

COMPARATIVE STATICS WITH ADJUSTMENT COSTS AND THE LE CHATELIER PRINCIPLE

EDDIE DEKEL

Department of Economics, Northwestern University and Department of Economics, Tel Aviv University

JOHN K.-H. QUAH

Department of Economics, National University of Singapore

LUDVIG SINANDER

Department of Economics and Nuffield College, University of Oxford

We develop a theory of monotone comparative statics for models with adjustment costs. We show that comparative-statics conclusions may be drawn under the usual ordinal complementarity assumptions on the objective function, assuming very little about costs: only a mild monotonicity condition is required. We use this insight to prove a general Le Chatelier principle: under the ordinal complementarity assumptions, if short-run adjustment is subject to a monotone cost, then the long-run response to a shock is greater than the short-run response. We extend these results to a fully dynamic model of adjustment over time: the Le Chatelier principle remains valid, and under slightly stronger assumptions, optimal adjustment follows a monotone path. We apply our results to models of saving, production, pricing, labor supply, and investment.

KEYWORDS: Adjustment costs, comparative statics, Le Chatelier.

1. INTRODUCTION

ADJUSTMENT COSTS PLAY A MAJOR ROLE in explaining a wide range of economic phenomena. Examples include the investment behavior of firms,¹ price stickiness,² trade in illiquid financial markets,³ aggregate consumption dynamics,⁴ labor supply,⁵ labor demand,⁶ and housing consumption and asset pricing.⁷

In this paper, we develop a theory of monotone comparative statics with adjustment costs. Our fundamental insight is that very little needs to be assumed about the cost function: comparative statics requires only that *not* adjusting be cheaper than adjusting, plus

Eddie Dekel: eddiedekel@gmail.com

John K.-H. Quah: ecsqkhj@nus.edu.sg

Ludvig Sinander: ludvig.sinander@economics.ox.ac.uk

We are grateful for comments from Gregorio Curello, Christian Ewerhardt, Alkis Georgiadis-Harris, Alex Kohlhas, Bart Lipman, Hamish Low, Konstantin Milbradt, Paul Milgrom, Sara Neff, Marek Pycia, Simon Quinn, Ronny Razin, Karthik Sastry, Jakub Steiner, Quitzé Valenzuela-Stookey, Nathaniel Ver Steeg, Jan Žemlička, Mu Zhang, three anonymous referees, and audiences at Berlin, Bonn, Bristol, Caltech, CERGE-EI, Carlos III Madrid, Collegio Carlo Alberto, Columbia, Edinburgh, LSE, Manchester, National University of Singapore, Oxford, Paris–Saclay, Purdue, Queen Mary, Singapore Management University, Stony Brook, Tsinghua, UCLA, Western, Yale, Zürich, and several conferences. Dekel acknowledges financial support from National Science Foundation Grant 1919494 and from the Foerder Institute for Economic Research at Tel Aviv University.

¹For example, Jorgenson (1963), Hayashi (1982), Cooper and Haltiwanger (2006).

²For example, Mankiw (1985), Caplin and Spulber (1987), Golosov and Lucas (2007), Midrigan (2011).

³For example, Kyle (1985), Back (1992).

⁴For example, Kaplan and Violante (2014), Berger and Vavra (2015), Chetty and Szeidl (2016).

⁵For example, Chetty, Friedman, Olsen, and Pistaferri (2011), Chetty (2012).

⁶For example, Hamermesh (1988), Bentolila and Bertola (1990).

⁷Grossman and Laroque (1990).

the usual ordinal complementarity assumptions on the objective function. We use this insight to show that Samuelson's (1947) *Le Chatelier principle* is far more general than previously claimed: it holds whenever adjustment is costly, given only minimal structure on costs. We extend our comparative-statics and Le Chatelier results to a fully dynamic model of adjustment.

We apply our results to models of factor demand, capital investment, pricing, labor supply, and saving by wishful thinkers. These models are typically studied only under strong functional-form assumptions, and the cases of convex and nonconvex costs are considered separately and handled very differently. Our general results yield robust comparative statics for these standard models, dispensing with auxiliary assumptions and handling convex and nonconvex costs in a unified fashion.

The abstract setting is as follows. An agent chooses an action x from a sublattice $L \subseteq \mathbb{R}^n$. Her objective $F(x, \theta)$ depends on a parameter θ . At the initial parameter $\underline{\theta}$, the agent chose $\underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta})$. The parameter now increases to $\bar{\theta} \geq \underline{\theta}$, and the agent may adjust her choice. Adjusting the action by $\varepsilon = x - \underline{x}$ costs $C(\varepsilon) \geq 0$, and the agent's new choice maximizes $G(x, \bar{\theta}) = F(x, \bar{\theta}) - C(x - \underline{x})$.

Our only assumption on the cost function C is *monotonicity*:

$$C(\varepsilon_1, \dots, \varepsilon_{i-1}, \varepsilon'_i, \varepsilon_{i+1}, \dots, \varepsilon_n) \leq C(\varepsilon) \quad \text{whenever } 0 \leq \varepsilon'_i \leq \varepsilon_i \text{ or } 0 \geq \varepsilon'_i \geq \varepsilon_i.$$

This means that cost falls whenever an adjustment vector ε is modified by shifting one of its entries closer to zero ("no adjustment"). An additively separable cost function $C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$ is monotone if and only if each dimension's cost function C_i is single-dipped and minimized at zero.

We allow some adjustments ε to be infeasible, as captured by a prohibitive cost $C(\varepsilon) = \infty$. In some of our results, the monotonicity assumption may be weakened to *minimal monotonicity*: cost falls whenever an adjustment vector ε is modified by replacing all of its positive entries with zero ($C(\varepsilon \wedge 0) \leq C(\varepsilon)$), and similarly for the negative entries ($C(\varepsilon \vee 0) \leq C(\varepsilon)$). In the additively separable case, this means that each dimension's cost C_i is minimized at zero.

We eschew restrictive *curvature* assumptions on costs, such as convexity; with multidimensional actions, even quasiconvexity is not needed. As we show, it is monotonicity-type properties, not convexity-type properties, which govern the direction of adjustment in response to a shock. (What curvature properties govern is the *speed* of adjustment—an important but distinct question.)

Our basic question is under what assumptions on the objective F and cost C the agent's choice increases, in the sense that $\hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$ (provided the argmax is not empty; such qualifiers are omitted throughout this introduction). Our fundamental result, Theorem 1, answers this question: nothing need be assumed about the cost C except minimal monotonicity, while F need only satisfy the ordinal complementarity conditions of *quasi-supermodularity* and *single-crossing differences* that feature in similar comparative-statics results absent adjustment costs (see Milgrom and Shannon (1994)). Thus costs need not even be monotone, and the objective need not satisfy any cardinal properties, such as supermodularity or increasing differences. We also provide a generalization (Theorem 1*) allowing for shifts of both the constraint set L and of the parameter θ , and we give a "V" variant (Proposition 1) showing that adding either of two mild assumptions yields the stronger conclusion that $\hat{x} \geq \underline{x}$ for every $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$.

We use our fundamental result to re-think Samuelson's (1947) *Le Chatelier principle*, which asserts that the response to a parameter shift is greater at longer horizons. Our

Theorem 2 provides that the Le Chatelier principle holds whenever short-run adjustment is subject to a monotone adjustment cost C , long-run adjustment is frictionless, and the objective F satisfies the ordinal complementarity conditions. Formally, the theorem states that under these assumptions, given any long-run choice $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfying $\bar{x} \geq \underline{x}$, we have $\bar{x} \geq \hat{x} \geq \underline{x}$ for some optimal short-run choice $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$.⁸ This substantially generalizes Milgrom and Roberts's (1996) Le Chatelier principle, in which short-run adjustment is assumed to be impossible for some dimensions i and costless for the rest: that is, $C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$, where some dimensions i have $C_i(\varepsilon_i) = \infty$ for all $\varepsilon_i \neq 0$, and the rest have $C_i \equiv 0$. We show that our Le Chatelier principle remains valid if long-run adjustment is also costly (Proposition 3), and we identify two weak assumptions under either of which $\bar{x} \geq \hat{x} \geq \underline{x}$ holds for every $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$ (Proposition 4).

We then extend our comparative-statics and Le Chatelier theorems to a fully dynamic, forward-looking model of costly adjustment over time. The parameter θ_t evolves over time $t \in \{1, 2, 3, \dots\}$, and the adjustment cost function C_t may also vary between periods. Starting at $x_0 = \underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta})$, the agent chooses a path $(x_t)_{t=1}^\infty$ to maximize the discounted sum of her period payoffs $F(x_t, \theta_t) - C_t(x_t - x_{t-1})$. Theorem 3 validates the Le Chatelier principle: under the same assumptions (ordinal complementarity of F and monotonicity of each C_t), if $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ in every period t , then given any $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ such that $\underline{x} \leq \bar{x}$, the agent's choices satisfy $\underline{x} \leq x_t \leq \bar{x}$ along some optimal path $(x_t)_{t=1}^\infty$. If the parameter and cost are time-invariant ($\theta_t = \bar{\theta}$ and $C_t = C$ for all periods t), then a stronger Le Chatelier principle holds (Theorem 4): under additional assumptions, the agent adjusts more at longer horizons, in the sense that $\underline{x} \leq x_t \leq x_T \leq \bar{x}$ holds at any dates $t < T$ along some optimal path $(x_t)_{t=1}^\infty$.

The Le Chatelier principle remains valid if decisions are instead made by a sequence of short-lived agents (Theorem 5): under the same assumptions as in Theorem 3, if $\underline{\theta} \leq \theta_t (\leq \theta_{t+1}) \leq \bar{\theta}$ in every period t , then $\underline{x} \leq \tilde{x}_t (\leq \tilde{x}_{t+1}) \leq \bar{x}$ in every period t along some equilibrium path $(\tilde{x}_t)_{t=1}^\infty$. Thus short-lived agents adjust in the same direction as a long-lived agent would. They may do so more sluggishly, however: Theorem 6 asserts that under stronger assumptions, $\underline{x} \leq \tilde{x}_t \leq x_t \leq \bar{x}$ holds along some short-lived equilibrium path $(\tilde{x}_t)_{t=1}^\infty$ and some long-lived optimal path $(x_t)_{t=1}^\infty$.

Several of our main results admit converses, which assert that monotonicity-type assumptions on adjustment costs are *necessary* (as well as sufficient) for drawing comparative-statics conclusions. In particular, the cost assumptions in Theorems 1 and 3 are necessary as well as sufficient, while a condition slightly weaker than monotonicity is necessary and sufficient for Theorem 2.

The rest of this paper is arranged as follows. In the next section, we describe the environment. We present our fundamental comparative-statics insight (Theorem 1) in Section 3, and apply it to saving. In Section 4, we develop a general Le Chatelier principle (Theorem 2), and apply it to pricing and factor demand. In Section 5, we introduce a dynamic, forward-looking adjustment model, derive two dynamic Le Chatelier principles (Theorems 3 and 4), and apply them to pricing, labor supply, and investment. In Section 6, we derive a Le Chatelier principle for short-lived agents (Theorem 5) and compare their behavior to that of a long-lived agent (Theorem 6). We conclude in Section 7 by establishing converses of several main results. The Appendix contains definitions of some standard terms, an extension to allow for uncertain adjustment costs, and all proofs omitted from the text.

⁸Furthermore, if \bar{x} is the largest element of $\arg \max_{x \in L} F(x, \underline{x})$, then $\bar{x} \geq \hat{x}$ for any short-run choice $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$.

2. SETTING

The agent's objective is $F(x, \theta)$, where x is the choice variable and $\theta \in \Theta$ is a parameter. The choice variable x belongs to a subset L of \mathbb{R}^n . (More generally, x could be infinite-dimensional; our results apply also in that case. We focus on the finite-dimensional case merely for simplicity.)

At the initial parameter $\theta = \underline{\theta}$, an optimal choice \underline{x} was made:

$$\underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta}).$$

(Note that we allow for a multiplicity of optimal actions.) This is the agent's "starting point," and we shall consider how she responds to a change in the parameter from $\underline{\theta}$ to $\bar{\theta}$, where $\bar{\theta} \geq \underline{\theta}$, when adjustment is costly.

Adjusting from \underline{x} to x costs $C(x - \underline{x})$. The cost function C is a map $\Delta L \rightarrow [0, \infty]$, where $\Delta L = \{x - y : x, y \in L\}$. Note that we allow some adjustments $\varepsilon \in \Delta L$ to have infinite cost $C(\varepsilon) = \infty$, meaning that they are infeasible. We assume throughout that $C(0) < \infty$.

The agent adjusts her action $x \in L$ to maximize

$$G(x, \bar{\theta}) = F(x, \bar{\theta}) - C(x - \underline{x}).$$

REMARK 1: Since \underline{x} is held fixed, our assumption that cost depends only on $\varepsilon = x - \underline{x}$ is without loss of generality. In particular, if the "true" cost has the general form $\tilde{C}(x, \underline{x})$, then we interpret C as $C(\varepsilon) = \tilde{C}(\underline{x} + \varepsilon, \underline{x})$.

2.1. Order Assumptions

Throughout, \mathbb{R}^n (and thus L) is endowed with the usual "product" order \geq , so " $x \geq y$ " means " $x_i \geq y_i$ for every dimension i ." We write " $x > y$ " whenever $x \geq y$ and $x \neq y$. We assume that the choice set L is a *sublattice* of \mathbb{R}^n , meaning that for any $x, y \in L$, the following two vectors also belong to L :

$$x \wedge y = (\min\{x_1, y_1\}, \dots, \min\{x_n, y_n\})$$

and $x \vee y = (\max\{x_1, y_1\}, \dots, \max\{x_n, y_n\}).$

Examples of sublattices include $L = \mathbb{R}^n$, "boxes" $L = \{x \in \mathbb{R}^n : y \leq x \leq z\}$ for $y, z \in \mathbb{R}^n$, "grids" such as $L = \mathbb{Z}^n$ (where \mathbb{Z} denotes the integers), and half-planes $L = \{(x_1, x_2) \in \mathbb{R}^2 : \alpha x_1 + \beta x_2 \geq k\}$ for $\alpha \leq 0 \leq \beta$ and $k \in \mathbb{R}$.

The parameter θ belongs to a partially ordered set Θ . We use the symbol " \geq " also for the partial order on Θ . In applications, the parameter θ is often a vector, in which case Θ is a subset of \mathbb{R}^n and \geq is the usual "product" order.

2.2. Monotonicity Assumptions on Costs

Most of our results assume that the cost C is *monotone*, but our first theorem requires only *minimal monotonicity*. We now define these two properties.

The cost function C is *monotone* if and only if, for any adjustment vector $\varepsilon \in \Delta L$ and any dimension i ,

$$C(\varepsilon_1, \dots, \varepsilon_{i-1}, \varepsilon'_i, \varepsilon_{i+1}, \dots, \varepsilon_n) \leq C(\varepsilon) \quad \text{whenever } 0 \leq \varepsilon'_i \leq \varepsilon_i \text{ or } 0 \geq \varepsilon'_i \geq \varepsilon_i.$$

In other words, modifying an adjustment vector by shifting one dimension's adjustment toward zero always reduces cost. An equivalent definition of monotonicity is that $C(\varepsilon') \leq C(\varepsilon)$ holds whenever ε' is "between 0 and ε " in the sense that in each dimension i , we have either $0 \leq \varepsilon'_i \leq \varepsilon_i$ or $0 \geq \varepsilon'_i \geq \varepsilon_i$. Clearly, monotonicity is an *ordinal* property: if C is monotone, then so is $\phi \circ C$ for any strictly increasing map $\phi : [0, \infty] \rightarrow [0, \infty]$.

If the choice variable is one-dimensional ($L \subseteq \mathbb{R}$), then monotonicity requires precisely that C be single-dipped and minimized at zero.⁹ More generally, if C has the *additively separable* form $C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$, then it is monotone if and only if each C_i is single-dipped and minimized at zero.

EXAMPLE 1: For a one-dimensional choice variable ($L \subseteq \mathbb{R}$), the following cost functions are monotone, for any values of the parameters $k, a \in (0, \infty)$: (a) Fixed cost: $C(\varepsilon) = k$ for $\varepsilon \neq 0$ and $C(0) = 0$. (b) Quadratic cost: $C(\varepsilon) = a\varepsilon^2$. (c) Quadratic with free disposal: $C(\varepsilon) = a\varepsilon^2$ if $\varepsilon \geq 0$ and $C(\varepsilon) = 0$ otherwise. (d) Quadratic with a constraint: $C(\varepsilon) = a\varepsilon^2$ if $\varepsilon \in \mathcal{E}$ and $C(\varepsilon) = \infty$ otherwise, where the constraint set $\mathcal{E} \subseteq \mathbb{R}$ is convex and contains 0.

EXAMPLE 2: The following cost functions are monotone: (a) Additively separable: $C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$, where each C_i is of one of the types in Example 1. (b) Euclidean: $C(\varepsilon) = \sqrt{\sum_{i=1}^n \varepsilon_i^2}$. (c) Cobb–Douglas: $C(\varepsilon) = \prod_{i=1}^n |\varepsilon_i|^{a_i}$, where $a_1, \dots, a_n \in (0, \infty)$.

Monotonicity is consistent with quite general nonconvexities, and even with failures of quasiconvexity: the cost function in Example 2(c) is monotone, but is not quasiconvex unless the choice variable is one-dimensional, that is, $L \subseteq \mathbb{R}$. (Monotonicity *does* imply quasiconvexity when the choice variable is one-dimensional, since then quasiconvexity is equivalent to single-dippedness.)

A cost function C is called *minimally monotone* if and only if

$$C(\varepsilon \wedge 0) \leq C(\varepsilon) \leq C(\varepsilon \vee 0) \quad \text{for any adjustment vector } \varepsilon \in \Delta L.$$

In other words, simultaneously *canceled* all upward adjustments, by replacing all of the positive entries of an adjustment vector ε with zeroes, reduces cost; similarly, canceling all downward adjustments reduces cost. Clearly, minimal monotonicity is implied by monotonicity, and clearly it is an ordinal property.

If the choice variable is one-dimensional ($L \subseteq \mathbb{R}$), then minimal monotonicity demands exactly that C be minimized at zero. If C is additively separable, then it is minimally monotone if and only if each C_i is minimized at zero.

EXAMPLE 3: Consider the cost function in Example 1(d), with a constraint set $\mathcal{E} \subseteq \mathbb{R}$ that contains 0 but is *not* convex. For instance, $\mathcal{E} = (-\infty, 0] \cup [I, \infty)$ for some $I > 0$, as in the recent literature on the investment behavior of entrepreneurs in developing countries (see Section 5.5 below). Or $\mathcal{E} = \mathbb{Z}$ due to an integer constraint. Such a cost function is minimally monotone, but not monotone.

EXAMPLE 4: If $C(x - \underline{x}) = 0$ whenever x is a permutation of \underline{x} and $C(x - \underline{x}) > 0$ otherwise, then C is not minimally monotone (hence not monotone). This may occur if x_1, \dots, x_n are prices at a firm's various establishments, swapping menus between establishments is costless, and printing new menus is costly.

