proof. Linear independence makes $v \mapsto (\langle v, p\rangle)_{p\in\mathcal{P}}$ a linear map from $\mathbb{R}^S$ onto $\mathbb{R}^{\mathcal{P}}$; pick v with $\langle v, p\rangle = \eta(p)$ for all p. Since u is onto, pick for each s a lottery $f(s)$ with $u(f(s)) = v(s)$. By affinity, $u(f^p) = \sum_s p(s)u(f(s)) = \langle v, p\rangle = \eta(p)$. □

Lemma A.2. *Let $\mathcal{P}$ be linearly independent, let u be affine and onto, and compute stakes with u.*

(i) (Profiles.) *Every $k \in \mathbb{R}^{\mathcal{P}}_+$ is the stakes profile of some situation.*

(ii) (Arrivals.) *For every situation $(\mathcal{F}, h)$ and act e, $\rho_{\mathcal{F}\cup\{e\},h} = \rho_{\mathcal{F},h} + \Delta_e$, where $\Delta_e(p) = \big(u(e^p) - b_{\mathcal{F}}(p)\big)^+$ does not depend on h. For all $k, k', \Delta \in \mathbb{R}^{\mathcal{P}}_+$ there are a menu $\mathcal{F}$, status quos $h, h' \in \mathcal{F}$ and an act e with $\rho_{\mathcal{F},h} = k$, $\rho_{\mathcal{F},h'} = k'$ and $\Delta_e = \Delta$.*

Proof. (i) Take by Lemma A.1 acts h, f with $u(h^q) = 0$ and $u(f^q) = k(q)$; the situation $(\{h, f\}, h)$ has profile k. (ii) The frontier of $\mathcal{F} \cup \{e\}$ is $\max\{b_{\mathcal{F}}, u(e^{\cdot})\} = b_{\mathcal{F}} + (u(e^{\cdot}) - b_{\mathcal{F}})^+$, and subtracting $u(h^{\cdot})$ gives the first claim. For the second, Lemma A.1 gives acts g, h, h', e with $u(g^q) = 0$, $u(h^q) = -k(q)$, $u(h'^q) = -k'(q)$ and $u(e^q) = \Delta(q)$. The menu $\mathcal{F} = \{g, h, h'\}$ has frontier 0, so $\rho_{\mathcal{F},h} = k$, $\rho_{\mathcal{F},h'} = k'$ and $\Delta_e = \Delta$. □

Proof of Proposition 1. (i) If $f \in \mathcal{F}$ has profile $b_{\mathcal{F}}$, then $V(f) = \sum_p Q(p) b_{\mathcal{F}}(p)$ and $D = 0$. Conversely, if $D = 0$, a V-best f has $\sum_p Q(p)[b_{\mathcal{F}}(p) - u(f^p)] = 0$ with nonnegative terms and Q of full support, so $u(f) = b_{\mathcal{F}}$. (ii) The acts g and $g_{x,y}$ are below the frontier of $\mathcal{F}$, so the enlarged menu has the frontier of $\mathcal{F}$ and the situation has the weights Q. Since $V(g) = \sum_p Q(p) b_{\mathcal{F}}(p) - \delta > \max_{\mathcal{F}} V$, the best element of the enlarged menu is g or $g_{x,y}$, and $V(g_{x,y}) \geq V(g)$ if and only if $Q(p)x \geq Q(p')y$. (iii) Under a scenario with $\rho_{\mathcal{F},h}(p) = 0$ the status quo is on the frontier, so the new menu has frontier $b_{\mathcal{F}}$. None of its options has profile $b_{\mathcal{F}}$: not h, since $\rho_{\mathcal{F},h} \neq 0$, and not e_p, which is below h under some scenario. (iv) Every act of such a menu is weakly worse than h under every scenario. A situation with stakes $c\mathbf{1}$ has weights proportional to $Q_{\text{base}}(p)e^{\alpha c}$, that is, Q_{base}. (v) The stakes of $(\{h, b_{\varepsilon,t}\}, h)$ are ε under p and zero elsewhere, so by (1) its odds of p against p' are $o_{\text{base}}(p, p')e^{\alpha\varepsilon}$, and $V(b_{\varepsilon,t}) - V(h) = Q^*(p)\varepsilon - Q^*(p')t\varepsilon$. □

A.2 Test pairs

For a profile b and $M > 0$, let $S_M(b)$ be a set of acts e_p, one for each scenario, with $u(e_p^p) = b(p)$ and $u(e_p^q) = b(q) - M$ for $q \neq p$. A menu that contains $S_M(b)$ and acts below b has frontier b, and under any full-support weights the expected utility of e_p tends to $-\infty$ as M grows.

The proofs use one device. Let h be an act, $k \in \mathbb{R}_+^{\mathcal{P}}$ nonzero, $b = u(h) + k$, let Q, Q' be full-support weights and $d \in \mathbb{R}^{\mathcal{P}}$. Choose $\delta > 0$ smaller than $\sum_p Q(p)k(p)$ and $\sum_p Q'(p)k(p)$, let g have profile $b - \delta\mathbf{1}$, and let f have profile $u(g) + \eta d$, with $\eta > 0$ so small that f is below b and both weight vectors rank f above h, as they rank g. Call (f, g) a *test pair*. Consider a menu made of h, f, g and of *props*: acts taken from sets $S_M(b)$ or $S_M(b + c\mathbf{1})$, or the act e of Lemma A.5. A prop is at a fixed level under one scenario and about M below b under the others, so for M large it is worth less than h under Q and under Q'. If the weights of the menu's situation are Q or Q' and M is large, Axiom 1 says that the choice from it is the set of best elements of $\{f, g\}$ under those weights: f is chosen if and only if the weights give d a nonnegative expectation.

Unique weights. Let $\rho_{\mathcal{F},h} \neq 0$, and suppose that $Q \neq \tilde{Q}$ both satisfy Axiom 1 for $(\mathcal{F}, h)$. Pick d with $\sum_p Q(p)d(p) > 0 > \sum_p \tilde{Q}(p)d(p)$ and a test pair (f, g) for h and $b = b_{\mathcal{F}}$. For M large, the choice from $\{h, f, g\} \cup S_M(b)$ would be both $\{f\}$ and $\{g\}$.

Revealed odds. Let σ have nonzero stakes and weights Q. If $f \in C(\mathcal{G}, h)$ and $g \in \mathcal{G}$, for a menu $\mathcal{G} \ni h$ with the frontier of σ, and $u(f) - u(g) = x(\mathbf{1}_p - t\mathbf{1}_{p'})$ with $x > 0$, Axiom 1 gives $Q(p) \geq tQ(p')$. Conversely, a test pair with $d = \mathbf{1}_p - \frac{Q(p)}{Q(p')}\mathbf{1}_{p'}$ has $f \in C(\{h, f, g\} \cup S_M(b), h)$ for M large. So the maximum defining $o_\sigma(p, p')$ is attained at $t = Q(p)/Q(p')$.

A.3 How the weights depend on the situation

Lemma A.3 (Only the stakes matter)**.** *Under Axioms 1 and 3, there is a map Γ from $\mathbb{R}_+^{\mathcal{P}} \setminus \{0\}$ to the full-support elements of $\Delta(\mathcal{P})$ such that $Q_{\mathcal{F},h} = \Gamma(\rho_{\mathcal{F},h})$ for every situation with nonzero stakes.*

Proof. Let $(\mathcal{F}, h)$ and $(\mathcal{G}, h')$ have the same nonzero profile k and weights $Q \neq Q'$. Pick d with $\sum_p Q(p)d(p) > 0 > \sum_p Q'(p)d(p)$ and a test pair (f, g) for h. Let $\mathcal{T} = \{h, f, g\} \cup S_M(b_{\mathcal{F}})$, let τ send h to h' and every other act of $\mathcal{T}$ to an act whose profile is its own plus $t = u(h') - u(h) = b_{\mathcal{G}} - b_{\mathcal{F}}$, and let $\mathcal{T}' = \tau(\mathcal{T})$. Then τ satisfies the premise of Axiom 3, $\mathcal{T}$ has the frontier of $\mathcal{F}$ and $\mathcal{T}'$ that

of $\mathcal{G}$. For M large, $C(\mathcal{T},h)=\{f\}$ and $C(\mathcal{T}',h')=\{\tau(g)\}$ by Axiom 1, which contradicts Axiom 3. Every nonzero $k\in\mathbb{R}^{\mathcal{P}}_{+}$ is the profile of a situation (Lemma A.2(i)). □

Lemma A.4 (Uniform improvements). *Under Axioms 1–3, $\Gamma(k+c\mathbf{1})=\Gamma(k)$ for every nonzero $k\in\mathbb{R}^{\mathcal{P}}_{+}$ and $c\geq 0$.*

Proof. Let $Q=\Gamma(k)\neq Q'=\Gamma(k+c\mathbf{1})$, pick d as before and a test pair (f,g) for an act h and $b=u(h)+k$. The menus $\mathcal{T}=\{h,f,g\}\cup S_M(b)$ and $\mathcal{T}'=\{h,f,g\}\cup S_M(b+c\mathbf{1})$ have parallel frontiers, profiles k and $k+c\mathbf{1}$, and for M large $C(\mathcal{T},h)=\{f\}$ and $C(\mathcal{T}',h)=\{g\}$. This contradicts Axiom 2. □

Lemma A.5 (Arrivals). *Under Axioms 1, 3 and 4, let $k\in\mathbb{R}^{\mathcal{P}}_{+}$ be nonzero, $p\in\mathcal{P}$ and $s>0$, and write $Q=\Gamma(k)$, $Q'=\Gamma(k+s\mathbf{1}_p)$. Then $Q'(q)/Q'(q')=Q(q)/Q(q')$ for all $q,q'\neq p$, and $Q'(p)/Q'(q)\geq Q(p)/Q(q)$ for all $q\neq p$.*

Proof. Let $d\in\mathbb{R}^{\mathcal{P}}$ with $d(p)\geq 0$ and $\sum_q Q(q)d(q)=0$, and take a test pair (f,g) for an act h and $b=u(h)+k$. Let $\mathcal{T}=\{h,f,g\}\cup S_M(b)$ and let e have profile $b(p)+s$ under p and $b(q)-M$ under $q\neq p$. Then e beats the frontier of $\mathcal{T}$ under p only, $(\mathcal{T},h)$ has profile k and $(\mathcal{T}\cup\{e\},h)$ has profile $k+s\mathbf{1}_p$. For M large, $C(\mathcal{T},h)=\{f,g\}$. If $\sum_q Q'(q)d(q)<0$, then $C(\mathcal{T}\cup\{e\},h)=\{g\}$ although $f^p\succsim g^p$, which contradicts Axiom 4. So $\sum_q Q'(q)d(q)\geq 0$. Apply this to $d=Q(q')\mathbf{1}_q-Q(q)\mathbf{1}_{q'}$ and to its opposite, for $q,q'\neq p$, and to $d=Q(q)\mathbf{1}_p-Q(p)\mathbf{1}_q$. □

A.4 Proof of Theorem 1

Necessity. Let C have a separable representation with nondecreasing φ_p, and write V and Q for its values and weights; a pragmatic representation is the case $\varphi_p(s)=\alpha s$. The weights have full support, and given the status quo they depend on the menu only through its frontier: this is Axiom 1. For τ as in Axiom 3, $(\mathcal{F},h)$ and $(\mathcal{G},\tau(h))$ have the same stakes, and the value of $\tau(f)$ differs from that of f by a constant: this gives the axiom. For Axiom 4, the weights after the arrival are proportional to $Q(q)\,w^{\mathbf{1}\{q=p\}}$ for some $w\geq 1$, so the value of f minus that of g is, up to a positive factor, $V(f)-V(g)+Q(p)(w-1)[u(f^p)-u(g^p)]\geq 0$; since g is chosen after the arrival, so is f.

If moreover $\varphi_p(s)=\alpha s$, situations with the same status quo and parallel frontiers have stakes that differ by a constant, hence the same weights: this gives Axiom 2. By Lemma A.2(ii) and (1), $\log o^*_{\mathcal{F}\cup\{e\},h}(p,p')-\log o^*_{\mathcal{F},h}(p,p')=\alpha[\Delta_e(p)-\Delta_e(p')]$, which does not depend on h: this is Axiom 5.

Sufficiency.

Step 1: separation. Let $|\mathcal{P}|\geq 3$, and let C satisfy Axioms 1, 3 and 4. Axiom 1 and Lemmas A.3 and A.5 give u, Γ and its properties. For distinct p,p' and nonzero k, let $L_{pp'}(k)=\log[\Gamma(k)(p)/\Gamma(k)(p')]$. By Lemma A.5, raising one coordinate q of a nonzero profile leaves $L_{pp'}$ unchanged if $q\notin\{p,p'\}$, raises it weakly if $q=p$ and lowers it weakly if $q=p'$. Two nonzero profiles that agree at p and p' are joined through nonzero profiles by raising coordinates, then lowering others. Hence $L_{pp'}(k)$ depends on k only through $(k(p),k(p'))$, is nondecreasing in the first and nonincreasing in the second. Since $|\mathcal{P}|\geq 3$, this defines $\Lambda_{pp'}$ on all of $\mathbb{R}^2_+$ with $L_{pp'}(k)=\Lambda_{pp'}(k(p),k(p'))$.

For p and $q\neq p$, let $\varphi_p(s)=\Lambda_{pq}(s,1)-\Lambda_{pq}(0,1)$. This does not depend on q: for distinct $q,q'\neq p$, evaluating $L_{pq'}=L_{pq}+L_{qq'}$ at $s\mathbf{1}_p+\mathbf{1}_q+\mathbf{1}_{q'}$ gives $\Lambda_{pq'}(s,1)=\Lambda_{pq}(s,1)+\Lambda_{qq'}(1,1)$. Moreover φ_p is nondecreasing and $\varphi_p(0)=0$. For distinct p,p' and $q\notin\{p,p'\}$, evaluating $L_{pp'}=L_{pq}+L_{qp'}$ at $s\mathbf{1}_p+t\mathbf{1}_{p'}+\mathbf{1}_q$ gives $\Lambda_{pp'}(s,t)=\Lambda_{pq}(s,1)-\Lambda_{p'q}(t,1)$, that is,

$$L_{pp'}(k)=c_{pp'}+\varphi_p\big(k(p)\big)-\varphi_{p'}\big(k(p')\big),\qquad c_{pp'}=\Lambda_{pp'}(0,0).$$

Since $L_{pp''} = L_{pp'} + L_{p'p''}$ and $L_{p'p} = -L_{pp'}$, the same relations hold for the constants c. Let $Q_{\text{base}}(p) \propto e^{c_{pp_0}}$ for a fixed p_0, with $c_{p_0p_0} = 0$; then $Q_{\text{base}}(p)/Q_{\text{base}}(p') = e^{c_{pp'}}$ and $\Gamma(k)(p) \propto Q_{\text{base}}(p)e^{\varphi_p(k(p))}$. So C maximizes expected utility under $Q^{\varphi}_{\mathcal{F},h}$ in every situation with nonzero stakes. With zero stakes, every act of the menu is weakly below h, so that the maximizers are the same under any full-support weights. Hence C has a separable representation with nondecreasing φ_p.

