Theoretical Economics 14 (2019), 1435–1482 1555-7561/20191435

Optimal dynamic contracting: The first-order approach and beyond

Marco Battaglini
Department of Economics, Cornell University

Rohit Lamba
Department of Economics, Pennsylvania State University

We explore the conditions under which the "first-order approach" (FO approach) can be used to characterize profit maximizing contracts in dynamic principal–agent models. The FO approach works when the resulting FO-optimal contract satisfies a particularly strong form of monotonicity in types, a condition that is satisfied in most of the solved examples studied in the literature. The main result of our paper is to show that except for nongeneric choices of the stochastic process governing the types' evolution, monotonicity and, more generally, incentive compatibility are necessarily violated by the FO-optimal contract if the frequency of interactions is sufficiently high (or, equivalently, if the discount factor, time horizon, and persistence in types are sufficiently large). This suggests that the applicability of the FO approach is problematic in environments in which expected continuation values are important relative to per period payoffs. We also present conditions under which a class of incentive compatible contracts that can be easily characterized is approximately optimal.

Keywords. Contract theory, dynamic contracts.

JEL classification. D82, D86.

1. Introduction

Most contractual relationships have a dynamic nature, involving long-term, non-anonymous interactions between a principal and an agent. Examples of these contractual relationships include income taxation, regulation, managerial compensation, and

Marco Battaglini: battaglini@cornell.edu
Rohit Lamba: rlamba@psu.edu
For useful comments and discussions we thank the three anonymous referees, Dilip Abreu, Dirk Bergemann, Carlo Cabrera, Sylvain Chassang, Ilia Krasikov, Stephen Morris, Wolfgang Pesendorfer, Roland Strausz, Juuso Valimaki, and seminar participants at the Einaudi Institute for Economics and Finance, New York University, Princeton University, University of Toronto, Yale University, University of Cambridge, the 9th Csef-Igier Symposium on Economics and Institutions, and the Tenth World Congress of the Econometric Society. We thank Nemanja Antic, Shomak Chakrabarti, Edoardo Grillo, and Derek Lougee for outstanding research assistance. Battaglini gratefully acknowledges the hospitality of the Einaudi Institute for Economics and Finance, and Lamba gratefully acknowledges the time spent at Cambridge INET Institute while working on this paper. An earlier version of this paper forms Chapter 3 of Lamba's Ph.D. dissertation at Princeton, and he is especially grateful to Stephen Morris for his guidance.

Available at http://econtheory.org. https://doi.org/10.3982/TE2355

a monopolist repeatedly selling a nondurable good to a buyer. In these environments, contracts can be made contingent on past realizations of the agent's type, allowing the principal to use the agent's revealed preferences to screen future types' realizations. This may be particularly useful in limiting asymmetric information and agency problems when the type is persistent over time.

Despite recent advances in contract theory, there is still a limited understanding about how to use this information to design optimal contracts. Dynamic contracts are difficult to study because they involve a large number of incentive compatibility constraints. The analysis has, therefore, been limited to economic environments in which a form of the "first-order approach" can be applied; that is, environments in which the optimal contract can be fully characterized using only the necessary conditions implied by local incentive compatibility constraints. While the first-order approach can be generally applied in static environments under standard regularity assumptions, in dynamic models, local incentive compatibility constraints have been shown to be sufficient only in certain specific economic environments.[1]

This leaves three sets of open questions. First, what is the general applicability of the first-order approach and what are its implications? Second, in environments in which the first-order approach does not hold, what do the optimal contracts look like? Finally, if characterizing the optimal contracts is complicated, can we approximate the optimal contracts with simpler contracts that guarantee a minimal loss in profits?

To address these questions, we consider a simple principal–agent model in which a monopolist repeatedly sells a nondurable good to a buyer. The "type" of the buyer, which parametrizes her utility, is private information, and it evolves over time according to a general N-state Markov process. Higher types are assumed to have higher marginal valuations, and their associated conditional distributions on future types first-order stochastically dominate the conditional distributions of lower types.

We start our analysis by exploring the applicability of the first-order approach. We show that if we ignore global incentive compatibility constraints, necessary local incentive constraints allow us to state a "dynamic envelope theorem" with discrete types through which the agent's equilibrium rent can be expressed just as a function of the expected allocation. The dynamic envelope theorem allows for a simple characterization of the profit maximizing contract. In keeping with the terminology of the static literature, this contract is referred to as the first-order-optimal (FO-optimal) contract. While this is not a completely new finding, as similar "envelope conditions" have been characterized in the literature, it allows us to characterize necessary and sufficient conditions for the applicability of the first-order (FO) approach in general environments with finite Markovian types.[2] These conditions tend to be very complicated functions of the underlying parameters of the problem. An easily verifiable sufficient (but not necessary) condition for the FO approach to work is that the associated FO-optimal contract

[1]We discuss the literature in greater detail in Section 7.

[2]Specifically, our formula is a straightforward generalization to N types of the formula in Battaglini (2005). For continuous types, see derivations by Baron and Besanko (1984), Besanko (1985), Laffont and Tirole (1996), Courty and Li (2000), and, more recently, Pavan et al. (2014).

satisfies a simple form of *monotonicity* that puts a natural partial order on the set of histories of types. We show that the solved examples that have been used in the literature to motivate the use of the FO approach satisfy this condition.

The main result of our paper is to show that, for a generic choice of the stochastic process governing the evolution of types, incentive compatibility is necessarily violated by the FO-optimal contract if the frequency of the principal–agent interactions is sufficiently high (or, equivalently, if the discount factor, time horizon, and types' persistence are sufficiently large). This suggests that the applicability of the FO approach is problematic in environments in which expected continuation values are important relative to per period payoffs, both in terms of magnitude and informational content.[3] These findings have significant implications for the literature on dynamic contracts, which has recently been applied to study problems ranging from optimal pricing to dynamic taxation and insurance.

If the first-order approach is not valid, what can a researcher do? A possibility is obviously to deal with all the incentive constraints, though this is, in general, very hard. We take a first step in this direction by characterizing the optimal contract in a simple two-period, three-type environment. While this is just an example, it provides a first look at the shape of the optimal contract in environments in which the FO approach is not applicable. It shows, among other things, that binding global incentive constraints lead to pooling across histories in the dynamic model.

A second possible approach is to give up on characterizing an exactly optimal contract and instead attempt to design an incentive compatible contract that is approximately optimal in environments of interest. We identify a particular class of allocations, which we term *monotonic contracts*, for which the optimal implementable contract can be easily characterized. We show that as the agent's discount factor converges to 1 and/or types converge to being perfectly persistent (independently of the order of the two limits), the average profit in the optimal monotonic contract converges to the average optimal profit. The optimal monotonic contracts is globally incentive compatible and numerical analysis reveal that it performs well even for intermediate value of persistence and discounting.

We proceed as follows. In Section 2, we present the model, and characterize the dynamic envelope formula and the associated first-order optimal contract. In Section 3, we characterize conditions for the validity of the first-order approach. In Section 4, we establish the limits of the first-order approach in the form of an impossibility result for a generic class of primitives. In Section 5, we briefly discuss the three-type, two-period example, which is completely characterized in the Supplemental Material, available in a supplementary file on the journal website, http://econtheory.org/supp/2355/supplement.pdf. In Section 6, we introduce and study monotonic contracts. In Section 7, we provide an overview of the literature. Finally, conclusions are presented in Section 4. Proofs can be found in the Appendices.

[3]Moreover, numerical examples presented here and in Battaglini and Lamba (2015) show that the FO approach is easily violated for intermediate levels of interaction and/or persistence as well.

2. Model

2.1 *Setup*

There are two players: a buyer (or consumer) and a seller (or monopolist). The buyer repeatedly buys a nondurable good from the seller. A consumer of type θ enjoys a per period utility $u(\theta, q) - p$ for q units of the good bought at a price p. In every period, the seller produces the good with a cost function $c(q)$. The utility and cost functions satisfy the usual conditions. The utility function $u(\theta, q)$ is increasing and differentiable in both arguments with $u(\theta, 0) = 0$, it is concave in q, and it satisfies the single-crossing condition.

Assumption 1. *We have $u_{\theta q}(\theta, q) > 0$ for any θ and q.*

The cost function $c(q)$ is increasing, convex, and differentiable with $c'(0) = 0$ and $\lim_{q\to\infty} c'(q) = \infty$. For future reference, let $s(\theta, q) = u(\theta, q) - c(q)$ be the instantaneous surplus generated by a contract that supplies quantity q to a buyer of type θ. In what follows, $s_q(\theta, q)$ and $u_q(\theta, q)$ denote the derivatives with respect to q. To illustrate some of the results, we will repeatedly use the classic version of this model proposed by Mussa and Rosen (1978) in which $u(\theta, q) = \theta q$ and $c(q) = (1/2) \cdot q^2$.