⁹Given $X \subseteq \mathbb{R}$, a function $\phi : X \rightarrow [0, \infty]$ is *single-dipped* if and only if there is an $x \in X$ such that ϕ is decreasing on $\{y \in X : y \leq x\}$ and increasing on $\{y \in X : y \geq x\}$.

2.3. Complementarity Assumptions on the Objective

We assume throughout that the objective function F satisfies the standard *ordinal complementarity conditions* of *quasi-supermodularity* and *single-crossing differences* (see [Milgrom and Shannon \(1994\)](#)), defined as follows.

The objective $F(x, \theta)$ has *single-crossing differences* in (x, θ) if and only if $F(y, \theta') - F(x, \theta') \geq (>) 0$ implies $F(y, \theta'') - F(x, \theta'') \geq (>) 0$ whenever $x \leq y$ and $\theta' \leq \theta''$. Economically, this means that a higher parameter implies a greater liking for higher actions: whenever a higher action is (strictly) preferred to a lower one, this remains true if the parameter increases. A sufficient condition is *increasing differences*, which requires that $F(y, \theta) - F(x, \theta)$ be increasing in θ whenever $x \leq y$. Related concepts, such as *log increasing differences*, are defined in the [Appendix](#).

A function $\phi : L \rightarrow \mathbb{R}$ is called *quasi-supermodular* if $\phi(x) - \phi(x \wedge y) \geq (>) 0$ implies $\phi(x \vee y) - \phi(y) \geq (>) 0$. A sufficient condition is *supermodularity*, which requires that $\phi(x) - \phi(x \wedge y) \leq \phi(x \vee y) - \phi(y)$ for any $x, y \in L$. If $L \subseteq \mathbb{R}$, then every function $\phi : L \rightarrow \mathbb{R}$ is automatically supermodular. See the [Appendix](#) for discussion and for definitions of related concepts, such as *submodularity*.

We say that $F(x, \theta)$ is *(quasi-)supermodular in x* if, for each parameter $\theta \in \Theta$, the function $F(\cdot, \theta) : L \rightarrow \mathbb{R}$ is (quasi-)supermodular. This captures complementarity between the different dimensions of the action.

REMARK 2: In applications to decision under uncertainty, the objective is typically $F(x, \theta) = \int_S f(x, s, \theta) H(ds, \theta)$, where $f(x, s, \theta)$ is the payoff of action $x \in L$ contingent on state $s \in S \subseteq \mathbb{R}^k$, and $H(\cdot, \theta)$ is the CDF from which the state is drawn. Clearly, $F(x, \theta)$ is supermodular in x if $f(x, s, \theta)$ is. If $S \subseteq \mathbb{R}$, $f(x, s, \theta)$ is independent of θ and has single-crossing differences in (x, s) , and $H(\cdot, \theta)$ increases with θ in the monotone likelihood ratio order, then $F(x, \theta)$ has single-crossing differences in (x, θ) ([Karlin and Rubin \(1956, Lemma 1\)](#)). For more conditions under which $F(x, \theta)$ is quasi-supermodular in x or has single-crossing differences in (x, θ) , see [Athey \(2002\)](#) and [Quah and Strulovici \(2012\)](#).

3. COMPARATIVE STATICS

Recall that the agent chooses $x \in L$ to maximize $G(x, \bar{\theta}) = F(x, \bar{\theta}) - C(x - \underline{x})$. Our fundamental comparative-statics result is the following.

THEOREM 1: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that the adjustment cost C is minimally monotone. If $\theta \geq \underline{\theta}$, then $\hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in L} G(x, \theta)$, provided the argmax is nonempty.*

In words, an increased parameter leads to a higher action (modulo tie-breaking). This parallels the basic comparative-statics result for costless adjustment (see [Milgrom and Shannon \(1994, Theorem 4\)](#)), one version of which states that under the same ordinal complementarity conditions on the objective F , we have $\bar{x} \geq \underline{x}$ for some $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$, provided the argmax is nonempty. [Theorem 1](#) shows that this basic result is strikingly robust to adjustment costs: the objective F need not satisfy any additional property, and the cost C need only be minimally monotone.

EXAMPLE 5: In many economic models, for example, neoclassical production (see [Sections 4.2 and 5.5](#) below), $F(x, \theta) = \phi(x) + \theta \cdot x$, where L and Θ are subsets of \mathbb{R}^n and

$\phi : L \rightarrow \mathbb{R}$ is supermodular. Then $F(x, \theta)$ is supermodular in x and has increasing differences in (x, θ) . Hence, provided the adjustment cost is minimally monotone, Theorem 1 guarantees that any coordinatewise increase of the parameter θ leads to an coordinatewise higher optimal choice of action x .

Theorem 1 does *not* follow from applying the basic comparative-statics result to the objective function $G(x, \bar{\theta})$, because its assumptions do not guarantee that $G(\cdot, \bar{\theta})$ is quasi-supermodular.¹⁰ A different argument is required.

PROOF: Let $x' \in \arg \max_{x \in L} G(x, \bar{\theta})$. We claim that $\hat{x} = \underline{x} \vee x'$ also maximizes $G(\cdot, \bar{\theta})$; obviously, $\hat{x} \geq \underline{x}$. We have $F(\underline{x}, \underline{\theta}) \geq F(\underline{x} \wedge x', \underline{\theta})$ by definition of \underline{x} . Thus $F(\underline{x} \vee x', \underline{\theta}) \geq F(x', \underline{\theta})$ by quasi-supermodularity, whence $F(\underline{x} \vee x', \bar{\theta}) \geq F(x', \bar{\theta})$ by single-crossing differences. Furthermore, by minimal monotonicity, $C(\underline{x} \vee x' - \underline{x}) = C((x' - \underline{x}) \vee 0) \leq C(x' - \underline{x})$. Thus,

$$G(\hat{x}, \bar{\theta}) = F(\underline{x} \vee x', \bar{\theta}) - C(\underline{x} \vee x' - \underline{x}) \geq F(x', \bar{\theta}) - C(x' - \underline{x}) = G(x', \bar{\theta}).$$

Since x' maximizes $G(\cdot, \bar{\theta})$ on L , it follows that \hat{x} does, too.

Q.E.D.

This proof illustrates the role played by the minimal-monotonicity assumption in delivering comparative statics. This assumption cannot be weakened: we show in Section 7 below that minimal monotonicity is necessary (as well as sufficient) for comparative statics to hold whatever the objective F .

In applications, it is often useful that Theorem 1 requires F to satisfy only the ordinal complementarity conditions, rather than the stronger *cardinal complementarity conditions* of supermodularity and increasing differences. In monopoly pricing, for example, the objective F has single-crossing differences, but not increasing differences—see Section 4.3 below.

Theorem 1 has a counterpart for parameter *decreases*: under the same assumptions, if $\bar{\theta} \leq \underline{\theta}$, then $\hat{x} \leq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$, provided the argmax is nonempty. The proof is exactly analogous.¹¹ All of the results in this paper have such counterparts for parameter decreases; we will not discuss them explicitly.

Theorem 1 also extends straightforwardly to the case in which adjustment costs are uncertain, even if the agent is risk-averse. To be precise, let all uncertainty be summarized by a random variable S , called “the state of the world.” The agent’s adjustment cost is $C_s(\cdot)$ in state $S = s$. Her ex ante payoff is

$$\tilde{G}(x, \theta) = \mathbb{E}[u(F(x, \theta) - C_s(x - \underline{x}))],$$

where u is an increasing function $\mathbb{R} \rightarrow \mathbb{R}$, whose curvature captures the agent’s risk attitude. Theorem 1 remains true verbatim, except with “ C is minimally monotone” replaced by “ C_s is minimally monotone for almost every realization s of the state S ,” and G replaced by \tilde{G} . Several subsequent results also generalize along these lines. These assertions are proved in Appendix B.

¹⁰Its second term $x \mapsto -C(x - \underline{x})$ need not be quasi-supermodular, and in any case, the sum of two quasi-supermodular functions is not quasi-supermodular in general.

¹¹The proof of Theorem 1 uses only one half of the minimal-monotonicity assumption: that $C(\varepsilon \vee 0) \leq C(\varepsilon)$ for every $\varepsilon \in \Delta L$. The proof of its parameter-decrease counterpart uses (only) the other half, namely, that $C(\varepsilon \wedge 0) \leq C(\varepsilon)$ for every $\varepsilon \in \Delta L$.

3.1. Shifts of the Constraint Set

We now generalize Theorem 1 to encompass shifts of the constraint set L as well as of the parameter θ . Write $\underline{L} \subseteq \mathbb{R}^n$ for the initial constraint set, $\underline{x} \in \arg \max_{x \in \underline{L}} F(x, \underline{\theta})$ for the agent's initial choice, and $\bar{L} \subseteq \mathbb{R}^n$ for the new constraint set. Recall that for two sets $X, Y \subseteq \mathbb{R}^n$, X is *higher in the strong set order* than Y , denoted $X \geq_{ss} Y$, if and only if, for any $x \in X$ and $y \in Y$, the vector $x \vee y$ belongs to X and the vector $x \wedge y$ belongs to Y .

THEOREM 1*: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that the adjustment cost C is minimally monotone. If $\bar{\theta} \geq \underline{\theta}$ and $\bar{L} \geq_{ss} \underline{L}$, then $\hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in \bar{L}} G(x, \bar{\theta})$, provided the argmax is nonempty.*

REMARK 3: Theorems 1 and 1* are phrased differently than the usual statement of the basic result (see Milgrom and Shannon (1994, Theorem 4)), which asserts that when $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , if $\bar{\theta} \geq \underline{\theta}$ and $\bar{L} \geq_{ss} \underline{L}$, then $\arg \max_{x \in \bar{L}} F(x, \bar{\theta}) \geq_{ss} \arg \max_{x \in \underline{L}} F(x, \underline{\theta})$. A version of Theorem 1* with this form also holds: under the same hypotheses, if $\bar{\theta} \geq \underline{\theta}$ and $\bar{L} \geq_{ss} \underline{L}$, then

$$\begin{aligned} \underline{x} \in \arg \max_{x \in \underline{L}} F(x, \underline{\theta}) \quad \text{and} \quad x' \in \arg \max_{x \in \bar{L}} G(x, \bar{\theta}) \\ \implies \underline{x} \wedge x' \in \arg \max_{x \in \underline{L}} F(x, \underline{\theta}) \quad \text{and} \quad \underline{x} \vee x' \in \arg \max_{x \in \bar{L}} G(x, \bar{\theta}). \end{aligned}$$

The latter claim (about $\underline{x} \vee x'$) is exactly what the proof of Theorem 1* shows. To see why $\underline{x} \wedge x' \in \arg \max_{x \in \underline{L}} F(x, \underline{\theta})$, suppose not; then $F(\underline{x}, \underline{\theta}) > F(\underline{x} \wedge x', \underline{\theta})$, so that replicating the steps in the proof of Theorem 1* delivers $G(\underline{x} \vee x', \bar{\theta}) > G(x', \bar{\theta})$, which contradicts the fact that x' maximizes $G(\cdot, \bar{\theta})$ on \bar{L} .

3.2. “ \forall ” Comparative Statics

We now provide a “ \forall ” counterpart to Theorem 1, giving two conditions under either of which $\hat{x} \geq \underline{x}$ holds for every optimal choice \hat{x} . The first of these conditions is *strict single-crossing differences* of the objective $F(x, \theta)$ in (x, θ) , which requires that $F(y, \theta') - F(x, \theta') \geq 0$ implies $F(y, \theta'') - F(x, \theta'') > 0$ whenever $x < y$ and $\theta' < \theta''$. The second is *strict minimal monotonicity* of the cost C , which demands that for any adjustment vector $\varepsilon \in \Delta L$,

$$C(\varepsilon \wedge 0) < C(\varepsilon) \quad \text{unless } \varepsilon \leq 0, \quad \text{and} \quad C(\varepsilon \vee 0) < C(\varepsilon) \quad \text{unless } \varepsilon \geq 0.^{12}$$

In other words, simultaneously canceling all upward adjustments, by replacing all of the strictly positive entries of an adjustment vector ε with zeroes, *strictly* reduces cost; and likewise for downward adjustments.

PROPOSITION 1: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x , and that either*

- (a) *the objective $F(x, \theta)$ has strict single-crossing differences in (x, θ) and the cost C is minimally monotone, or*

¹²Equivalently: $C(\varepsilon \wedge 0) < C(\varepsilon)$ unless $\varepsilon \wedge 0 = \varepsilon$, and $C(\varepsilon \vee 0) < C(\varepsilon)$ unless $\varepsilon \vee 0 = \varepsilon$.

(b) the objective $F(x, \theta)$ has single-crossing differences in (x, θ) and the cost C is strictly minimally monotone.
 If $\bar{\theta} > \underline{\theta}$, then $\hat{x} \geq \underline{x}$ for any $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$.¹³

Proposition 1 is the costly-adjustment analog of the standard “ \forall ” comparative-statics result (see Milgrom and Shannon (1994, Theorem 4’)), which states that given any $\underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta})$, if $F(x, \theta)$ is quasi-supermodular in x and has strict single-crossing differences in (x, θ) , then $\bar{x} \geq \underline{x}$ for any $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$. Part (a) directly extends this result to the costly-adjustment case. Part (b) shows that the “strictness” in the hypotheses required to obtain a “ \forall ” comparative-statics conclusion can come from the cost C rather than the objective F : in particular, strict minimal monotonicity ensures that even if some action $x \not\geq \underline{x}$ maximizes $F(\cdot, \bar{\theta})$, it will not be chosen on account of its cost.

3.3. Application to Saving by Wishful Thinkers

An extensive literature in psychology and economics documents the prevalence of *motivated reasoning*: believing (to some extent) what it is convenient to believe. Theoretical work has modeled this as an agent (perhaps subconsciously) *choosing* her belief, balancing convenience against the costs of inaccuracy.¹⁴

Wishful thinking is motivated reasoning driven by a desire to be optimistic. Caplin and Leahy (2019) studied the economic implications of wishful thinking, showing (among other things) that wishful thinking may suppress saving in a standard consumption–saving model. In this section, we show how Theorem 1 delivers this result without the authors’ functional-forms assumptions.

An agent starts period 1 with wealth $w > 0$. She consumes $c \in [0, w]$ and saves the rest. Savings accrue interest at rate $r > 0$. The agent’s period-2 income is uncertain, drawn from a finite set $\mathcal{Y} \subseteq \mathbb{R}_+$. Given her belief G (a CDF on \mathcal{Y}) about her period-2 income, the agent’s expected lifetime payoff is

$$U(c, G) = u_1(c) + \int_{\mathcal{Y}} u_2((1+r)(w-c) + y)G(dy),$$

where u_1 and u_2 are continuous, concave, and strictly increasing. (A natural special case is when $u_2 = \delta u_1$ for some $\delta \in (0, 1)$.)

In deciding what to believe about period-2 income, the agent contemplates a set \mathcal{G} of beliefs (CDFs on \mathcal{Y}). We compare beliefs according to optimism, formalized by first-order stochastic dominance: $G \leq_1 H$ if and only if $G(y) \geq H(y)$ for every $y \in \mathcal{Y}$. We assume that \mathcal{G} is a sublattice: if $G, H \in \mathcal{G}$, then $G \wedge_1 H$ and $G \vee_1 H$ also belong to \mathcal{G} , where $(G \wedge_1 H)(y) = \max\{G(y), H(y)\}$ and $(G \vee_1 H)(y) = \min\{G(y), H(y)\}$ for each $y \in \mathcal{Y}$. We further assume that \mathcal{G} has a most optimistic element \bar{G} (namely, $\bar{G}(y) = \inf_{G \in \mathcal{G}} G(y)$ for each $y \in \mathcal{Y}$).

A *realist* holds belief $G_0 \in \mathcal{G}$, so consumes $c_0 \in \arg \max_{c \in [0, w]} U(c, G_0)$. A *wishful thinker* chooses both what to believe and how much to consume:

$$(\hat{c}, \hat{G}) \in \arg \max_{(c, G) \in [0, w] \times \mathcal{G}} [U(c, G) - C(G - G_0)],$$

¹³A variant of Proposition 1 can be obtained by mixing the “strictness” properties (a) and (b): if $x = (y, z)$, where $F(y, z, \theta)$ has strict single-crossing differences in (y, θ) for any fixed z , and $C(y - \underline{y}, \cdot)$ is strictly minimally monotone for any fixed y , then the conclusion goes through, with essentially the same proof.

¹⁴This literature has been surveyed by Bénabou and Tirole (2016) and Bénabou (2015).

where $C : \Delta\mathcal{G} \rightarrow [0, \infty]$ is a minimally monotone cost function. An example is when $C(G - G_0)$ is the Kullback–Leibler divergence of G from G_0 ;¹⁵ this is the functional form assumed by [Caplin and Leahy \(2019\)](#).

We claim that wishful thinkers save less than realists: $\bar{c} \geq \hat{c} \geq c_0$ and $\bar{G} \geq_1 \hat{G} \geq_1 G_0$, where $\bar{c} \in \arg \max_{c \in [0, w]} U(c, \bar{G})$ is how much the agent would consume if she were to hold the most optimistic belief \bar{G} . Formally, we have the following:

PROPOSITION 2: *In the wishful-thinking application, under the stated assumptions, $\bar{c} \geq \hat{c} \geq c_0$ and $\bar{G} \geq_1 \hat{G} \geq_1 G_0$ for some*

$$(\bar{c}, \bar{G}) \in \arg \max_{(c, G) \in [0, w] \times \mathcal{G}} [U(c, G) - C(G - G_0)] \quad \text{and} \quad \bar{c} \in \arg \max_{c \in [0, w]} U(c, \bar{G}),$$

provided the argmaxes are nonempty.

The proof hinges on Theorem 1*. It is important here that Theorem 1* demands only *minimal* monotonicity, because Caplin and Leahy’s (2019) Kullback–Leibler functional form does *not* satisfy full-blown monotonicity.¹⁶

PROOF: Observe that $[0, w] \times \mathcal{G}$ is a sublattice.

CLAIM 1: *U is supermodular: $U(c, G) - U(c \wedge c', G \wedge_1 H) \leq U(c \vee c', G \vee_1 H) - U(c', H)$ for any $c, c' \in [0, w]$ and any beliefs $G, H \in \mathcal{G}$.*

The proof of Claim 1 (Appendix E) turns on the concavity of u_2 .

Assume that the argmaxes are nonempty. By Claim 1, $U(c, G)$ has increasing differences in (c, G) , that is, $c \mapsto U(c, H) - U(c, G)$ is increasing if $G \leq_1 H$. Hence, by the basic comparative-statics result (see [Milgrom and Shannon \(1994, Theorem 4\)](#)), we may choose a $\bar{c} \in \arg \max_{c \in [0, w]} U(c, \bar{G})$ such that $\bar{c} \geq c_0$. By the basic result again, if $\bar{G} \geq_1 \hat{G} \geq_1 G_0$, then $\bar{c} \geq \hat{c} \geq c_0$ for some

$$\hat{c} \in \arg \max_{c \in [0, w]} U(c, \hat{G}) = \arg \max_{c \in [0, w]} [U(c, \hat{G}) - C(\hat{G} - G_0)].$$

It remains only to show that $\bar{G} \geq_1 \hat{G} \geq_1 G_0$ for some

$$\hat{G} \in \arg \max_{G \in \mathcal{G}} [F(G) - C(G - G_0)],$$

where $F : \mathcal{G} \rightarrow \mathbb{R}$ is given by $F(G) = \max_{c \in [0, w]} U(c, G)$ for each $G \in \mathcal{G}$.