Step 2: linearity. Let C also satisfy Axiom 2. By Lemma A.4, evaluated at $k = s\mathbf{1}_p + t\mathbf{1}_{p'} + \mathbf{1}_q$, for distinct p, p' and all $s, t, c \geq 0$,

$$\varphi_p(s + c) - \varphi_p(s) = \varphi_{p'}(t + c) - \varphi_{p'}(t).$$

So this difference depends neither on s, nor on t, nor on the scenario; call it $\psi(c)$. Then ψ is additive and nondecreasing on $\mathbb{R}_+$, hence $\psi(c) = \alpha c$ with $\alpha \geq 0$, and $\varphi_p(s) = \psi(s) = \alpha s$ for every p. The separable representation is a pragmatic one.

Step 3: two scenarios. Let $\mathcal{P} = \{p, p'\}$ and let C satisfy Axioms 1–5. By Remark 1, $L_{pp'}(k) = \log o_{\text{base}}(p, p') + \ell(k(p) - k(p'))$ for nonzero k. Let $x, \delta \in \mathbb{R}$. Take a nonzero k with $k(p) - k(p') = x$, $k' = \mathbf{1}$, and $\Delta(p) = \delta^+$, $\Delta(p') = \delta^-$. Lemma A.2(ii) gives a menu, two status quos and an act that realize k, k' and Δ, and Axiom 5 gives $\ell(x + \delta) - \ell(x) = \ell(\delta) - \ell(0) = \ell(\delta)$. So ℓ is additive on $\mathbb{R}$ and nondecreasing, hence $\ell(x) = \alpha x$ with $\alpha \geq 0$, and C has a pragmatic representation.

Uniqueness. Two representations have $u' = \gamma u + \delta$, since both represent the choices among lotteries. Since weights are unique, they give the same weights to every situation with nonzero stakes: $Q_{\text{base}}(p)e^{\alpha k(p)} \propto Q'(p)e^{\alpha'\gamma k(p)}$ for every nonzero k. Take $k = \mathbf{1}$, then $k = \mathbf{1}_p$. The converse is immediate. For separable representations, with u fixed and distinct p, p', $q \notin \{p, p'\}$, the odds of p against p' are $Q_{\text{base}}(p)/Q_{\text{base}}(p')$ at the profile $\mathbf{1}_q$ and $e^{\varphi_p(s)}Q_{\text{base}}(p)/Q_{\text{base}}(p')$ at $s\mathbf{1}_p + \mathbf{1}_q$, which pins down Q_{base} and φ_p.

The case $\alpha = 0$. If $\alpha = 0$, all weights equal Q_{base} and $C(\mathcal{F}, h)$ does not depend on h. If $\alpha > 0$, take acts h, h' and a profile b with $b - u(h) = \mathbf{1}_p$ and $b - u(h') = \mathbf{1}_{p'}$, for $p \neq p'$. By (1), the weights Q and Q' of the two situations with status quos h, h' and frontier b differ. Pick d with $\sum_q Q(q)d(q) > 0 > \sum_q Q'(q)d(q)$, and a test pair (f, g) with δ and η so small that both weight vectors rank f and g above h and h'. For M large, the menu $\{h, h', f, g\} \cup S_M(b)$ has $C(\cdot, h) = \{f\}$ and $C(\cdot, h') = \{g\}$. □

A.5 Proofs of Remarks 1, 2 and 3 and of Proposition 2

Remark 1. With two scenarios, the lemmas give $L_{pp'}(k) = \Lambda(k(p), k(p'))$ for $k \neq 0$, monotone as in Step 1, with $\Lambda(s + c, t + c) = \Lambda(s, t)$. Hence $\Lambda(s, t) = \bar{\ell}(s - t)$, where $\bar{\ell}(x) = \Lambda(x, 0)$ for $x > 0$, $\bar{\ell}(x) = \Lambda(0, -x)$ for $x < 0$ and $\bar{\ell}(0) = \Lambda(c, c)$ for any $c > 0$; $\bar{\ell}$ is nondecreasing, since $\Lambda(x, 0) = \Lambda(x + c, c) \geq \Lambda(c, c)$ for $x > 0$, and similarly on the other side. Set $\ell = \bar{\ell} - \bar{\ell}(0)$ and $Q_{\text{base}}(p)/Q_{\text{base}}(p') = e^{\bar{\ell}(0)}$. With zero stakes, Step 1 applies. The converse is checked as in the necessity part of the theorem. □

Remark 2. Since $u(g) \leq u(h) \leq b_{\mathcal{F}}$, the frontier and the stakes do not change, nor do the weights, and $V(g) \leq V(h)$. □

Remark 3. In each model the weights have full support and, given the status quo, depend on the menu only through its frontier, so that Axiom 1 holds. In the first model, the weights after an arrival are proportional to $Q(q)w^{\mathbf{1}\{q=p\}}$ with $w < 1$ for ρ raised at p; stakes that differ by a constant give the same weights, and the situations that Axiom 3 relates have the same stakes, so that Axioms 2 and 3 hold. The menus of Lemma A.5, with $d(p) > 0$, $\sum_q Q(q)d(q) = 0$ and $\sum_q Q'(q)d(q) < 0$, violate Axiom 4. In the second and third models, $w \geq 1$ and the situations that Axiom 3 relates have the same stakes, so that Axioms 4 and 3 hold. Some profile k and $c > 0$

have $\Gamma(k+c\mathbf{1}) \neq \Gamma(k)$: otherwise, as in Step 2, the α_p would be equal or φ affine. The menus of Lemma A.4 then violate Axiom 2. In the fourth model, for a fixed status quo the factor $e^{\gamma u(h^p)}$ is a constant of the situation, so Axioms 2 and 4 hold as in the necessity part. Two situations with the same stakes and status quos whose profiles differ by a nonconstant vector have different weights, and the menus of Lemma A.3 violate Axiom 3. □

Proof of Proposition 2. Necessity is that of Theorem 1, with (1). For sufficiency, Pragmatism applied in both directions makes the weights a function Γ of the nonzero stakes profile, and $L_{pp'}(k)$ a function of $(k(p), k(p'))$, monotone as in Step 1. By Lemma A.2(ii), for nonzero k, k' and $\Delta \in \mathbb{R}^{\mathcal{P}}_{+}$ there are a menu, two status quos and an act that realize the profiles k, k', $k+\Delta$ and $k'+\Delta$, all nonzero. Common reception gives $L_{pp'}(k+\Delta) - L_{pp'}(k) = L_{pp'}(k'+\Delta) - L_{pp'}(k')$; call this $\xi(\Delta)$. Then $\xi(\Delta+\Delta') = [L_{pp'}(k+\Delta+\Delta') - L_{pp'}(k+\Delta)] + [L_{pp'}(k+\Delta) - L_{pp'}(k)] = \xi(\Delta') + \xi(\Delta)$, and ξ depends on $(\Delta(p), \Delta(p'))$ only and is monotone in each. Hence $\xi(\Delta) = a\Delta(p) - a'\Delta(p')$ with $a, a' \geq 0$. For nonzero k, k', $L_{pp'}(k+k') = L_{pp'}(k) + \xi(k') = L_{pp'}(k') + \xi(k)$, so $L_{pp'}(k) - \xi(k)$ is a constant $c_{pp'}$: this is the log of the baseline odds. Equal odds at all profiles $c\mathbf{1}$ give $a = a' =: \alpha_{pp'}$. Since odds are ratios of weights, $c_{pp''} = c_{pp'} + c_{p'p''}$, which defines the baseline. If $|\mathcal{P}| = 2$, let α be $\alpha_{pp'}$. If $|\mathcal{P}| \geq 3$, the same relation for L evaluated at unit profiles gives $\alpha_{pp''} = \alpha_{pp'} = \alpha_{p'p''}$ for distinct p, p', p'', and any two pairs of scenarios are linked through such triples. With zero stakes, the maximizers are the same under any full-support weights, as in Step 1. □

A.6 Proof of Proposition 3 and Corollary 1

(i) By definition, $Q^{*\ell}_{\mathcal{F},h}(p) \propto Q^{\ell}_{\text{base}}(p)\, e^{\alpha\rho_{\mathcal{F},h}(p)}$, and $Q^{\ell}_{\text{base}}(p) \propto Q_{\text{base}}(p)\, \ell(p)$. (ii) Take the ratio of the expressions in (i) at p and p'; the normalizations cancel. (iii) Let the stakes profiles of the two situations be k and $k+\mathbf{1}_p$. By (i), the weights in the second are proportional to $Q_{\text{base}}(q)\, \ell(q)\, e^{\alpha k(q)}\, \ell'(q)$, where $\ell'(p) = e^{\alpha}$ and $\ell'(q) = 1$ for $q \neq p$; this is the expression in (i) for the first situation after the datum $\ell\ell'$. (iv) By (ii), the left side is $\alpha|\rho_{\mathcal{F},h}(p) - \rho_{\mathcal{F},h}(p')|$, and since stakes are nonnegative, $-\rho_{\mathcal{F},h}(p') \leq \rho_{\mathcal{F},h}(p) - \rho_{\mathcal{F},h}(p') \leq \rho_{\mathcal{F},h}(p)$.

Corollary 1. (a) By Proposition 3(iv), $\log o^{*\ell_n}_{\mathcal{F},h}(p,p') \leq \log o_{\text{base}}(p,p') + \log\big(\ell_n(p)/\ell_n(p')\big) + \alpha K$ for every situation with $\max\{\rho_{\mathcal{F},h}(p), \rho_{\mathcal{F},h}(p')\} \leq K$, and the right side tends to $-\infty$. (b) By Proposition 3(ii), the term $\log(\ell_n(p)/\ell_n(p'))$ vanishes. □

A.7 Proofs of Propositions 4, 5 and 6

Proof of Proposition 4. (i) Since $x \in \mathcal{F}$ and $u(x) \geq u(f^p)$ for all $f \in \mathcal{F}$ and p, the frontier is $b_{\mathcal{F}}(p) = u(x)$ for every p, so that $\rho_{\mathcal{F},h}(p) = u(x) - u(h^p)$. The factor $e^{\alpha u(x)}$ is common to all scenarios and cancels in the normalization. (ii) By (1), for all p, p',

$$\log o^{*}_{\mathcal{F},h}(p,p') - \log o^{*}_{\mathcal{G},h}(p,p') = \alpha\Big[\big(b_{\mathcal{F}}(p) - b_{\mathcal{G}}(p)\big) - \big(b_{\mathcal{F}}(p') - b_{\mathcal{G}}(p')\big)\Big],$$

since the terms in $u(h^{\cdot})$ cancel. Two weight vectors coincide if and only if all their odds coincide. Since $\alpha > 0$, this holds if and only if $b_{\mathcal{F}} - b_{\mathcal{G}}$ takes the same value at every scenario. □

Proof of Proposition 5. (i) Let $Z = \sum_q Q_{\text{base}}(q) e^{\alpha\rho(q)}$, so that $Q^{\alpha}(p) = Q_{\text{base}}(p) e^{\alpha\rho(p)}/Z$. Since Q_{base} has full support, $\mathrm{KL}(Q \,\|\, Q_{\text{base}})$ is finite for every $Q \in \Delta(\mathcal{P})$, and

$$\sum_p Q(p)\rho(p) - \tfrac{1}{\alpha}\mathrm{KL}(Q \,\|\, Q_{\text{base}}) = \tfrac{1}{\alpha}\sum_p Q(p)\Big[\alpha\rho(p) - \log\tfrac{Q(p)}{Q_{\text{base}}(p)}\Big] = \tfrac{1}{\alpha}\Big[\log Z - \mathrm{KL}(Q \,\|\, Q^{\alpha})\Big].$$

Relative entropy is nonnegative and vanishes only at $Q = Q^\alpha$, so Q^α is the unique maximizer and the value is $\frac{1}{\alpha}\log Z = R_\alpha$. (ii) Let $\Lambda(\alpha) = \log\sum_q Q_{\text{base}}(q)e^{\alpha\rho(q)}$, so that $R_\alpha = \Lambda(\alpha)/\alpha$ and $\Lambda'(\alpha) = \sum_p Q^\alpha(p)\rho(p)$. Then

$$\frac{dR_\alpha}{d\alpha} = \frac{\alpha\Lambda'(\alpha) - \Lambda(\alpha)}{\alpha^2} = \frac{\mathrm{KL}(Q^\alpha \,\|\, Q_{\text{base}})}{\alpha^2} \geq 0,$$

because $\mathrm{KL}(Q^\alpha \,\|\, Q_{\text{base}}) = \sum_p Q^\alpha(p)\big[\alpha\rho(p) - \Lambda(\alpha)\big] = \alpha\Lambda'(\alpha) - \Lambda(\alpha)$. As $\alpha \to 0$, $\Lambda(0) = 0$ and $\Lambda'(0) = \sum_p Q_{\text{base}}(p)\rho(p)$, so $R_\alpha \to \sum_p Q_{\text{base}}(p)\rho(p)$. Let $\bar\rho = \max_p \rho(p)$. Then $\Lambda(\alpha) = \alpha\bar\rho + \log\big(Q_{\text{base}}(M) + \sum_{q\notin M} Q_{\text{base}}(q)e^{\alpha(\rho(q)-\bar\rho)}\big)$, and the logarithm tends to $\log Q_{\text{base}}(M)$, which is finite; hence $R_\alpha \to \bar\rho$. As for $\sum_p Q^\alpha(p)\rho(p) = \Lambda'(\alpha)$, it exceeds R_α by $\mathrm{KL}(Q^\alpha \,\|\, Q_{\text{base}})/\alpha \geq 0$, it is nondecreasing because $\Lambda''(\alpha)$ is the variance of ρ under Q^α, it tends to $\Lambda'(0) = \sum_p Q_{\text{base}}(p)\rho(p)$ as $\alpha \to 0$, and it tends to $\bar\rho$ as $\alpha \to \infty$ by (iii). (iii) For $p \in M$ and $q \notin M$, $Q^\alpha(q)/Q^\alpha(p) = \big(Q_{\text{base}}(q)/Q_{\text{base}}(p)\big)e^{\alpha(\rho(q)-\bar\rho)} \to 0$, while for $p, p' \in M$ the ratio $Q^\alpha(p)/Q^\alpha(p')$ is $Q_{\text{base}}(p)/Q_{\text{base}}(p')$. Hence $Q^\alpha \to Q_{\text{base}}(\cdot \mid M)$. □

Proof of Proposition 6. The stakes of $(\{h, f\}, h)$ are $\max(u(f^p) - u(h^p), 0)$, and the value of f minus that of h is proportional to $\sum_p Q_{\text{base}}(p)e^{\alpha\rho(p)}[u(f^p) - u(h^p)]$. □

A.8 Proofs of Section 5

Proof of Proposition 7. (i) Under every scenario, $\max\{b_{\mathcal{F}}(p), u(e^p)\} - u(h^p) = \rho_{\mathcal{F},h}(p) + \Delta(p)$. Hence the exponent of $Q^1(p)$ is that of $Q^0(p)$ plus $\alpha\Delta(p)$. (ii) Δ is defined from $b_{\mathcal{F}}$ and e. (iii) Since Q^0 has full support, $Q^1 = Q^0$ if and only if the factor $e^{\alpha\Delta}$ is constant across scenarios. An entrant with $u(e^p) \leq b_{\mathcal{F}}(p)$ for all p has $\Delta \equiv 0$. (iv) Let $f \in \mathcal{F}$, $d(p) = u(e^p) - u(f^p)$ and $\varpi(p) = e^{\alpha\Delta(p)} \geq 1$. Where $\Delta(p) > 0$, $u(e^p) > b_{\mathcal{F}}(p) \geq u(f^p)$, so $d(p) > 0$; elsewhere $\varpi(p) = 1$. Hence $\sum_p Q^0(p)\varpi(p)d(p) = \sum_p Q^0(p)d(p) + \sum_{p:\Delta(p)>0} Q^0(p)(\varpi(p)-1)d(p) \geq \sum_p Q^0(p)d(p)$. Write $V^0 = V_{\mathcal{F},h}$ and $V^1 = V_{\mathcal{F}\cup\{e\},h}$. By (i), $V^1(e) - V^1(f) = \sum_p Q^0(p)\varpi(p)d(p) \big/ \sum_p Q^0(p)\varpi(p)$, whose sign is that of the numerator, while $V^0(e) - V^0(f) = \sum_p Q^0(p)d(p)$. So a nonnegative (positive) difference before the arrival remains nonnegative (positive) after it. □