We assume that time is discrete, and that the relationship between the buyer and the seller lasts for $T \geq 2$ periods. The type θ^t evolves over time according to a Markov process. There are $N + 1$ possible types, $\Theta = \{\theta_0, \theta_1, \dots, \theta_N\}$, with $\theta_i - \theta_{i+1} = \Delta\theta > 0$ for any $i = 0, \dots, N - 1$. Let $\mathcal{N} = \{0, 1, 2, \dots, N\}$ denote the set of all indices of types, noting that the indices uniquely identify the types. The probability of next period type being k if the current type is i is given by $f(\theta_k|\theta_i) = f_{ik}$. Let F be the conditional cumulative distribution function (CDF) defined by $F(\theta_j|\theta_i) = \sum_{k=0}^{N-j} f_{i(j+k)}$, with a shorthand F_{ij}. The distribution of next period's types conditional on current type being i today is denoted by $\mathbf{f}_i = (f_{i0}, f_{i1}, \dots, f_{iN})$, where we assume that $\mathbf{f}_i$ has full support (i.e., $f_{ij} > 0$ for any i, j), and $\mathbf{f}_i$ first-order stochastically dominates $\mathbf{f}_j$ for any i and any $j > i$. Given that higher indices imply lower values, first-order stochastic dominance can be stated as follows.

Assumption 2. *We have $F(\theta_j|\theta_i) \leq F(\theta_j|\theta_k)$ for any j and $i \leq k$.*

In each period, the consumer observes the realization of his own type; the seller, in contrast, can only observe past allocations. At date 0 the seller has a prior $\mu = (\mu_0, \dots, \mu_N)$ on the agent's type. For convenience in most of what follows, we assume the prior has full support, so $\mu_i > 0$ for any i.

In static models, standard concavity assumptions on the objective and distributional assumptions such as monotone hazard rate on the prior ensure the validity of the first-order approach; see, for example, Stole (2001). We require the former assumption, but we do not need the latter. Define

$$\Phi(\theta_i, q) = s(\theta_i, q) - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \big[u(\theta_{i-1}, q) - u(\theta_i, q)\big].$$

Assumption 3. *The function* Φ *is concave in* q *for all* θ_i *and has a a unique interior maximum over* q $\forall i$.

This assumption rules out situations in which even in the static model, the optimal solution is the zero-supply corner solution.

In period 1 the seller offers a supply contract to the buyer. The buyer can reject the offer or accept it. In the latter case, the buyer can walk away from the relationship at any time $t \geq 1$ if the expected continuation utility offered by the contract falls below the reservation value $\underline{U} = 0$. In line with the standard model of price discrimination, the monopolist commits to the contract that is offered. The common discount factor is $\delta \in (0, 1)$.[4]

It is easy to show that in this environment a form of the revelation principle is valid (Myerson (1986)), which allows us to consider, without loss of generality, only contracts that depend on the history of type revelations; i.e., the contract can be written as $\langle \mathbf{p}, \mathbf{q} \rangle = (p(\theta^t|h^{t-1}), q(\theta^t|h^{t-1}))_{t=1}^T$, where h^{t-1} and θ^t are, respectively, the public history up to period $t-1$ and the type revealed at time t.[5] In general, h^t can be defined recursively as $h^t = \{h^{t-1}, \theta^t\}$, $h^0 = \varnothing$. The set of possible histories at time t is denoted H^t (for simplicity $H = H^T$). Let κ_t be the mapping that projects the first t elements of a vector. The set of full histories that follow h^t until time t is given by $H(h^t) = \{h \in H | \kappa_t(h) = h^t\}$. It is also useful to define the set $\widehat{H}(h^t) = \{h \in H(h^t) | h_\tau < \theta_0 \ \forall \tau > t\}$. This is the set of histories following h^t in which all realizations after t are lower than θ_0, the highest type.

A strategy for a seller consists of offering a direct mechanism $\langle \mathbf{p}, \mathbf{q} \rangle$ as described above. The strategy of a consumer is, at least potentially, contingent on a richer history $h_A^t = \{h_A^{t-1}, \theta^t, \widehat{\theta}^{t-1}\}$, starting from $h_A^0 = \theta^1$, where θ^t is the actual type every period and $\widehat{\theta}^t$ is the revealed type. For a given contract, a strategy for the consumer is simply a function that maps a history h_A^t into a revealed type: $h_A^t \mapsto s(h_A^t)$.

2.2 *The principal's problem*

The seller's problem consists of choosing a contract $\langle \mathbf{p}, \mathbf{q} \rangle$ that maximizes profits under two sets of constraints: incentive compatibility constraints, which guarantee that an agent of type i does not want to report being of type j after any history h^t, and individual rationality constraints, which guarantee that all types expect to receive at least their reservation utility $\underline{U} = 0$ after any history h^t. Since the choice of prices and quantities corresponds to the choice of utilities and quantities for the buyer, this problem can be conveniently represented as a choice of $\langle \mathbf{U}, \mathbf{q} \rangle = (U(\theta^t|h^{t-1}), q(\theta^t|h^{t-1}))_{t=1}^T$, where $U(\theta^t|h^{t-1})$ is the expected utility of type θ^t after history h^{t-1}.

By the one-shot deviation principle, the incentive compatibility constraint $IC_{i,j}(h^{t-1})$ requires $U(\theta_i|h^{t-1}) \geq U(\theta_j; \theta_i|h^{t-1}) \ \forall \theta_i, \theta_j, h^{t-1}$, where $U(\theta_j; \theta_i|h^{t-1})$ is the

[4]So by "static model," we mean the model described above with $T = 1$ or $\delta = 0$.

[5]Note that the superscript on θ signifies time period and subscript signifies the type: θ_i^t. Often we write just one of them and the other is clear from the context.

expected utility of type θ_i reporting to be of type θ_j at time t after history h^{t-1} and truthfully reporting thereafter. This constraint can easily be rewritten in terms of $\langle \mathbf{U}, \mathbf{q} \rangle$ as

$$U(\theta_i|h^{t-1}) \geq U(\theta_j|h^{t-1}) + \delta \sum_{k=0}^{N} (f_{ik} - f_{jk}) U(\theta_k|h^{t-1}, \theta_j) + u(\theta_i, q(\theta_j|h^{t-1})) - u(\theta_j, q(\theta_j|h^{t-1})).$$

The individual rationality constraint for type i at history h^{t-1}, $IR_i(h^{t-1})$, is a simple nonnegativity constraint,

$$U(\theta_i|h^{t-1}) \geq 0.$$

For future reference, we call *local downward constraints* the incentive constraints that are associated with a deviation to a contiguous lower type (i.e., $IC_{i,i+1}(h^{t-1})$), and *local upward constraints* the incentive constraints that are associated with a deviation to a contiguous higher type (i.e., $IC_{i+1,i}(h^{t-1})$). We refer to all the other constraints as *global*. A contract that satisfies all incentive and individual rationality constraints is said to be *implementable*.

Let $E[S(\mathbf{q})]$ denote the expected discounted surplus across time and types. The monopolist's problem is to maximize expected surplus net of the buyer's expected equilibrium rents:

$$\max_{\langle \mathbf{U}, \mathbf{q} \rangle} \left\{ \begin{array}{c} E[S(\mathbf{q})] - \sum_{i=0}^{N} \mu_i U(\theta_i|h^0) \\ \text{s.t. } \mathbf{q} \geq 0 \text{ and } IR_i(h^{t-1}), IC_{i,j}(h^{t-1}) \\ \text{for any } i, j, t \text{ and } h^{t-1} \in H^{t-1} \end{array} \right\}. \tag{1}$$

This is a standard maximization problem of a concave function under a system of nonlinear constraints. As T and N increase, the number of variables and constraints becomes prohibitively large, making it unclear whether (1) is analytically tractable.

3. The first-order approach

3.1 *Characterization*

The typical approach in the literature to attack problem (1) is first to study a relaxed problem in which only the individual rationality constraint of the lowest type and the local downward constraints $IC_{i,i+1}(h^t)$ are considered. The remaining constraints can be verified ex post after the solution of the relaxed problem has been characterized.

Definition 1. A contract is first-order optimal if it maximizes expected profit under the constraints $IR_N(h^{t-1})$ and $IC_{i,i+1}(h^{t-1})$ $\forall i \in \mathcal{N} \backslash \{N\}$, $\forall h^{t-1} \in H^{t-1}$, $\forall t$.

Interest in FO-optimal contracts is based on the fact that in many environments they coincide with the optimal contracts. Under standard assumptions, the FO-optimal contract coincides with the optimal contract in a static environment, both with finite and continuous type spaces (see, for example, Stole (2001)).[6] This approach has also been used in all papers that have extended the principal–agent model to dynamic environments; for example, the first-order autoregressive environment (Besanko (1985)) and the Markov environment with two types (Battaglini (2005)).[7] A significant applied literature follows this approach, using numerical approximations to verify a sample of the remaining constraints.[8]

It is easy to show that when we consider the relaxed problem, the downward incentive compatibility constraints can be assumed to hold with equality.[9] This allows us to simplify the the optimization problem drastically by eliminating the vector **U**. Define

$$\Delta F(\theta_j|\theta_i) = F(\theta_j|\theta_i) - F(\theta_j|\theta_{i-1})$$

to be the effect on the conditional distribution of a marginal change in type in the previous period. It is important to note that first-order stochastic dominance implies $\Delta F(\theta_j|\theta_i) \geq 0$ for all i and j. Recalling that $\widehat{H}(h^t)$ is the set of histories following h^t in which all realizations after t are lower than θ_0, and representing by h_k the kth element of history h, we have the following characterization of the agent's utility as a function of **q** only.[10]

Lemma 1. *For a FO-optimal contract,*

$$\begin{aligned}
&\frac{U^*(\theta_i|h^{t-1}) - U^*(\theta_{i+1}|h^{t-1})}{\Delta\theta} \\
&\quad = \frac{\displaystyle\int_{\theta_{i+1}}^{\theta_i} u_\theta\big(x, q^*(\theta_{i+1}|h^{t-1})\big)\,dx}{\Delta\theta} \\
&\qquad + \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+1})} \sum_{\tau>t} \delta^{\tau-t} \left[\begin{array}{c} \displaystyle\prod_{k=t+1}^{\tau} \Delta F(\hat{h}_k|\hat{h}_{k-1}) \\ \cdot \dfrac{\displaystyle\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta\big(x, q^*(\hat{h}_\tau|\hat{h}^{\tau-1})\big)\,dx}{\Delta\theta} \end{array} \right]
\end{aligned} \tag{2}$$

for any $i \in \mathcal{N}\backslash\{N\}$, $h^{t-1} \in H^{t-1}$, *and* $t = 1, \ldots, T$.