Note that F is increasing: $F(G) \leq F(H)$ whenever $G \leq_1 H$. Hence, $G_0 \in \arg \max_{G \in \mathcal{G}_0} F(G)$, where $\mathcal{G}_0 = \{G \in \mathcal{G} : G \leq_1 G_0\}$. Claim 1 implies that F is supermodular (see [Topkis \(1998, Theorem 2.7.6\)](#)). C is minimally monotone, and clearly $\mathcal{G} \geq_{ss} \mathcal{G}_0$.

¹⁵That is, $C(\varepsilon) = D(G_0 + \varepsilon, G_0)$ for every $\varepsilon \in \Delta\mathcal{G}$, where D is the Kullback–Leibler divergence. C is minimally monotone since for any G , $D(G, G_0) = D(G, G_0) + D(G_0, G_0) \geq D(G_0 \wedge_1 G, G_0) + D(G_0 \vee_1 G, G_0)$ since $D(\cdot, G_0)$ is submodular (see [Dziewulski and Quah \(2024, Example 16\)](#)), which since $D \geq 0$ implies $D(G, G_0) \geq \max\{D(G_0 \wedge_1 G, G_0), D(G_0 \vee_1 G, G_0)\}$.

¹⁶Let $\mathcal{Y} = \{1, 2, 3\}$, and consider beliefs G_0, G, H given by $(G_0(1), G_0(2), G_0(3)) = (\frac{1}{3}, \frac{2}{3}, 1)$, $(G(1), G(2), G(3)) = (\frac{1}{4}, \frac{1}{4}, 1)$ and $(H(1), H(2), H(3)) = (\frac{1}{8}, \frac{1}{4}, 1)$. Then $G_0 \leq_1 G \leq_1 H$, but $C(G - G_0) - C(H - G_0) = \frac{1}{4} \ln(\frac{3}{4}) - 2 \times \frac{1}{8} \ln(\frac{3}{8}) = \frac{1}{4} \ln 2 > 0$.

Hence, by Theorem 1*, there is a

$$\widehat{G} \in \arg \max_{G \in \mathcal{G}} [F(G) - C(G - G_0)]$$

such that $\widehat{G} \geq_1 G_0$. We have $\bar{G} \geq_1 \widehat{G}$ by definition of \bar{G} . Q.E.D.

4. THE LE CHATELIER PRINCIPLE

The Le Chatelier principle asserts that an agent will adjust less (in every dimension) if subjected to an adjustment friction. A common interpretation equates frictional adjustment with the “short run” and frictionless adjustment with the “long run,” making the principle a claim about how the action response to parameter changes varies with the horizon.

In this section, we show that the Le Chatelier principle is far more general than previously claimed: it arises whenever the friction takes the form of a monotone adjustment cost. The classic formalization, which models friction as a constraint whereby some dimensions cannot be adjusted at all, is the special case in which each dimension has an adjustment cost that is either prohibitively high or equal to zero.

We shall compare the agent’s response to a shift of the parameter from $\underline{\theta}$ to $\bar{\theta}$ in two cases: the case in which adjustment is costly, so that the agent’s choice \widehat{x} maximizes $G(\cdot, \bar{\theta})$, and the case in which adjustment is costless, so the agent chooses a frictionless optimum $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$. We call these cases “short run” and “long run,” respectively.¹⁷

Recall from Section 2.2 the definition of a *monotone* cost function C .

THEOREM 2—Le Chatelier principle: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that the adjustment cost C is monotone. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.¹⁸ Then*

- $\bar{x} \geq \widehat{x} \geq \underline{x}$ for some $\widehat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$, provided the argmax is nonempty, and
- if \bar{x} is the largest element of $\arg \max_{x \in L} F(x, \bar{\theta})$, then $\bar{x} \geq \widehat{x}$ for any $\widehat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$.

Theorem 2 nests the Le Chatelier principle of Milgrom and Roberts (1996), in which it is assumed that only some dimensions x_i of the choice variable can be adjusted in the short run, and that such adjustments are costless. This is the special case of our model in which $C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$, where some dimensions i have $C_i \equiv 0$, and the other dimensions i have $C_i(\varepsilon_i) = \infty$ for every $\varepsilon_i \neq 0$.

Like Theorem 1, Theorem 2 requires F only to satisfy ordinal complementarity properties, not cardinal ones. This greatly extends its applicability, allowing it to be used to study pricing, for example (see Section 4.3 below).

PROOF: For the first part, assume that $\arg \max_{x \in L} G(x, \bar{\theta})$ is nonempty. By Theorem 1, we may choose an $x' \in \arg \max_{x \in L} G(x, \bar{\theta})$ such that $x' \geq \underline{x}$. We claim that $\widehat{x} = \bar{x} \wedge x'$ also maximizes $G(\cdot, \bar{\theta})$; this suffices since $\bar{x} \geq \widehat{x} \geq \underline{x}$. We have $F(\bar{x} \vee x', \bar{\theta}) \leq F(\bar{x}, \bar{\theta})$ by

¹⁷Even in settings with richer dynamics, comparing adjustment between these two cases can be useful: we give an infinite-horizon example in Remark 5 below (Section 5.1).

¹⁸Such an \bar{x} must exist, provided the argmax is nonempty. This follows from the basic comparative-statics result (see Milgrom and Shannon (1994, Theorem 4)).

definition of \bar{x} , which by quasi-supermodularity implies that $F(x', \bar{\theta}) \leq F(\bar{x} \wedge x', \bar{\theta})$. Since C is monotone and $x' \geq \bar{x} \wedge x' \geq \underline{x}$, we have $C(x' - \underline{x}) \geq C(\bar{x} \wedge x' - \underline{x})$. Thus

$$G(x', \bar{\theta}) = F(x', \bar{\theta}) - C(x' - \underline{x}) \leq F(\bar{x} \wedge x', \bar{\theta}) - C(\bar{x} \wedge x' - \underline{x}) = G(\bar{x}, \bar{\theta}),$$

which since x' maximizes $G(\cdot, \bar{\theta})$ on L implies that \hat{x} does, too.

For the second part, let \bar{x} be the largest element of $\arg\max_{x \in L} F(x, \bar{\theta})$, and let $\hat{x} \in \arg\max_{x \in L} G(x, \bar{\theta})$; we will show that $\bar{x} \geq \hat{x}$. The optimality of \hat{x} implies that $G(\hat{x}, \bar{\theta}) \geq G(\bar{x} \wedge \hat{x}, \bar{\theta})$. It furthermore holds that $C(\hat{x} - \underline{x}) \geq C(\bar{x} \wedge \hat{x} - \underline{x})$, by the monotonicity of C and the fact that in each dimension i , either $\bar{x}_i \leq \hat{x}_i$ so $0 \leq (\bar{x} \wedge \hat{x} - \underline{x})_i \leq (\hat{x} - \underline{x})_i$, or $\bar{x}_i > \hat{x}_i$ in which case $(\bar{x} \wedge \hat{x} - \underline{x})_i = (\hat{x} - \underline{x})_i$. Hence, $F(\hat{x}, \bar{\theta}) \geq F(\bar{x} \wedge \hat{x}, \bar{\theta})$, which implies $F(\bar{x} \vee \hat{x}, \bar{\theta}) \geq F(\bar{x}, \bar{\theta})$ by quasi-supermodularity. Since \bar{x} is the largest maximizer of $F(\cdot, \bar{\theta})$, it follows that $\bar{x} \geq \bar{x} \vee \hat{x}$, which is to say that $\bar{x} \geq \hat{x}$. Q.E.D.

The monotonicity assumption in Theorem 2 cannot be dropped: if the cost C were merely minimally monotone, then $\bar{x} \geq \hat{x}$ would not necessarily hold.¹⁹ Monotonicity is not quite necessary, however: a somewhat weaker property is necessary and sufficient for the Le Chatelier principle to hold whatever the objective F , as we show in Section 7 below.

4.1. Extensions

Theorem 2 remains true if adjustment is costly also in the long run: that is, if in addition to the short-run cost $C_1(x_1 - \underline{x})$ of moving from the initial choice \underline{x} to her short-run choice x_1 , the agent incurs a further cost $C_2(x_2 - x_1)$ of moving from her short-run choice x_1 to her long-run choice x_2 .

PROPOSITION 3: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that the adjustment costs C_1 and C_2 are monotone. If $\bar{\theta} \geq \underline{\theta}$, then $x_2 \geq x_1 \geq \underline{x}$ for some $x_1 \in \arg\max_{x \in L} [F(x, \bar{\theta}) - C_1(x - \underline{x})]$ and $x_2 \in \arg\max_{x \in L} [F(x, \bar{\theta}) - C_2(x - x_1)]$, provided the argmaxes are nonempty.²⁰*

To interpret this result, note that when long-run adjustment is costly, it matters whether or not the agent is forward-looking when making her short-term choice x_1 , because x_1 now enters her long-run payoff $F(x_2, \bar{\theta}) - C_2(x_2 - x_1)$. Proposition 3 describes an agent who is myopic, taking no account of the long-run implications of her short-run choice x_1 . Forward-looking behavior is studied in Section 5 below. In Section 6, we revisit myopic behavior, proving a general result (Theorem 5) of which Proposition 3 is a special case.

Theorem 2 also has a “ \forall ” counterpart. Say that the cost function C is *strictly monotone* if and only if $C(\varepsilon') < C(\varepsilon)$ holds whenever $\varepsilon' \neq \varepsilon$ and ε' is “between 0 and ε ” in the sense that in each dimension i , either $0 \leq \varepsilon'_i \leq \varepsilon_i$ or $0 \geq \varepsilon'_i \geq \varepsilon_i$. Strict monotonicity implies monotonicity and strict minimal monotonicity.

PROPOSITION 4: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that the cost C is strictly monotone. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg\max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.²¹ Then $\bar{x} \geq \hat{x} \geq \underline{x}$ for any $\hat{x} \in \arg\max_{x \in L} G(x, \bar{\theta})$.*

¹⁹For example, if $L = \mathbb{R}$, $F(x, \underline{\theta}) = -x^2$, $F(x, \bar{\theta}) = -(x - 2)^2$, and $C(\varepsilon) = \infty$ if $0 < \varepsilon < 3$ and $C(\varepsilon) = 0$ otherwise, then $G(\cdot, \bar{\theta})$ is uniquely maximized by $\hat{x} = 3$, and $\bar{x} = 2 < \hat{x}$.

²⁰In fact, x_1 and x_2 may be chosen so that $\bar{x} \geq x_2 \geq x_1 \geq \underline{x}$ holds for any $\bar{x} \in \arg\max_{x \in L} F(x, \bar{\theta})$ that satisfies $\bar{x} \geq \underline{x}$.

²¹Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

4.2. Application to Factor Demand

Consider a stylized model of production, following [Milgrom and Roberts \(1996\)](#). A firm uses capital k and labor ℓ to produce output $f(k, \ell)$. Profit at real factor prices (r, w) is $F(k, \ell, -w) = f(k, \ell) - rk - w\ell$. The adjustment cost C is monotone, but otherwise unrestricted.

If the production function f is supermodular, meaning that capital and labor are complements, then profit $F(k, \ell, -w)$ is supermodular in $x = (k, \ell)$. By inspection, the profit function $F(k, \ell, -w)$ has increasing differences in $(x, \theta) = ((k, \ell), -w)$. So by Theorem 2, any drop in the wage w precipitates a short-run increase of both k and ℓ , and a further increase in the long run.

If f is instead submodular, meaning that capital and labor are substitutes in production, then we may apply Theorem 2 to the choice variable $(x_1, x_2) = (-k, \ell)$, since profit $F^\dagger(x_1, x_2, -w) = f(-x_1, x_2) + rx_1 + (-w)x_2$ is then supermodular in $x = (x_1, x_2)$ and has increasing differences in $(x, -w)$.²² The conclusion is that ℓ still increases in the short run and further increases in the long run, whereas k now *decreases*.

[Milgrom and Roberts \(1996\)](#) were the first to use the theory of monotone comparative statics to obtain such a result. They assumed that labor adjustments are costless and that capital cannot be adjusted at all in the short run: in other words, $C(\varepsilon_k, \varepsilon_\ell) = C_k(\varepsilon_k) + C_\ell(\varepsilon_\ell)$, where $C_\ell \equiv 0$ and $C_k(\varepsilon_k) = \infty$ for every $\varepsilon_k \neq 0$. Our analysis reveals that much weaker assumptions suffice. It turns out not to matter whether labor is cheap to adjust relative to capital. What matters is, rather, that short-run adjustments are costly.

4.3. Application to Pricing

The central plank of new Keynesian macroeconomic models is price stickiness, and the oldest and most important microfoundation for this property is (nonconvex) adjustment costs (e.g., [Mankiw \(1985\)](#), [Caplin and Spulber \(1987\)](#), [Golosov and Lucas \(2007\)](#), [Midrigan \(2011\)](#)). These may be real costs of updating what prices are displayed: empirically, such “menu costs” can be nonnegligible (see, e.g., [Levy, Bergen, Dutta, and Venable \(1997\)](#)). Or they may arise from consumers reacting adversely to price hikes by temporarily reducing demand (as in [Antić and Salant \(2020\)](#)).

To study pricing, we consider the simplest model, following [Milgrom and Roberts \(1990\)](#): a monopolist with constant marginal cost $c \geq 0$ faces a decreasing demand curve $D(\cdot, \eta)$ parameterized by η , thus earning a profit of $F(p, (c, -\eta)) = (p - c)D(p, \eta)$ if she prices at $p \in \mathbb{R}_+$. We assume that demand $D(p, \eta)$ is always strictly positive, and that η is an elasticity shifter: when it increases, so does the absolute elasticity of demand at every price p . Then profit $F(p, \theta) = F(p, (c, -\eta))$ has increasing differences in (p, c) and has log increasing differences in $(p, -\eta)$, so it has single-crossing differences in $(p, \theta) = (p, (c, -\eta))$. Furthermore, profit $F(p, \theta)$ is automatically quasi-supermodular in p since this choice variable is one-dimensional ($L \subseteq \mathbb{R}$).

Adjusting the price by ε incurs a cost of $C(\varepsilon) \geq 0$. We assume nothing about C except that it is minimized at zero. In many macroeconomic models, it is a pure fixed cost: $C(\varepsilon) = k > 0$ for every $\varepsilon \neq 0$. When adjustment costs arise from price-hike-averse consumers, we have $C(\varepsilon) = 0$ for $\varepsilon \leq 0$ and $C(\varepsilon) > 0$ for $\varepsilon > 0$. If consumers are inattentive to small price changes, then $C(\varepsilon) = 0$ if $\varepsilon \in [\underline{\varepsilon}, \bar{\varepsilon}]$ and $C(\varepsilon) > 0$ otherwise, where $\underline{\varepsilon} < 0 < \bar{\varepsilon}$.

By Theorem 1, the familiar comparative-statics properties of the monopoly problem are robust to the introduction of adjustment costs: it remains true that the monopolist

²²This trick is due to [Milgrom and Roberts \(1996\)](#).

raises her price whenever her marginal cost c rises and whenever demand becomes less elastic (i.e., η falls). No assumptions on the adjustment cost C are required except that it be minimized at zero.

Under the mild additional assumption that C is single-dipped, Theorem 2 yields a dynamic prediction: in response to a shock that increases her marginal cost or decreases the elasticity of demand, the monopolist initially raises her price, and then increases it further over the longer run. Thus, one-off permanent cost and demand-elasticity shocks lead, quite generally, to price increases in both the short and long run.

A key reason why we can draw such general conclusions about pricing is that Theorems 1 and 2 require F to satisfy only ordinal (not cardinal) complementarity conditions. Specifically, we used the fact that the monopolist's profit undergoes a "single-crossing differences" shift when demand becomes less elastic (i.e., when η falls). A result which assumed the cardinal property of *increasing* differences would have been inapplicable, since elasticity shifts do not generally cause profit to shift in an "increasing differences" fashion.²³

5. DYNAMIC ADJUSTMENT

The Le Chatelier principle takes a classical, "reduced-form" approach to dynamics, following Samuelson and Milgrom–Roberts. In this section and the next, we consider a fully-fledged dynamic model of adjustment. We show that the Le Chatelier principle remains valid: in the short run, the agent's choices exceed the initial choice \underline{x} and do not overshoot the new frictionless optimum \bar{x} . We furthermore show that under additional assumptions, the path of adjustment is monotone, so that the agent adjusts more over longer horizons.

In this section, we assume that the agent is long-lived and forward-looking. The alternative case in which each period t 's choice x_t is made by a short-lived agent (or equivalently, by a myopic long-lived agent) is studied in Section 6.

5.1. Setting

The agent faces an infinite-horizon decision problem in discrete time. In each period $t \in \mathbb{N} = \{1, 2, 3, \dots\}$, she takes an action $x_t \in L$, and earns a payoff of $F(x_t, \theta_t)$. Adjusting from x_{t-1} to x_t in period t costs $C_t(x_t - x_{t-1})$.

The agent's initial choice $x_0 = \underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta})$ is given, as are the parameter sequence $(\theta_t)_{t=1}^\infty$ and the sequence $(C_t)_{t=1}^\infty$ of adjustment cost functions. The simplest example is a one-off parameter shift ($\theta_t = \bar{\theta}$ for all $t \in \mathbb{N}$) with a time-invariant cost ($C_t = C$ for all $t \in \mathbb{N}$).

The agent is forward-looking, and discounts future payoffs by a factor of $\delta \in (0, 1)$. Given her period-0 choice $x_0 \in L$, the agent's payoff from a sequence $(x_t)_{t=1}^\infty$ in L is

$$\mathcal{G}((x_t)_{t=1}^\infty, x_0) = \mathcal{F}((x_t)_{t=1}^\infty) - \mathcal{C}(x_0, (x_t)_{t=1}^\infty), \quad \text{where}$$

$$\mathcal{F}((x_t)_{t=1}^\infty) = \sum_{t=1}^{\infty} \delta^{t-1} F(x_t, \theta_t) \quad \text{and} \quad \mathcal{C}(x_0, (x_t)_{t=1}^\infty) = \sum_{t=1}^{\infty} \delta^{t-1} C_t(x_t - x_{t-1}).$$

²³This applied advantage of requiring only ordinal complementarity was pointed out by Milgrom and Roberts (1990) and Milgrom and Shannon (1994) in the context of models with costless adjustment.

REMARK 4: As we describe below (Remark 4, continued), our results apply also in the finite-horizon case.