Proof of Proposition 8. (i) In the situation $(\{h, e\}, h)$, the frontier is h under p_1 and e under p_2. So $\rho = (0, D)$, and the odds of p_2 against p_1 are $(q_2/q_1)e^{\alpha D}$. The entrant is adopted if and only if $Q^*(p_2)D > Q^*(p_1)L$, that is, $(q_2/q_1)e^{\alpha D}D > L$. For a Bayesian, $\alpha = 0$. (ii) For e_s the condition reads $sL < (q_2/q_1)\, sD\, e^{\alpha sD}$, that is, $L < (q_2/q_1)\, D\, e^{\alpha sD}$. If $\alpha > 0$, the right side is increasing in s and tends to ∞. For a Bayesian it is free of s. □

Proof of Proposition 9. For a status quo $h \in \mathcal{F}$, $\log Q^*_{\mathcal{F},h}(p) = \log Q_{\text{base}}(p) + \alpha\big(b_{\mathcal{F}}(p) - u(h^p)\big) - \log Z_h$, where Z_h does not depend on p. The terms in $b_{\mathcal{F}}$ are common to h and h', and the normalizations cancel in the odds. What remains is $-\alpha$ times the difference in the status quo's utility, which gives the formula. □

Proof of Proposition 10. From the status quo h, the profile is $\rho = (0, D)$ and the odds of p_2 against p_1 are $(q_2/q_1)e^{\alpha D}$. The DM keeps h if and only if $Q^*(p_2)D \leq Q^*(p_1)L$, that is, $(q_2/q_1)e^{\alpha D} \leq L/D$, that is, $\alpha D \leq \kappa$. From the status quo e, the frontier is h under p_1 and e under p_2, so the profile is $\rho = (L, 0)$ and the odds are $(q_2/q_1)e^{-\alpha L}$. The DM keeps e if and only if $Q^*(p_1)L \leq Q^*(p_2)D$, that is, $\kappa \leq -\alpha L$. With two technologies, a DM who is not at a rest point switches to the other one. (i) If $\alpha > 0$, both technologies would be rest points only if $\alpha D \leq \kappa \leq -\alpha L$, which is impossible; and if $-\alpha L < \kappa < \alpha D$, neither is, so the DM switches in every period. (ii) If $\alpha = 0$, h is a rest point

if and only if $\kappa \geq 0$ and e if and only if $\kappa \leq 0$; when $\kappa \neq 0$ exactly one is, and a DM holding the other switches to it in the first period and stays. □

Proof of Proposition 11. Let $d = u(h) - u(h')$ and $\mu = Q^*_{\mathcal{F},h}$. Since the two situations share their frontier, $Q^*_{\mathcal{F},h'}(p) \propto \mu(p)e^{\alpha d(p)}$. That h is chosen when held gives $\sum_p \mu(p)d(p) \geq 0$; that h' is chosen when held gives $\sum_p \mu(p)e^{\alpha d(p)}d(p) \leq 0$. Subtracting, $\sum_p \mu(p)d(p)\big(e^{\alpha d(p)} - 1\big) \leq 0$. Every term is nonnegative, and vanishes only if $d(p) = 0$. □

Proof of Proposition 12. Fix a Markov strategy σ and the self who holds x, and write $v(y) = \sum_p Q_x(p)u(y^p)$, so that $G_x = v(x') - v(x)$. For $z \in \mathcal{F}$, let $B(z)$ be the value, for this self, of the stream that starts at a date at which z is held and continues according to σ:

$$B(z) = v\big(\sigma(z)\big) - c\,\mathbf{1}\{\sigma(z) \neq z\} + \delta B\big(\sigma(z)\big).$$

Keeping x is worth $v(x)+\delta B(x)$ and switching is worth $v(x')-c+\delta B(x')$, so the gain from switching is $\Gamma_x = G_x - c + \delta\big[B(x') - B(x)\big]$. There are four Markov strategies. If σ keeps both technologies, $B(z) = v(z)/(1-\delta)$ and $\Gamma_x = G_x/(1-\delta) - c$. If σ always chooses x, then $B(x) = v(x) + \delta B(x)$ and $B(x') = v(x) - c + \delta B(x) = B(x) - c$, so $\Gamma_x = G_x - (1+\delta)c$. If σ always chooses x', symmetrically $B(x) = B(x') - c$ and $\Gamma_x = G_x - (1-\delta)c$. If σ is the cycle, $B(x) = v(x') - c + \delta B(x')$ and $B(x') = v(x) - c + \delta B(x)$, so $(1+\delta)\big[B(x') - B(x)\big] = -G_x$ and $\Gamma_x = G_x/(1+\delta) - c$. The strategy σ is an equilibrium if and only if, for each x, $\sigma(x) = x'$ when $\Gamma_x > 0$ and $\sigma(x) = x$ when $\Gamma_x \leq 0$. (i) If $c = 0$, Γ_x is a positive multiple of G_x under every strategy, so whatever σ is, the best reply of each self is the myopic choice. This proves the first two claims. The myopic choice is that of Proposition 10: the self holding h keeps it if and only if $G_h \leq 0$, that is, $\kappa \geq \alpha D$, and the self holding e keeps it if and only if $\kappa \leq -\alpha L$. If $-\alpha L < \kappa < \alpha D$, both switch, and the equilibrium is the cycle. (ii) Read the four cases with $c > 0$. Keeping both is an equilibrium if and only if $\Gamma_x = G_x/(1-\delta) - c \leq 0$ for both x. Always choosing x is an equilibrium if and only if the self holding x keeps it, $\Gamma_x = G_x - (1+\delta)c \leq 0$, and the self holding x' switches; for her the strategy always chooses the other technology, so $\Gamma_{x'} = G_{x'} - (1-\delta)c > 0$. The cycle is an equilibrium if and only if $\Gamma_x = G_x/(1+\delta) - c > 0$ for both x. □

A.9 Proofs of Section 6

Proof of Lemma 6.1. Let $D = D(m)$. Under I, $u(h^I_m) = u_0 \geq u_0 - L = u(e^I)$, and under T, $u(e^T) = u_0 + \theta \geq u_0 - \lambda m = u(h^T_m)$. From the status quo h_m, the stakes are therefore 0 under I and D under T, and the odds of T against I are $r\,e^{\alpha D}$. The firm adopts if and only if $V(e) > V(h_m)$, that is, $Q^*(T)D > Q^*(I)L$, that is, $r\,D\,e^{\alpha D} > L$. If $D = 0$, then $L^*(m) = 0$ and no firm adopts. Firms with $L = L^*(m)$ are indifferent and keep the old technology; they have mass zero because G is continuous. □

Proof of Proposition 13. (i) $L^*(m) = rD(m)e^{\alpha D(m)}$ is continuous and nondecreasing in m, and so is $\Phi = G \circ L^*$. Let $\mathcal{U} = \{m \in [0,1] : \Phi(m) \geq m\}$, which is closed and contains 0, and let $\overline{m} = \max \mathcal{U}$. Then $\Phi(\overline{m}) \geq \overline{m}$, and by monotonicity $\Phi(\Phi(\overline{m})) \geq \Phi(\overline{m})$, so $\Phi(\overline{m}) \in \mathcal{U}$ and $\Phi(\overline{m}) \leq \overline{m}$. Hence $\overline{m}$ is a fixed point, and it is the greatest since every fixed point lies in $\mathcal{U}$. Symmetrically, $\underline{m} = \min\{m : \Phi(m) \leq m\}$ is the least fixed point. (ii) Since $D(m) \geq 0$, $L^*(m)$ is nondecreasing in α, r and $D(m)$, and $D(m)$ is nondecreasing in θ and λ. Hence Φ is pointwise nondecreasing in $\alpha, \theta, \lambda, r$. If $\Phi' \geq \Phi$ pointwise, then $\{m : \Phi'(m) \geq m\} \supseteq \{m : \Phi(m) \geq m\}$, so the greatest fixed point of Φ' is at least that of Φ; and $\{m : \Phi'(m) \leq m\} \subseteq \{m : \Phi(m) \leq m\}$, so the least fixed point of Φ' is at least that of Φ. A Bayesian market is the case $\alpha = 0$. (iii) From

$m_0 = 0$, $m_1 = \Phi(0) \geq m_0$, and by induction the sequence is nondecreasing; it is bounded by $\underline{m}$, since $m_t \leq \underline{m}$ implies $m_{t+1} = \Phi(m_t) \leq \Phi(\underline{m}) = \underline{m}$. It converges to some $m_\infty \leq \underline{m}$, which is a fixed point by continuity, hence equal to $\underline{m}$. The argument from $m_0 = 1$ is symmetric. □

Proof of Proposition 14. Taking logarithms, $\log L^* = \log r + \log D + \alpha D$, so $d\log L^*/d\log D = 1 + \alpha D$. By the chain rule, $dL^*/dm = (L^*/D)(1+\alpha D)D'(m) = rD'(m)e^{\alpha D}(1+\alpha D)$. For $\alpha = 0$, $L^* = rD$ and $dL^*/dm = rD'(m)$; dividing gives M. □

Proof of Proposition 15. (i) With $\alpha = 0$, $\Phi(m) = G(r(\theta + \lambda m))$ is the composition of a concave function with a nondecreasing affine one, hence concave, and so is ξ. A concave cdf with $G(0) = 0$ is positive on $(0, \infty)$: if $G(x_0) = 0$ with $x_0 > 0$, concavity gives $G(x) \leq 0$ for $x > x_0$. Since $\theta > 0$, $\xi(0) = G(r\theta) > 0$. A concave function positive at 0 has at most one zero on $[0, 1]$: if $\xi(m_1) \leq 0$ with $m_1 > 0$, then for $m > m_1$, writing $m_1 = (1-t)\cdot 0 + t\, m$ with $t \in (0,1)$, concavity gives $\xi(m_1) \geq (1-t)\xi(0) + t\,\xi(m)$, so $\xi(m) \leq \big(\xi(m_1) - (1-t)\xi(0)\big)/t < 0$. Since an equilibrium exists, it is unique.

(ii) Write $\psi(D) = De^{\alpha D}$, so that $L^*(m) = r\psi(\theta + \lambda m)$. If $L^*(0) \geq \bar{L}$, then $\Phi \equiv 1$, $\xi > 0$ on $[0, 1)$, and the only equilibrium is $m = 1$. Otherwise let $m_c = \sup\{m \in [0,1] : L^*(m) \leq \bar{L}\} \in (0, 1]$. On $[0, m_c]$, $\Phi(m) = L^*(m)/\bar{L}$, and on $[m_c, 1]$, $\Phi = 1$. Since $\psi''(D) = \alpha(2 + \alpha D)e^{\alpha D} > 0$ and $\lambda > 0$, Φ and ξ are strictly convex on $[0, m_c]$. Moreover $\xi(0) = L^*(0)/\bar{L} > 0$, and $\xi(m) = 1 - m > 0$ on $[m_c, 1)$. Suppose $L^*(1) > \bar{L}$, so that $m_c < 1$, and $\xi(\hat{m}) < 0$. Then $\hat{m} < m_c$. The strictly convex ξ is positive at 0, negative at $\hat{m}$ and positive at m_c, so it has exactly one zero $m_L \in (0, \hat{m})$ and exactly one zero $m_T \in (\hat{m}, m_c)$, and no other zero on $[0, m_c]$, since a strictly convex function has at most two zeros. On $(m_c, 1]$ its only zero is 1. Hence the equilibria are exactly $m_L < m_T < 1$, with $\xi > 0$ on $[0, m_L)$, $\xi < 0$ on (m_L, m_T) and $\xi > 0$ on $(m_T, 1)$. Adjustment: if $m_0 < m_L$, the sequence increases and is bounded by m_L, as in Proposition 13(iii). If $m_0 \in (m_L, m_T)$, then $\xi(m_0) < 0$ and the sequence decreases; it is bounded below by m_L, since $m_t \geq m_L$ implies $\Phi(m_t) \geq \Phi(m_L) = m_L$, so it converges to a fixed point in $[m_L, m_0]$, which is m_L. If $m_0 > m_T$, the sequence increases, since $\xi > 0$ on $(m_T, 1)$, and converges to a fixed point in $(m_T, 1]$, which is 1. If $L^*(1) < \bar{L}$, then $m_c = 1$ and $\xi(1) < 0$: the strictly convex ξ is positive at 0 and negative at 1, and has exactly one zero. If $\xi > 0$ on $[0, 1)$, the only equilibrium is $m = 1$. □

Proof of Proposition 16. With ψ as above, $\psi'(D) = (1 + \alpha D)e^{\alpha D}$ increases continuously from 1 to ∞, and $\bar{L}/(r\lambda) > 1$, so $\hat{D}$ exists and is unique. Using $(1 + \alpha\hat{D})e^{\alpha\hat{D}} = \bar{L}/(r\lambda)$, we get $r\psi(\hat{D})/\bar{L} = \hat{D}/\big(\lambda(1 + \alpha\hat{D})\big) = \mu^*$ and $\hat{D} - \lambda\mu^* = \alpha\hat{D}^2/(1 + \alpha\hat{D}) = \theta^*$. Fix θ and write Φ_θ, ξ_θ and $m_c(\theta)$ as in the proof of Proposition 15. Where $\Phi_\theta < 1$, $\xi_\theta'(m) = (r\lambda/\bar{L})\,\psi'(\theta + \lambda m) - 1$, so $\xi_\theta' = 0$ exactly where $\theta + \lambda m = \hat{D}$, that is, at $\mu(\theta) = (\hat{D} - \theta)/\lambda$. There $\Phi_\theta(\mu(\theta)) = r\psi(\hat{D})/\bar{L} = \mu^* < 1$, so $\mu(\theta)$, when it lies in $[0, 1]$, lies in the region where $\Phi_\theta < 1$, and $\xi_\theta(\mu(\theta)) = \mu^* - (\hat{D} - \theta)/\lambda$. (i) Let $\theta \in (\max\{0, \hat{D} - \lambda\}, \theta^*)$; the interval is nonempty because $\theta^* > 0$ and $\theta^* - (\hat{D} - \lambda) = \lambda(1 - \mu^*) > 0$. Then $\mu(\theta) \in (\mu^*, 1)$ and $\xi_\theta(\mu(\theta)) < 0$, because $\theta < \theta^*$ means $(\hat{D} - \theta)/\lambda > \mu^*$. Since $\xi_\theta(0) > 0$ when $\theta > 0$, ξ_θ has a zero in $(0, \mu(\theta))$, so $\underline{m}(\theta) < \mu(\theta)$. As $\theta \uparrow \theta^*$, $\mu(\theta) \downarrow \mu^*$. (ii) Let $\theta > \theta^*$. By assumption and monotonicity, $r(\theta + \lambda)e^{\alpha(\theta+\lambda)} > \bar{L}$, so $m_c(\theta) < 1$ and $\xi_\theta = 1 - m > 0$ on $[m_c(\theta), 1)$. On $[0, m_c(\theta)]$, ξ_θ is convex and its minimum is either at $\mu(\theta)$, if $\mu(\theta) \in [0, m_c(\theta)]$, where it equals $\mu^* - (\hat{D} - \theta)/\lambda > 0$; or, if $\mu(\theta) < 0$, at 0, where it equals $\xi_\theta(0) > 0$. The case $\mu(\theta) > m_c(\theta)$ cannot occur: $\theta > \theta^* > \hat{D} - \lambda$ gives $\mu(\theta) < 1$, and $\Phi_\theta(\mu(\theta)) = \mu^* < 1$ places $\mu(\theta)$ below $m_c(\theta)$. Hence $\xi_\theta > 0$ on $[0, 1)$ and full adoption is the unique equilibrium. (iii) $\Phi_\theta(1) = \min\{1, r(\theta + \lambda)e^{\alpha(\theta+\lambda)}/\bar{L}\} = 1$. (iv) With $\alpha = 0$ and concave G, the equilibrium is unique for every $\theta > 0$ by Proposition 15(i). Let $\theta_n \to \theta > 0$ and $m_n = m^B(\theta_n)$. Every limit point $\tilde{m}$ of (m_n) satisfies $\Phi_\theta(\tilde{m}) = \tilde{m}$, by joint continuity of $(\theta, m) \mapsto \Phi_\theta(m)$, so $\tilde{m} = m^B(\theta)$, and m^B is continuous. With uniform G,