[6] A sufficient condition for the FO-optimal contract to be optimal in a static environment is that the prior μ satisfies the monotone hazard rate condition and $u_\theta(\theta, q)$ is non-increasing in θ—conditions satisfied, for example, by a uniform prior and $u(\theta, q) = \theta q$. See Stole (2001) for discussion of these results.

[7] Another class of models in which the allocation takes place only in the final period, called sequential screening, has produced interesting cases in which the FO-optimal contract is indeed optimal. See Courty and Li (2000) and the discussion in Section 7.

[8] See Section 7 for a discussion of this literature.

[9] The details of the statements made in this section are formally proven in the Appendix.

[10] To interpret (2), note that, given a history $\hat{h}$, $\hat{h}^{\tau-1} = (\hat{h}_1, \ldots, \hat{h}_{\tau-1})$ is the realized history at time $\tau - 1$, and $\hat{h}_\tau$ is the realization of the type at time τ. It follows that $q(\hat{h}_\tau|\hat{h}^{\tau-1})$ is the quantity at time τ when the realized history is $\hat{h}^{\tau-1}$.

Lemma 1 presents a straightforward dynamic extension of the envelope formula introduced by Myerson (1981). (This can be seen by taking δ to zero, in which case the second term on the right-hand side vanishes and (2) coincides with the classic static formula.[11]) Although the formula is a complicated function of conditional probabilities and allocations, in specific environments it is quite tractable, as we will see in the examples of Section 3.2.

We can express the utility vector solely as a function of $\mathbf{q}$ using Lemma 1. Define

$$U^*(\theta_i|h^{t-1};\mathbf{q}) = \sum_{n=1}^{N-i}\left[\begin{array}{l} \displaystyle\int_{\theta_{i+n}}^{\theta_{i+n-1}} u_\theta\big(x, q^*(\theta_{i+n}|h^{t-1})\big)\,dx \\ \displaystyle + \sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n})}\sum_{\tau>t}\delta^{\tau-t}\left[\begin{array}{l}\displaystyle\prod_{k=t+1}^{\tau}\Delta F(\hat{h}_k|\hat{h}_{k-1}) \\ \displaystyle\cdot\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta\big(x, q^*(\hat{h}_\tau|\hat{h}^{\tau-1})\big)\,dx\end{array}\right]\end{array}\right] \tag{3}$$

for all $i < N$, and define $U^*(\theta_N|h^{t-1};\mathbf{q}) = 0$. Corollary 1 immediately follows from (2).

Corollary 1. *For a FO-optimal contract, we have $U^*(\theta_i|h^{t-1}) = U^*(\theta_i|h^{t-1};\mathbf{q})$ for all $i \in \mathcal{N}$, $h^{t-1} \in H^{t-1}$, and t.*

The FO-optimal contract can now be characterized as the solution of the program

$$\max_{\mathbf{q}\geq 0}\left\{E\big[S(\mathbf{q})\big] - \sum_{i=0}^{N}\mu_i U^*\big(\theta_i|h^0;\mathbf{q}\big)\right\}. \tag{4}$$

This problem can be solved to obtain a closed-form solution. Let $D(h^t)$ be equal to 1 at $t = 1$ and for $t > 1$, define

$$D(h^t) = \begin{cases} 0 & \text{if } h_\tau^t = \theta_0 \text{ for any } \tau \leq t, \\ \displaystyle\prod_{\tau=1}^{t-1}\left(\frac{\Delta F(h_{\tau+1}^t|h_\tau^t)}{f(h_{\tau+1}^t|h_\tau^t)}\right) & \text{otherwise.}\end{cases} \tag{5}$$

These are the dynamic distortions associated with the FO-optimal contract. Recall that for any θ_i, $s(\theta_i, q)$ is the per period surplus (i.e., $u(\theta_i, q) - c(q)$) and $s_q(\theta_i, q)$ is its derivative with respect to q. From the first-order necessary conditions of (4) we can easily characterize the FO-optimal contract as follows.[12]

[11]A continuous-type version of the formula is presented by Baron and Besanko (1984) for the case in which $T = 3$ and by Besanko (1985) for an infinite-horizon model with first-order autoregressive types in which shocks have independent realizations. Battaglini (2005) states the formula for a Markov process with two states; (2) is a direct, but more involved extension of this result for the case with $|\Theta| \geq 2$. Pavan et al. (2014) present a general version of the formula for a continuous-type space and other stochastic processes.

[12]Note that in the following expression, $D(h^{t-1}, \theta_i)$ corresponds to $D(h^t)$ for $h^t = \{h^{t-1}, \theta_i\}$. Also, θ_{-1} is any dummy type.

Proposition 1. *For a FO-optimal contract, we have*

$$s_q\big(\theta_i, q^*\big(\theta_i|h^{t-1}\big)\big) \leq \frac{1-\sum_{k=j}^{N}\mu_k}{\mu_j} \cdot D\big(h^{t-1}, \theta_i\big) \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta q}\big(x, q^*\big(\theta_i|h^{t-1}\big)\big)\,dx \tag{6}$$

for all $i \in \mathcal{N}$, $h^{t-1} \in H^{t-1}$ and t, where $\theta_j = h_1^t$, and the above is satisfied with equality if $q^*(\theta_i|h^{t-1}) > 0$.

It is customary in the literature to assume that the objective function in (4) is concave (see Stole (2001) for example). In this case, (6) is necessary and sufficient, and so it uniquely defines the FO-optimal contract. Although this assumption is not required for the following results, it is always verified if we assume preferences à la Mussa and Rosen (1978); i.e., $u(\theta, q) = \theta q$ and $c(q) = (1/2)q^2$. In this case, at an interior solution,

$$q^*\big(\theta_i|h^{t-1}\big) = \theta_i - \frac{1-\sum_{k=j}^{N}\mu_k}{\mu_j} D\big(h^{t-1}, \theta_i\big)\Delta\theta, \tag{7}$$

where $\theta_j = h_1^t$.

Some distinct characteristics easily emerge from (6) even without assuming that it admits a unique solution. Since the right-hand side of (6) is nonnegative, the contract is always distorted downward, at least weakly. So, analogous to the static case, we never have overprovision, but we can have under provision. Moreover, the right-hand side becomes zero when the type is θ_0, the highest type. In this case the contract is efficient in all following periods, a phenomenon that has been called *generalized no-distortion at the top* (Battaglini (2005)). For any other history, the quantities are distorted strictly below the efficient level. The distortion (or wedge) is exactly equal to $[\sum_{k=0}^{j-1}\mu_k/\mu_i]D(h^{t-1}, \theta_i)\Delta\theta$, which is state contingent and depends on the entire history.

3.2 *When does the first-order approach work?*

Given the (relatively) simple characterization of Proposition 1, the imperative question is, "when is it without loss of generality to focus on the first-order approach?" To verify the validity of the FO approach, we need to establish that the solution of (4) satisfies the full set of constraints in (1). It is not difficult to define a necessary and sufficient condition for the FO approach to work. Let $\Gamma_{ij}(h^{t-1}) = U(\theta_i|h^{t-1}) - U(\theta_j; \theta_i|h^{t-1})$ be the marginal rent received by type i for truthfully reporting her type rather than reporting type j. Naturally, a contract is incentive compatible if and only if these rents are nonnegative. For a FO-optimal contract with $q^* = \{q^*(\theta_i|h^{t-1})\ \forall i, h^{t-1}\}$ that solves problem (4), $\Gamma_{ij}(h^{t-1})$ can be written as

$$\Gamma_{ij}^*\big(h^{t-1}\big) = U^*\big(\theta_i|h^{t-1}; \mathbf{q}^*\big) - U^*\big(\theta_j|h^{t-1}; \mathbf{q}^*\big) - \int_{\theta_j}^{\theta_i} u_\theta\big(x, q^*\big(\theta_j|h^{t-1}\big)\big)\,dx$$

$$-\sum_{k=0}^{N}(f_{ik}-f_{jk})U^*\big(\theta_k|h^{t-1},\theta_j;\mathbf{q}^*\big), \tag{8}$$

where $U^*(\theta_k|h^{t-1},\theta_i;\mathbf{q}^*)$ is the rent associated with the FO-optimal allocation $\mathbf{q}^*$ as given by (3). Note that Γ^* is purely a function of the parameters of the model: the agent's utilities u, the discount factor δ, the types' distribution F, and the prior μ.