REMARK 5: In some applications, the adjustment cost in each period t is calculated relative to the *initial* choice \underline{x} , not (as assumed in this section) relative to the previous period's choice x_{t-1} . For example, the initial choice \underline{x} could be a norm, default, or reference point that was formed through custom, bargaining, or other processes. In such cases, the agent incurs a cost in each period of deviating from \underline{x} . The long-lived agent's problem is then the same in every period: choose $x \in L$ to maximize $(1 - \delta)[F(x, \bar{\theta}) - C(x - \underline{x})]$. This is formally equivalent to the “static” model studied in Sections 3 and 4 above.

A more general model would allow the prevailing norm y_t to evolve sluggishly, for example, $y_t = \lambda x_{t-1} + (1 - \lambda)\underline{x}$, where $\lambda \in [0, 1]$. (The above discussion concerns $\lambda = 1$ and $\lambda = 0$.) We do not study this more general model in the present paper, but view it as a potentially interesting avenue for future work.

5.2. Dynamic Le Chatelier Principles

The following result shows that our Le Chatelier principle (Theorem 2) remains valid when the agent can adjust over time and is forward-looking: for any new frictionless optimum $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$, the agent's “short-run” actions x_t satisfy $\underline{x} \leq x_t \leq \bar{x}$ along some optimal path $(x_t)_{t=1}^\infty$.

THEOREM 3—Dynamic Le Chatelier: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that each adjustment cost C_t is monotone. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.²⁴ If $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$, then provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq \bar{x}$ for every period $t \in \mathbb{N}$.*

The proof (Appendix H) is a direct extension of the arguments used to prove Theorems 1 and 2. The straightforwardness of this extension is perhaps surprising, since the dynamic adjustment problem is superficially quite different from the one-shot problem: the agent chooses a *sequence* of actions, and her objective $\mathcal{F}(\cdot)$ need not be quasi-supermodular (since the sum of quasi-supermodular functions is not quasi-supermodular in general).

Monotonicity cannot be weakened in Theorem 3: it is necessary as well as sufficient for the dynamic Le Chatelier principle to hold whatever the objective F and parameter sequence $(\theta_t)_{t=1}^\infty$, as we show in Section 7 below.

The next result shows that under stronger assumptions, a stronger dynamic Le Chatelier principle holds: $\underline{x} \leq x_t \leq x_T \leq \bar{x}$ for any periods $t < T$, which is to say that the agent adjusts more at longer horizons. Let us use “BCS” as shorthand for “bounded on compact sets.”

THEOREM 4—Strong dynamic Le Chatelier: *Suppose that the objective $F(x, \theta)$ is supermodular and BCS in x and has single-crossing differences in (x, θ) , and that $C_t = C$ for every period t , where the adjustment cost C is monotone and additively separable. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.²⁵ Let the parameter shift once and for all: $\theta_t = \bar{\theta}$ for every $t \in \mathbb{N}$. Then, provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq x_{t+1} \leq \bar{x}$ for every period $t \in \mathbb{N}$.*

²⁴Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

²⁵Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

The assumptions of Theorem 4 strengthen those of Theorem 3 in two directions: (1) the parameter θ_t and cost function C_t must not vary over time, and (2) the payoff F and cost $C_t = C$ must satisfy additional properties, namely, supermodularity and BCS of $F(\cdot, \theta)$ and additive separability of C . The BCS requirement is mild (continuity is a sufficient condition). The other two requirements, supermodularity and additive separability, are substantial assumptions. However, they are both automatically satisfied when the choice variable is one-dimensional ($L \subseteq \mathbb{R}$), as in our applications to pricing, labor supply, and capital investment (Sections 5.3 to 5.5 below).

The proof (Appendix I) runs as follows. Any sequence $(x_t)_{t=1}^\infty$ can be made increasing by replacing its t th entry x_t with the cumulative maximum $x_1 \vee x_2 \vee \cdots \vee x_{t-1} \vee x_t$, for each $t \in \mathbb{N}$. It suffices to show that such “monotonization” preserves optimality, since then the optimal sequence delivered by Theorem 3 may be monotonized to yield an optimal sequence with all of the desired properties. To prove that monotonization preserves optimality, it suffices (by BCS and a limit argument) to show that an optimal sequence $(x_t)_{t=1}^\infty$ which satisfies $x_1 \leq x_2 \leq \cdots \leq x_{k-1} \leq x_k$ remains optimal if its t th entry x_t is replaced by $x_{t-1} \vee x_t$ for each $t \geq k+1$. We prove this using the supermodularity of $F(\cdot, \bar{\theta})$ and the monotonicity and additive separability of C .

REMARK 4—Continued: If there is a finite horizon K , so the agent chooses a length- K sequence $(x_t)_{t=1}^K$ to maximize $\sum_{t=1}^K \delta^{t-1}[F(x_t, \bar{\theta}) - C(x_t - x_{t-1})]$, then Theorem 3 remains directly applicable (set $C_t(\varepsilon) = \infty$ for all $t > K$ and $\varepsilon \neq 0$). Theorem 4 also remains true as stated (and BCS can be dropped); this is shown in the Appendix, immediately after the proof of Theorem 4.

REMARK 6: There is a familiar way of obtaining comparative statics in dynamic problems, via the Bellman equation (see Hopenhayn and Prescott (1992)). This approach can be used to obtain an increasing optimal path $(x_t)_{t=1}^\infty$, as in Theorem 4, but only under stronger assumptions. The Bellman equation is

$$V(x) = \max_{x' \in L} [F(x', \bar{\theta}) - C(x' - x) + \delta V(x')] \quad \text{for every } x \in L.$$

Suppose we find conditions which guarantee that

$$\Phi(x) = \arg \max_{x' \in L} [F(x', \bar{\theta}) - C(x' - x) + \delta V(x')]$$

is a nonempty compact sublattice for each $x \in L$, and that the correspondence Φ is increasing in the strong set order. Then $\Phi(x)$ has a greatest element $\phi(x)$ for each $x \in L$, and ϕ is an increasing function. Define an optimal sequence $(x_t)_{t=1}^\infty$ by $x_t = \phi(x_{t-1})$ for each $t \in \mathbb{N}$. We have $x_1 = \phi(x_0) \geq x_0$ by Theorem 3, and by repeatedly applying ϕ on both sides of this inequality, we obtain $x_2 = \phi(x_1) \geq x_1$, then $x_3 = \phi(x_2) \geq x_2$, and so on; thus $(x_t)_{t=1}^\infty$ is increasing.

The usual sufficient conditions for the above argument are that $G(x', x) = F(x', \bar{\theta}) - C(x' - x)$ is supermodular in x' and has increasing differences in (x', x) . These cardinal complementarity conditions, together with ancillary assumptions, ensure that the value function V is supermodular (see Hopenhayn and Prescott (1992)), allowing us to apply the basic comparative-statics result (see Milgrom and Shannon (1994, Theorem 4)) to the objective $G(x', x) + \delta V(x')$ to conclude that Φ is increasing in the strong set order. The cardinal complementarity conditions effectively require that $-C(x' - x)$ have increasing differences (x', x) . This is a very restrictive assumption; in the additively separable case

$C(\varepsilon) = \sum_{i=1}^n C_i(\varepsilon_i)$, it demands that each C_i be convex. Theorem 4 avoids this assumption, requiring merely that the cost C be monotone.

5.3. Application to Pricing, Continued

An active literature in macroeconomics (e.g., Golosov and Lucas (2007), Midrigan (2011)) examines the price stickiness central to the new Keynesian paradigm by studying forward-looking dynamic models of pricing subject to adjustment costs (usually called “menu costs” in this context—see Section 4.3 above). The basic mechanism is that non-convexities in adjustment costs give rise to price stickiness.

Theorem 4 delivers comparative statics for such pricing models, without any of the parametric assumptions that are typically placed on adjustment costs.²⁶ Consider again the monopoly pricing problem described in Section 4.3. Assume that the adjustment cost C is monotone. It is automatically true that the cost C is additively separable and that profit $F(\cdot, \eta)$ is supermodular, because the choice variable $p \in \mathbb{R}_+$ is one-dimensional. Thus, by Theorem 4, supply shocks cause inflation at every horizon: a one-off permanent increase of marginal cost c leads prices to increase monotonically over time. The same is true of demand shocks that make the demand curve less elastic (lower η).

Theorem 3 furthermore provides that the path of prices remains always above the original frictionless monopoly price, and never overshoots the new frictionless monopoly price. This conclusion is more general, holding even when marginal cost c_t , the demand elasticity parameter η_t , and the adjustment cost function $C_t(\cdot)$ vary over time.

Although we phrased these findings in terms of a monopolist’s pricing problem, they apply equally to the typical new Keynesian setting of monopolistic competition between many firms selling differentiated goods (see, e.g., Galí (2015)). In that case, the demand curve in our analysis above is to be understood as *residual* demand, taking into account the other firms’ pricing.

5.4. Application to Labor Supply

A recent literature attempts to reconcile “micro” and “macro” estimates of labor supply elasticities by appeal to adjustment costs (see Chetty et al. (2011), Chetty (2012)). The idea is that since job design in firms is often inflexible, with no scope for big changes in hours, a worker wishing to adjust her labor supply may have to find a new job, which entails costly search.

The workhorse model of this literature features a worker choosing consumption $c \in \mathbb{R}_+$ and labor supply $x \in L \subseteq \mathbb{R}_+$. Per-period utility has the quasilinear form $c - \kappa(x)$. Given quasilinearity, it is without loss of generality to assume that the worker cannot save or borrow.²⁷ She then consumes whatever she earns, so her per-period utility equals $F(x, T) = wx - T(wx) - \kappa(x)$, where $w > 0$ is the wage and $T(\cdot)$ is the tax schedule, and we have normalized the price of consumption to 1. Labor supply is subject to adjustment costs (arising from search), which are assumed to be single-dipped. We depart from the literature by eschewing functional-form restrictions on the effort disutility κ , tax schedule T , and adjustment cost function.

²⁶Common functional forms include quadratic (Rotemberg (1982)) and pure fixed cost (many papers, e.g., Caplin and Spulber (1987), Golosov and Lucas (2007), Midrigan (2011)).

²⁷Given the standard assumption that the worker’s discount rate is equal to the interest rate, the marginal utility of a dollar is equal across all periods.

We consider tax reforms which reduce marginal rates. Formally, we write $\bar{T} \geq_{\text{flat}} T$ (“ \bar{T} is flatter than T ”) if and only if $\bar{T}(y') - \bar{T}(y) \leq T(y') - T(y)$ for all $y' \geq y \geq 0$. The per-period utility $F(x, T)$ has single-crossing differences in (x, T) ,²⁸ and is automatically supermodular in x .

Our Le Chatelier principles imply that when marginal tax rates are cut, labor supply increases at every horizon. More specifically, in the “classical” environment of Sections 2 to 4, labor supply increases in the short run and further increases in the long run (Theorem 2), while in the infinite-horizon model of Sections 5.1 and 5.2, labor supply rises monotonically over time (Theorem 4). The adjustment may be gradual or abrupt, depending on functional forms.

5.5. Application to Capital Investment

In the neoclassical theory of investment (originating with Jorgenson (1963)), a firm adjusts its capital stock over time subject to adjustment costs. In the simplest such model, the profit of a firm with capital stock $k_t \in \mathbb{R}_+$ is $F(k_t, (p, \eta, -r)) = pf(k_t, \eta) - rk_t$, where (p, r) are the prices of output and capital and $f(\cdot, \eta)$ is an increasing production function. Capital is subject to an adjustment cost: investing $i_t = k_t - k_{t-1}$ costs $C(i_t) \geq 0$, where $C(0) = 0$.

We assume that f has increasing differences, so that the parameter η shifts the marginal product of capital. Then $F(k, \theta)$ has increasing differences (and hence single-crossing differences) in (k, θ) , where $\theta = (p, \eta, -r)$. Profit $F(k, \theta)$ is automatically supermodular in k , since $k \in \mathbb{R}_+$ is one-dimensional. Our discussion below may be extended to richer variants of this model featuring, for example, depreciation and time-varying prices.

The early literature assumed a convex adjustment cost $C(\cdot)$, which yields gradual capital accumulation and an equivalence of the neoclassical theory with Tobin’s (1969) “ q ” theory of investment (see Hayashi (1982)). Later work focused on the “lumpy” investment behavior that arises when adjustment costs are nonconvex. “Lumpiness” is empirically well-documented (see Cooper and Haltiwanger (2006)), and has implications for, among other things, business cycles (e.g., Thomas (2002), Bachmann, Caballero, and Engel (2013), Winberry (2021)) and the effects of microfinance programs on entrepreneurship in developing countries (e.g., Field, Pande, Papp, and Rigol (2013), Bari, Malik, Meki, and Quinn (2024)).

Our comparative-statics theory handles both the convex case and rich forms of non-convexity. Our Le Chatelier principles (the “classical,” reduced-form Theorem 2 and the dynamic, forward-looking Theorems 3 and 4) are applicable provided merely that adjustment costs are single-dipped. Investment then increases at every horizon, and by more at longer horizons, whenever the marginal profitability of capital increases, whether due to a drop in its price r , a rise in the price p of output, or an increase of the marginal product of capital (an increase of η).

The aforementioned papers on microfinance consider models in which adjustment costs fail even to be single-dipped: there is a minimum investment size $I > 0$, meaning that investing $i \in (0, I)$ costs $C(i) = \infty$ (whereas investing $i \geq I$ has finite cost). Our fundamental result, Theorem 1, can accommodate such failures of single-dippedness: it remains true that a rise in the marginal profitability of capital increases investment, just as would

²⁸For $x' \geq x$ and $\bar{T} \geq_{\text{flat}} T$, $F(x', T) - F(x, T) \geq (>) 0$ is equivalent to $T(wx') - T(wx) \leq (<) wx' - wx - k$, where $k = \kappa(x') - \kappa(x)$, which implies $\bar{T}(wx') - \bar{T}(wx) \leq (<) wx' - wx - k$ by definition of \geq_{flat} , which is equivalent to $F(x', \bar{T}) - F(x, \bar{T}) \geq (>) 0$.

be the case if adjustment were costless. Our remaining results (Theorems 2 to 4) cannot be applied in this case, however.

All of these results generalize to multiple factors of production, on the pattern of Section 4.2. It suffices to assume that the factors $x = (x_1, \dots, x_n)$ are complements in production, meaning that the production function $f(x, \eta)$ is supermodular in x . Then, denoting factor prices by $r = (r_1, \dots, r_n)$, the profit function $F(x, (p, \eta, -r)) = pf(x, \eta) - r \cdot x$ is supermodular in x , and has increasing differences in $(x, (p, \eta, -r))$ as before, so that all of our general results remain applicable. In case there are just $n = 2$ factors of production, the complementarity hypothesis may be replaced with substitutability (submodularity of $f(x, \eta)$ in x), using the trick described in Section 4.2.

6. DYNAMIC ADJUSTMENT BY SHORT-LIVED AGENTS

In this section, we continue our study of the dynamic adjustment model introduced in the previous section, under a different behavioral assumption: that each period's decision is made by a short-lived agent. We recover the Le Chatelier principle, and provide conditions under which short-lived agents adjust more sluggishly than a long-lived agent would.

Recall the model from Section 5.1. We now assume that there is one agent per period. The period- t agent takes her predecessor's choice $x_{t-1} \in L$ as given, and chooses $x_t \in L$ to maximize the period- t payoff $G_t(\cdot, x_{t-1})$, where $G_t(x, y) = F(x, \theta_t) - C_t(x - y)$. A (*short-lived*) *equilibrium sequence* is a sequence $(x_t)_{t=1}^\infty$ in L such that $x_t \in \arg \max_{x \in L} G_t(x, x_{t-1})$ for every $t \in \mathbb{N}$, where (as before) $x_0 = \underline{x} \in \arg \max_{x \in L} F(x, \underline{\theta})$ is given.

THEOREM 5—Short-lived dynamic Le Chatelier: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that each adjustment cost C_t is monotone. Assume that $\arg \max_{x \in L} G_t(x, y)$ is nonempty for all $t \in \mathbb{N}$ and $y \in L$. Fix $\underline{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.²⁹*

- *If $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$, then there is an equilibrium sequence $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq \bar{x}$ for every $t \in \mathbb{N}$.*
- *If $\underline{\theta} \leq \theta_t \leq \theta_{t+1} \leq \bar{\theta}$ for every $t \in \mathbb{N}$, then there is an equilibrium sequence $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq x_{t+1} \leq \bar{x}$ for every $t \in \mathbb{N}$.*

By inspection, Proposition 3 (Section 4) is the special case of Theorem 5 in which there are only two periods of adjustment ($C_t(\varepsilon) = \infty$ for all $t \geq 3$ and $\varepsilon \neq 0$) and the parameter shifts once and for all ($\theta_1 = \theta_2 = \bar{\theta}$).

Taken together, our three dynamic Le Chatelier principles (Theorems 3 to 5) tell us that long- and short-lived agents adjust in the same direction. The *speed* of adjustment generally differs, however. The following result gives additional assumptions under which short-lived agents adjust more sluggishly.

THEOREM 6—Short- versus long-lived: *Suppose that the objective $F(x, \theta)$ is supermodular in x and has single-crossing differences in (x, θ) , and that each adjustment cost C_t is monotone, additively separable, and convex. Let $\{F(\cdot, \theta)\}_{\theta \in \Theta}$ and $\{C_t\}_{t \in \mathbb{N}}$ be equi-BCS.³⁰ Fix $\bar{\theta} \geq \underline{\theta}$, let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$,³¹ and assume that $\underline{\theta} \leq \theta_t \leq \theta_{t+1} \leq \bar{\theta}$ for every*

²⁹Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

³⁰Given $X \subseteq \mathbb{R}^n$, a collection $\{\phi_k\}_{k \in \mathcal{K}}$ of functions $\phi_k : X \rightarrow (-\infty, \infty]$ is equi-BCS if and only if the map $x \mapsto \sup_{k \in \mathcal{K}} |\phi_k(x)|$ is BCS.

³¹Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

$t \in \mathbb{N}$. Fix a short-lived equilibrium sequence $(\tilde{x}_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq \tilde{x}_t \leq \tilde{x}_{t+1} \leq \bar{x}$ for every $t \in \mathbb{N}$.³² Then, provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\tilde{x}_t \leq x_t \leq \bar{x}$ for every period $t \in \mathbb{N}$.

The intuition is straightforward: short-lived agents adjust more sluggishly because they do not take into account that the less they adjust today, the more adjustment will be required tomorrow. Note that this result relies on stronger assumptions than our earlier results: in particular, convexity of costs.

The proof of Theorem 6 (Appendix L) establishes that if $(x_t)_{t=1}^\infty$ solves the long-lived agent's problem and satisfies $\underline{x} \leq x_t \leq \bar{x}$ for every $t \in \mathbb{N}$,³³ then $(\tilde{x}_t \vee x_t)_{t=1}^\infty$ also solves the long-lived agent's problem. By equi-BCS and a limit argument, it suffices to show that for every $T \in \{0, 1, 2, \dots\}$, the sequence $(\tilde{x}_{\min\{t, T\}} \vee x_t)_{t=1}^\infty$ solves the long-lived agent's problem, where $\tilde{x}_0 = \underline{x}$. We show this by induction on $T \in \{0, 1, 2, \dots\}$; the base case $T = 0$ is immediate, and the induction step uses the supermodularity of the objective $F(\cdot, \theta)$ and the monotonicity, additive separability, and convexity of each cost C_t .