$\Phi_\theta(m) = \min\{1, r(\theta+\lambda m)/\bar{L}\}$. If $r\theta \le \bar{L} - r\lambda$, then $m = r\theta/(\bar{L}-r\lambda) \le 1$ solves $m = r(\theta+\lambda m)/\bar{L}$; otherwise $r(\theta+\lambda) > \bar{L}$ and $\Phi_\theta(1) = 1$. □

Proof of Proposition 17. The odds are those of Lemma 6.1. (i) $Q^*(T)$ is an increasing function of $\alpha D(m^*)$, which is strictly increasing in m^* when $\alpha\lambda > 0$; at $\alpha D = 0$ it equals $r/(1+r) = \pi_T$. (ii) At the true probabilities, a firm's expected gain from adopting is $\pi_T D(m^*) - \pi_I L$, negative if and only if $L > rD(m^*)$. The firm adopts if and only if $L < L^*(m^*) = rD(m^*)e^{\alpha D(m^*)}$. □

Proof of Proposition 18. At the true probabilities, a firm with downside L that adopts at share m gets $\pi_I(u_0 - L) + \pi_T(u_0 + \theta) = u_0 + \pi_T\theta - \pi_I L$, independent of m; a firm that does not adopt gets $\pi_I u_0 + \pi_T(u_0 - \lambda m) = u_0 - \pi_T\lambda m$. Since L^* is nondecreasing, adopters at m_1 adopt at m_2. Firms that adopt in both get the same payoff. Firms that adopt in neither get $u_0 - \pi_T\lambda m_1 \ge u_0 - \pi_T\lambda m_2$, strictly if $\lambda > 0$. A firm that switches has $L \ge L^*(m_1) \ge rD(m_1)$, and its payoff at m_1 minus its payoff at m_2 is $\pi_I L - \pi_T(\theta + \lambda m_1) = \pi_I\big(L - rD(m_1)\big) \ge 0$, strictly unless $L = rD(m_1)$, a single downside. □

Proof of Corollary 2. With $\theta = 0$, $D(0) = 0$, so $L^*(0) = 0$ and $\Phi(0) = G(0) = 0$. Full adoption is an equilibrium if and only if $\Phi(1) = 1$, that is, $L^*(1) = r\lambda e^{\alpha\lambda} \ge \bar{L}$; for $\alpha = 0$ it reads $r\lambda \ge \bar{L}$. At $m = 1$ every firm adopts and gets $u_0 - \pi_I L$; at $m = 0$ no firm adopts and every firm gets u_0. □

Proof of Proposition 19. Relative to u_0, at the true probabilities, a firm with downside L that adopts at the share m gets $\pi_T(\theta + \beta m) - \pi_I L$, and one that does not adopt gets $-\pi_T\lambda m$. Since D is nondecreasing, so is L^*, and adopters at m_1 adopt at m_2. (i) follows. (ii) A firm that adopts at m_2 only gets, at m_2 minus at m_1, $\pi_T(\theta + \beta m_2) - \pi_I L + \pi_T\lambda m_1 = \pi_I\big[r(\theta + \beta m_2 + \lambda m_1) - L\big]$. (iii) Let $\lambda = 0$. Firms that adopt in neither equilibrium are indifferent, and firms that adopt in both are better off at m_2 by (i). A firm that adopts at m_2 only has $L^*(m_1) \le L < L^*(m_2)$ and, by (ii), is worse off at m_2 if and only if $L > r(\theta + \beta m_2) = rD(m_2)$. If $\alpha = 0$, $L^*(m_2) = rD(m_2)$, so no such firm exists. For the last part, let $\lambda = 0$. Firms that adopt at neither share are indifferent. A Bayesian adopter at m_B has $L < rD(m_B) \le rD(m_F) \le L^*(m_F)$, so it also adopts at m_F and is better off there by $\pi_T\beta(m_F - m_B) \ge 0$. A firm that adopts at m_F only gains $\pi_I\big[rD(m_F) - L\big]$, which is nonnegative because $L \le \sup\operatorname{supp} G \le rD(m_F)$. □

A.10 Proofs of Section 7

Proof of Lemma 7.1. (i) is Lemma 6.1. (ii) Consider a firm with downside $L > 0$ that holds e, facing the menu $\{h_m, e\}$, and let $D = D(m)$. Under I the frontier is h_m, and the stakes are $u_0 - (u_0 - L) = L$; under T the frontier is e, and the stakes are 0. The odds of T against I are $r\,e^{-\alpha L}$. The firm keeps e if and only if $Q^*(T)D \ge Q^*(I)L$, that is, $L\,e^{\alpha L} \le rD$. Since $L \mapsto Le^{\alpha L}$ is continuous and strictly increasing from 0 to ∞, this is $L \le \tilde{L}(m)$; a firm with $L = 0$ keeps e as well. If $\alpha > 0$, then $\tilde{L} > 0$ and $\tilde{L} = rD\,e^{-\alpha\tilde{L}} < rD < rD\,e^{\alpha D} = L^*$. □

Proof of Theorem 2. Fix a profile x with share m, and write $L^* = L^*(m)$ and $\tilde{L} = \tilde{L}(m)$. Since $\nu > 0$, x is stationary if and only if $\mathbf{1}\{L < L^*\}(1 - x(L)) = \mathbf{1}\{L > \tilde{L}\}\,x(L)$ for almost every L. The points L^* and $\tilde{L}$ have G-measure zero, since G is continuous, and can be ignored.

(i) If $\alpha = 0$, then $L^* = \tilde{L} = rD(m)$. For $L < L^*$ the condition reads $1 - x = 0$, and for $L > L^*$ it reads $0 = x$. So $x = \mathbf{1}\{L < L^*(m)\}$ almost everywhere, and $m = m(x) = G(L^*(m)) = \Phi(m)$. Conversely, the launch profile of any $m = \Phi(m)$ satisfies the condition. At such a profile, no holder of the old technology has $L < L^*$ and no holder of the new one has $L > \tilde{L}$, so no firm switches.

(ii) If $\alpha > 0$, then $\tilde{L} < L^*$ by Lemma 7.1. For $L < \tilde{L}$ the condition reads $1 - x = 0$; for $\tilde{L} < L < L^*$ it reads $1 - x = x$; for $L > L^*$ it reads $0 = x$. This gives the description, and the share

is $G(\tilde{L}) + \frac{1}{2}[G(L^*) - G(\tilde{L})] = \Psi(m)$. Conversely, every $m = \Psi(m)$ yields a stationary profile. The map Ψ is continuous and nondecreasing from $[0,1]$ into itself, so it has a least and a greatest fixed point, by the argument of Proposition 13(i). In a period, the revising holders of the old technology in the restless band adopt, a mass $\frac{\nu}{2}[\Phi(m) - \tilde{\Phi}(m)]$, and the revising holders of the new technology in the band abandon it, a mass $\frac{\nu}{2}[\Phi(m) - \tilde{\Phi}(m)]$; no other firm switches. Since $\tilde{\Phi} \leq \Psi \leq \Phi$ pointwise, the comparison argument of Proposition 13(ii) orders the least and the greatest fixed points of the three maps. Finally, $\Psi(m) \leq \frac{1}{2}[1 + \tilde{\Phi}(m)] \leq \frac{1}{2}[1 + \tilde{\Phi}(1)]$, since $\tilde{L}$ is nondecreasing; and $\Psi(1) = 1$ if and only if $\Phi(1) = \tilde{\Phi}(1) = 1$, that is, since $\tilde{\Phi} \leq \Phi$, if and only if $\tilde{\Phi}(1) = 1$. □

Proof of Proposition 20. (i) Write $(\mathcal{T}_\nu x)(L) = \chi\big(x(L), \iota_A(L), \iota_K(L)\big)$, with $\chi(z,a,b) = z(1 - \nu a - \nu b) + \nu a$, $\iota_A(L) = \mathbf{1}\{L < L^*(m)\}$ and $\iota_K(L) = \mathbf{1}\{L > \tilde{L}(m)\}$. The function χ is nondecreasing in a and nonincreasing in b, and it is nondecreasing in z whenever $\nu(a+b) \leq 1$. This holds if $\nu \leq 1/2$; it also holds if $\alpha = 0$, because then $L^* = \tilde{L}$ and the two indicators cannot both equal one. Let $x \leq y$ almost everywhere, so that $m(x) \leq m(y)$, and write (a_x, b_x) and (a_y, b_y) for the indicators at $m(x)$ and $m(y)$. Since L^* and $\tilde{L}$ are nondecreasing, $a_x \leq a_y$ and $b_x \geq b_y$. Hence $\chi(x(L), a_x, b_x) \leq \chi(y(L), a_x, b_x) \leq \chi(y(L), a_y, b_y)$ for almost every L.

(ii) Let $x_0 \equiv 0$ and $x_{n+1} = \mathcal{T}_\nu x_n$. Then $x_1 \geq x_0$, and by (i) and induction the sequence is nondecreasing; it converges pointwise to some x_∞, and $m_n = m(x_n)$ increases to $m_\infty = m(x_\infty)$ by monotone convergence. Since L^* and $\tilde{L}$ are continuous and nondecreasing, $L^*(m_n) \uparrow L^*(m_\infty)$ and $\tilde{L}(m_n) \uparrow \tilde{L}(m_\infty)$. Hence $\mathbf{1}\{L < L^*(m_n)\} \to \mathbf{1}\{L < L^*(m_\infty)\}$ for every L, and $\mathbf{1}\{L > \tilde{L}(m_n)\} \to \mathbf{1}\{L > \tilde{L}(m_\infty)\}$ for every $L \neq \tilde{L}(m_\infty)$. Since χ is continuous in z, $x_{n+1}(L) \to (\mathcal{T}_\nu x_\infty)(L)$ for almost every L, so x_∞ is stationary. If y is stationary, then $y \geq x_0$, so $y = \mathcal{T}_\nu^n y \geq x_n$ for every n and $y \geq x_\infty$. Every stationary share is therefore at least m_∞, which is the least stationary share.

(iii) In the launch profile, the holders of the old technology have $L \geq L^*(m^S)$ and do not adopt, and the holders of the new one have $L < L^*(m^S)$; those that leave when they revise are the ones with $L > \tilde{L}(m^S)$, a set of mass $\Phi(m^S) - \tilde{\Phi}(m^S)$. If this mass is zero, $\mathcal{T}_\nu x^S = x^S$ almost everywhere. Otherwise $\alpha > 0$, since $\Phi = \tilde{\Phi}$ when $\alpha = 0$, and therefore $\nu \leq 1/2$. Then $\mathcal{T}_\nu x^S \leq x^S$, with $m(\mathcal{T}_\nu x^S) = m^S - \nu[\Phi(m^S) - \tilde{\Phi}(m^S)]$. By (i) and induction the sequence $x_n = \mathcal{T}_\nu^n x^S$ is nonincreasing, and it converges pointwise to some x_∞ with $m(x_n) \downarrow m(x_\infty)$. Now $L^*(m_n) \downarrow L^*(m_\infty)$ and $\tilde{L}(m_n) \downarrow \tilde{L}(m_\infty)$, so $\mathbf{1}\{L < L^*(m_n)\} \to \mathbf{1}\{L < L^*(m_\infty)\}$ for every $L \neq L^*(m_\infty)$ and $\mathbf{1}\{L > \tilde{L}(m_n)\} \to \mathbf{1}\{L > \tilde{L}(m_\infty)\}$ for every L. As in (ii), x_∞ is stationary, and if y is stationary with $y \leq x^S$, then $y = \mathcal{T}_\nu^n y \leq x_n$ for every n, so $y \leq x_\infty$: x_∞ is the greatest stationary profile below x^S. To identify its share, let x_m denote the stationary profile of Theorem 2(ii) with share $m = \Psi(m)$. If $m \leq m'$, then $x_m \leq x_{m'}$: where $x_m = 1$, $L \leq \tilde{L}(m) \leq \tilde{L}(m')$; where $x_m = \frac{1}{2}$, $L > \tilde{L}(m)$ and $L < L^*(m) \leq L^*(m')$, so $x_{m'} \geq \frac{1}{2}$. Moreover $x_m \leq x^S$ if and only if $m \leq m^S$: if $m \leq m^S$, x_m vanishes above $L^*(m) \leq L^*(m^S)$, while $x^S = 1$ below $L^*(m^S)$; conversely, $x_m \leq x^S$ gives $m \leq m^S$ by integration. Since $x_\infty = x_{m_\infty}$ with $m_\infty \leq m^S$, and $x_m \leq x_\infty$ for every fixed point $m \leq m^S$ of Ψ, m_∞ is the greatest fixed point of Ψ in $[0, m^S]$.

(iv) Here $\nu \leq 1/2$. Let x be any profile. By (i), $\mathcal{T}_\nu^n 0 \leq \mathcal{T}_\nu^n x \leq \mathcal{T}_\nu^n 1$ for every n. By (ii), $\mathcal{T}_\nu^n 0$ converges to the least stationary profile. Since $\mathcal{T}_\nu 1 \leq 1$, the argument of (iii) shows that $\mathcal{T}_\nu^n 1$ decreases to the greatest stationary profile. When Ψ has a unique fixed point m^*, Theorem 2(ii) gives a unique stationary profile up to null sets, so both bounds converge to it, and $m(\mathcal{T}_\nu^n x) \to m^*$ by dominated convergence. □

Proof of Corollary 3. Write $\tilde{L}_\theta$ for the keeping threshold at the gain θ. Then $\tilde{L}_\theta(1)$ solves $\tilde{L} e^{\alpha \tilde{L}} = r(\theta + \lambda)$, so $\tilde{\Phi}_\theta(1) = 1$, that is, $\tilde{L}_\theta(1) \geq \bar{L}$, holds if and only if $\bar{L} e^{\alpha \bar{L}} \leq r(\theta + \lambda)$, that is, $\theta \geq \theta^{**}$. Let $\theta \in (\max\{0, \hat{D} - \lambda\}, \theta^*)$, an interval that is nonempty by the proof of Proposition 16(i). Suppose that $\theta \geq \theta^{**}$. Then $r(\theta + \lambda) \geq \bar{L} e^{\alpha \bar{L}} > \bar{L}$, so $r\theta/(\bar{L} - r\lambda) > 1$, and by Proposition 16(iv)

the Bayesian market has the unique equilibrium 1. By Proposition 13(ii) the least equilibrium of the FOMO market is at least that of the Bayesian market, hence equal to 1. This contradicts Proposition 16(i), which gives a least equilibrium below $(\hat{D}-\theta)/\lambda < 1$. So every θ in the interval is below θ^{**}, and $\theta^{**} \geq \theta^{*}$. Finally, if full adoption is a launch equilibrium at some $\theta \in (0, \theta^{**})$, then $\Phi_\theta(1) = 1 > \tilde{\Phi}_\theta(1)$: in the first period the share falls by $\nu[1-\tilde{\Phi}_\theta(1)]$, and with $\nu \leq 1/2$ Proposition 20(iii) applies with $m^S = 1$. The greatest fixed point of Ψ_θ is below one because $\Psi_\theta(1) = \frac{1}{2}[1+\tilde{\Phi}_\theta(1)] < 1$. □

Proof of Proposition 21. The inequalities are Lemma 7.1.