Lemma 2. *A FO-optimal contract is incentive compatible if and only if* $\Gamma^*_{i,j}(h^{t-1})\geq 0$ *for all* $i, j\in\mathcal{N}$, $h^{t-1}\in H^{t-1}$, *and* t.

As for all necessary and sufficient conditions for incentive compatibility, Lemma 2 is just a reformulation of the definition, which is useful to the extent that it characterizes a nonempty set of environments and to the extent that it allows us to better understand the features of the parameter space for which incentive compatibility holds. While this typically reduces to requiring a monotone hazard rate in the static model, there is no such hope in the dynamic model. In Section 4, we use Lemma 2 to characterize general properties of the parameter space in which the FO-optimal contract is incentive compatible. We show that when interactions are frequent or when private information is highly persistent and expected continuation values are much greater than per period payoffs, this parameter space is nongeneric.

In the remainder of this section, first we characterize a simpler sufficient condition that can easily be checked, and then apply this condition to the cases for which the FO approach is known to work. Let $q(h^t)=q(h^t_t|h^{t-1})$ be an allocation after history h^t and let $h^t\succeq\widehat{h}^t$ if $h^t_s\geq\widehat{h}^t_s$ $\forall s\leq t$.

Definition 2. An allocation is monotonic if $q(h^t)\geq q(\widehat{h}^t)$ for any $h^t\succeq\widehat{h}^t$.

A simple sufficient condition for the validity of the first-order approach can now be stated.[13]

Proposition 2. *The envelope formula (3) and monotonicity of the FO-optimal contract are sufficient for implementability.*

Proposition 2 directly parallels the well known results in static environments that show that local incentive compatibility (i.e., the envelope formula) and monotonicity of the allocation are necessary and sufficient for implementability. The result is, however, weaker for two reasons: first, the monotonicity condition is stronger than in a static environment, since it compares quantities along all histories as opposed to just the current quantities; second, the result is only sufficient. There are a number of applications in which the FO-optimal contract is indeed monotonic. The advantage of Proposition 2 with respect to Lemma 2 is that it is easy to check.

[13]This result generalizes, to an environment with N types, the method used in Battaglini (2005) to establish the sufficiency of (3) for $N=2$. The proof of Claim 2 in Battaglini (2005) employs a weaker monotonicity condition—the marginal expected utilities are nondecreasing in the current type—and shows that it is sufficient for implementability. Analogous monotonicity results for continuous types and more general stochastic processes are presented by Pavan et al. (2014).

3.3 *Examples*

We now apply Proposition 2 to the environments in which the FO approach has been proven to work and that have been suggested in the literature to motivate its use. These examples illustrate why the approach works when it does; they will also be important in the next section to illustrate why they should be seen as special cases that cannot be easily generalized. For simplicity, in the remainder of this section we assume $u(\theta, q) = \theta q$.

Example 1 (The independent and identically distributed (i.i.d.) case.). When types are i.i.d., we have $f(\theta_i|\theta_j) = f(\theta_i|\theta_k)$ for all i, j, k, so for all histories $\Delta F(\hat{h}_k|\hat{h}_{k-1}) = 0$. Applying (2), we have

$$[U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})]/\Delta\theta = q(\theta_{i+1}|h^{t-1}).$$

We have that type i's expected rent at $t = 1$, $U(\theta_i|h^0)$, depends only on quantities in the first period and is the same as in the static model. The agent has no private information about future realizations beyond period 1 when the contract is signed, so she is unable to extract any rents for $t \geq 2$. From (5) and (7) we can see that when types are i.i.d., it is optimal to offer the optimal static contract in the first period and the efficient contract in all following periods; i.e., $q^*(\theta_i) = \theta_i - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i}\Delta\theta$ and $q^*(\theta_i|h^{t-1}) = \theta_i$ for $t > 1$. The contract is history independent and monotonic in all periods $t > 1$ (since it coincides with the efficient allocation). The contract is also monotonic in the type at $t = 1$ if the optimal static contract is monotonic. We conclude that under standard conditions the FO contract is monotonic and, therefore, optimal for the seller. ◊

Example 2 (Constant types). Assume that types are constant, i.e., $f(\theta_i|\theta_i) = 1$ for all $i = 0, \ldots, N$, as in Baron and Besanko (1984). In this case, after (h^{t-1}, θ_{i+1}), only history $\hat{h} = \{h^{t-1}, \theta_{i+1}, \ldots, \theta_{i+1}\}$ (in which the type remains equal to θ_{i+1}) has positive probability and $\Delta F(\theta_i|\theta_i) = 1$ for all i. Applying (2), it follows that

$$[U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})]/\Delta\theta = \sum_{\tau \geq t} \delta^{\tau-t} \cdot q(\hat{h}_\tau|\hat{h}^{\tau-1})$$

for all $i \in \mathcal{N}\backslash\{N\}$, where $\hat{h} \in H(h^{t-1}, \theta_{i+1})$ is the history that has all realizations following period t equal to θ_{i+1}.

The expected rents are thus a discounted sum of quantities along the constant histories; in particular, $U(\theta_i|h^0)$ is simply a function of quantities along the constant history $\{\theta_{i+1}, \theta_{i+1}, \theta_{i+1}, \ldots\}$. From (7), it follows that when types are constant, it is optimal to offer the same quantities $q^*(\theta_i) = \theta_i - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i}\Delta\theta$ in all periods. To see this, note that for histories in which types remain constant, we have $D(h^{t-1}, \theta_i) = 1$, so (7) is equal to $\theta_i - \frac{1-\sum_{k=i}^N \mu_k}{\mu_i}\Delta\theta$. For histories in which types are not constant, any quantity is optimal. Since these quantities affect neither the surplus nor the rents of the agent, they do not enter the objective function (4).[14] Once again, the FO contract is monotonic and, therefore, it maximizes the seller's expected profit. ◊

[14]In the rest of the paper, we assume that types have full support, so (5) is always well defined. With perfect persistence, for histories in which types change, $D(h^t)$ is indeterminate, since in this case both

Example 3 (Two Markovian types). Assume, as in Battaglini (2005), that there are two types, $\theta_0 = \theta_H$ and $\theta_1 = \theta_L$, evolving according to a stationary Markov process and imperfectly correlated over time. In this case, all histories except the "lowest history" (in which all the type realizations are always θ_L) disappear from (2). Given this, we obtain

$$\big[U(\theta_H|h^{t-1}) - U(\theta_L|h^{t-1})\big]/\Delta\theta = \sum_{\tau\geq t} \delta^{\tau-t} \cdot \big[F(\theta_L|\theta_L) - F(\theta_L|\theta_H)\big]^{\tau-t} \cdot q\big(\widehat{h}_\tau|\widehat{h}^{\tau-1}\big)),$$

where $\widehat{h} = \{h^{t-1}, \theta_L, \dots, \theta_L\}$ is the history following h^{t-1} in which all realizations after $t-1$ are θ_L. In this case the rent of the agent at $t = 1$ depends only on the quantities in the lowest history, in which the realizations are always θ_L. With two types, (7) implies that $q^*(\theta_i|h^{t-1}) = \theta_i$ if $\theta_i = \theta_H$ and/or θ_H is a realization in h^{t-1}. For the remaining history, $\widetilde{h}^{t-1}$, in which the type is always θ_L, we have $q^*(\theta_L|\widetilde{h}^{t-1}) = \theta_L - \frac{\mu_H}{\mu_L}(\frac{F(\theta_L|\theta_L)-F(\theta_L|\theta_H)}{F(\theta_L|\theta_L)})^{t-1}\Delta\theta$. In this case the FO-optimal contract is efficient for all histories except the lowest in which the type is always θ_L. Since $\frac{F(\theta_L|\theta_L)-F(\theta_L|\theta_H)}{F(\theta_L|\theta_L)} < 1$, distortions along the low history vanish as $t \to \infty$. The FO-optimal contract is monotonic and so it maximizes the seller's expected profit. ◊

Example 4 (AR(k) models). Besanko (1985) assumes an AR(1) model in which $\theta^t = \gamma\theta^{t-1} + \varepsilon_t$, where ε_t is the realization of an i.i.d. random variable and $\gamma \in (0, 1)$. While he assumes continuous types, the logic of his results easily extends to our environment if we assume that in the first period, the type is $\theta^1 = \theta_i$ with prior probability μ_i and support $\theta_0, \dots, \theta_N$, and in the subsequent periods, $\varepsilon_t = \theta_j$ with some probability f_j (independent of the history up to t) and support again equal to $\theta_0, \dots, \theta_N$. This AR(1) environment can be seen as a simple extension of the i.i.d. case of Example 1. To see this, note that the evolution of a type can be divided into two components. There is a i.i.d. shock ε_t, for which at time $t-1$ there is symmetric information between the principal and the agent—as in Example 1, this term has no effect on the agent's dynamic distortions. There is then a deterministic component (the term $\gamma\theta^{t-1}$) that affects all following periods in a uniform and perfectly predictable way.[15] It can indeed be verified that the FO-optimal quantity at time t is

$$\theta^t - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \gamma^{t-1}\Delta\theta \tag{9}$$

when θ_i is the realization in the first period. Since (9) is monotonic if $(1 - \sum_{k=i}^N \mu_k)/\mu_i$ is monotonic in i, the FO-optimal contract is then incentive compatible. This case works because distortions depend *only* on the first realization θ_i and the length of time through γ^{t-1}. This type of logic can easily be extended to allow for the kth-order autocorrelation case (i.e., $\theta^t = \sum_{j=0}^k \gamma_j\theta^{t-j} + \varepsilon_t$) or other variations in which, given the information at time $t-1$, the distribution of the shock is independent of the type at time t. ◊

the numerator and the denominator of $D(h^t)$ are zero. These histories occur with zero probability, so the associated quantities are irrelevant.