7. NECESSITY OF MONOTONICITY ASSUMPTIONS

The basic comparative-statics result for costless adjustment (see Milgrom and Shannon (1994, Theorem 4)) includes a converse, which asserts that the ordinal complementarity assumptions on the objective are not only sufficient for comparative statics, but also necessary in a natural sense. When adjustment is costly, ordinal complementarity is still necessary.³⁴ In this section, we show that the same is true of monotonicity-type assumptions on the cost function: they are necessary (in the same sense) for comparative statics.³⁵

First, Theorem 1 has a direct converse:

THEOREM 1[†]: Assume that there exist distinct parameters $\underline{\theta}, \bar{\theta} \in \Theta$ such that $\underline{\theta} \leq \bar{\theta}$. For a cost function C , the following are equivalent:

- (a) C is minimally monotone.
- (b) $F : X \times \Theta \rightarrow \mathbb{R}$ such that $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) ,
 - $\bar{\theta} \geq \underline{\theta}$ implies $\hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in X} G(x, \bar{\theta})$, provided the argmax is nonempty, and
 - $\bar{\theta} \leq \underline{\theta}$ implies $\hat{x} \leq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in X} G(x, \bar{\theta})$, provided the argmax is nonempty.

Minimal monotonicity is *not* equivalent to the “ $\bar{\theta} \geq \underline{\theta}$ ” half of property (b) alone. Rather, as noted in footnote 11, the “ $\bar{\theta} \geq \underline{\theta}$ ” half of (b) is implied by one half of minimal monotonicity ($C(\varepsilon \vee 0) \leq C(\varepsilon)$ for every $\varepsilon \in \Delta L$), while the “ $\bar{\theta} \leq \underline{\theta}$ ” half follows from the other half of minimal monotonicity.

PROOF: (a) implies (b) by Theorem 1. To show that (b) implies (a), we prove the contrapositive: suppose that C fails to be minimally monotone, meaning that there is an adjustment vector $\varepsilon \in \Delta L$ such that $C(\varepsilon \wedge 0) > C(\varepsilon)$ or $C(\varepsilon \vee 0) > C(\varepsilon)$; we will

³²Such a sequence exists by Theorem 5, provided there is an equilibrium sequence.

³³Such a solution exists by Theorem 3, provided there is a solution.

³⁴Since ordinal complementarity is necessary for comparative statics in the costless case $C \equiv 0$, it is a fortiori necessary for comparative statics to hold whatever the cost C .

³⁵Thanks to a referee and the editor for encouraging us to develop this material.

show that (b) fails. Assume that $C(\varepsilon \vee 0) > C(\varepsilon)$; the other case is analogous. Let $\alpha = C(\varepsilon \vee 0) - C(\varepsilon) > 0$ and $\beta = 2\alpha + \max\{0, C(\varepsilon) - C(\varepsilon \wedge 0), C(\varepsilon) - C(0)\}$.

Choose $\underline{x}, \widehat{x} \in L$ such that $\widehat{x} - \underline{x} = \varepsilon$, and note that $\widehat{x} \not\leq \underline{x}$ (else $\varepsilon \vee 0 = \varepsilon$, which would imply $\alpha = 0$). Let $X = \{\underline{x} \wedge \widehat{x}, \underline{x}, \widehat{x}, \underline{x} \vee \widehat{x}\}$; X is a sublattice. Fix distinct $\underline{\theta} \leq \bar{\theta}$ in Θ (possible by hypothesis). Let $F(\underline{x} \wedge \widehat{x}, \underline{\theta}) = F(\underline{x} \wedge \widehat{x}, \bar{\theta}) = \alpha/2$, $F(\underline{x}, \underline{\theta}) = F(\underline{x}, \bar{\theta}) = \alpha$, $F(\widehat{x}, \underline{\theta}) = 0$ and $F(\widehat{x}, \bar{\theta}) = \beta$, $F(\underline{x} \vee \widehat{x}, \underline{\theta}) = \alpha/2$, and $F(\underline{x} \vee \widehat{x}, \bar{\theta}) = \alpha/2 + \beta$.³⁶ Then $F(x, \theta)$ is supermodular in x and has increasing differences in (x, θ) , and we have

$$\arg \max_{x \in X} F(x, \underline{\theta}) = \{\underline{x}\} \quad \text{and} \quad \arg \max_{x \in X} F(x, \bar{\theta}) = \{\widehat{x}\}.$$

Hence, (b) fails, as $\bar{\theta} \geq \underline{\theta}$ and $\widehat{x} \not\leq \underline{x}$.

Q.E.D.

The converse of Theorem 2 is not true; instead, a property slightly weaker than monotonicity is necessary and sufficient for the Le Chatelier principle. Say that $C : \Delta L \rightarrow [0, \infty]$ is *weakly monotone* if and only if $C(\varepsilon') \leq C(\varepsilon)$ whenever $0 \leq \varepsilon' \leq \varepsilon \vee 0$ or $0 \geq \varepsilon' \geq \varepsilon \wedge 0$. Equivalently, C is weakly monotone if and only if it is minimally monotone and $C(\varepsilon') \leq C(\varepsilon)$ whenever $0 \leq \varepsilon' \leq \varepsilon$ or $0 \geq \varepsilon' \geq \varepsilon$. The latter requirement is monotonicity, except applied only to “unidirectional” adjustments (up in every dimension, $\varepsilon \geq 0$, or down in every dimension, $\varepsilon \leq 0$).

THEOREM 2[†]: *Assume that there exist distinct parameters $\underline{\theta}, \bar{\theta} \in \Theta$ such that $\underline{\theta} \leq \bar{\theta}$. For a cost function C , the following are equivalent:*

- (a) *The adjustment cost C is weakly monotone.*
- (b) *For any sublattice $X \subseteq L$ and any objective $F : X \times \Theta \rightarrow \mathbb{R}$ such that $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) ,*
 - $\bar{\theta} \geq \underline{\theta}$ *implies that for any $\bar{x} \in \arg \max_{x \in X} F(x, \bar{\theta})$ such that $\bar{x} \geq \underline{x}$,³⁷ we have $\bar{x} \geq \widehat{x} \geq \underline{x}$ for some $\widehat{x} \in \arg \max_{x \in X} G(x, \bar{\theta})$, provided the argmax is nonempty, and*
 - $\bar{\theta} \leq \underline{\theta}$ *implies that for any $\bar{x} \in \arg \max_{x \in X} F(x, \bar{\theta})$ such that $\bar{x} \leq \underline{x}$,³⁸ we have $\bar{x} \leq \widehat{x} \leq \underline{x}$ for some $\widehat{x} \in \arg \max_{x \in X} G(x, \bar{\theta})$, provided the argmax is nonempty.*

Theorem 3 also has a converse: monotonicity is necessary as well as sufficient for the dynamic Le Chatelier principle. Define $\underline{C}, \bar{C} : \Delta L \rightarrow [0, \infty]$ by $\underline{C} \equiv 0$, $\bar{C}(0) = 0$, and $\bar{C}(\varepsilon) = \infty$ for all $\varepsilon \neq 0$, and call a set \mathcal{C} of cost functions *admissible* if and only if it contains both \underline{C} and \bar{C} .

THEOREM 3[†]: *Assume that there exist distinct parameters $\underline{\theta}, \theta^1, \theta^2, \bar{\theta} \in \Theta$ such that $\underline{\theta} \leq \theta^1 \leq \bar{\theta} \geq \theta^2 \geq \underline{\theta}$ and $\theta^1 \not\leq \theta^2 \not\leq \theta^1$. For an admissible set \mathcal{C} of cost functions, the following are equivalent:*

- (a) *Every member of \mathcal{C} is monotone.*
- (b) *For any sublattice $X \subseteq L$, any objective $F : X \times \Theta \rightarrow \mathbb{R}$ such that $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , any parameter sequence $(\theta_i)_{i=1}^\infty \subseteq \Theta$, and any cost-function sequence $(C_i)_{i=1}^\infty \subseteq \mathcal{C}$, it holds that*

³⁶For $\theta \in \Theta \setminus \{\underline{\theta}, \bar{\theta}\}$, let $F(\cdot, \theta) = F(\cdot, \underline{\theta})$ if $\theta \leq \underline{\theta}$ and $F(\cdot, \theta) = F(\cdot, \bar{\theta})$ otherwise.

³⁷Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

³⁸Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

- for any $\underline{\theta} \leq \bar{\theta}$ in Θ and any $\underline{x} \in \arg \max_{x \in X} F(x, \underline{\theta})$ and $\bar{x} \in \arg \max_{x \in X} F(x, \bar{\theta})$ such that $\underline{x} \leq \bar{x}$,³⁹ if $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$, then provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq \bar{x}$ for every period $t \in \mathbb{N}$, and
- for any $\underline{\theta} \geq \bar{\theta}$ in Θ and any $\underline{x} \in \arg \max_{x \in X} F(x, \underline{\theta})$ and $\bar{x} \in \arg \max_{x \in X} F(x, \bar{\theta})$ such that $\underline{x} \geq \bar{x}$,⁴⁰ if $\underline{\theta} \geq \theta_t \geq \bar{\theta}$ for every $t \in \mathbb{N}$, then provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \geq x_t \geq \bar{x}$ for every period $t \in \mathbb{N}$.

The role of the assumption that there exist incomparable parameters ($\theta^1, \theta^2 \in \Theta$ such that $\theta^1 \not\leq \theta^2 \not\leq \theta^1$) is to avoid conflating property (b) with the weaker property that quantifies only over “one-dimensional” parameter sequences $(\theta_t)_{t=1}^\infty$, meaning those such that either $\theta_t \leq \theta_T$ or $\theta_T \leq \theta_t$ holds for all $t, T \in \mathbb{N}$. We do not know whether (a) is necessary also for the weaker “one-dimensional” property. However, *weak* monotonicity is necessary, by Theorem 2[†] (see footnote 41 below).

The proof of Theorem 3[†] (Appendix N) has three parts. First, (a) implies (b) by Theorem 3. Second, if there is a $C \in \mathcal{C}$ that fails to be *weakly* monotone, then (b) fails by Theorem 2[†].⁴¹ Third, suppose there is a $C \in \mathcal{C}$ that is weakly monotone but not monotone. In this case, we construct suitable $X \subseteq L$ and $F : X \times \Theta \rightarrow \mathbb{R}$ such that if $\theta_1 = \theta^1$, $\theta_t = \theta^2$ for all $t \geq 2$, $C_1 = \underline{C}$, $C_2 = C$, and $C_t = \bar{C}$ for every $t \geq 3$, then in every solution $(x_t)_{t=1}^\infty$ of the long-lived agent's problem, x_2 fails to lie between \underline{x} and \bar{x} .

We do not have a converse for Theorem 4. We conjecture that the strong dynamic Le Chatelier principle remains true under somewhat weaker assumptions, but showing this would seem to require a different proof strategy, since our strategy relies heavily on the stated assumptions. Identifying the necessary and sufficient conditions seems challenging; we leave this problem open.

APPENDIX A: STANDARD DEFINITIONS

Given $X \subseteq \mathbb{R}$, a function $\psi : X \rightarrow (-\infty, \infty]$ is *single-dipped* if and only if there is an $x \in X$ such that ψ is decreasing on $\{y \in X : y \leq x\}$ and increasing on $\{y \in X : y \geq x\}$. Given $X \subseteq \mathbb{R}^n$, a function $\psi : X \rightarrow (-\infty, \infty]$ is *bounded on compact sets (BCS)* if and only if, for each compact $Y \subseteq X$, there is a constant $K_Y > 0$ such that $|\psi(y)| \leq K_Y$ for every $y \in Y$.

Fix a sublattice L of \mathbb{R}^n . A function $\phi : L \rightarrow \mathbb{R}$ is called *supermodular* if $\phi(x) - \phi(x \wedge y) \leq \phi(x \vee y) - \phi(y)$ for any $x, y \in L$, *quasi-supermodular* if $\phi(x) - \phi(x \wedge y) \geq (>) 0$ implies $\phi(x \vee y) - \phi(y) \geq (>) 0$, and *(quasi-)submodular* if $-\phi$ is (quasi-)supermodular. Clearly, supermodularity implies quasi-supermodularity. If $n = 1$, then every function $\phi : L \rightarrow \mathbb{R}$ is automatically supermodular.

Fix a partially ordered set Θ . A function $F : L \times \Theta \rightarrow \mathbb{R}$ has *(strict) increasing differences* if $F(y, \theta) - F(x, \theta)$ is (strictly) increasing in θ whenever $x \leq y$, has *single-crossing differences* if $F(y, \theta') - F(x, \theta') \geq (>) 0$ implies $F(y, \theta'') - F(x, \theta'') \geq (>) 0$ whenever

³⁹Such \underline{x} and \bar{x} must exist, provided the argmaxes are nonempty (refer to footnote 18).

⁴⁰Such \underline{x} and \bar{x} must exist, provided the argmaxes are nonempty (refer to footnote 18).

⁴¹By Theorem 2[†], there exist suitable $X \subseteq L$ and $\tilde{F} : X \times \Theta \rightarrow \mathbb{R}$ such that every $\hat{x} \in M = \arg \max_{x \in X} [\tilde{F}(x, \bar{\theta}) - C(x - \underline{x})]$ fails to lie between \underline{x} and \bar{x} . Hence, if $F = (1 - \delta)\tilde{F}$, $\theta_t = \bar{\theta}$ for every $t \in \mathbb{N}$, $C_1 = C$ and $C_t = \bar{C}$ for every $t \geq 2$, then every solution $(x_t)_{t=1}^\infty$ of the long-lived agent's problem has $x_1 \in M$, so x_1 fails to lie between \underline{x} and \bar{x} .

$x \leq y$ and $\theta' \leq \theta''$, has *strict single-crossing differences* if $F(y, \theta') - F(x, \theta') \geq 0$ implies $F(y, \theta'') - F(x, \theta'') > 0$ whenever $x < y$ and $\theta' < \theta''$, and has (strict) *decreasing differences* if $-F$ has (strict) increasing differences. A function $F : L \times \Theta \rightarrow \mathbb{R}_{++}$ has (strict) *log increasing differences* if and only if $\ln F$ has (strict) increasing differences. (Strict) increasing differences and (strict) log increasing differences each imply (strict) single-crossing differences.

Quasi-supermodularity and single-crossing differences are *ordinal* properties: they are preserved by strictly increasing transformations.⁴² By contrast, supermodularity and increasing differences are in general preserved only by strictly increasing *affine* transformations: in other words, they are *cardinal*.

The sum of quasi-supermodular functions need not be quasi-supermodular. Likewise, single-crossing differences is not preserved by summation. By contrast, the sum of supermodular functions is supermodular, and the sum of functions with increasing differences also has increasing differences.

APPENDIX B: EXTENSION TO UNCERTAIN ADJUSTMENT COST

Several of our results are robust to uncertainty about adjustment costs. For Theorems 1 and 2, augment the setting from Section 2 as follows. Let all uncertainty be summarized by a random variable S , called “the state of the world.” The agent’s adjustment cost is $C_s(\cdot)$ in state $S = s$. Her ex ante payoff is

$$\tilde{G}(x, \theta) = \mathbb{E}[u(F(x, \theta) - C_s(x - \underline{x}))],$$

where u is an increasing function $\mathbb{R} \rightarrow \mathbb{R}$.

THEOREM 1’: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that at almost every realization s of the state S , the adjustment cost C_s is minimally monotone. If $\bar{\theta} \geq \underline{\theta}$, then $\hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in L} \tilde{G}(x, \bar{\theta})$, provided the argmax is nonempty.*

PROOF: Let $x' \in \arg \max_{x \in L} \tilde{G}(x, \bar{\theta})$, and let $\hat{x} = \underline{x} \vee x'$. For almost every realization s , applying the proof of Theorem 1 yields $F(\hat{x}, \bar{\theta}) - C_s(\hat{x} - \underline{x}) \geq F(x', \bar{\theta}) - C_s(x' - \underline{x})$. Hence, $\tilde{G}(\hat{x}, \bar{\theta}) \geq \tilde{G}(x', \bar{\theta})$ as u is increasing, which since x' maximizes $\tilde{G}(\cdot, \bar{\theta})$ on L implies that \hat{x} does, too. Clearly, $\hat{x} \geq \underline{x}$. Q.E.D.

THEOREM 2’: *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that at almost every realization s of the state S , the adjustment cost C_s is monotone. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.⁴³ Then*

- $\bar{x} \geq \hat{x} \geq \underline{x}$ for some $\hat{x} \in \arg \max_{x \in L} \tilde{G}(x, \bar{\theta})$, provided the argmax is nonempty, and
- if \bar{x} is the largest element of $\arg \max_{x \in L} F(x, \bar{\theta})$, then $\bar{x} \geq \hat{x}$ for any $\hat{x} \in \arg \max_{x \in L} \tilde{G}(x, \bar{\theta})$.

⁴²That is, for any strictly increasing $f : \mathbb{R} \rightarrow \mathbb{R}$, if ϕ is quasi-supermodular, then so is $f \circ \phi$, and if $F(x, \theta)$ has increasing differences in (x, θ) , then so does $\tilde{F}(x, \theta) = f(F(x, \theta))$.

⁴³Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

The proof is that of Theorem 2, modified along the lines of the proof of Theorem 1' (above). We omit the details.

To extend our dynamic results, augment the dynamic model of Section 5.1 as follows. Let the period- t adjustment cost function be $C_{S_t,t}$, where S_t is a random variable. The agent's period- t expected payoff is $\tilde{G}_t(x_t, x_{t-1}) = \mathbb{E}[u(F(x_t, \theta_t) - C_{S_t,t}(x_t - x_{t-1}))]$, where $u: \mathbb{R} \rightarrow \mathbb{R}$ is increasing. Given her period-0 choice $x_0 \in L$, the long-lived agent's problem is to choose a sequence $(x_t)_{t=1}^\infty$ in L to maximize $\tilde{G}((x_t)_{t=1}^\infty, x_0) = \sum_{t=1}^\infty \delta^{t-1} \tilde{G}_t(x_t, x_{t-1})$. We assume that the states $(S_t)_{t=1}^\infty$ are independent across periods.⁴⁴

THEOREM 3': *Suppose that the objective $F(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , and that in each period t , at almost every realization s_t of the state S_t , the adjustment cost $C_{s_t,t}$ is monotone. Fix $\bar{\theta} \geq \underline{\theta}$, and let $\bar{x} \in \arg \max_{x \in L} F(x, \bar{\theta})$ satisfy $\bar{x} \geq \underline{x}$.⁴⁵ If $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$, then provided the long-lived agent's problem admits a solution, there is a solution $(x_t)_{t=1}^\infty$ that satisfies $\underline{x} \leq x_t \leq \bar{x}$ for every period $t \in \mathbb{N}$.*

Again, a simple "a.s." modification to the proof Theorem 3 delivers this result. By contrast, the proof of Theorem 4 is not easily modified to accommodate uncertain cost, except in the risk-neutral case (when u is affine).