(i) With $\alpha = 0$, $\Phi(m) = G(rD(m)) \geq G(r\theta) = 1$ for every m, so the only stationary share of the Bayesian market is 1, by Theorem 2(i). The bound for the FOMO market is Theorem 2(ii).

(ii) Let $D = D(m)$. Since $r \leq 1$, $\tilde{L} < rD \leq D$. If $L^*(m) < \bar{L}$, then $\Psi(m) = (L^* + \tilde{L})/(2\bar{L})$, and, since $\tilde{L} = rDe^{-\alpha\tilde{L}}$,

$$L^* + \tilde{L} - 2rD = rD\big(e^{\alpha D} + e^{-\alpha\tilde{L}} - 2\big) \geq rD\big(e^{\alpha\tilde{L}} + e^{-\alpha\tilde{L}} - 2\big) > 0,$$

so $\Psi(m) > rD/\bar{L} = G(rD(m))$, the last equality because $rD(m) \leq rD(1) \leq \bar{L}/2$. If $L^*(m) \geq \bar{L}$, then $\Psi(m) = \frac{1}{2}[1+G(\tilde{L})] > \frac{1}{2} \geq rD(1)/\bar{L} \geq G(rD(m))$, since $\tilde{L} > 0$. The Bayesian response $G(rD(m)) = r(\theta+\lambda m)/\bar{L}$ has slope $r\lambda/\bar{L} \leq \frac{1}{2}$, so the Bayesian market has a unique equilibrium m^B, and $G(rD(m)) < m$ if and only if $m > m^B$. A stationary share m of the FOMO market satisfies $m = \Psi(m) > G(rD(m))$, hence $m > m^B$. □

Proof of Proposition 22. Let $D = D(m) > 0$. (i) A holder of the old technology compares h_m with the act "switch to e", whose utilities are those of e minus c. If $D \leq c$, the switch is not better under T and worse under I, and the firm does not adopt. Otherwise, from the status quo h_m, the stakes are 0 under I and $D-c$ under T. The odds of T against I are $r\,e^{\alpha(D-c)}$, and the firm adopts if and only if $r\,(D-c)e^{\alpha(D-c)} > L+c$, that is, $L < L^*(D-c) - c$. A holder of the new technology compares e with "switch to h_m", whose utilities are those of h_m minus c. If $L \leq c$, the switch is not better under I and worse under T, and the firm keeps e. Otherwise, from the status quo e, the stakes are $L-c$ under I and 0 under T. The odds are $r\,e^{-\alpha(L-c)}$, and the firm keeps e if and only if $r(D+c) \geq (L-c)e^{\alpha(L-c)}$, that is, $L \leq c + \tilde{L}(D+c)$. Both thresholds are continuous and nondecreasing in m; A_c is nonincreasing and K_c increasing in c.

(ii) For $c > D$, $A_c = 0$ and the band is empty. For $c \in [0, D]$ let $f(c) = 2c - L^*(D-c) + \tilde{L}(D+c)$. If $A_c > 0$, the band is empty if and only if $K_c \geq A_c$, that is, $f(c) \geq 0$. If $A_c = 0$, the band is empty, and $L^*(D-c) \leq c$ gives $f(c) \geq c + \tilde{L}(D+c) > 0$. The function f is continuous and strictly increasing; $f(0) = \tilde{L}(D) - L^*(D) < 0$ by Lemma 7.1; and $f(D) = 2D + \tilde{L}(2D) > 0$. So f has a unique zero $\bar{c} \in (0, D)$, and the band is empty if and only if $c \geq \bar{c}$. Let $c_0 = \frac{1}{2}[L^*(D) - \tilde{L}(D)]$. If $c_0 \geq D$, then $\bar{c} < c_0$. Otherwise $f(c_0) = [L^*(D) - L^*(D-c_0)] + [\tilde{L}(D+c_0) - \tilde{L}(D)] \geq 0$, so $\bar{c} \leq c_0$. □

Proof of Remark 4. Fix a stationary profile with share m, a firm with downside L that revises while holding x, and a Markov strategy σ of its future selves. Let $v(y)$ be the expected one-period payoff of y under the weights of a firm with downside L that holds x at the share m, computed from the stakes net of c, and $G_x = v(x') - v(x)$. A technology chosen now is held, and yields v, until the next revision. With $a = 1/\big(1-\delta(1-\nu)\big)$ and $b = \delta\nu/\big(1-\delta(1-\nu)\big) \in [0,1)$, choosing y is worth $a\,v(y) - c\,\mathbf{1}\{y \neq x\} + b\,B(y)$, where $B(z) = a\,v(\sigma(z)) - c\,\mathbf{1}\{\sigma(z) \neq z\} + b\,B(\sigma(z))$ is the value of a revision at which z is held. The computation in the proof of Proposition 12, with (a, b) in place of $(1, \delta)$, gives the gain from switching under the four Markov strategies: $aG_x/(1-b) - c$, $aG_x - (1+b)c$, $aG_x - (1-b)c$ and $aG_x/(1+b) - c$. With $c = 0$, each is a positive multiple

of G_x, so the best reply of a revising firm is the myopic one whatever the Markov strategy σ; hence a profile is stationary for forward-looking firms if and only if it is stationary in the sense of Theorem 2. With $c > 0$, alternating is an equilibrium if and only if $aG_x/(1+b) > c$ for both x, and $a/(1+b) = 1/(1-\delta+2\delta\nu)$. The myopic firm of Section 7.5 compares one-period payoffs, and switches if and only if $G_x > c$. □

B Foundations of Axiom 1

Axiom 1 has a representation in it. This appendix gives conditions on C that do not, the *core axioms*. For a situation $\sigma = (\mathcal{F}, h)$ and acts f, g, write $f\ R_\sigma\ g$ if some menu $\mathcal{G}$ that contains h, f and g and has the frontier of $\mathcal{F}$ satisfies $f \in C(\mathcal{G}, h)$, and $f\ P_\sigma\ g$ if moreover $g \notin C(\mathcal{G}, h)$. In words, f is revealed weakly, or strictly, preferred to g while the status quo and the frontier stay put. Both relations hold only between acts below the frontier of $\mathcal{F}$. Frontiers are computed with the index u of Axiom B.1, and Axioms B.4 and B.5 are required of every situation σ.

Axiom B.1 (Tastes)**.** There is a complete and transitive relation $\succsim$ on lotteries such that, in every situation whose menu contains only lotteries, $C(\mathcal{F}, h)$ is the set of $\succsim$-best elements of $\mathcal{F}$. Moreover, for all lotteries x, y, z: (i) if $x \succ y \succ z$, there are $a, b \in (0,1)$ such that $ax+(1-a)z \succ y \succ bx+(1-b)z$; (ii) for every $a \in (0,1]$, $x \succsim y$ if and only if $ax+(1-a)z \succsim ay+(1-a)z$; (iii) there are lotteries $x \succ y$ such that, for every $a \in (0,1)$, there are lotteries z, z' with $az+(1-a)y \succ x$ and $y \succ az'+(1-a)x$.

Axiom B.2 (Scenario consistency)**.** Let $f, g \in \mathcal{F}$ with $f^p \succsim g^p$ for every $p \in \mathcal{P}$. If $g \in C(\mathcal{F}, h)$, then $f \in C(\mathcal{F}, h)$. If moreover $f^p \succ g^p$ for some p, then $g \notin C(\mathcal{F}, h)$.

Axiom B.3 (Frontier WARP, equal frontiers)**.** Let $\mathcal{F}$ and $\mathcal{G}$ contain h and have the same frontier. If $f, g \in \mathcal{F} \cap \mathcal{G}$, $f \in C(\mathcal{F}, h)$ and $g \in C(\mathcal{G}, h)$, then $f \in C(\mathcal{G}, h)$.

Axiom B.4 (Independence below the frontier)**.** Let f, g, k be acts and $\lambda \in (0,1)$, and write $f' = \lambda f + (1-\lambda)k$ and $g' = \lambda g + (1-\lambda)k$. If $f\ R_\sigma\ g$, then not $g'\ P_\sigma\ f'$. If $f\ P_\sigma\ g$, then not $g'\ R_\sigma\ f'$.

Axiom B.5 (Continuity below the frontier)**.** If $f\ P_\sigma\ g$ and $g\ P_\sigma\ k$, there are $\lambda, \mu \in (0,1)$ such that $\lambda f + (1-\lambda)k\ P_\sigma\ g$ and $g\ P_\sigma\ \mu f + (1-\mu)k$.

When nothing depends on the scenario, the choice is a matter of taste: Axiom B.1 says this, with the unboundedness axiom of Maccheroni et al. (2006) as its part (iii). Axiom B.2 is the consistency axiom of Cerreia-Vioglio et al. (2013), read on choices. Axiom B.3 says that an option below the frontier is an irrelevant alternative; it is an axiom of minimax regret (Hayashi, 2008; Stoye, 2011). Axioms B.4 and B.5 are the independence and Archimedean conditions, on the relations that choices reveal.

Proposition 23. *Let $\mathcal{P}$ be finite and linearly independent. A choice correspondence C satisfies Axiom 1 if and only if it satisfies Axioms B.1–B.5.*

Proof. Axiom 1 implies the core axioms. Axiom B.1 holds because the values equal u on lotteries and u is affine and onto. Axiom B.2 holds because the weights have full support. Situations with the same status quo and the same frontier maximize the same function V: this gives Axiom B.3. Within such a class, $f\ R_\sigma\ g$ implies $V(f) \geq V(g)$ and $f\ P_\sigma\ g$ implies $V(f) > V(g)$; conversely, if f and g are below the frontier, $V(f) > V(g)$ and $V(f) \geq V(h)$, then $f\ P_\sigma\ g$, as the menu $\{h, f, g\} \cup S_M(b_\mathcal{F})$ shows for M large. Axiom B.4 follows from the linearity of V. For Axiom B.5, let $f\ P_\sigma\ g$ and $g\ P_\sigma\ k$. Then $V(f) > V(g) > V(k)$, and $V(g) \geq V(h)$, since g is chosen from a menu

that contains h. The acts f and k are below the frontier, and so are their mixtures. For λ near one, $V(\lambda f + (1-\lambda)k) > V(g)$, which gives $\lambda f + (1-\lambda)k \; P_\sigma \; g$. For μ near zero, adding $\mu f + (1-\mu)k$ to a menu from which g is chosen gives $g \; P_\sigma \; \mu f + (1-\mu)k$.

The core axioms imply Axiom 1. By Axiom B.1(i)–(ii) and the von Neumann–Morgenstern theorem, $\succsim$ has an affine representation u. By (iii), for $x \succ y$ and every $a \in (0,1)$ there are z, z' with $u(z) > u(y) + (u(x) - u(y))/a$ and $u(z') < u(x) - (u(x) - u(y))/a$, so that u is unbounded above and below. Its range is convex, hence $\mathbb{R}$. Fix a situation $(\mathcal{F}, h)$. Write $b = b_{\mathcal{F}}$, let $\mathcal{K}$ be the set of menus that contain h and have frontier b, and let B be the set of acts below b. The set $\mathcal{K}$ is closed under unions and under the addition of acts of B. Write R and P for R_σ and P_σ, and $C(\mathcal{G})$ for $C(\mathcal{G}, h)$.

Zero stakes. If $u(h) = b$, every act of a menu $\mathcal{G} \in \mathcal{K}$ is weakly worse than h under every scenario. By Axiom B.2, h is chosen with any chosen act, hence is chosen; the acts with its profile are chosen with it, and the others are not. These are the maximizers under any full-support weights.

Let now $u(h) \neq b$.

Step 1: order. For $\mathcal{G}_1 \subseteq \mathcal{G}$ in $\mathcal{K}$,

$$C(\mathcal{G}) \cap \mathcal{G}_1 \text{ is empty or equal to } C(\mathcal{G}_1). \tag{2}$$

Indeed, let $z \in C(\mathcal{G}) \cap \mathcal{G}_1$ and $f \in C(\mathcal{G}_1)$. Axiom B.3 applied to $(\mathcal{G}_1, \mathcal{G})$ gives $f \in C(\mathcal{G})$, and applied to $(\mathcal{G}, \mathcal{G}_1)$ it gives $z \in C(\mathcal{G}_1)$. Let $A = \{f : f \; R \; f\}$ be the set of acts chosen from some menu of $\mathcal{K}$. Then: (a) $f \; R \; g$ implies not $g \; P \; f$, by Axiom B.3. (b) R is transitive. Let $f \in C(\mathcal{G}_1)$, $g \in \mathcal{G}_1 \cap C(\mathcal{G}_2)$ and $k \in \mathcal{G}_2$, and $\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2$. If $C(\mathcal{G})$ did not meet $\mathcal{G}_1$, it would meet $\mathcal{G}_2$, and (2) would give $g \in C(\mathcal{G}) \cap \mathcal{G}_1$. So it meets $\mathcal{G}_1$, (2) gives $f \in C(\mathcal{G})$, and $f \; R \; k$. (c) R is complete on A, by the same argument: for $f \in C(\mathcal{G}_1)$ and $g \in C(\mathcal{G}_2)$, $C(\mathcal{G}_1 \cup \mathcal{G}_2)$ contains f or g. (d) For $f, g \in A$, $f \; P \; g$ if and only if $f \; R \; g$ and not $g \; R \; f$. If $f \; P \; g$, then not $g \; R \; f$ by (a). If $f \in C(\mathcal{G}_1)$, $g \in \mathcal{G}_1$ and not $g \; R \; f$, then $g \notin C(\mathcal{G}_1)$.