[15]See Battaglini and Lamba (2015) for the envelope formula in the AR(1) model.

The examples presented above show that the first-order approach can be extended to study quite complex dynamic environments. All the examples, however, can be reduced to two basic assumptions. The environment studied in Besanko (1985) allows for many possible types (in fact a continuum) and arbitrary persistence, but assumes that types change because of linearly additive stochastic shocks uncorrelated with the agent's type. In this environment the shocks are irrelevant for the equilibrium distortions, which are independent of the history of realized types (except for the first). The environment studied in Battaglini (2005) allows the conditional distributions of the future types to depend on the current type, but limits the analysis to two types only. In this case, the optimal contract is history dependent. The two environments have a common feature: the FO-optimal allocation is monotonic. In the next section, however, we show that this is not a general property of FO-optimal contracts.

4. The limits of the first-order approach

Is monotonicity a property to be generally expected for optimal dynamic contracts? Consider an example with two periods, three types $\{\theta_H, \theta_M, \theta_L\}$ with $\theta_H > \theta_M > \theta_L$, and Mussa and Rosen (1978) preferences $u(\theta, q) = \theta q$. Using (7), we have

$$\begin{aligned} q^*(\theta_M|\theta_M) &= \theta_M - \frac{\mu_H}{\mu_M}\frac{F(\theta_M|\theta_M) - F(\theta_M|\theta_H)}{f(\theta_M|\theta_M)}\Delta\theta, \\ q^*(\theta_M|\theta_L) &= \theta_M - \frac{1-\mu_L}{\mu_L}\frac{F(\theta_M|\theta_L) - F(\theta_M|\theta_M)}{f(\theta_M|\theta_L)}\Delta\theta. \end{aligned} \tag{10}$$

There is no reason to expect that $q^*(\theta_M|\theta_M) > q^*(\theta_M|\theta_L)$. For example, if we assume $f(\theta_i|\theta_i) = \alpha$ and $f(\theta_i|\theta_j) = (1-\alpha)/2$ for $i \neq j$ (a simple transition function that satisfies first-order stochastic dominance), we have $q^*(\theta_M|\theta_L) = \theta_M$ but $q^*(\theta_M|\theta_M) = \theta_M - \frac{\mu_H}{\mu_M}\frac{3\alpha-1}{2\alpha}\Delta\theta < q^*(\theta_M|\theta_L)$.

This failure of monotonicity is problematic for incentive compatibility. From (8), we have

$$\begin{aligned} \Gamma^*_{HL}(h^0) &= \Delta\theta\big(q^*(\theta_M) - q^*(\theta_L)\big) + \delta \sum_{k=H,M,L} (f_{H,k} - f_{M,k})\Delta U^*(\theta_k|\theta_M) \\ &= \Delta\theta\big(q^*(\theta_M) - q^*(\theta_L)\big) \\ &\quad + \delta\Delta\theta \sum_{k=M,L} \big[\Delta F(\theta_k|\theta_M)\big(q^*(\theta_k|\theta_M) - q^*(\theta_k|\theta_L)\big)\big]. \end{aligned} \tag{11}$$

From Lemma 2, $\Gamma^*_{HL}(h^0) \geq 0$ is a necessary condition for incentive compatibility.[16] As types become persistent, that is, as $\Delta F(\theta_L|\theta_M) \to 0$ and $\Delta F(\theta_M|\theta_M) \to 1$, we have

$$\Gamma^*_{HL}(h^0) \to \Delta\theta\big(q^*(\theta_M) - q^*(\theta_L)\big) + \delta\Delta\theta\big[\big(q^*(\theta_M|\theta_M) - q^*(\theta_M|\theta_L)\big)\big].$$

[16] The inequality $\Gamma_{HL}(h^0) < 0$ corresponds to a failure of the incentive constraint $IC_{HL}(h^0)$.

It follows that when $q^*(\theta_M|\theta_M) < q^*(\theta_M|\theta_L)$, there is a δ^* (possibly larger than 1) such that for $\delta > \delta^*$, incentive compatibility of the first-order optimal contract fails if types are sufficiently persistent.[17]

In the remainder of this section, we show that the monotonicity problem illustrated in this example is a generic property of the FO approach when the "future" is sufficiently important in the dynamic interaction between the principal and the agent. A natural way to make the future matter, in terms of both private information and payoffs, is to allow the frequency of interactions between the principal and the agent to increase.[18] We therefore extend the model presented above by allowing the length of a period to be τ, where $0 \le \tau \le 1$, and the discount factor to be $\delta = e^{-r\tau}$. The evolution of types is naturally defined in this context if we assume that the stochastic process for θ is derived by sampling, at intervals of length τ, a given continuous-time Markov chain with $N+1$ possible states, $\Theta = \{\theta_0, \theta_1, \dots, \theta_N\}$.[19] The model presented in the previous sections can be seen as a special case of this model in which $\tau = 1$.

The continuous Markov chain is uniquely defined by an exponential transition probability with parameters described by an $N+1$ dimensional row vector $\Lambda = (\lambda_0, \dots, \lambda_n)$ and an $N+1$ dimensional matrix P. Here λ_i is the rate of switching out of the current state i, and P_{ij} is the conditional transition rate of switching to state j from i, with $P_{ii} = 0$ and $\sum_j P_{ij} = 1$ for all i. For a general τ, the transition probabilities sampling the process at τ intervals can be written as

$$f_\tau(\theta_j|\theta_i) = e^{\tau Q} = \sum_{k=0}^{\infty} \frac{\tau^k Q^k}{k!}, \tag{12}$$

where $Q = \Lambda(P - I)$, so $Q_{ii} = -\lambda_i$ and $Q_{ij} = \lambda_i P_{ij}$. Note that $f(\theta_j|\theta_i)$ is simply equal to $f_1(\theta_j|\theta_i)$.[20] To keep notation simple, in what follows we suppress the dependence of f on τ unless explicitly required by the context.

4.1 *Failure of monotonicity in the FO-optimal contract*

To grasp the intuition of why the FO-optimal contract is not generically monotonic, let us consider the simple preference $u(\theta, q) = \theta q$ and focus on two histories $h_i = \{\theta_i, \theta_i\}$ and $h_i' = \{\theta_i, \theta_{i+1}\}$, where $i \in \{0, 1, 2, \dots, N\}$. Assuming an interior solution, from (7), we have

$$q^*(\theta_i|h_i) = \theta_i - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \left[\frac{\Delta F(\theta_i|\theta_i)}{f(\theta_i|\theta_i)} \cdot \frac{\Delta F(\theta_i|\theta_i)}{f(\theta_i|\theta_i)} \right] \Delta\theta. \tag{13}$$

[17]We may need a $\delta^* > 1$ in this simple example only because we are assuming two periods. We show that, in general, for frequent interactions or highly persistent Markov processes, monotonicity and, in turn, incentive compatibility fail for the first-order optimal contract for large enough $\delta < 1$.

[18]Another way to do it is to consider sequences of environments with increasing T, α, and δ. We discuss this equivalent approach in Section 4.3.

[19]Given this, if the total length of interaction is T, we have T/τ periods. Conversely, the length of interaction for t periods is τt.

[20]It is easy to verify that first-order stochastic dominance for the process $f_\tau(\theta_j|\theta_i)$ is preserved for any τ if $f_1(\theta_j; \theta_i)$ satisfies first-order stochastic dominance.