APPENDIX C: PROOF OF THEOREM 1*

Let $x' \in \arg \max_{x \in \bar{L}} G(x, \bar{\theta})$. Define $\hat{x} = \underline{x} \vee x'$; obviously $\hat{x} \geq \underline{x}$, and \hat{x} belongs to \bar{L} since $\bar{L} \geq_{ss} \underline{L} \ni \underline{x}$. We claim that \hat{x} maximizes $G(\cdot, \bar{\theta})$ on \bar{L} . We have $F(\underline{x}, \underline{\theta}) \geq F(\underline{x} \wedge x', \underline{\theta})$ by definition of \underline{x} since $\underline{x} \wedge x' \in \underline{L}$ (by $\bar{L} \geq_{ss} \underline{L} \ni \underline{x}$). Thus, $F(\underline{x} \vee x', \underline{\theta}) \geq F(x', \underline{\theta})$ by quasi-supermodularity, whence $F(\underline{x} \vee x', \bar{\theta}) \geq F(x', \bar{\theta})$ by single-crossing differences. Furthermore, by minimal monotonicity, $C(\underline{x} \vee x' - \underline{x}) = C((x' - \underline{x}) \vee 0) \leq C(x' - \underline{x})$. Thus

$$G(\hat{x}, \bar{\theta}) = F(\underline{x} \vee x', \bar{\theta}) - C(\underline{x} \vee x' - \underline{x}) \geq F(x', \bar{\theta}) - C(x' - \underline{x}) = G(x', \bar{\theta}).$$

Since x' maximizes $G(\cdot, \bar{\theta})$ on \bar{L} , it follows that \hat{x} does, too.

Q.E.D.

APPENDIX D: PROOF OF PROPOSITION 1

Let $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$, and suppose toward a contradiction that $\hat{x} \not\geq \underline{x}$. Then $\underline{x} \vee \hat{x} > \hat{x}$. The proof of Theorem 1 yields $F(\underline{x} \vee \hat{x}, \bar{\theta}) \geq F(\hat{x}, \bar{\theta})$ and $C(\underline{x} \vee \hat{x} - \underline{x}) \leq C(\hat{x} - \underline{x})$, where the first inequality is strict if the single-crossing differences of F is strict, and the second inequality is strict if the minimal monotonicity of C is strict. In either case, we have

$$G(\underline{x} \vee \hat{x}, \bar{\theta}) = F(\underline{x} \vee \hat{x}, \bar{\theta}) - C(\underline{x} \vee \hat{x} - \underline{x}) > F(\hat{x}, \bar{\theta}) - C(\hat{x} - \underline{x}) = G(\hat{x}, \bar{\theta}),$$

which contradicts the fact that \hat{x} maximizes $G(\cdot, \bar{\theta})$ on L .

Q.E.D.

⁴⁴Theorem 3' is true without independence, but our *interpretation* of it hinges on independence. The result is about sequences $(x_t)_{t=1}^\infty$ that maximize the time-0 expected payoff $\tilde{G}(\cdot, x_0)$, not about optimal real-time choice of actions. The time-0 and real-time problems are equivalent if and only if the agent does not learn over time about the realizations of future states, and this is ensured by independence.

⁴⁵Such an \bar{x} must exist, provided the argmax is nonempty (refer to footnote 18).

APPENDIX E: PROOF OF CLAIM 1 IN SECTION 3.3

Label the elements of \mathcal{Y} as $\mathcal{Y} = \{y_1, \dots, y_N\}$, where $y_1 < \dots < y_N$. Let $\phi(c, y) = u_1(c) + u_2((1+r)(w-c) + y)$ for each $c \in [0, w]$ and $y \in \mathcal{Y}$. The function ϕ is supermodular since u_2 is concave. For all $c \in [0, w]$ and $G \in \mathcal{G}$,

$$U(c, G) = \int_{\mathcal{Y}} \phi(c, y) G(dy) = \phi(c, y_N) + \sum_{n=1}^{N-1} [\phi(c, y_{n+1}) - \phi(c, y_n)] [-G(y_n)]$$

from telescoping / integration by parts. Thus, for any $c, c' \in [0, w]$ and $G, H \in \mathcal{G}$, letting $k = \phi(c, y_N) - \phi(c \wedge c', y_N) = \phi(c \vee c', y_N) - \phi(c', y_N)$, we have

$$\begin{aligned} & U(c, G) - U(c \wedge c', G \wedge_1 H) \\ &= k + \sum_{n=1}^{N-1} [\phi(c, y_{n+1}) - \phi(c \wedge c', y_n)] [(G \wedge_1 H)(y_n) - G(y_n)] \\ &= k + \sum_{n=1}^{N-1} [\phi(c, y_{n+1}) - \phi(c \wedge c', y_n)] [H(y_n) - (G \vee_1 H)(y_n)] \\ &\leq k + \sum_{n=1}^{N-1} [\phi(c \vee c', y_{n+1}) - \phi(c', y_n)] [H(y_n) - (G \vee_1 H)(y_n)] \\ &= U(c \vee c', G \vee_1 H) - U(c', H), \end{aligned}$$

where the inequality holds since ϕ is supermodular and $H \leq_1 G \vee_1 H$. Q.E.D.

APPENDIX F: PROOF OF PROPOSITION 4

Fix any $\hat{x} \in \arg \max_{x \in L} G(x, \bar{\theta})$. We have $\hat{x} \geq \underline{x}$ by Proposition 1. To show that $\bar{x} \geq \hat{x}$, suppose toward a contradiction that $\bar{x} \not\geq \hat{x}$. The proof of Theorem 2 yields $F(\hat{x}, \bar{\theta}) \leq F(\bar{x} \wedge \hat{x}, \bar{\theta})$ and $C(\hat{x} - \underline{x}) < C(\bar{x} \wedge \hat{x} - \underline{x})$, where the latter inequality is strict since $\hat{x} - \underline{x} \neq \bar{x} \wedge \hat{x} - \underline{x}$ and C is *strictly* monotone. Thus

$$G(\bar{x} \wedge \hat{x}, \bar{\theta}) = F(\bar{x} \wedge \hat{x}, \bar{\theta}) - C(\bar{x} \wedge \hat{x} - \underline{x}) > F(\hat{x}, \bar{\theta}) - C(\hat{x} - \underline{x}) = G(\hat{x}, \bar{\theta}),$$

which contradicts the fact that \hat{x} maximizes $G(\cdot, \bar{\theta})$ on L . Q.E.D.

APPENDIX G: A MONOTONICITY LEMMA

The following lemma will be used below in the proofs of Theorems 3 and 6.

LEMMA 1: *If $C : \Delta L \rightarrow [0, \infty]$ is monotone, then $C(z \vee x - z \vee y) \leq C(x - y) \leq C(z \wedge x - z \wedge y)$ for any $x, y, z \in L$.*

PROOF: We shall prove the first inequality; the second follows similarly. By monotonicity, it suffices to show that for each i , one of the following holds:

- (a) $0 \leq (z \vee x - z \vee y)_i \leq (x - y)_i$
- (b) $0 \geq (z \vee x - z \vee y)_i \geq (x - y)_i$.

If $x_i \geq z_i \geq y_i$, then (a) holds by inspection; if $x_i \leq z_i \leq y_i$, then (b) holds by inspection; if $x_i \leq z_i \geq y_i$, then (a) or (b) holds since $(z \vee x - z \vee y)_i = 0$; and if $x_i \geq z_i \leq y_i$, then (a) or (b) holds since $(z \vee x - z \vee y)_i = (x - y)_i$. Q.E.D.

The following corollary will be used below to prove Theorems 5 and 6.

COROLLARY 1: *Let $C : \Delta L \rightarrow [0, \infty]$ be monotone, and consider $x, y, z \in L$.*

- *If $z \leq y$, then $C(z \vee x - y) \leq C(x - y)$.*
- *If $z \geq y$, then $C(z \wedge x - y) \leq C(x - y)$.*

APPENDIX H: PROOF OF THEOREM 3

Let $(x_t)_{t=1}^\infty$ maximize $\mathcal{G}(\cdot, x_0)$. We shall show that $(\bar{x} \wedge (\underline{x} \vee x_t))_{t=1}^\infty$ also maximizes $\mathcal{G}(\cdot, x_0)$; this suffices since $\underline{x} \leq \bar{x} \wedge (\underline{x} \vee x_t) \leq \bar{x}$ for each t .

We first show that $(\hat{x}_t)_{t=1}^\infty = (\underline{x} \vee x_t)_{t=1}^\infty$ maximizes $\mathcal{G}(\cdot, x_0)$. For every $t \in \mathbb{N}$, we have $F(\underline{x}, \underline{\theta}) \geq F(\underline{x} \wedge x_t, \underline{\theta})$ by definition of \underline{x} , which by quasi-supermodularity implies that $F(\underline{x} \vee x_t, \underline{\theta}) \geq F(x_t, \underline{\theta})$, whence $F(\underline{x} \vee x_t, \theta_t) \geq F(x_t, \theta_t)$ by single-crossing differences and $\theta_t \geq \underline{\theta}$. Thus, $\mathcal{F}((\underline{x} \vee x_t)_{t=1}^\infty) \geq \mathcal{F}((x_t)_{t=1}^\infty)$. Furthermore, for every $t \in \mathbb{N}$, we have $C_t(\underline{x} \vee x_t - \underline{x} \vee x_{t-1}) \leq C_t(x_t - x_{t-1})$ by Lemma 1 (Appendix G) since C_t is monotone, so $\mathcal{C}(x_0, (\underline{x} \vee x_t)_{t=1}^\infty) \leq \mathcal{C}(x_0, (x_t)_{t=1}^\infty)$. So $\mathcal{G}((\underline{x} \vee x_t)_{t=1}^\infty, x_0) \leq \mathcal{G}((x_t)_{t=1}^\infty, x_0)$, which since $(x_t)_{t=1}^\infty$ maximizes $\mathcal{G}(\cdot, x_0)$ implies that $(\hat{x}_t)_{t=1}^\infty = (\underline{x} \vee x_t)_{t=1}^\infty$ does, too.

It remains to show that $(\bar{x} \wedge \hat{x}_t)_{t=1}^\infty$ also maximizes $\mathcal{G}(\cdot, x_0)$. For every $t \in \mathbb{N}$, we have $F(\bar{x} \vee \hat{x}_t, \bar{\theta}) \leq F(\bar{x}, \bar{\theta})$ by definition of \bar{x} , which by quasi-supermodularity implies that $F(\hat{x}_t, \bar{\theta}) \leq F(\bar{x} \wedge \hat{x}_t, \bar{\theta})$, whence $F(\hat{x}_t, \theta_t) \leq F(\bar{x} \wedge \hat{x}_t, \theta_t)$ by single-crossing differences and $\theta_t \leq \bar{\theta}$. Thus, $\mathcal{F}((\hat{x}_t)_{t=1}^\infty) \leq \mathcal{F}((\bar{x} \wedge \hat{x}_t)_{t=1}^\infty)$. Furthermore, for every $t \in \mathbb{N}$, we have $C_t(\hat{x}_t - \hat{x}_{t-1}) \geq C_t(\bar{x} \wedge \hat{x}_t - \bar{x} \wedge \hat{x}_{t-1})$ by Lemma 1 since C_t is monotone, so $\mathcal{C}(x_0, (\hat{x}_t)_{t=1}^\infty) \geq \mathcal{C}(x_0, (\bar{x} \wedge \hat{x}_t)_{t=1}^\infty)$. So $\mathcal{G}((\hat{x}_t)_{t=1}^\infty, x_0) \leq \mathcal{G}((\bar{x} \wedge \hat{x}_t)_{t=1}^\infty, x_0)$, which since $(\hat{x}_t)_{t=1}^\infty$ maximizes $\mathcal{G}(\cdot, x_0)$ implies that $(\bar{x} \wedge \hat{x}_t)_{t=1}^\infty$ does, too. Q.E.D.

APPENDIX I: PROOF OF THEOREM 4

For any sequence $\mathbf{x} = (x_t)_{t=1}^\infty$ in L and any $T \in \mathbb{N}$, let $M_T \mathbf{x}$ denote the sequence in L whose t th entry is x_t for $t < T$ and $x_{t-1} \vee x_t$ for $t \geq T$.

Assume that the agent's problem admits a solution. Let $\mathbf{x}^1 = (x_t^1)_{t=1}^\infty$ be a solution satisfying $\underline{x} \leq x_t^1 \leq \bar{x}$ in every period t ; such a solution exists by Theorem 3. Define $X_t = x_1^1 \vee x_2^1 \vee \dots \vee x_{t-1}^1 \vee x_t^1$ for $t \in \mathbb{N}$, and $X_0 = \underline{x}$.

Write $\mathbf{x}^T = M_T M_{T-1} \dots M_3 M_2 \mathbf{x}^1$ for $T \geq 2$. By inspection, the first T entries of \mathbf{x}^T are $X_1, X_2, \dots, X_{T-1}, X_T$. Clearly, $\underline{x} \leq X_t \leq X_{t+1} \leq \bar{x}$ for any period $t \in \mathbb{N}$. To prove the theorem, we need only show that $\mathbf{x}^\infty = (X_1, X_2, X_3, \dots)$ is optimal.

It suffices to show for each $T \in \mathbb{N}$ that \mathbf{x}^T is optimal. For then, letting V be the optimal value and noting that both $\mathbf{x}^T = (x_t)_{t=1}^\infty$ and \mathbf{x}^∞ have X_1, \dots, X_T as their first T entries, we have

$$\begin{aligned}
 0 &\geq \mathcal{G}(\mathbf{x}^\infty, x_0) - V = \mathcal{G}(\mathbf{x}^\infty, x_0) - \mathcal{G}(\mathbf{x}^T, x_0) \\
 &= \delta^T [\mathcal{G}((X_t)_{t=T+1}^\infty, X_T) - \mathcal{G}((x_t)_{t=T+1}^\infty, X_T)] \\
 &= \delta^T [\mathcal{F}((X_t)_{t=T+1}^\infty) - \mathcal{F}((x_t)_{t=T+1}^\infty)] \\
 &\quad - \delta^T [\mathcal{C}(X_T, (X_t)_{t=T+1}^\infty) - \mathcal{C}(X_T, (x_t)_{t=T+1}^\infty)] \\
 &\geq \delta^T [\mathcal{F}((X_t)_{t=T+1}^\infty) - \mathcal{F}((x_t)_{t=T+1}^\infty)], \tag{1}
 \end{aligned}$$

where the final inequality holds since

$$C(X_t - X_{t-1}) \leq C(x_t - x_{t-1}) \quad \text{for every } t \geq T + 2$$

and $C(X_{T+1} - X_T) \leq C(x_{T+1} - X_T)$

by the monotonicity of C . (For $t \geq T + 2$, for each dimension i , if $X_{t,i} = X_{t-1,i}$, then $0 = (X_t - X_{t-1})_i$, while if $X_{t,i} > X_{t-1,i}$, then $x_{t,i} = X_{t,i} \geq X_{t-1,i} \geq x_{t-1,i}$, so $0 \leq (X_t - X_{t-1})_i \leq (x_t - x_{t-1})_i$. For the $t = T + 1$ inequality, if $X_{T+1,i} = X_{T,i}$, then $0 = (X_{T+1} - X_T)_i \geq (x_{T+1} - X_T)_i$, while if $X_{T+1,i} > X_{T,i}$, then $X_{T+1,i} = x_{T+1,i}$, so $(X_{T+1} - X_T)_i = (x_{T+1} - X_T)_i$.) Since $F(\cdot, \bar{\theta})$ is BCS, the “[\cdot]” expression in (1) is bounded below uniformly over $T \in \mathbb{N}$,⁴⁶ so letting $T \rightarrow \infty$ yields $0 \leq \mathcal{G}(x^\infty, x_0) - V \geq 0$, which is to say that x^∞ is optimal.

To show that x^T is optimal for each $T \in \mathbb{N}$, we employ induction on $T \in \mathbb{N}$. The base case $T = 1$ is immediate.

For the induction step, fix any $T \in \mathbb{N}$, and suppose that $x^T = (x_t)_{t=1}^\infty$ is optimal; we will show that $x^{T+1} = M_{T+1}x^T$ is also optimal. Let $(\tilde{x}_t)_{t=1}^\infty$ be the sequence with t th entry x_t for $t < T$ and $x_t \wedge x_{t+1}$ for $t \geq T$. Since $x^T = (x_t)_{t=1}^\infty$ is optimal, and $(\tilde{x}_t)_{t=1}^\infty$ shares its first $T - 1$ entries X_1, \dots, X_{T-1} , we have $\mathcal{G}((x_t)_{t=T}^\infty, X_{T-1}) \geq \mathcal{G}((\tilde{x}_t)_{t=T}^\infty, X_{T-1})$, which may be written in full as

$$\sum_{t=T}^{\infty} \delta^{t-T} ([F(x_t, \bar{\theta}) - F(x_t \wedge x_{t+1}, \bar{\theta})] - [C(x_t - x_{t-1}) - C(x_t \wedge x_{t+1} - x_{t-1} \wedge x_t)]) \geq 0. \quad (2)$$

(Note that since $x_t = X_t$ for every $t \leq T$, we have $x_{T-1} \wedge x_T = X_{T-1} = x_{T-1}$.) Since $F(\cdot, \bar{\theta})$ is supermodular, it holds for every $t \geq T$ that

$$F(x_t \vee x_{t+1}, \bar{\theta}) - F(x_{t+1}, \bar{\theta}) \geq F(x_t, \bar{\theta}) - F(x_t \wedge x_{t+1}, \bar{\theta}). \quad (3)$$

We furthermore claim that for each $t \geq T$,

$$C(x_t \vee x_{t+1} - x_{t-1} \vee x_t) - C(x_{t+1} - x_t) \leq C(x_t - x_{t-1}) - C(x_t \wedge x_{t+1} - x_{t-1} \wedge x_t); \quad (4)$$

we shall prove this shortly. Combining (2), (3), and (4), and changing variables in the sum, we obtain

$$\sum_{t=T+1}^{\infty} \delta^{t-(T+1)} ([F(x_{t-1} \vee x_t, \bar{\theta}) - F(x_t, \bar{\theta})] - [C(x_{t-1} \vee x_t - x_{t-2} \vee x_{t-1}) - C(x_t - x_{t-1})]) \geq 0.$$

By inspection, this says precisely that $(\hat{x}_t)_{t=1}^\infty = x^{T+1} = M_{T+1}x^T$ satisfies

$$\mathcal{G}((\hat{x}_t)_{t=T+1}^\infty, X_T) \geq \mathcal{G}((x_t)_{t=T+1}^\infty, X_T).$$

⁴⁶Since X_t and x_t belong to the compact set $[\underline{x}, \bar{x}]$ for every $t \in \mathbb{N}$, there is a $K > 0$ such that $F(X_t, \bar{\theta}) - F(x_t, \bar{\theta}) \geq -2K$ for all t , so “[\cdot]” is bounded below by $-2K/(1 - \delta)$.