Step 2: A is closed upward and convex. (i) Let $f \in A$ and $f' \in B$ with $u(f') \geq u(f)$. Then $f' \in A$ and $f' \; R \; f$; if $u(f') \neq u(f)$, then $f' \; P \; f$. Indeed, let $f \in C(\mathcal{G}_1)$ and $\mathcal{G} = \mathcal{G}_1 \cup \{f'\}$. If $u(f') \neq u(f)$, Axiom B.2 gives $f \notin C(\mathcal{G})$, so $C(\mathcal{G}) \cap \mathcal{G}_1$ is empty by (2), and $C(\mathcal{G}) = \{f'\}$. If $u(f') = u(f)$, either $C(\mathcal{G}) = \{f'\}$, or $f \in C(\mathcal{G})$ by (2) and then $f' \in C(\mathcal{G})$ by Axiom B.2. (ii) Let $f, g \in A$ with $f \; R \; g$, and $m = \lambda f + (1-\lambda)g$ with $\lambda \in (0,1)$. Axiom B.4 with $k = g$ gives not $g \; P \; m$. Let $g \in C(\mathcal{G}_1)$ and $\mathcal{G} = \mathcal{G}_1 \cup \{m\}$, which is in $\mathcal{K}$ since $m \in B$. Either $C(\mathcal{G}) = \{m\}$, or $g \in C(\mathcal{G})$ by (2), and then $m \in C(\mathcal{G})$ since not $g \; P \; m$. Hence $m \in A$.

Step 3: A contains an act strictly below the frontier. No act of $\mathcal{G}_0 = \{h\} \cup S_1(b)$ has profile b. Let $a_0 \in C(\mathcal{G}_0)$, let $\bar{f}$ have profile b and k have profile $u(a_0) - \mathbf{1}$. By Step 2(i), $\bar{f} \; P \; a_0$. Also $a_0 \; P \; k$: by Axiom B.2, $k \notin C(\mathcal{G}_0 \cup \{k\})$, so $C(\mathcal{G}_0 \cup \{k\}) = C(\mathcal{G}_0)$ by (2). By Axiom B.5 there is $\lambda \in (0,1)$ with $m = \lambda \bar{f} + (1-\lambda)k \; P \; a_0$. Then $m \in A$ and $u(m) = \lambda b + (1-\lambda)(u(a_0) - \mathbf{1}) < b$ under every scenario. By Step 2(i), A contains every act with profile between $u(m)$ and b.

Step 4: representation on A. By Steps 1 and 2, R is a complete preorder on the convex set A. For $f, g, k \in A$ with $f \; P \; g$, Axiom B.4 gives not $g' \; R \; f'$, hence $f' \; P \; g'$ by (c) and (d); and Axiom B.5 is the Archimedean condition. These are the hypotheses of the expected utility theorem for mixture sets (Fishburn, 1970, Chapter 8), which goes back to Herstein and Milnor (1953): there is an affine $U : A \to \mathbb{R}$ that represents R on A. By Step 2(i), $U(f)$ depends on f only through its profile and is strictly increasing in it. So $U(f) = \bar{U}(u(f))$ for an affine and strictly increasing $\bar{U}$ on the convex set $\bar{A} = \{u(f) : f \in A\}$, which has a nonempty interior by Step 3. Hence $\bar{U}(x) = \sum_p w_p x(p) + w_0$ for a unique w, and every $w_p > 0$. Let $Q = w / \sum_p w_p$ and $V(g) = \sum_p Q(p) u(g^p)$.

Step 5: $C = \arg\max V$ on $\mathcal{K}$. Fix $a^* \in A$ with $u(a^*)$ in the interior of $\bar{A}$, and write $g_\lambda = \lambda g + (1-\lambda)a^*$. For $g \in B$ and $\lambda > 0$ small, $u(g_\lambda) \in \bar{A}$, so $g_\lambda \in A$ by Step 2(i), and $U(g_\lambda) - U(f_\lambda)$

has the sign of $V(g) - V(f)$. Let $\mathcal{G} \in \mathcal{K}$, $f \in C(\mathcal{G})$ and $g \in \mathcal{G}$. Then $f \; R \; g$, so not $g_\lambda \; P \; f_\lambda$ by Axiom B.4, so $U(f_\lambda) \geq U(g_\lambda)$ by (c) and (d), and $V(f) \geq V(g)$. Conversely, let f maximize V on $\mathcal{G}$ and $z \in C(\mathcal{G})$; then $V(z) = V(f)$. If $f \notin C(\mathcal{G})$, then $z \; P \; f$, so not $f_\lambda \; R \; z_\lambda$ by Axiom B.4, which contradicts $U(f_\lambda) = U(z_\lambda)$.

□

B.1 The axioms without the utility index

Under the core axioms, u comes from the choices among lotteries of Axiom B.1, so that all the axioms are statements about C. Their wording can also do without u. By Axiom B.1, the gain from a lottery y to a lottery x is at least the gain from y' to x' exactly when $\frac{1}{2}x + \frac{1}{2}y' \succsim \frac{1}{2}x' + \frac{1}{2}y$. For a menu $\mathcal{F}$ and a scenario q, let $b^q_{\mathcal{F}}$ be a $\succsim$-best element of $\{f^q : f \in \mathcal{F}\}$; nothing below depends on which one is chosen.

Frontiers. An act e is below the frontier of $\mathcal{F}$ if $b^q_{\mathcal{F}} \succsim e^q$ for every q, and it beats the frontier under p if $e^p \succ b^p_{\mathcal{F}}$. Two menus have the same frontier if $b^q_{\mathcal{F}} \sim b^q_{\mathcal{G}}$ for every q, and parallel frontiers if $\frac{1}{2}b^p_{\mathcal{F}} + \frac{1}{2}b^q_{\mathcal{G}} \sim \frac{1}{2}b^p_{\mathcal{G}} + \frac{1}{2}b^q_{\mathcal{F}}$ for all p, q. Each of these reads on u as in the text, by affinity. The premise of Axiom 3 is already in this form.

Stakes. For situations $\sigma = (\mathcal{F}, h)$ and $\sigma' = (\mathcal{G}, h')$, σ has at least as high stakes under q as σ' if $\frac{1}{2}\, b^q_{\mathcal{F}} + \frac{1}{2}\, h'^q \succsim \frac{1}{2}\, b^q_{\mathcal{G}} + \frac{1}{2}\, h^q$, that is, if $\rho_\sigma(q) \geq \rho_{\sigma'}(q)$.

Odds. For σ with nonzero stakes and lotteries $x \succ x'$ and y, y', write $(x, x') \approx^{p,p'}_{\sigma} (y, y')$ if some acts f, g and some $a \in (0, 1]$ satisfy the following: f and g are chosen together from a menu that contains h and has the frontier of σ; they yield equally good reductions under every scenario other than p, p'; and the gain from g^p to f^p is a times the gain from x' to x, while the change from $g^{p'}$ to $f^{p'}$ is a times the change from y to y'. In σ, a gain from x' to x if p is true exactly compensates a change from y to y' if p' is. By Axiom 1 and the test pairs of Appendix A, this holds if and only if $o_\sigma(p, p')[u(x) - u(x')] = u(y) - u(y')$.

Common reception. Let $h, h' \in \mathcal{F}$ with nonzero stakes, let e be an act, and let $a \succ a'$, b, b', c, c' and d, d' be lotteries. If

$$(a, a') \approx^{p,p'}_{\mathcal{F},h} (b, b'), \qquad (a, a') \approx^{p,p'}_{\mathcal{F}\cup\{e\},h} (c, c'), \qquad (c, c') \approx^{p,p'}_{\mathcal{F},h'} (d, d'),$$

then $(b, b') \approx^{p,p'}_{\mathcal{F}\cup\{e\},h'} (d, d')$. (The premises imply $b \succ b'$ and $c \succ c'$, so the relation is defined where it is used.) The first two premises measure the factor by which the arrival of e moves the odds of the user who holds h. The third premise and the conclusion say that it moves the odds of the user who holds h' by the same factor. To see the equivalence with Axiom 5, write o_1, o_2, o_3, o_4 for the odds of p against p' in $(\mathcal{F}, h)$, $(\mathcal{F} \cup \{e\}, h)$, $(\mathcal{F}, h')$ and $(\mathcal{F} \cup \{e\}, h')$, and $\bar{a}, \bar{b}, \bar{c}, \bar{d}$ for the differences in utility of the four pairs. The premises read $o_1\bar{a} = \bar{b}$, $o_2\bar{a} = \bar{c}$ and $o_3\bar{c} = \bar{d}$, and the conclusion reads $o_4\bar{b} = \bar{d}$. If $o_4o_1 = o_2o_3$, the conclusion follows. Conversely, since u is onto, choose lotteries with $\bar{a} = 1$, $\bar{b} = o_1$, $\bar{c} = o_2$ and $\bar{d} = o_3o_2$; the conclusion gives $o_4o_1 = o_3o_2$.

C Further results

This appendix collects results that complete Sections 6 and 7: which distributions of downsides allow tipping, the welfare cost of churn, the long run with general thresholds, and when staggered revisions rule out cycles.

C.1 When fear of missing out creates tipping

Proposition 15 shows tipping with uniform downsides. For which distributions of downsides can a market tip, and what exactly does fear of missing out add? The answer turns on one elasticity.

Fix r and α, and let the market vary over $\theta > 0$ and $\lambda \geq 0$. A Bayesian firm adopts when its downside is below $y = rD$. A pragmatic firm adopts when its downside is below

$$\tau(y) = y\, e^{\alpha y/r},$$

the pragmatic threshold of the firm whose Bayesian threshold is y. So the launch response is $\Phi(m) = F\big(r(\theta + \lambda m)\big)$ with $F = G \circ \tau$: a pragmatic market behaves like a Bayesian market whose downsides are distributed according to F. The elasticity of τ is the *FOMO elasticity*

$$\mathcal{E}(L) = 1 + \alpha D, \qquad D = \tau^{-1}(L)/r,$$

of the firm whose pragmatic threshold is L. We say that a market is *immune to tipping* if the launch equilibrium is unique for every $\theta > 0$ and $\lambda \geq 0$.

Theorem 3 (When fear of missing out creates tipping)**.** *(i) The market is immune to tipping if and only if $F(y)/y$ is nonincreasing on $(0, \infty)$. Otherwise, some market with $\theta > 0$ and $\lambda > 0$ has at least three launch equilibria.*

(ii) If $\alpha > 0$, the ratio $\tau(y)/y$ is strictly increasing.

(iii) Let G have a derivative G' that is continuous and positive on $[0, \bar{L})$, where $\bar{L} \leq \infty$ and $G(\bar{L}) = 1$ if $\bar{L} < \infty$. The market is immune to tipping if and only if $\varepsilon_G(L)\, \mathcal{E}(L) \leq 1$ for every $L \in (0, \bar{L})$, where $\varepsilon_G(L) = LG'(L)/G(L)$ is the elasticity of G.

In words, when downsides have a density, a market is immune to tipping exactly when the elasticity of adoption with respect to the stakes never exceeds one. For a Bayesian market this elasticity is that of the distribution of downsides. Fear of missing out multiplies it by $1 + \alpha D > 1$. The level effect, which raises every threshold, does not decide whether some market tips; it only decides which markets do. What creates tipping is that fear of missing out makes adoption more responsive to the stakes: when others' adoption raises the upside by one percent, the share of firms that want to adopt may rise by more than one percent, and adoption then feeds on itself.

Corollary 4 (Bayesian and pragmatic tipping)**.** *(a) A Bayesian market is immune to tipping if and only if $G(L)/L$ is nonincreasing, that is, if the average density of downsides on $[0, L]$ does not increase with L. This holds for every concave G.*

(b) Let $\alpha > 0$. With uniform downsides, the Bayesian market is immune to tipping and the pragmatic market is not. With exponential downsides of mean $\bar{\ell}$, the Bayesian market is immune, and the pragmatic market is immune if and only if $\bar{\ell} \leq r/(2\alpha)$. More generally, under the assumptions of Theorem 3(iii), if G' is differentiable at 0, the pragmatic market is not immune when $-G''(0)/G'(0) < 2\alpha/r$, and it is immune when $\varepsilon_G(L) \leq \exp(-\alpha L/r)$ for every L.

Part (a) extends Proposition 15(i) from concave distributions to distributions with a nonincreasing average density: a Bayesian market tips only when large downsides are relatively frequent. Part (b) says what fear of missing out adds. It creates tipping where Bayesian beliefs cannot, and the less likely the transformation, the less fear it takes: what matters is α/r. With exponential downsides, $r = 1/9$ and $\alpha = 1/2$, fear of missing out creates tipping unless the mean downside is at most $1/9$.

Lemma C.1 (One-dimensional responses)**.** *Let $F : [0,\infty) \to [0,1]$ be continuous and nondecreasing, with $F(0) = 0$ and $F(y) \to 1$ as $y \to \infty$. For $a > 0$ and $b \geq 0$, consider the fixed points of $m \mapsto F(a + bm)$ on $[0,1]$.*

(i) If $F(y)/y$ is nonincreasing on $(0,\infty)$, there is exactly one for every $a > 0$ and $b \geq 0$.

(ii) Otherwise, there are $a > 0$ and $b > 0$ with at least three.

Proof. A fixed point exists, since $m \mapsto F(a + bm)$ is a continuous map of $[0,1]$ into itself. For $b = 0$ it is unique. Let $b > 0$ and $H_a(y) = F(y)/(y - a)$ for $y > a$. A share $m \in (0,1]$ is a fixed point if and only if $y = a + bm$ satisfies $H_a(y) = 1/b$: then $m = (y - a)/b = F(y) \leq 1$, so every solution $y > a$ of $H_a(y) = 1/b$ gives a fixed point, and conversely. The share $m = 0$ is a fixed point if and only if $F(a) = 0$. (i) If $F(y_0) = 0$ for some $y_0 > 0$, then $F(y)/y \leq F(y_0)/y_0 = 0$ for $y \geq y_0$, which contradicts $F \to 1$. So $F > 0$ on $(0,\infty)$, $m = 0$ is not a fixed point, and for $a < y < y'$, $F(y')/(y' - a) \leq \frac{F(y)}{y} \cdot \frac{y'}{y'-a} < \frac{F(y)}{y} \cdot \frac{y}{y-a} = H_a(y)$, since $t \mapsto t/(t - a)$ is decreasing on (a,∞). So H_a is strictly decreasing, and the fixed point is unique. (ii) Let $0 < y_0 < y_1$ with $F(y_1)/y_1 > F(y_0)/y_0$. Since $H_a(y) \to F(y)/y$ as $a \downarrow 0$, there is $a \in (0, y_0)$ with $H_a(y_0) < H_a(y_1)$. Pick $c \in (H_a(y_0), H_a(y_1))$ and $b = 1/c$. Since H_a is continuous on (a,∞) and tends to 0 at infinity, the equation $H_a(y) = c$ has a solution in (y_0, y_1) and one in (y_1,∞). If $F(a) > 0$, then $H_a(y) \to \infty$ as $y \downarrow a$, and there is a third solution in (a, y_0); if $F(a) = 0$, $m = 0$ is a third fixed point. □

Proof of Theorem 3. (i) With $y = rD(m)$, $L^*(m) = y\,e^{\alpha y/r} = \tau(y)$, so $\Phi(m) = F(r\theta + r\lambda m)$. The map τ is a continuous increasing bijection of $[0,\infty)$, so $F = G \circ \tau$ satisfies the assumptions of Lemma C.1, which applies with $a = r\theta$ and $b = r\lambda$. (ii) $\tau(y)/y = e^{\alpha y/r}$. (iii) Since $\log \tau(y) = \log y + \alpha y/r$, the elasticity of τ is $y\tau'(y)/\tau(y) = 1 + \alpha y/r$, which is $\mathcal{E}(L)$ at $L = \tau(y)$. Under the assumptions, $G > 0$ on $(0,\infty)$, G is absolutely continuous, and so is F on compact intervals, with $F'(y) = G'(\tau(y))\tau'(y)$ except at the point where $\tau(y) = \bar{L}$, if $\bar{L} < \infty$. So $F(y)/y$ is nonincreasing if and only if $yF'(y) \leq F(y)$ almost everywhere, that is, $\varepsilon_G(\tau(y))\, y\tau'(y)/\tau(y) \leq 1$, that is, $\varepsilon_G(L)\,\mathcal{E}(L) \leq 1$ at $L = \tau(y)$. For $L > \bar{L}$, $G' = 0$. Since ε_G and $\mathcal{E}$ are continuous on $(0,\bar{L})$, the inequality holds almost everywhere there if and only if it holds everywhere there. □