Equation (7) also tells us that after a "mixed" history $h_i' = \{\theta_i, \theta_{i+1}\}$, we have

$$q^*(\theta_i|h_i') = \theta_i - \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \left[\frac{\Delta F(\theta_{i+1}|\theta_i)}{f(\theta_{i+1}|\theta_i)} \cdot \frac{\Delta F(\theta_i|\theta_{i+1})}{f(\theta_i|\theta_{i+1})} \right] \Delta\theta. \tag{14}$$

The only difference between (13) and (14) is the term in the square brackets. The former, as is easy to verify, converges to 1 for frequent interactions or highly persistent types. As we formally prove in the Appendix, however, as $\tau \to 0$, we have

$$\frac{\Delta F(\theta_{i+1}|\theta_i)}{f(\theta_{i+1}|\theta_i)} \cdot \frac{\Delta F(\theta_i|\theta_{i+1})}{f(\theta_i|\theta_{i+1})} \to \frac{\sum_{k=i+1}^{N} [Q_{i,k} - Q_{i-1,k}]}{Q_{i,i+1}} \cdot \frac{\sum_{k=i}^{N} [Q_{i+1,k} - Q_{i,k}]}{Q_{i+1,i}}, \tag{15}$$

where $Q_{ii} = -\lambda_i$ and $Q_{ij} = \lambda_i P_{ij}$ for $i \neq j$. The right-hand side of (15) depends only on the primitives of the stochastic process (Λ, P), and it is easy to see that for a generic choice of (Λ, P), this term is different from 1. Denote D to be the limit of this term as $\tau \to 0$. Suppose first that $D < 1$. In this case, from the discussion above, we have $q^*(\theta_i|h_i) < q^*(\theta_i|h_i')$ for sufficiently frequent interactions. Since $(h_i, \theta_i) \succ (h_i', \theta_i)$, we have a failure of monotonicity. Suppose instead that $D > 1$. In this case, we consider histories $h_{i+1} = \{\theta_{i+1}, \theta_{i+1}\}$ and $h_{i+1}' = \{\theta_{i+1}, \theta_i\}$. We now have

$$q^*(\theta_{i+1}|h_{i+1}') = \theta_{i+1} - \frac{1 - \sum_{k=i+1}^{N} \mu_k}{\mu_{i+1}} \left[\frac{\Delta F(\theta_i|\theta_{i+1})}{f(\theta_i|\theta_{i+1})} \cdot \frac{\Delta F(\theta_{i+1}|\theta_i)}{f(\theta_{i+1}|\theta_i)} \right] \Delta\theta \tag{16}$$

and

$$q^*(\theta_{i+1}|h_{i+1}) = \theta_{i+1} - \frac{1 - \sum_{k=i+1}^{N} \mu_k}{\mu_{i+1}} \left[\frac{\Delta F(\theta_{i+1}|\theta_{i+1})}{f(\theta_{i+1}|\theta_{i+1})} \cdot \frac{\Delta F(\theta_{i+1}|\theta_{i+1})}{f(\theta_{i+1}|\theta_{i+1})} \right] \Delta\theta. \tag{17}$$

As $\tau \to 0$, the term in square brackets in (16) converges to 1, and that in (17) converges to D. Since $D > 1$, we have $q^*(\theta_{i+1}|h_{i+1}) > q^*(\theta_{i+1}|h_{i+1}')$, but $(h_{i+1}', \theta_{i+1}) \succ (h_{i+1}, \theta_{i+1})$, culminating in a failure of monotonicity. We conclude that no matter how we choose D, except for the nongeneric case in which D is *exactly* equal to 1, we can find some history after which monotonicity fails.

On an intuitive level, what happens can be described as follows. From (13), (14), (16), and (17), the ratio between the distortion after a history $\{\theta_i, \theta_{i+1}, \theta_i\}$ and a history $\{\theta_i, \theta_i, \theta_i\}$ is the same as the ratio between the distortion after a history $\{\theta_{i+1}, \theta_i, \theta_{i+1}\}$ and a history $\{\theta_{i+1}, \theta_{i+1}, \theta_{i+1}\}$.[21] This is because from an ex ante perspective, the likelihood of moving from i to $i+1$ and back is the same as moving from $i+1$ to i and back.

[21]That is, formally, $(\theta_i - q^*(\theta_i|h_i'))/(\theta_i - q^*(\theta_i|h_i))$ and $(\theta_{i+1} - q^*(\theta_i|h_{i+1}'))/(\theta_{i+1} - q^*(\theta_{i+1}|h_{i+1}))$ are the same.

Therefore, in terms of the reduction at $t=0$ of the agent's expected rent, the marginal benefits of distortion at $\{\theta_i, \theta_{i+1}, \theta_i\}$ and $\{\theta_{i+1}, \theta_i, \theta_{i+1}\}$ are the same, relative to the cases with the constant histories at, respectively, i and at $i+1$. In both cases, as $\tau \to 0$, the ratio is given by

$$D = \frac{\sum_{k=i+1}^{N} [Q_{i,k} - Q_{i-1,k}]}{Q_{i,i+1}} \cdot \frac{\sum_{k=i}^{N} [Q_{i+1,k} - Q_{i,k}]}{Q_{i+1,i}},$$

which is exogenous and depends only on the shape of the transition matrix of the stochastic process. If $D < 1$, then we have that the distortion after $\{\theta_i, \theta_i, \theta_i\}$ is higher than after $\{\theta_i, \theta_{i+1}, \theta_i\}$, a failure of monotonicity; if $D > 1$, we have that the distortion after $\{\theta_{i+1}, \theta_i, \theta_{i+1}\}$ is higher than after $\{\theta_{i+1}, \theta_{i+1}, \theta_{i+1}\}$, again a failure of monotonicity. This makes it impossible for a generic choice of $Q = \{Q_{i,j}\}_{i\in N, j\in N}$ to satisfy monotonicity: there are too many constraints to be satisfied. This argument is general because it depends neither on the specific types i and $i+1$ nor on the length of the histories that we are considering.

To formalize this observation, we first precisely define what it means for a Markov chain to be generic. Let $\mathcal{Q}$ be the space of continuous-time Markov chains with typical element (Λ, P) such that first-order stochastic dominance is satisfied for all $\tau \le 1$.[22]

Definition 3. A subset $\mathcal{Q}^*$ of the space of Markov chains $\mathcal{Q}$ is said to be generic in $\mathcal{Q}$ if it is open and dense in $\mathcal{Q}$.[23]

This is the standard definition of genericity in our environment.[24] We say that a set $\mathcal{Q}'$ is nongeneric if $\mathcal{Q}\backslash\mathcal{Q}'$ is generic. A property holds for any generic Markov chain if it holds for all $Q \in \mathcal{Q}$, except at most for a nongeneric subset of $\mathcal{Q}$. Our first result proves that for a generic Markov chain, the optimal contract is nonmonotonic when the frequency of interactions between the principal and the agent is sufficiently high.

Proposition 3. *Let $|\Theta| > 2$. For a generic Markov chain Q, there exists a τ^* such that the FO-optimal contract is not monotonic for $\tau \le \tau^*$.*

We stress that the argument of Proposition 3 requires τ to be small only because it aims at a very strong result: the *generic* failure of monotonicity. It is easy to show that a small τ is not at all necessary for the failure of monotonicity. For example, in the three-type example presented above with $f(\theta_i|\theta_i) = \alpha$ and $f(\theta_j|\theta_i) = (1-\alpha)/2$, monotonicity fails for any τ.[25]

[22]The requirement of first-order stochastic dominance is assumed here to be consistent with the previous analysis; it is not necessary for the arguments in Propositions 3 and 4 presented below.

[23]The subset $\mathcal{Q}^*$ is open if for any $Q \in \mathcal{Q}^*$, there exists a $Q' \in \mathcal{Q}^*$ such that Q' is arbitrarily close to Q. The subset $\mathcal{Q}^*$ is dense if for any $Q \in \mathcal{Q}$, that may or may not be in $\mathcal{Q}^*$, we can find a $Q' \in \mathcal{Q}^*$ arbitrarily close to Q.

[24]Note that $\mathcal{Q}$ is a complete metric space. The complement of an open and dense set in $\mathcal{Q}$ is a set of first category. The Baire category theorem guarantees that these sets have empty interiors and, therefore, are topologically small (Royden (1988, Chapter 7.8)).

[25]See the discussion following (10).

4.2 *From monotonicity to incentive compatibility*

A failure of monotonicity of the contract is not in itself sufficient for a failure of incentive compatibility. This is obvious from (11): if δ is close to zero, the problem is essentially a static one, even if $T > 1$. A failure of incentive compatibility, however, becomes inevitable when the future increases in importance relative to current payoffs.

To see this point, consider the three-type example considered above, now assuming $T > 2$ periods. A necessary condition for the FO-optimal contract to be incentive compatible is that $\Gamma^*_{HL}(\theta_M) \geq 0$ (see Lemma 2), which requires that the constraint $IC_{HL}(\theta_M)$ be satisfied. This condition can be written as

$$\Delta\theta\big(q^*(\theta_M|\theta_M) - q^*(\theta_L|\theta_M)\big) + \delta \sum_{k=H,M,L} (f_{H,k} - f_{M,k}) \begin{pmatrix} U^*(\theta_k|\theta_M, \theta_M) \\ -\, U^*(\theta_k|\theta_M, \theta_L) \end{pmatrix} \geq 0. \tag{18}$$

In comparison to (11), we are evaluating the constraint after history $h^1 = \theta_M$, and now the terms $U(\theta_k|\theta_M, \theta_M)$ and $U(\theta_k|\theta_M, \theta_L)$ are expected continuation values, which are typically complicated functions of future quantities and transition probabilities.