(Note that since $x_t = X_t$ for every $t \leq T$, we have $x_{(T+1)-2} \vee x_{(T+1)-1} = X_{T-1} \vee X_T = X_T = x_{(T+1)-1}$.) Since $\mathbf{x}^{T+1} = (\hat{x}_t)_{t=1}^\infty$ and $\mathbf{x}^T = (x_t)_{t=1}^\infty$ agree in their first T entries, and \mathbf{x}^T is optimal, it follows that \mathbf{x}^{T+1} is optimal, too.

It remains to show that (4) holds. It suffices to prove for each i that

$$C_i(y \vee z - x \vee y) + C_i(y \wedge z - x \wedge y) \leq C_i(y - x) + C_i(z - y) \quad \text{for any } x, y, z. \quad (5)$$

(We have renamed $x_{t-1,i} = x$, $x_{t,i} = y$, and $x_{t+1,i} = z$.) When y is not extreme (neither least nor greatest), (5) holds trivially because the left-hand side is equal to the right-hand side. When y is extreme, (5) reads

$$C_i(0) + C_i(z - x) \leq C_i(y - x) + C_i(z - y),$$

and we have either

$$\begin{aligned} & \text{(i) } 0 \leq z - x \leq y - x \quad \text{or} \quad \text{(ii) } 0 \leq z - x \leq z - y \\ & \text{or (iii) } 0 \geq z - x \geq y - x \quad \text{or} \quad \text{(iv) } 0 \geq z - x \geq z - y. \end{aligned}$$

Since C_i is single-dipped and minimized at zero, we have $C_i(z - x) \leq C_i(y - x) \leq C_i(y - x) + C_i(z - y) - C_i(0)$ in the first and third cases, and $C_i(z - x) \leq C_i(z - y) \leq C_i(y - x) + C_i(z - y) - C_i(0)$ in the second and fourth. *Q.E.D.*

The above proof applies nearly unchanged if there is a finite horizon K , so that the agent chooses a length- K sequence $(x_t)_{t=1}^K$ in L to maximize $\sum_{t=1}^K \delta^{t-1} [F(x_t, \bar{\theta}) - C(x_t - x_{t-1})]$. The limit argument early on is superfluous, so BCS is not needed. Let $x_{K+1} = x_K$ and follow the same steps to obtain

$$\begin{aligned} & \sum_{t=T}^K \delta^{t-T} ([F(x_t, \bar{\theta}) - F(x_t \wedge x_{t+1}, \bar{\theta})] \\ & \quad - [C(x_t - x_{t-1}) - C(x_t \wedge x_{t+1} - x_{t-1} \wedge x_t)]) \geq 0. \end{aligned}$$

This together with (3) and (4) and a change of variable delivers

$$\begin{aligned} & \sum_{t=T+1}^{K+1} \delta^{t-(T+1)} ([F(x_{t-1} \vee x_t, \bar{\theta}) - F(x_t, \bar{\theta})] \\ & \quad - [C(x_{t-1} \vee x_t - x_{t-2} \vee x_{t-1}) - C(x_t - x_{t-1})]) \geq 0. \end{aligned}$$

The sum's final term equals $\delta^{K-(T+1)}(0 - [C(x_K - x_{K-1} \vee x_K) - C(0)])$, which is nonpositive since C is minimized at 0 by monotonicity. Hence, the inequality is preserved when the final term is dropped:

$$\begin{aligned} & \sum_{t=T+1}^K \delta^{t-(T+1)} ([F(x_{t-1} \vee x_t, \bar{\theta}) - F(x_t, \bar{\theta})] \\ & \quad - [C(x_{t-1} \vee x_t - x_{t-2} \vee x_{t-1}) - C(x_t - x_{t-1})]) \geq 0. \end{aligned}$$

The remainder of the argument now applies unchanged.

APPENDIX J: PROOF OF THEOREM 5

For the first part, fix a sequence $(\theta_t)_{t=1}^\infty$ in Θ such that $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$. Call a finite sequence $(x_t)_{t=1}^T$ *equilibrium caged* if and only if $x_t \in \arg \max_{x \in L} G_t(x, x_{t-1})$ and $\underline{x} \leq x_t \leq \bar{x}$ for every $t \in \{1, \dots, T\}$. By Theorem 2, there exists an equilibrium caged sequence of length $T = 1$. Given this, it suffices to prove that for each $T \geq 2$, any length- $(T - 1)$ equilibrium caged sequence $(x_t)_{t=1}^{T-1}$ may be extended to a length- T equilibrium caged sequence $(x_t)_{t=1}^T$ (by appropriately choosing x_T). To that end, fix an arbitrary $T \geq 2$, and let $(x_t)_{t=1}^{T-1}$ be equilibrium caged. We shall prove two claims:

CLAIM 2: *There is an $x' \in \arg \max_{x \in L} G_T(x, x_{T-1})$ such that $x' \geq \underline{x}$.*

CLAIM 3: *$\bar{x} \wedge x'$ belongs to $\arg \max_{x \in L} G_T(x, x_{T-1})$ whenever x' does.*

These claims suffice because $x_T = \bar{x} \wedge x'$ satisfies $\underline{x} \leq x_T \leq \bar{x}$.

PROOF OF CLAIM 2: Fix any $x'' \in \arg \max_{x \in L} G_T(x, x_{T-1})$. We will show that $x' = \underline{x} \vee x''$ also maximizes $G_T(\cdot, x_{T-1})$; obviously $x' \geq \underline{x}$. We have $F(\underline{x}, \underline{\theta}) \geq F(\underline{x} \wedge x'', \underline{\theta})$ by definition of \underline{x} . Thus, $F(\underline{x} \vee x'', \underline{\theta}) \geq F(x'', \underline{\theta})$ by quasi-supermodularity, whence $F(\underline{x} \vee x'', \theta_T) \geq F(x'', \theta_T)$ by single-crossing differences and $\theta_T \geq \underline{\theta}$. Furthermore, since C_T is monotone and $\underline{x} \leq x_{T-1}$, we have $C_T(\underline{x} \vee x' - x_{T-1}) \leq C_T(x' - x_{T-1})$ by Corollary 1 (Appendix G). Thus

$$\begin{aligned} G_T(x', x_{T-1}) &= F(\underline{x} \vee x'', \theta_T) - C_T(\underline{x} \vee x'' - x_{T-1}) \\ &\geq F(x'', \theta_T) - C_T(x'' - x_{T-1}) = G_T(x'', x_{T-1}). \end{aligned}$$

Since x'' maximizes $G_T(\cdot, x_{T-1})$ on L , it follows that x' does, too.

Q.E.D.

PROOF OF CLAIM 3: Let x' belong to $\arg \max_{x \in L} G_T(x, x_{T-1})$; we claim that $\hat{x} = \bar{x} \wedge x'$ also maximizes $G_T(\cdot, x_{T-1})$. We have $F(\bar{x} \vee x', \bar{\theta}) \leq F(\bar{x}, \bar{\theta})$ by definition of \bar{x} , whence $F(x', \bar{\theta}) \leq F(\bar{x} \wedge x', \bar{\theta})$ by quasi-supermodularity, so that $F(x', \theta_T) \leq F(\bar{x} \wedge x', \theta_T)$ by single-crossing differences and $\theta_T \leq \bar{\theta}$. Since C_T is monotone and $\bar{x} \geq x_{T-1}$, we have $C_T(x' - x_{T-1}) \geq C_T(\bar{x} \wedge x' - x_{T-1})$ by Corollary 1. Thus

$$\begin{aligned} G_T(x', x_{T-1}) &= F(x', \theta_T) - C_T(x' - x_{T-1}) \\ &\leq F(\bar{x} \wedge x', \theta_T) - C_T(\bar{x} \wedge x' - x_{T-1}) = G_T(\hat{x}, x_{T-1}), \end{aligned}$$

which since x' maximizes $G_T(\cdot, x_{T-1})$ on L implies that \hat{x} does, too.

Q.E.D.

To prove the second part of Theorem 5, fix a sequence $(\theta_t)_{t=1}^\infty$ in Θ such that $\underline{\theta} \leq \theta_t \leq \theta_{t+1} \leq \bar{\theta}$ for every $t \in \mathbb{N}$. Recall that $x_0 = \underline{x}$. Call a finite sequence $(x_t)_{t=1}^T$ *equilibrium monotone* if and only if $x_t \in \arg \max_{x \in L} G_t(x, x_{t-1})$ and $x_{t-1} \leq x_t \leq \bar{x}$ for every $t \in \{1, \dots, T\}$. By Theorem 2, there exists an equilibrium monotone sequence of length $T = 1$. Given this, it suffices to prove that for every $T \geq 2$, any length- $(T - 1)$ equilibrium monotone sequence $(x_t)_{t=1}^{T-1}$ may be extended to a length- T equilibrium monotone sequence $(x_t)_{t=1}^T$ (by an appropriate choice of x_T).

To that end, fix an arbitrary $T \geq 2$, and let $(x_t)_{t=1}^{T-1}$ be equilibrium monotone; we shall show that for every $t \in \{1, \dots, T\}$, there is an $x' \in L$ which belongs to $\arg \max_{x \in L} G_T(x, x_{T-1})$ and satisfies $x_{t-1} \leq x' \leq \bar{x}$. We proceed by induction on $t \in$

$\{1, \dots, T\}$. The base case $t = 1$ follows from Claims 2 and 3. For the induction step, suppose that there is an $x'' \in \arg \max_{x \in L} G_T(x, x_{T-1})$ that satisfies $x_{t-2} \leq x'' \leq \bar{x}$; we claim that $x' = x_{t-1} \vee x''$ also maximizes $G_T(\cdot, x_{T-1})$. This suffices since $x_{t-1} \leq x' \leq \bar{x}$, where the latter inequality holds because $x_{t-1} \leq \bar{x}$ (as $(x_s)_{s=1}^{T-1}$ is equilibrium monotone) and $x'' \leq \bar{x}$.

We have $G_{t-1}(x_{t-1}, x_{t-2}) \geq G_{t-1}(x_{t-1} \wedge x'', x_{t-2})$ by definition of x_{t-1} . Since C_{t-1} is monotone and $x'' \geq x_{t-2}$ by the induction hypothesis, we have $C_{t-1}(x_{t-1} - x_{t-2}) \geq C_{t-1}(x_{t-1} \wedge x'' - x_{t-2})$ by Corollary 1 (Appendix G). It follows that $F(x_{t-1}, \theta_{t-1}) \geq F(x_{t-1} \wedge x'', \theta_{t-1})$. Thus, $F(x_{t-1} \vee x'', \theta_{t-1}) \geq F(x'', \theta_{t-1})$ by quasi-supermodularity, whence $F(x_{t-1} \vee x'', \theta_T) \geq F(x'', \theta_T)$ by single-crossing differences and $\theta_T \geq \theta_{t-1}$. We have $C_T(x_{t-1} \vee x'' - x_{T-1}) \leq C_T(x'' - x_{T-1})$ by Corollary 1 since C_T is monotone and $x_{t-1} \leq x_{T-1}$, where the latter holds since $(x_s)_{s=1}^{T-1}$ is equilibrium monotone. Thus

$$\begin{aligned} G_T(x', x_{T-1}) &= F(x_{t-1} \vee x'', \theta_T) - C_T(x_{t-1} \vee x'' - x_{T-1}) \\ &\geq F(x'', \theta_T) - C_T(x'' - x_{T-1}) = G_T(x'', x_{T-1}), \end{aligned}$$

which since x'' maximizes $G_T(\cdot, x_{T-1})$ on L implies that x' does, too. Q.E.D.

APPENDIX K: PROOF OF PROPOSITION 3

Define a sequence $(\theta_t)_{t=1}^\infty$ in Θ by $\theta_t = \bar{\theta}$ for every t . For each $t \geq 3$, define $C_t : \Delta L \rightarrow [0, \infty]$ by $C_t(\varepsilon) = \infty$ for every $\varepsilon \neq 0$ and $C_t(0) = 0$. Now apply Theorem 5. Q.E.D.

APPENDIX L: PROOF OF THEOREM 6

For any sequence $\mathbf{x} = (x_t)_{t=1}^\infty$ in L and any $T \in \{0, 1, 2, \dots\}$, let $R_T \mathbf{x}$ denote the sequence in L whose t th entry is x_t for $t < T$ and $\tilde{x}_T \vee x_t$ for $t \geq T$.

The long-lived agent's problem is to maximize $\mathcal{G}(\cdot, x_0)$. Assume that it admits a solution. Let $\mathbf{x}' = (x'_t)_{t=1}^\infty$ be a solution satisfying $\underline{x} \leq x'_t \leq \bar{x}$ in every period t ; such a solution exists by Theorem 3. Define $X_t = \tilde{x}_t \vee x'_t$ for each $t \in \mathbb{N}$, and $X_0 = \underline{x}$.

Write $\mathbf{x}^T = R_T R_{T-1} \dots R_2 R_1 R_0 \mathbf{x}'$ for $T \in \{0, 1, 2, \dots\}$. The sequence \mathbf{x}^T has t th entry X_t for $t \leq T$ and $\tilde{x}_T \vee x_t$ for $t > T$, since $(\tilde{x}_t)_{t=1}^\infty$ is increasing. Clearly, $\tilde{x}_t \leq X_t \leq \bar{x}$ for any period $t \in \mathbb{N}$. To prove the theorem, we need only show that $\mathbf{x}^\infty = (X_1, X_2, X_3, \dots)$ maximizes $\mathcal{G}(\cdot, x_0)$.

It suffices to show for each $T \in \{0, 1, 2, \dots\}$ that \mathbf{x}^T maximizes $\mathcal{G}(\cdot, x_0)$. For then, letting V be the long-lived agent's optimal value and noting that both $\mathbf{x}^T = (x_t)_{t=1}^\infty$ and \mathbf{x}^∞ have X_1, \dots, X_T as their first T entries, we have

$$\begin{aligned} 0 &\geq \mathcal{G}(\mathbf{x}^\infty, x_0) - V = \mathcal{G}(\mathbf{x}^\infty, x_0) - \mathcal{G}(\mathbf{x}^T, x_0) \\ &= \delta^T [\mathcal{G}((X_t)_{t=T+1}^\infty, X_T) - \mathcal{G}((x_t)_{t=T+1}^\infty, X_T)]. \end{aligned}$$

By equi-BCS, the right-hand "[\cdot]" is bounded below uniformly over $T \in \mathbb{N}$,⁴⁷ so letting $T \rightarrow \infty$ yields $0 \geq \mathcal{G}(\mathbf{x}^\infty, x_0) - V \geq 0$, meaning that \mathbf{x}^∞ is optimal.

To show that \mathbf{x}^T is optimal for each $T \in \{0, 1, 2, \dots\}$, we employ induction on $T \in \{0, 1, 2, \dots\}$. The base case $T = 0$ is immediate, since $\mathbf{x}^0 = R_0 \mathbf{x}' = \mathbf{x}'$ is optimal.

⁴⁷Since X_t and x_t belong to the compact set $[\underline{x}, \bar{x}]$ for every $t \in \mathbb{N}$, there are constants $A, B > 0$ such that $F(X_t, \theta_t) - F(x_t, \theta_t) \geq -2A$ and $-[C(X_t - X_{t-1}) + C(x_t - x_{t-1})] \geq -2B$ for all $t \in \mathbb{N}$, so the right-hand "[\cdot]" is bounded below by $-2(A+B)/(1-\delta)$.

For the induction step, fix any $T \in \mathbb{N}$, and suppose that $\mathbf{x}^{T-1} = (x_t)_{t=1}^\infty$ is optimal; we will show that $\mathbf{x}^T = R_T \mathbf{x}^{T-1}$ is also optimal. Since \mathbf{x}^{T-1} and \mathbf{x}^T have the same first $T-1$ entries (namely, X_1, X_2, \dots, X_{T-1}), and since for $t \geq T$ the t th entry of \mathbf{x}^T is $\tilde{x}_T \vee x_t$, it suffices to show that

$$G_T(X_T, X_{T-1}) \geq G_T(x_T, X_{T-1}), \quad (6)$$

$$G_t(\tilde{x}_T \vee x_t, \tilde{x}_T \vee x_{t-1}) \geq G_t(x_t, x_{t-1}) \quad \text{for all } t \geq T+1. \quad (7)$$

For (6), since C_T is convex and $X_T \geq x_T$ and $X_{T-1} \geq \tilde{x}_{T-1}$, we have

$$C_T(X_T - X_{T-1}) - C_T(x_T - X_{T-1}) \leq C_T(X_T - \tilde{x}_{T-1}) - C_T(x_T - \tilde{x}_{T-1}).$$

It follows that

$$\begin{aligned} G_T(X_T, X_{T-1}) - G_T(x_T, X_{T-1}) &\geq G_T(X_T, \tilde{x}_{T-1}) - G_T(x_T, \tilde{x}_{T-1}) \\ &\geq G_T(\tilde{x}_T, \tilde{x}_{T-1}) - G_T(\tilde{x}_T \wedge x_T, \tilde{x}_{T-1}) \geq 0, \end{aligned}$$

where the second inequality holds since

$$\begin{aligned} F(X_T, \theta_T) - F(x_T, \theta_T) &\geq F(\tilde{x}_T, \theta_T) - F(\tilde{x}_T \wedge x_T, \theta_T) \quad \text{and} \\ C_T(X_T - \tilde{x}_{T-1}) - C_T(x_T - \tilde{x}_{T-1}) &= C_T(\tilde{x}_T - \tilde{x}_{T-1}) - C_T(\tilde{x}_T \wedge x_T - \tilde{x}_{T-1}) \end{aligned}$$

by the supermodularity of $F(\cdot, \theta_T)$ and the additive separability of C_T , and the final inequality holds since \tilde{x}_T maximizes $G_T(\cdot, \tilde{x}_{T-1})$ on L by definition.

It remains to establish (7). Fix an arbitrary $t \geq T+1$. It suffices to show that $F(\tilde{x}_T \vee x_t, \theta_t) \geq F(x_t, \theta_t)$ and $C_t(\tilde{x}_T \vee x_t - \tilde{x}_T \vee x_{t-1}) \leq C_t(x_t - x_{t-1})$. The latter holds by Lemma 1 (Appendix G) since C_t is monotone. To show the former, begin by noting that $G_T(\tilde{x}_T, \tilde{x}_{T-1}) \geq G_T(\tilde{x}_T \wedge x_t, \tilde{x}_{T-1})$ since \tilde{x}_T maximizes $G_T(\cdot, \tilde{x}_{T-1})$ on L by definition. We have $C_T(\tilde{x}_T - \tilde{x}_{T-1}) \geq C_T(\tilde{x}_T \wedge x_t - \tilde{x}_{T-1})$ by Corollary 1 (Appendix G) since $\tilde{x}_T \geq \tilde{x}_{T-1}$. It follows that $F(\tilde{x}_T, \theta_T) \geq F(\tilde{x}_T \wedge x_t, \theta_T)$. Thus, $F(\tilde{x}_T \vee x_t, \theta_T) \geq F(x_t, \theta_T)$ by supermodularity, whence $F(\tilde{x}_T \vee x_t, \theta_t) \geq F(x_t, \theta_t)$ by single-crossing differences and the fact that $\theta_t \geq \theta_T$ (since $t > T$ and $(\theta_s)_{s=1}^\infty$ is increasing). Q.E.D.