Proof of Corollary 4. (a) For $\alpha = 0$, τ is the identity and $F = G$; Theorem 3(i) applies. A concave G with $G(0) = 0$ has a nonincreasing ratio $G(L)/L$. (b) *Uniform downsides on $[0,\bar{L}]$.* $G(L)/L$ equals $1/\bar{L}$ on $(0,\bar{L}]$ and $1/L$ beyond, so it is nonincreasing; and $F(y)/y = \tau(y)/(\bar{L}y)$ as long as $\tau(y) \leq \bar{L}$, which is strictly increasing by Theorem 3(ii). *General densities.* Let $y = \tau^{-1}(L)$. Since $\alpha > 0$, $\tau(y) > y$ for $y > 0$, so $y < L$ and $\mathcal{E}(L) = 1 + \alpha y/r \leq 1 + \alpha L/r \leq \exp(\alpha L/r)$. Moreover $y/L = e^{-\alpha y/r} \to 1$ as $L \to 0$, so $\mathcal{E}(L) = 1 + (\alpha/r)L + o(L)$. If G' is differentiable at 0, then $G(L) = G'(0)L + G''(0)L^2/2 + o(L^2)$ and $LG'(L) = G'(0)L + G''(0)L^2 + o(L^2)$, so $\varepsilon_G(L) = 1 + \frac{G''(0)}{2G'(0)}L + o(L)$ and $\varepsilon_G(L)\,\mathcal{E}(L) = 1 + \left(\frac{\alpha}{r} + \frac{G''(0)}{2G'(0)}\right)L + o(L)$, which exceeds one for small L when $-G''(0)/G'(0) < 2\alpha/r$; Theorem 3(iii) applies. If $\varepsilon_G(L) \leq \exp(-\alpha L/r)$ for every L, then $\varepsilon_G\,\mathcal{E} \leq 1$ everywhere. *Exponential downsides.* $G(L) = 1 - \exp(-L/\bar{\ell})$ is concave, so the Bayesian market is immune by (a). Here $-G''(0)/G'(0) = 1/\bar{\ell}$, so the pragmatic market is not immune if

$\bar{\ell} > r/(2\alpha)$. If $\bar{\ell} \le r/(2\alpha)$, then, with $x = L/\bar{\ell}$, $\varepsilon_G(L) = x/(e^x - 1) \le e^{-x/2} \le \exp(-\alpha L/r)$, since $x \le e^{x/2} - e^{-x/2}$ for $x \ge 0$, and the market is immune. □

Remark 6 (Tipping with memory). When $\alpha > 0$, the stationary shares are the fixed points of $\Psi(m) = F_{\text{mem}}\big(r(\theta + \lambda m)\big)$, where $F_{\text{mem}} = \frac{1}{2}\big(G \circ \tau + G \circ \tilde{\tau}\big)$, τ is the map of Section C.1, and $\tilde{\tau}(y)$ solves $\tilde{\tau}\, e^{\alpha\tilde{\tau}} = y$. By Lemma C.1, the stationary share is unique for every $\theta > 0$ and $\lambda \ge 0$ if and only if $F_{\text{mem}}(y)/y$ is nonincreasing. With uniform downsides and $r < 1$ this fails near the origin, since $\big(\tau(y) + \tilde{\tau}(y)\big)/y = 2 + \alpha(1/r - 1)\, y + o(y)$. Memory therefore does not remove tipping in general, although in Example 4 it removes it at the parameters of the example.

Proof of Remark 6. With $y = rD(m)$, $L^*(m) = \tau(y)$ and $\tilde{L}(m) = \tilde{\tau}(y)$, so $\Psi(m) = F_{\text{mem}}(r\theta + r\lambda m)$. The function F_{mem} is continuous and nondecreasing, with $F_{\text{mem}}(0) = 0$ and $F_{\text{mem}} \to 1$, and Lemma C.1 applies with $a = r\theta$ and $b = r\lambda$. With uniform downsides on $[0, \bar{L}]$ and $\tau(y) \le \bar{L}$, $F_{\text{mem}}(y)/y = \big(\tau(y) + \tilde{\tau}(y)\big)/(2\bar{L}y)$. Here $\tau(y)/y = e^{\alpha y/r} = 1 + \alpha y/r + o(y)$ and $\tilde{\tau}(y)/y = e^{-\alpha\tilde{\tau}(y)} = 1 - \alpha y + o(y)$, since $\tilde{\tau}(y)/y \to 1$. The sum is $2 + \alpha(1/r - 1)y + o(y)$, which increases near 0 when $r < 1$. □

C.2 Welfare with memory

Proposition 18 ranks launch equilibria: with competitive pressure, more adoption leaves no firm better off. With memory, a new cost appears, and the ranking acquires an exception. We evaluate expected payoffs per period at the true probabilities, relative to u_0: holding the new technology yields $\pi_T\theta - \pi_I L$, holding the old one yields $-\pi_T\lambda m$, and a restless firm, which holds each half of the time, yields the average.

Proposition 24 (Welfare with memory)**.** *Let $\alpha > 0$ and $\lambda > 0$.*

(i) At a stationary profile with share m, loyal adopters and non-adopters make the choice that a Bayesian firm would make given m, and restless firms make it half of the time. The average payoff falls short of that of Bayesian choices at the same share by

$$\frac{1}{2}\int_{(\tilde{L}(m),L^*(m))} \big|\pi_T D(m) - \pi_I L\big|\, dG(L).$$

(ii) Let $m_1 < m_2$ be stationary shares. Every firm weakly prefers m_1, except the firms with $\tilde{L}(m_1) < L < rD(m_1)$ and $L \le \tilde{L}(m_2)$, a nondegenerate interval, which strictly prefer m_2. Firms that do not adopt at m_1, and firms that are restless at both shares, strictly prefer m_1.

The cost of churn is a cost of indecision, and it is borne by the restless firms only. Part (ii) says that the rat race of Proposition 18 survives memory with one exception. Some firms would be better off with the new technology at the true probabilities, $L < rD(m_1)$, yet fear of missing out makes them restless, because each time they hold it they wonder whether they should have stayed. More pressure makes them loyal, and they gain from it. In Example 3, the shortfall is 0.0071 per firm and period at the stationary share 0.052, where the average payoff is -0.036 and that of Bayesian choices at the same share is -0.028. In the market of Remark 7 with three stationary shares, 0.018, 0.339 and 0.552, the shortfall is 0.00004, 0.30 and 0.85. Only 0.05% of the firms, those with $L \in (0.073, 0.075)$, prefer the high share to the low one, and the average payoff falls from -0.017 to -1.36.

Proof of Proposition 24. Payoffs are relative to u_0. Holding the new technology yields $\pi_T\theta - \pi_I L$ and holding the old one yields $-\pi_T\lambda m$; the first exceeds the second by $\pi_T D(m) - \pi_I L$, which is positive if and only if $L < rD(m)$, the Bayesian choice. By Lemma 7.1, $\tilde{L}(m) < rD(m) < L^*(m)$.

(i) Loyal adopters have $L \leq \tilde{L} < rD$ and hold the new technology; non-adopters have $L \geq L^* > rD$ and hold the old one; both choose as a Bayesian firm would. Half of the restless firms with a given downside hold each technology, so on average each loses $\frac{1}{2}|\pi_T D(m) - \pi_I L|$ relative to the Bayesian choice. Integrating gives the formula.

(ii) Since L^* and $\tilde{L}$ are nondecreasing, a firm's status moves, from m_1 to m_2, from non-adopter to non-adopter, restless or loyal; from restless to restless or loyal; or from loyal to loyal. Write Δ for the payoff at m_1 minus the payoff at m_2. For a non-adopter at both, $\Delta = \pi_T\lambda(m_2 - m_1) > 0$. From non-adopter to restless, $\Delta = \frac{1}{2}\pi_I(L - rD(m_1)) + \frac{1}{2}\pi_T\lambda(m_2 - m_1) > 0$, since $L \geq L^*(m_1) > rD(m_1)$. From non-adopter to loyal, $\Delta = \pi_I(L - rD(m_1)) > 0$. For a restless firm at both, $\Delta = \frac{1}{2}\pi_T\lambda(m_2 - m_1) > 0$. From restless to loyal, $\Delta = \frac{1}{2}\pi_I(L - rD(m_1))$. For a loyal firm at both, $\Delta = 0$. So $\Delta < 0$ exactly for the firms restless at m_1 and loyal at m_2 with $L < rD(m_1)$, that is, $\tilde{L}(m_1) < L < rD(m_1)$ and $L \leq \tilde{L}(m_2)$, since $rD(m_1) < L^*(m_1)$. Since $\lambda > 0$, $D(m_2) > D(m_1)$, and $\tilde{L}$ is strictly increasing in the upside, so $\tilde{L}(m_2) > \tilde{L}(m_1)$; with $\tilde{L}(m_1) < rD(m_1)$, the set is a nondegenerate interval. □

C.3 General thresholds

Theorem 2 uses only two properties of the thresholds: each is continuous and nondecreasing in the share, and the adoption threshold exceeds the keeping threshold. A cost of switching can reverse the second property, and the next lemma covers that case. Let a firm that holds the old technology adopt if and only if $L < A(m)$, and one that holds the new technology keep it if and only if $L \leq K(m)$, where A and K are continuous and nondecreasing. Write $\Phi_A = G \circ A$, $\Phi_K = G \circ K$ and $\Psi_{AK} = \frac{1}{2}(\Phi_A + \Phi_K)$.

Lemma C.2 (General thresholds). *(i) The stationary shares are the $m \in [0,1]$ such that*

$$\min\{\Phi_A(m), \Psi_{AK}(m)\} \leq m \leq \max\{\Phi_K(m), \Psi_{AK}(m)\}.$$

At a stationary profile with share m, the switching mass is $\nu[\Phi_A(m) - \Phi_K(m)]^+$.

(ii) Let $\nu \leq 1/2$, and let $m^S = \Phi_A(m^S)$ be a launch equilibrium, with launch profile $\mathbf{1}\{L < A(m^S)\}$. From it, the share of adopters is nonincreasing and converges to the greatest $m \in [0, m^S]$ such that $m = \Upsilon(m)$, where

$$\Upsilon(m) = G\big(\min\{K(m), A(m^S)\}\big) + \tfrac{1}{2}\big[\Phi_A(m) - \Phi_K(m)\big]^+.$$

When $K(m) < A(m)$, part (i) gives $m = \Psi_{AK}(m)$, as in Theorem 2: the firms in $(K(m), A(m))$ are restless. When $A(m) \leq K(m)$, the firms in $[A(m), K(m)]$ form a *loyal band*: they keep whatever they hold, and every m with $\Phi_A(m) \leq m \leq \Phi_K(m)$ is a stationary share. The market then remembers what happened to it.

Proof of Lemma C.2. (i) Fix a profile x with share m. It is stationary if and only if $\mathbf{1}\{L < A(m)\}(1 - x(L)) = \mathbf{1}\{L > K(m)\}\,x(L)$ almost everywhere. If $K(m) < A(m)$, this gives $x = 1$ on $[0, K(m)]$, $\frac{1}{2}$ on $(K(m), A(m))$ and 0 on $[A(m), \infty)$ almost everywhere, as in Theorem 2(ii), so that $m(x) = \Psi_{AK}(m)$; and since $\Phi_K \leq \Psi_{AK} \leq \Phi_A$, the two bounds of the statement both equal $\Psi_{AK}(m)$. If $A(m) \leq K(m)$, it gives $x = 1$ on $[0, A(m))$, $x = 0$ on $(K(m), \infty)$ and leaves x free in between, so that $m(x)$ ranges over $[\Phi_A(m), \Phi_K(m)]$; and since $\Phi_A \leq \Psi_{AK} \leq \Phi_K$, the bounds of the statement

are $\Phi_A(m)$ and $\Phi_K(m)$. The switching mass is computed as in Theorem 2(ii) when $K(m) < A(m)$, and is zero otherwise.

(ii) Since $\nu \le 1/2$, $\mathcal{T}_\nu$ is order-preserving, by the argument of Proposition 20(i). At the launch profile x^S, holders of the old technology have $L \ge A(m^S)$ and do not adopt, so $\mathcal{T}_\nu x^S \le x^S$; as in Proposition 20(iii), the sequence decreases to the greatest stationary profile below x^S. For $m \le m^S$, let $X_m = 1$ on $[0, \min\{K(m), A(m^S)\}]$, $X_m = \frac{1}{2}$ on $(K(m), A(m))$ if $K(m) < A(m)$, and $X_m = 0$ elsewhere. Its share is $\Upsilon(m)$. Let $y \le x^S$ be stationary with share m; then $m \le m^S$. If $K(m) < A(m)$, y is the profile of (i), which equals X_m since $K(m) < A(m) \le A(m^S)$. If $A(m) \le K(m)$, $y = 0$ above $K(m)$, and $y \le x^S$ forces $y = 0$ above $A(m^S)$, so $y \le X_m$. In both cases $y \le X_m$, and in particular $m \le \Upsilon(m)$. The profile X_m is nondecreasing in m. Let $m \le m' \le m^S$. Where $X_m = 1$, $L \le \min\{K(m'), A(m^S)\}$, so $X_{m'} = 1$. Where $X_m = \frac{1}{2}$, $K(m) < L < A(m)$; if $X_{m'}(L) < 1$, then $L > K(m')$, since $L < A(m) \le A(m^S)$, and $L < A(m) \le A(m')$, so $X_{m'}(L) = \frac{1}{2}$. Hence Υ is nondecreasing. It is continuous, with $\Upsilon(0) \ge 0$ and $\Upsilon(m^S) \le G(A(m^S)) = m^S$, so by the argument of Proposition 13(i) it has a greatest fixed point $\bar{m}$ in $[0, m^S]$, and every $m \le m^S$ with $m \le \Upsilon(m)$ satisfies $m \le \bar{m}$. The profile $X_{\bar{m}}$ is stationary: its share is $\bar{m}$; if $K(\bar{m}) < A(\bar{m})$ it is the profile of (i); and if $A(\bar{m}) \le K(\bar{m})$ it equals 1 on $[0, A(\bar{m}))$, since $A(\bar{m}) \le \min\{K(\bar{m}), A(m^S)\}$, and 0 above $K(\bar{m})$. It lies below x^S, which equals 1 on $[0, A(m^S))$. Every stationary $y \le x^S$, with share m, satisfies $y \le X_m \le X_{\bar{m}}$, since $m \le \bar{m}$. So $X_{\bar{m}}$ is the greatest stationary profile below x^S, and its share is $\bar{m}$. □

C.4 Staggered revisions

Write $\zeta(\nu) = \max\{1 - \nu, 2\nu - 1\}$, and let Lip denote the smallest Lipschitz constant on $[0, 1]$.

Proposition 25 (Staggered revisions). *Let $\alpha > 0$ and $\nu \in (0, 1)$.*

(i) If $\zeta(\nu) + \nu\big[\mathrm{Lip}(\Phi) + \mathrm{Lip}(\tilde{\Phi})\big] < 1$, then $\mathcal{T}_\nu$ is a contraction of the set of profiles, with the distance of $L^1(G)$. The stationary profile is unique, and the share of adopters converges to it geometrically from every initial profile.