To show that (18) is violated, let us normalize payoffs, multiplying (18) by $(1 - e^{-r\tau})$, and denote $\overline{U}(\theta_k|h^{t-1}) = (1 - e^{-r\tau})U(\theta_k|h^{t-1})$ as the average rent.[26] As formally shown in the proof of Proposition 4, as τ decreases, the first term of the normalized left-hand side of (18) and the expected rents along histories in which the type changes (i.e., $f_\tau(\theta_k|\theta_j)\overline{U}(\theta_k|h^{t-1})$) converge to zero and so become irrelevant. For a small τ, the sign of (18) depends on

$$\left[f^\tau_{H,H} \begin{pmatrix} \overline{U}(\theta_H|\theta_M, \theta_M) \\ -\, \overline{U}(\theta_H|\theta_M, \theta_L) \end{pmatrix} - f^\tau_{M,M} \begin{pmatrix} \overline{U}(\theta_M|\theta_M, \theta_M) \\ -\, \overline{U}(\theta_M|\theta_M, \theta_L) \end{pmatrix}\right] + z_\tau,$$

where f^τ_{jk} is a shorthand for $f_\tau(\theta_k|\theta_j)$ and z_τ is a term that converges to zero as $\tau \to 0$. Since $f^\tau_{H,H} \to 1$ and $f^\tau_{M,M} \to 1$ as $\tau \to 0$, we can reshuffle this as

$$\big(\overline{U}(\theta_H|\theta_M, \theta_M) - \overline{U}(\theta_M|\theta_M, \theta_M)\big) - \big(\overline{U}(\theta_H|\theta_M, \theta_L) - \overline{U}(\theta_M|\theta_M, \theta_L)\big) + z_\tau. \tag{19}$$

The first term of (19) is the average rent of type θ_H over θ_M after history (θ_M, θ_M), and the second term is the average rent of type θ_H over type θ_M after history (θ_M, θ_L). The incentive constraint $IC_{HL}(\theta_M)$ is violated if (19) is negative.

To complete our argument, we now make two observations: first, the difference in average rents between types θ_H and θ_M after any history is increasing in the quantities offered thence; second, the contract is generically not monotonic for small τ, so the quantities after (θ_M, θ_M) are smaller than the quantities offered after (θ_M, θ_L).[27] This implies that (19) is generically negative for a sufficiently small τ.

[26]We are indulging in a slight abuse of notation here to keep notation simple. Technically speaking each time period is now divided into $\frac{1}{\tau}$ intervals. So after time t, the principal and agent have interacted $\frac{t}{\tau}$ times. The normalized expected utility should, therefore, be represented by $\overline{U}(\theta_k|h^{\frac{t-1}{\tau}})$.

[27]That is, $q(\theta_M|\theta_M, \theta_L) > q(\theta_M|\theta_M, \theta_M)$ and $q(\theta_i|\theta_M, \theta_L, \theta_M, h^{t-3}) > q(\theta_i|\theta_M, \theta_M, \theta_M, h^{t-3})$. If $q(\theta_M|\theta_M, \theta_L,) < q(\theta_M|\theta_M, \theta_M)$, then we can select two other histories after which quantities are nonmonotonic as in the discussion before Proposition 3.

The formal result now follows. The proof, presented in the Appendix, exploits the aforementioned intuition for the general model.

Proposition 4. *Let $|\Theta| > 2$. For a generic Markov chain Q, there exists a τ^* such that the first-order approach fails to be verified for $\tau \leq \tau^*$.*

Again, as we stated for monotonicity, we require a small τ because we are looking for a very strong result, one that holds generically. The simple example in which $|\Theta| = 3$, the prior is uniform, and the Markov matrix is given by $f_\tau(\theta_i|\theta_i) = \alpha(\tau)$ and $f_\tau(\theta_j|\theta_i) = (1-\alpha(\tau))/2$ illustrates this point.[28] In this case, (18) simplifies to

$$1 - \sum_{t=0}^{\frac{T}{\tau}-2} \delta^t \left(\frac{3\alpha(\tau)-1}{2}\right)^t \left(\frac{3\alpha(\tau)-1}{2\alpha(\tau)}\right)^{t+1}$$

$$= 1 - \frac{3\alpha(\tau)-1}{2\alpha(\tau)} \frac{1-\left(\delta\left(\frac{3\alpha(\tau)-1}{2}\right)^2 \frac{1}{\alpha(\tau)}\right)^{\frac{T}{\tau}-1}}{1-\delta\left(\frac{3\alpha(\tau)-1}{2}\right)^2 \frac{1}{\alpha(\tau)}} \geq 0. \tag{20}$$

We calibrate using $r = 0.105$ (so that $\delta = e^{-r} = 0.9$), $T = 10$, and $\lambda = 1$. The top two quadrants in Figure 1 plot the values of α and $V = \frac{\Gamma^*_{HL}(\theta_M)}{\Delta\theta}$; i.e., the right-hand side of (20) as functions of τ. We can observe that even for moderate frequencies of interactions, $\tau \leq 0.95$ (that implies a level of persistence α greater than 0.64), the constraint $IC_{HL}(\theta_M)$ is violated.

4.3 *Discussion*

We conclude this section with a few remarks on Propositions 3 and 4.

Frequent interactions versus high persistence and long time horizons. In the previous section, we considered economies in which the principal and the agent interact frequently so as to study environments in which future payoffs are important. An alternative is to fix the frequency of interactions (by setting $\tau = 1$, for instance) and consider economies in which types are highly persistent (that is, economies with low λ_is). A similar result can then be established.

Corollary 2. *Let $\lambda = \max \Lambda$. Assume that $\tau = 1$, $|\Theta| > 2$, and $\delta\frac{1-\delta^T}{1-\delta}$ is sufficiently high (i.e., higher than a threshold ζ^*, independent of Λ, and specified in the Appendix). Then, for a generic Markov transition function P, there exists a λ^* such that the first-order approach fails to be verified for $\lambda \leq \lambda^*$.*

[28] These probabilities, $\alpha(\tau)$, are computed using (12) by setting $P_{ij} = 1/2$ for $i \neq j$ and specifying values for λ and τ.

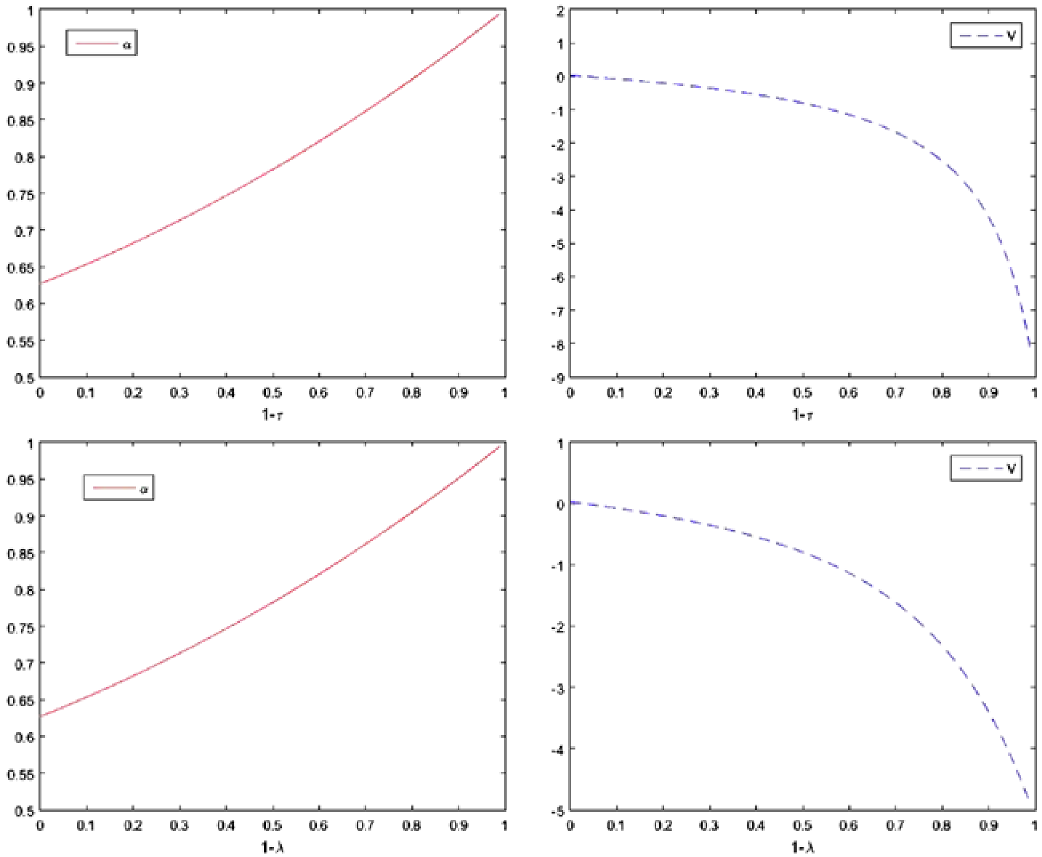

FIGURE 1. The top two quadrants illustrate the change in α and $V = \Gamma^*_{HL}(\theta_M)/\Delta\theta$ as functions of τ in the example of Section 4.2. The bottom two quadrants illustrate the same variables as functions of λ.

This result makes it clear that a high frequency of interactions is not essential for the main result as long as there is a sufficiently large intertemporal horizon and types are sufficiently persistent. These are economically relevant primitives because if $\delta(1-\delta^T)/(1-\delta)$ is small and/or types have low persistence, then the incentive structure is close to that of a static model. For the three-type example mentioned before, fixing $\tau = 1$, $r = 0.105$ (or $\delta \approx 0.9$), and $T = 10$, the bottom two panels in Figure 1 plot α and $\frac{\Gamma^*_{HL}(\theta_M)}{\Delta\theta}$ as functions of λ. It is clear that for even moderately high levels of persistence, $IC_{HL}(\theta_M)$ is violated.

The types of environments with high type persistence described by Corollary 2 are especially important in applications of the theory. Recent empirical evidence shows that in important applications of dynamic principal–agent models (including the study of optimal taxation), the key variable for which agents have private information is highly persistent.[29] These are precisely the environments where our results establish that the use of the first-order approach is particularly problematic.