APPENDIX M: PROOF OF THEOREM 2[†]

The proof of Theorem 2 (Section 4) shows that (a) implies (b). To show that (b) implies (a), we prove the contrapositive. If C fails to be minimally monotone, then (b) fails by Theorem 1[†]. Suppose instead that there are adjustment vectors $\varepsilon, \varepsilon' \in \Delta L$ such that $C(\varepsilon') > C(\varepsilon)$ and either $0 \leq \varepsilon' \leq \varepsilon$ or $0 \geq \varepsilon' \geq \varepsilon$; we must show that (b) fails. Assume that $0 \leq \varepsilon' \leq \varepsilon$; the other case is analogous. Let $\alpha = C(\varepsilon') - C(\varepsilon) > 0$ and $\beta = \alpha + \max\{0, C(\varepsilon) - C(0)\}$. Note that $\varepsilon' < \varepsilon$ (else $\varepsilon' = \varepsilon$, which would imply $\alpha = 0$).

Choose $\underline{x}, \bar{x}, \hat{x} \in L$ such that $\bar{x} - \underline{x} = \varepsilon$ and $\bar{x} - \underline{x} = \varepsilon'$, and note that $\underline{x} \leq \bar{x} < \hat{x}$. Let $X = \{\underline{x}, \bar{x}, \hat{x}\}$; X is a sublattice. Fix distinct $\underline{\theta} \leq \bar{\theta}$ in Θ (possible by hypothesis). Let $F(\underline{x}, \underline{\theta}) = 2$, $F(\bar{x}, \underline{\theta}) = 1$, and $F(\hat{x}, \underline{\theta}) = 0$, and let $F(\underline{x}, \bar{\theta}) = 0$, $F(\bar{x}, \bar{\theta}) = \beta$, and $F(\hat{x}, \bar{\theta}) = \beta - \alpha/2$.⁴⁸ Then $F(x, \theta)$ is supermodular in x and has single-crossing differences in (x, θ) , and we have

$$\arg \max_{x \in X} F(x, \underline{\theta}) = \{\underline{x}\}, \quad \arg \max_{x \in X} G(x, \bar{\theta}) = \{\hat{x}\}, \quad \arg \max_{x \in X} F(x, \bar{\theta}) = \{\bar{x}\}.$$

Hence, (b) fails, as $\bar{\theta} \geq \underline{\theta}$ and $\hat{x} > \bar{x}$. Q.E.D.

⁴⁸For $\theta \in \Theta \setminus \{\underline{\theta}, \bar{\theta}\}$, let $F(\cdot, \theta) = F(\cdot, \underline{\theta})$ if $\theta \leq \underline{\theta}$ and $F(\cdot, \theta) = F(\cdot, \bar{\theta})$ otherwise.

APPENDIX N: PROOF OF THEOREM 3[†]

Part (a) implies (b) by Theorem 3. To show that (b) implies (a), we prove the contrapositive: suppose that there is a $C \in \mathcal{C}$ that is not monotone, meaning that there are adjustment vectors $\varepsilon, \varepsilon' \in \Delta L$ such that $C(\varepsilon') > C(\varepsilon)$ and for each dimension i , either $0 \leq \varepsilon'_i \leq \varepsilon_i$ or $0 \geq \varepsilon'_i \geq \varepsilon_i$; we will show that (b) fails.

Suppose first that $\varepsilon \geq 0$ or $\varepsilon \leq 0$. Then C fails to be *weakly* monotone, so by Theorem 2[†], there exist a sublattice $X \subseteq L$, an objective $\tilde{F} : X \times \Theta \rightarrow \mathbb{R}$ such that $\tilde{F}(x, \theta)$ is quasi-supermodular in x and has single-crossing differences in (x, θ) , parameters $\underline{\theta}, \bar{\theta} \in \Theta$, and actions $\underline{x} \in \arg\max_{x \in X} \tilde{F}(x, \underline{\theta})$ and $\bar{x} \in \arg\max_{x \in X} \tilde{F}(x, \bar{\theta})$, such that $M = \arg\max_{x \in X} [\tilde{F}(x, \bar{\theta}) - C(x - \underline{x})]$ is nonempty and either (i) $\underline{\theta} \leq \bar{\theta}$, $\underline{x} \leq \bar{x}$ and $M \cap \{x \in X : \underline{x} \leq x \leq \bar{x}\} = \emptyset$ or (ii) $\underline{\theta} \geq \bar{\theta}$, $\underline{x} \geq \bar{x}$ and $M \cap \{x \in X : \underline{x} \geq x \geq \bar{x}\} = \emptyset$. Let $F = (1 - \delta)\tilde{F}$, $C_1 = C$, $C_t = \bar{C}$ for every $t \geq 2$, and $\theta_t = \bar{\theta}$ for every $t \in \mathbb{N}$. Then the long-lived agent's problem admits a solution, and every solution $(x_t)_{t=1}^\infty$ has $x_1 \in M$, so (b) fails. Assume for the remainder that $0 \not\leq \varepsilon \not\leq 0$.

Note that ε may satisfy $\varepsilon \wedge \varepsilon' = \varepsilon'$ or $\varepsilon \vee \varepsilon' = \varepsilon'$, but not both (else $\varepsilon' = \varepsilon$, which would imply $C(\varepsilon') = C(\varepsilon)$). Hence, $E = \{\varepsilon \wedge \varepsilon', \varepsilon, \varepsilon', \varepsilon \vee \varepsilon'\}$ has either two or four elements, and $(E \setminus \{\varepsilon'\}) \cap \{\tilde{\varepsilon} \in \Delta L : \varepsilon' \wedge 0 \leq \tilde{\varepsilon} \leq \varepsilon' \vee 0\} = \emptyset$.⁴⁹

Choose $\hat{x} \in L$ such that $\hat{x} + \varepsilon \in L$. Let $X = \{\hat{x}\} + (E \cup \{0, \varepsilon \wedge 0, \varepsilon \vee 0, \varepsilon' \wedge 0, \varepsilon' \vee 0\})$, where “+” denotes Minkowski addition;⁵⁰ X is a sublattice. Let $\underline{x} = \hat{x} + \varepsilon' \wedge 0 \in X$ and $\bar{x} = \hat{x} + \varepsilon' \vee 0 \in X$, and note that

$$((\{\hat{x}\} + E) \setminus \{\hat{x} + \varepsilon'\}) \cap \{x \in X : \underline{x} \leq x \leq \bar{x}\} = \emptyset. \quad (8)$$

Arrange the elements of X in a matrix as

$$\hat{x} + \begin{pmatrix} \varepsilon & \varepsilon \vee \varepsilon' & \varepsilon \vee 0 \\ \varepsilon \wedge \varepsilon' & \varepsilon' & \varepsilon' \vee 0 \\ \varepsilon \wedge 0 & \varepsilon' \wedge 0 & 0 \end{pmatrix} = \begin{pmatrix} \hat{x} + \varepsilon & \hat{x} + \varepsilon \vee \varepsilon' & \hat{x} + \varepsilon \vee 0 \\ \hat{x} + \varepsilon \wedge \varepsilon' & \hat{x} + \varepsilon' & \bar{x} \\ \hat{x} + \varepsilon \wedge 0 & \underline{x} & \hat{x} \end{pmatrix}.$$

The top-left 2×2 submatrix is $\{\hat{x}\} + E$. Here and below, ignore the leftmost column if $\varepsilon \wedge \varepsilon' = \varepsilon'$, since then its entries equal those of the middle column, and similarly ignore the top row if $\varepsilon \vee \varepsilon' = \varepsilon'$, the rightmost column if $\varepsilon' \wedge 0 = 0$, and the bottom row if $\varepsilon' \vee 0 = 0$.

Fix distinct $\underline{\theta}, \theta_1, \theta_2, \bar{\theta} \in \Theta$ such that $\underline{\theta} \leq \theta_1 \leq \bar{\theta} \geq \theta_2 \geq \underline{\theta}$ and $\theta_1 \not\leq \theta_2 \not\leq \theta_1$ (possible by hypothesis). Using the matrix notation from above, let

$$F(\cdot, \underline{\theta}) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 2 & 1 \\ 2 & 3 & 2 \end{pmatrix}, \quad F(\cdot, \theta_1) = \gamma \begin{pmatrix} 0 & 1 & 2 \\ 1 & 2 & 3 \\ 2 & 3 & 4 \end{pmatrix}, \quad F(\cdot, \bar{\theta}) = \begin{pmatrix} 0 & 1 & 2 \\ 1 & 2 & 3 \\ 0 & 1 & 2 \end{pmatrix},$$

$$\text{and } F(\cdot, \theta_2) = (1 - \delta) \begin{pmatrix} 0 & \alpha/3 & -\beta \\ \alpha/3 & 2\alpha/3 & \alpha/3 - \beta \\ -\beta & \alpha/3 - \beta & -2\beta \end{pmatrix},$$

where $\alpha = C(\varepsilon') - C(\varepsilon) > 0$, $\beta = 2[\max_{\tilde{\varepsilon} \in \Delta X} C(\tilde{\varepsilon}) - \min_{\tilde{\varepsilon} \in \Delta X} C(\tilde{\varepsilon})] \geq 2\alpha$, and $\gamma = 3\beta$.⁵¹ Then $F(x, \theta)$ is supermodular in x and has single-crossing differences in (x, θ) , and we

⁴⁹If $\varepsilon \wedge \varepsilon' \neq \varepsilon'$, then $0 \geq \varepsilon'_i > \varepsilon_i$ for some dimension i , so $(\varepsilon \wedge \varepsilon')_i = \varepsilon_i < \varepsilon'_i = (\varepsilon' \wedge 0)_i$, whence $\varepsilon \wedge \varepsilon' \not\leq \varepsilon' \wedge 0 \not\leq \varepsilon$. Similarly, if $\varepsilon \vee \varepsilon' \neq \varepsilon'$, then $\varepsilon \vee \varepsilon' \not\leq \varepsilon' \vee 0 \not\leq \varepsilon$.

⁵⁰That is, $A + B = \{a + b : a \in A, b \in B\}$ for any sets $A, B \subseteq \mathbb{R}^n$.

⁵¹For $\theta \in \Theta \setminus \{\underline{\theta}, \theta_1, \theta_2, \bar{\theta}\}$, let $F(\cdot, \theta) = F(\cdot, \underline{\theta})$ if $\theta_1 \leq \theta \leq \theta_2$, $F(\cdot, \theta) = F(\cdot, \theta_1)$ if $\theta \leq \theta_1$ and $\theta_2 \not\leq \theta \not\leq \theta_2$, $F(\cdot, \theta) = F(\cdot, \theta_2)$ if $\theta \leq \theta_2$ and $\theta_1 \not\leq \theta \not\leq \theta_1$, and $F(\cdot, \theta) = F(\cdot, \bar{\theta})$ if $\theta > \theta_1$ or $\theta > \theta_2$.

have

$$\arg \max_{x \in X} F(x, \underline{\theta}) = \{\underline{x}\}, \quad \arg \max_{x \in X} F(x, \theta_1) = \{\hat{x}\}, \quad \arg \max_{x \in X} F(x, \bar{\theta}) = \{\bar{x}\},$$

and

$$\arg \max_{x \in X} \left[\left(\sum_{t=2}^{\infty} \delta^{t-1} \right) F(x, \theta_2) - \delta C(x - \hat{x}) \right] \subseteq (\{\hat{x}\} + E) \setminus \{\hat{x} + \varepsilon'\}.$$

The final inclusion holds because β is large enough that $\hat{x} + \varepsilon$ is strictly better than every $x \in X \setminus (\{\hat{x}\} + E)$ and because $\hat{x} + \varepsilon$ is strictly better than $\hat{x} + \varepsilon'$.

Let $C_1 = \underline{C}$, $C_2 = C$, and $C_t = \bar{C}$ and $\theta_t = \theta_2$ for every $t \geq 3$. Then the long-lived agent's problem admits a solution, and every solution $(x_t)_{t=1}^{\infty}$ has $x_1 = \hat{x}$ and $x_2 \subseteq (\{\hat{x}\} + E) \setminus \{\hat{x} + \varepsilon'\}$. Here $x_1 = \hat{x}$ holds since γ is large enough that it is optimal to choose myopically in period $t = 1$. Hence, (b) fails, as $\underline{\theta} \leq \theta_t \leq \bar{\theta}$ for every $t \in \mathbb{N}$ and $x_2 \notin \{x \in X : \underline{x} \leq x \leq \bar{x}\}$ by (8). Q.E.D.

REFERENCES

- ANTIĆ, NEMANJA, AND YUVAL SALANT (2020): "Monopoly Theory With Price-Increase Aversion," Work in progress. [0673]
- ATHEY, SUSAN (2002): "Monotone Comparative Statics Under Uncertainty," *Quarterly Journal of Economics*, 117 (1), 187–223. [0666]
- BACHMANN, RÜDIGER, RICARDO J. CABALLERO, AND EDUARDO M. R. A. ENGEL (2013): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *American Economic Journal: Macroeconomics*, 5 (4), 29–67. [0678]
- BACK, KERRY (1992): "Insider Trading in Continuous Time," *Review of Financial Studies*, 5 (3), 387–409. [0661]
- BARI, FAISAL, KASHIF MALIK, MUHAMMAD MEKI, AND SIMON QUINN (2024): "Asset-Based Microfinance for Microenterprises: Evidence From Pakistan," *American Economic Review*, 114 (2), 534–574. [0678]
- BÉNABOU, ROLAND (2015): "The Economics of Motivated Beliefs," *Revue d'économie politique*, 125 (5), 665–685. [0669]
- BÉNABOU, ROLAND, AND JEAN TIROLE (2016): "Mindful Economics: The Production, Consumption, and Value of Beliefs," *Journal of Economic Perspectives*, 30 (3), 141–164. [0669]
- BENTOLILA, SAMUEL, AND GIUSEPPE BERTOLA (1990): "Firing Costs and Labour Demand: How Bad Is Eursclerosis?" *Review of Economic Studies*, 57 (3), 381–402. [0661]
- BERGER, DAVID, AND JOSEPH VAVRA (2015): "Consumption Dynamics During Recessions," *Econometrica*, 83 (1), 101–154. [0661]
- CAPLIN, ANDREW, AND JOHN LEAHY (2019): "Wishful Thinking," NBER working paper 25707. [0669,0670]
- CAPLIN, ANDREW, AND DANIEL F. SPULBER (1987): "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102 (4), 703–725. [0661,0673,0677]
- CHETTY, RAJ (2012): "Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80 (3), 969–1018. [0661,0677]
- CHETTY, RAJ, AND ADAM SZEIDL (2016): "Consumption Commitments and Habit Formation," *Econometrica*, 84 (2), 855–890. [0661]
- CHETTY, RAJ, JOHN N. FRIEDMAN, TORE OLSEN, AND LUIGI PISTAFERRI (2011): "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence From Danish Tax Records," *Quarterly Journal of Economics*, 126 (2), 749–804. [0661,0677]
- COOPER, RUSSELL W., AND JOHN C. HALTIWANGER (2006): "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 73 (3), 611–633. [0661,0678]
- DZIEWULSKI, PAWEŁ, AND JOHN K.-H. QUAH (2024): "Comparative Statics With Linear Objectives: Normality, Complementarity, and Ranking Multi-Prior Beliefs," *Econometrica*, 92 (1), 167–200. [0670]
- FIELD, ERICA, ROHINI PANDE, JOHN PAPP, AND NATALIA RIGOL (2013): "Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence From India," *Journal of Political Economy*, 103 (6), 2196–2226. [0678]
- GALI, JORDI (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications* (Second Ed.). Princeton, NJ: Princeton University Press. [0677]

- GOLOSOV, MIKHAIL, AND ROBERT E. JR. LUCAS (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115 (2), 171–199. [0661,0673,0677]
- GROSSMAN, SANFORD J., AND GUY LAROQUE (1990): "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58 (1), 25–51. [0661]
- HAMERMESH, DANIEL S. (1988): "Labor Demand and the Structure of Adjustment Costs," NBER working paper 2572. [0661]
- HAYASHI, FUMIO (1982): "Tobin's Marginal q and Average q : A Neoclassical Interpretation," *Econometrica*, 50 (1), 213–224. [0661,0678]
- HOPENHAYN, HUGO A., AND EDWARD C. PRESCOTT (1992): "Stochastic Monotonicity and Stationary Distributions for Dynamic Economies," *Econometrica*, 60 (6), 1387–1406. [0676]
- JORGENSON, DALE W. (1963): "Capital Theory and Investment Behavior," *American Economic Review*, 53 (2), 247–259. [0661,0678]
- KAPLAN, GREG, AND GIOVANNI L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82 (4), 1199–1239. [0661]
- KARLIN, SAMUEL, AND HERMAN RUBIN (1956): "The Theory of Decision Procedures for Distributions With Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 27 (2), 272–299. [0666]
- KYLE, ALBERT S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53 (6), 1315–1335. [0661]
- LEVY, DANIEL, MARK BERGEN, SHANTANU DUTTA, AND ROBERT VENABLE (1997): "The Magnitude of Menu Costs: Direct Evidence From Large U.S. Supermarket Chains," *Quarterly Journal of Economics*, 112 (3), 791–824. [0673]
- MANKIW, N. GREGORY (1985): "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly," *Quarterly Journal of Economics*, 100 (2), 529–538. [0661,0673]
- MIDRIGAN, VIRGILIU (2011): "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, 79 (4), 1139–1180. [0661,0673,0677]
- MILGROM, PAUL, AND JOHN ROBERTS (1990): "Rationalizability, Learning, and Equilibrium in Games With Strategic Complementarities," *Econometrica*, 58 (6), 1255–1277. [0673,0674]
- (1996): "The LeChatelier Principle," *American Economic Review*, 86 (1), 173–179. [0663,0671,0673]
- MILGROM, PAUL, AND CHRIS SHANNON (1994): "Monotone Comparative Statics," *Econometrica*, 62 (1), 157–180. [0662,0666,0668–0671,0674,0676,0680]
- QUAH, JOHN K.-H., AND BRUNO STRULOVICI (2012): "Aggregating the Single Crossing Property," *Econometrica*, 80 (5), 2333–2348. [0666]
- ROTEMBERG, JULIO J. (1982): "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies*, 49 (4), 517–531. [0677]
- SAMUELSON, PAUL A. (1947): *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press. [0662]
- THOMAS, JULIA K. (2002): "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy*, 110 (3), 508–534. [0678]
- TOBIN, JAMES (1969): "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit and Banking*, 1 (1), 15–29. [0678]
- TOPKIS, DONALD M. (1998): *Supermodularity and Complementarity*. Princeton, NJ: Princeton University Press. [0670]
- WINBERRY, THOMAS (2021): "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 111 (1), 364–396. [0678]

Co-editor Marina Halac handled this manuscript.

Manuscript received 11 March, 2024; final version accepted 13 December, 2024; available online 16 December, 2024.