(ii) Let m^* be a stationary share at which Φ and $\tilde{\Phi}$ are differentiable. If $\zeta(\nu) + \nu[\Phi'(m^*) + \tilde{\Phi}'(m^*)] < 1$, the stationary profile is locally stable: every profile close enough to it in $L^1(G)$ converges to it geometrically.

The condition says that the launch response and the keeping response are not too steep. For uniform downsides and $D(m) = \theta + \lambda m$, $\Phi'(m) = (r\lambda/\bar{L})\, M(m)$ wherever $L^*(m) < \bar{L}$, where M is the FOMO multiplier of Proposition 14. Writing $\mathcal{L} = \mathrm{Lip}(\Phi) + \mathrm{Lip}(\tilde{\Phi})$, the condition holds if and only if $\mathcal{L} < 1$ and $\nu < 2/(2 + \mathcal{L})$; for $\mathcal{L} > 0$ it fails near full synchronization, where cycles can exist. With $\lambda = 1$ instead of 6 in Example 3, the condition of part (i) holds for every $\nu < 0.93$. In Example 3 itself, $\mathrm{Lip}(\Phi) \approx 4.1$ and part (i) does not apply, but $\Phi' + \tilde{\Phi}' \approx 0.51$ at the stationary share, so part (ii) gives local stability for every $\nu < 0.79$. Beyond these bounds, we have only computations. We simulated 300 random markets with uniform downsides and a linear upside, with revision probabilities between 0.6 and 0.99 and two initial profiles, for a total of 3,000 runs of 3,000 periods. We found no sustained cycle. The bounds of Proposition 25 are crude because they ignore a sign. The cycles we found with synchronized revisions have period two, and for these the sign settles the matter.

Proposition 26 (No two-period cycles). *Let $\nu \in (0, 1)$, and let the thresholds be any continuous nondecreasing functions A and K. Every profile x with $\mathcal{T}_\nu^2 x = x$ is stationary.*

The proof is short. Suppose the share were higher in even periods than in odd ones. Then every firm would face weakly higher thresholds in even periods, and a direct computation shows that it would then hold the new technology at least as much in odd periods, which contradicts the higher share in even periods. With $\nu = 1$ the argument fails, because a restless firm that revises at every date simply alternates, and two-period cycles exist (Remark 5). So cycles of period two, the only ones we found, require full synchronization. Longer cycles are not excluded by this argument; we have neither excluded them in general nor found one.

Proof of Proposition 25. (i) Let x, y be profiles with shares m_x, m_y. Since $\alpha > 0$, every firm is, at each share, in one of three regions: it keeps or adopts ($L \leq \tilde{L}$), it is restless ($\tilde{L} < L < L^*$), or it stays out ($L \geq L^*$). Call a downside L *in a layer* if it lies between $\tilde{L}(m_x)$ and $\tilde{L}(m_y)$ or between $L^*(m_x)$ and $L^*(m_y)$. Outside the layers, L is in the same region for both profiles, and $|(\mathcal{T}_\nu x)(L) - (\mathcal{T}_\nu y)(L)|$ equals $(1-\nu)|x(L) - y(L)|$ in the first and third regions and $|1-2\nu|\,|x(L) - y(L)|$ in the second; in all cases it is at most $\zeta(\nu)|x(L) - y(L)|$. In a layer, writing $\mathcal{T}_\nu$ with the function χ of the proof of Proposition 20, $|(\mathcal{T}_\nu x)(L) - (\mathcal{T}_\nu y)(L)| \leq \zeta(\nu)|x(L) - y(L)| + \nu$, because χ has slope at most $\zeta(\nu)$ in z in each region and changing the region at a fixed z moves χ by at most ν. The layers have G-measure at most $|\Phi(m_x) - \Phi(m_y)| + |\tilde{\Phi}(m_x) - \tilde{\Phi}(m_y)| \leq [\mathrm{Lip}(\Phi) + \mathrm{Lip}(\tilde{\Phi})]\,\|x - y\|_1$, since $|m_x - m_y| \leq \|x - y\|_1$. Hence $\|\mathcal{T}_\nu x - \mathcal{T}_\nu y\|_1 \leq \{\zeta(\nu) + \nu[\mathrm{Lip}(\Phi) + \mathrm{Lip}(\tilde{\Phi})]\}\|x - y\|_1$, and Banach's fixed-point theorem applies on the complete space of profiles. (ii) Apply the same estimate with y the stationary profile x^*, whose share is m^*. For every $\varepsilon > 0$ there is $\delta > 0$ such that $|\Phi(m) - \Phi(m^*)| + |\tilde{\Phi}(m) - \tilde{\Phi}(m^*)| \leq (\Phi'(m^*) + \tilde{\Phi}'(m^*) + \varepsilon)|m - m^*|$ when $|m - m^*| \leq \delta$. If $\|x - x^*\|_1 \leq \delta$, then $|m_x - m^*| \leq \delta$ and $\|\mathcal{T}_\nu x - x^*\|_1 \leq \big[\zeta(\nu) + \nu(\Phi'(m^*) + \tilde{\Phi}'(m^*) + \varepsilon)\big]\,\|x - x^*\|_1$, where the factor is below one for ε small. The ball of radius δ is then mapped into itself, and iterates converge to x^* geometrically. □

Proof of Proposition 26. Let $x_0 = x$ and $x_1 = \mathcal{T}_\nu x_0$, so that $\mathcal{T}_\nu x_1 = x_0$, with shares m_0 and m_1. For a firm with downside L and a share m, let $\iota_A = \mathbf{1}\{L < A(m)\}$ and $\iota_K = \mathbf{1}\{L > K(m)\}$, so that the firm's holding moves by $\chi(z) = z + \nu[\iota_A(1-z) - \iota_K z]$. The pair (ι_A, ι_K) defines four *regions*: region 1, $(1,0)$, with $\chi_1(z) = z + \nu(1-z)$; region 2, $(1,1)$, the restless band, with $\chi_2(z) = (1-2\nu)z + \nu$; region 3, $(0,1)$, with $\chi_3(z) = (1-\nu)z$; and region 4, $(0,0)$, the loyal band, with $\chi_4(z) = z$. Then $x_1(L) = \chi_{R_0}(x_0(L))$ and $x_0(L) = \chi_{R_1}(x_1(L))$, where R_t is the region at m_t. Suppose $m_0 > m_1$. Since the thresholds are nondecreasing, ι_A is weakly larger and ι_K weakly smaller at m_0 than at m_1. Checking the four regions, either $R_0 = R_1$, or $R_0 = 1$, or $R_1 = 3$. If $R_0 = R_1 = R$, then $x_0(L)$ is a fixed point of $\chi_R \circ \chi_R$. The maps χ_1 and χ_3 have slope $1 - \nu < 1$, and χ_2 has slope $1 - 2\nu \in (-1, 1)$ because $\nu < 1$; so for $R \neq 4$, $\chi_R \circ \chi_R$ is a contraction whose fixed point is that of χ_R, and $x_1(L) = x_0(L)$; for $R = 4$, $x_1(L) = x_0(L)$ directly. If $R_0 = 1$, $x_1(L) = \chi_1(x_0(L)) \geq x_0(L)$. If $R_1 = 3$, $x_0(L) = \chi_3(x_1(L)) \leq x_1(L)$. Hence $x_1 \geq x_0$ almost everywhere, and $m_1 \geq m_0$, a contradiction. Exchanging the roles of x_0 and x_1, $m_0 < m_1$ is impossible too. So $m_0 = m_1$, every firm is in the same region at both dates, and $x_1 = x_0$. □

C.5 Temporary policies

Remark 7 (Temporary policies). Take Example 3 with $\nu = 1/2$, let the market settle, and raise θ from 2 to 6 for 30 periods, for instance by a subsidy. Adoption goes from 0.052 to 0.54 during the subsidy, and returns to 0.052 after it. The return follows from Proposition 20(iv), since the stationary share is unique (as computed). Fear of missing out thus amplifies a temporary policy and forgets it. This is not a general law: a churning market can keep a temporary effect when it has several stationary shares. With $r = 1/9$, $\theta = 1/2$, $\lambda = 10$, downsides uniform on $[0, 5]$ and $\alpha = 1/2$,

Ψ has the fixed points 0.018, 0.339 and 0.552; raising θ to 2 for 30 periods moves adoption from 0.018 to 0.56, and it settles at 0.552 for good, with 90% of the firms restless.

D Computations

Every computed number stated in the paper, apart from the elementary arithmetic of Section 5, is reproduced by the script `fomo_only_checks.py` (Python 3 with numpy and scipy; it runs in about a minute). It covers Example 1 and the discussions of Propositions 3 and 10; Example 3, Figure 2, Table 1 and the discussions of Propositions 14–18 and Corollary 2; the exponential case of Corollary 4 and the expansion of Remark 6; and Example 4, Figure 3 and every number of Section 7. The threshold $\tilde{L}$ is computed with the Lambert function, $\tilde{L}(m) = W(\alpha r D(m))/\alpha$. Fixed points of Φ and Ψ are located by a sign-change search on a grid of 2×10^5 shares, refined by bisection. Paths are computed by iterating $\mathcal{T}_\nu$ on a grid of 10^4 downsides, and welfare integrals on a grid of 2×10^5 downsides. Lipschitz constants are maxima of difference quotients on a grid of 2×10^4 shares. The cycle search draws 300 markets with uniform downsides and $D(m) = \theta + \lambda m$, with $r \in [0.05, 1]$, $\theta \in [0.05, 5]$, $\lambda \in [0, 15]$, $\bar{L} \in [1, 20]$ and $\alpha \in [0.1, 2]$ (seed 2026), and runs each for 3,000 periods from the arrival profile and from full adoption, with $\nu \in \{0.6, 0.8, 0.9, 0.95, 0.99\}$, on a grid of 400 downsides. A run counts as a sustained cycle if the range of the share over its last 200 periods exceeds 10^{-3}; the largest range found is below 10^{-14}.

A second script, `D6_checks.py`, draws random situations and checks Proposition 1, the reversal of Section 4.4, Propositions 6 and 11, the necessity of Axioms 2, 3 and 4 and of one pattern of Axiom B.4, and the examples of Remark 3. Of 20,000 random menus, 30% have no rest point and none has two.

These checks concern the stated parameters only; the results are proved above. The claim that staggering restores convergence beyond the bounds of Propositions 25 and 26 rests on these computations alone.

References

Anscombe, F. J., and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1), 199–205.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99(394), 116–131.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797–817.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5), 961–981.

Berk, R. H. (1966). Limiting behavior of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics*, 37(1), 51–58.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *Quarterly Journal of Economics*, 127(3), 1243–1285.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5), 803–843.

Brunnermeier, M. K., and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4), 1092–1118.

Bursztyn, L., Handel, B., Jiménez-Durán, R., and Roth, C. (2025). When product markets become collective traps: The case of social media. *American Economic Review*, 115(12), 4105–4136.

Bushong, B., Rabin, M., and Schwartzstein, J. (2021). A model of relative thinking. *Review of Economic Studies*, 88(1), 162–191.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. (2013). Classical subjective expected utility. *Proceedings of the National Academy of Sciences*, 110(17), 6754–6759.

Clifford, W. K. (1877). The ethics of belief. *Contemporary Review*, 29, 289–309.

Engelmann, J. B., Lebreton, M., Salem-Garcia, N. A., Schwardmann, P., and van der Weele, J. J. (2024). Anticipatory anxiety and wishful thinking. *American Economic Review*, 114(4), 926–960.

Esponda, I., and Pouzo, D. (2016). Berk–Nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(3), 1093–1130.

Fantl, J., and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.

Farrell, J., and Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16(1), 70–83.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*, 129(4), 1661–1710.

Gilboa, I., and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 141–153.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1420–1443.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4), 501–522.

Guimaraes, B., Machado, C., and Pereira, A. E. (2025). Dynamic network effects with fear of missing out. Documento de Trabajo 582, Instituto de Economía, Pontificia Universidad Católica de Chile.

Hansen, L. P., and Sargent, T. J. (2008). *Robustness*. Princeton University Press.

Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1), 242–268.

Herstein, I. N., and Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21(2), 291–297.

James, W. (1896). The will to believe. *The New World*, 5, 327–347.

Katz, M. L., and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3), 424–440.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849–1892.

Kőszegi, B., and Szeidl, A. (2013). A model of focusing in economic choice. *Quarterly Journal of Economics*, 128(1), 53–104.

Kovach, M. (2020). Twisting the truth: Foundations of wishful thinking. *Theoretical Economics*, 15(3), 989–1022.

Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement, Vol. I.* Academic Press.

Lanzani, G. (2025). Dynamic concern for misspecification. *Econometrica*, 93(4), 1333–1370.

Loomes, G., and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92(368), 805–824.

Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6), 1447–1498.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6), 1022–1100.

Masatlioglu, Y., and Ok, E. A. (2005). Rational choice with status quo bias. *Journal of Economic Theory*, 121(1), 1–29.

Masatlioglu, Y., and Ok, E. A. (2014). A canonical model of choice with initial endowments. *Review of Economic Studies*, 81(2), 851–883.

Matějka, F., and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1), 272–298.

Mayraz, G. (2011a). Priors and desires. CEP Discussion Paper 1047, London School of Economics.

Mayraz, G. (2011b). Wishful thinking. CEP Discussion Paper 1092, London School of Economics.

Milgrom, P., and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6), 1255–1277.

Milgrom, P., and Roberts, J. (1994). Comparing equilibria. *American Economic Review*, 84(3), 441–459.

Novák, V., Matveenko, A., and Ravaioli, S. (2024). The status quo and belief polarization of inattentive agents: Theory and experiment. *American Economic Journal: Microeconomics*, 16(4), 1–39.

Ortoleva, P. (2010). Status quo bias, multiple priors and uncertainty aversion. *Games and Economic Behavior*, 69(2), 411–424.

Peirce, C. S. (1878). How to make our ideas clear. *Popular Science Monthly*, 12, 286–302. Page references are to the reprint in N. Houser and C. Kloesel (eds.), *The Essential Peirce: Selected Philosophical Writings, Volume 1 (1867–1893)*, Indiana University Press, 1992, pp. 124–141.

Piermont, E. (2017). Context dependent beliefs. *Journal of Mathematical Economics*, 71, 63–73.

Przybylski, A. K., Murayama, K., DeHaan, C. R., and Gladwell, V. (2013). Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior*, 29(4), 1841–1848.

Ramsey, F. P. (1926). Truth and probability. In R. B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*, Routledge and Kegan Paul, 1931, pp. 156–198.

Riella, G., and Teper, R. (2014). Probabilistic dominance and status quo bias. *Games and Economic Behavior*, 87, 288–304.

Rogers, E. M. (1962). *Diffusion of Innovations*. Free Press.

Samuelson, W., and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1(1), 7–59.

Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46(253), 55–67.

Schaal, E., and Taschereau-Dumouchel, M. (2023). Herding through booms and busts. *Journal of Economic Theory*, 210, 105669.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690.

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.

Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6), 2226–2251.

Strzalecki, T. (2011). Axiomatic foundations of multiplier preferences. *Econometrica*, 79(1), 47–73.

Suryanarayanan, R. (2006). A model of anticipated regret and endogenous beliefs. CSEF Working Paper 161, University of Naples (revised 2008).

Wakker, P. P. (1989). *Additive Representations of Preferences*. Kluwer.