Perfectly persistent types. As we have seen in the previous sections, the first-order approach *always* works when types are perfectly persistent; Proposition 4 and Corollary 2, however, show that the FO approach does not generically work when interactions are very frequent or types are highly persistent. How is this possible? The key to understanding this apparent contradiction is to realize that when types are constant, the repetition of the optimal static contract is only one of the many possible solutions: in histories that occur with *exactly* zero probability, the quantities are irrelevant and so they can be set to any arbitrary number, for example, equal to the static optimum. To the contrary, when types are highly persistent but probabilities off the main diagonal are not exactly zero, quantities cannot be set arbitrarily in these histories. Typically, the quantities are uniquely defined along all histories. As persistence converges to 1, these quantities along the nonconstant histories converge to values that are different from the static optimum and are nonmonotonic. The problem is that there is a lack of *lower hemicontinuity* in the limit with constant types, and some of the limit solutions (including the repetition of the static optimum) cannot be seen as the limit of the optimal solution as persistence converges to 1.

On serially independent shocks. Esö and Szentes (2017) suggest that AR(1) models can be seen as an example of a more general class of environments for which the first-order approach works. They suggest that any model with correlated and continuous types can be transformed into an equivalent model with i.i.d. shocks, what they call the "orthogonalization of information." Since, as discussed in Section 3.3, with i.i.d. shocks the first-order approach works under some general conditions, it is tempting to believe that this transformation may help in characterizing a sufficient condition for the validity of the first-order approach. This is, however, not the case, since the change of variables that they suggest is just an alternative representation of the same environment. Their idea is that fixing a cumulative distribution $F(\theta^t|\theta^{t-1})$, observing θ^t is equivalent to observing a random variable $v_t = F(\theta^t|\theta^{t-1})$, since F is increasing and invertible in θ^t. In addition, v_t is uniformly distributed on $[0, 1]$. Assume the utility is $u(\theta, q) = \theta q$. To make the equivalent transformation, we need to substitute $\theta^t = F^{-1}(v_t; \theta^{t-1})$, so we have $u(v_1, q_1) = F^{-1}(\mathbf{v}_1) \cdot q_1$, $u(\mathbf{v}^2, q_2) = F^{-1}(v_2; F^{-1}(v_1)) \cdot q_2$, and, iterating,

$$u(\mathbf{v}^t, q_t) = F^{-1}\big(v_t; F^{-1}\big(v_{t-1}; F^{-1}\big(v_{t-2}; [\cdots]\big)\big)\big) \cdot q_t, \tag{21}$$

where $\mathbf{v}^t = (v_1, \dots, v_t)$. It is clear from (21) that, even starting from the simplest utility function, the per period utility of the "equivalent transformation" is a complicated, time-inseparable function of the entire history of the shocks $\mathbf{v}^t$. The change of variables from θ^t to v_t allows us to get rid of serial correlation in the types; the correlation, however, does not disappear—it must be incorporated into the transformed utility function.

[29]Using a recent large data set, Guvenen et al. (2014, 2016) show that individual income in the United States is very persistent and the empirical distribution of income changes has extremely high kurtosis. This suggests that in applications where income is the key variable, it is appropriate to assume that types are highly persistent.

All the problems that induce a failure of the first-order approach in the original problem are just shifted from the distribution function to the transformed per period utility function. The benefit of having independent shocks is compensated by the complications of having these new utilities.

Utilities that are not quasilinear. In the preceding analysis, both for the sufficient conditions of Section 3 and the impossibility result of Section 4, we have focused on quasilinear preferences of the form $u(q, \theta) - p$, where p is a monetary transfer. We have focused on this case because it is the most widely studied in contract theory and because little is known about the cases with utilities that are not quasilinear. Battaglini and Coate (2008) have characterized the necessary and sufficient conditions for a two-type model with quasilinear preferences and then extended these conditions to a model with isoelastic preferences of the form $u(q, \theta) - \frac{1}{\alpha} p^{\alpha}$, showing that when α is sufficiently small, then the structure of the binding incentive constraint of the quasilinear model is the same as in the isoelastic case (thus extending the sufficient conditions of the quasilinear to the isoelastic for small α). In more applied work, however, the solution of dynamic contracts with non-quasilinear preferences has relied on numerical methods. Kapička (2013), Farhi and Werning (2013), and Golosov et al. (2016) also adopt separable preferences like Battaglini and Coate (2008).[30] These papers allow for more than two types and verify the validity of the FO approach numerically in specific calibrated examples in which types are subject to i.i.d. innovations as in a random walk or in a geometric random walk. The focus of these papers is not the study of conditions under which the FO approach is valid, so they do not explore whether the applicability of the FO approach relies on the parametrization or on the specific stochastic processes. In future research it would be interesting to provide a formal extension of the results of Section 4 to these and more general environments beyond quasilinear utility.

From discrete to continuous types. While we have focused the analysis on the case with discrete types, there is a clear connection between models with discrete and continuous types, and the same issues discussed above arise in continuous-type models as well. Consider a continuous-type model with type set $\Theta = [\underline{\theta}, \overline{\theta}] \subset R^{+}$, prior distribution $\mu(\theta)$, and transition distribution $F(\theta'|\theta)$. We can define an associated discrete model by defining the type space as $\Theta^N = \{\theta_0, \dots, \theta_N\}$ with $\theta_0 = \overline{\theta}$, $\theta_N = \underline{\theta}$, and $\theta_i = \theta_{i+1} + \Delta\theta_N$, and the prior as $\mu^N(\theta_i) = \mu(\theta_i)$ and the transition matrix as $F^N(\theta_j|\theta_i) = F(\theta_j|\theta_i)$. In the Supplemental Material, we show that the envelope formula and the FO-optimal contracts of the continuous model can be obtained as limits of the discrete formulas (2) and (7) as $N \to \infty$ and $\Delta\theta_N \to 0$.

5. A solved example

As we have seen in Section 4, even with two periods and three types the FO-optimal contract fails to be monotonic and the FO approach generally cannot be applied. In the Supplemental Material, we fully characterize the optimal contract in the motivating example of Section 4 in which $f(\theta|\theta) = \alpha$ and $f(\theta|\theta') = (1 - \alpha)/2$ for any $\theta, \theta' \in$

[30] Golosov et al. (2016) also considered a nonseparable variant of the form $\frac{1}{\beta}(u(q, \theta) - p)^{\beta}$.

$\{\theta_H, \theta_M, \theta_L\}$, $\theta \neq \theta'$, and $\alpha > 1/3$.[31] Other than the types, there are four parameters here: $\{\mu_H, \mu_M, \alpha, \delta\}$. This example provides novel insights on the structure of the binding constraints.

We find that the principal delays pooling as much as possible. We may have pooling of types in the second period and not in the first; we may even have cases in which both IC_{HL} and IC_{HM} are binding, yet the principal offers a separating contract in the first period (which is not possible in the static model). If the principal does indeed pool types θ_M and θ_L in the first period (this happens when μ_M is very low), then we have pooling in the second period, both within and across histories: $q(\theta_M) = q(\theta_L) \Rightarrow q(\theta_M|\theta_M) = q(\theta_L|\theta_M) = q(\theta_M|\theta_L) = q(\theta_L|\theta_L)$. In future research, it will be interesting to study whether these properties are generally valid or are specific to our solved example.

6. Approximate optimality and implementability

Without the first-order approach, we have no systematic way to simplify the constraint set with many types and periods. This may make the analysis extremely complicated, even from a numerical point of view. In this section we show that there is a class of contracts that is relatively easy to characterize and that induces a minimal loss (if any) in the principal's payoff precisely when the first-order approach fails; that is, when the agent's type is highly persistent. This class consists of contracts that are monotonic in the sense of Definition 2. For simplicity, in the remainder of this section we assume the standard Mussa and Rosen (1978) preferences $u(\theta^t, q) = \theta^t q$.

Define $\mathcal{M}$ as the set of *monotonic contracts,*

$$\mathcal{M} = \left\{ \mathbf{q} \,\middle|\, \begin{array}{c} q(\theta_i|h^{t-1}) \geq q(\theta_{i+1}|h^{t-1}),\, i < N, \text{ and } q(\theta_i|h^{t-1}) \geq q(\theta_i|\widehat{h}^{t-1}), \\ i = 1, \dots, N, \forall h^{t-1} \text{ and } h^{t-1} \succeq \widehat{h}^{t-1} \end{array} \right\},$$

where, as before, $h^t \succeq \widehat{h}^t$ if $h^t_s \geq \widehat{h}^t_s$ $\forall s \leq t$.

In static environments, the envelope formula plus monotonicity are necessary and sufficient for a contract to be implementable: if the FO-optimal contract fails to satisfy the monotonicity constraint, then the contract must be ironed out to make it monotonic; otherwise implementability fails (see Myerson (1981)). In a dynamic environment, monotonicity is not necessary. If we impose monotonicity in the seller's problem, we guarantee implementability even if we ignore the global constraints, but we may obtain a suboptimal contract.[32]

It follows immediately from Proposition 2 that the optimal monotonic contract can be characterized by solving

$$\max_{\mathbf{q} \in \mathcal{M}} \left\{ \mathbb{E}[S(\mathbf{q})] - \sum_{i=0}^{N} \mu_i U^*(\theta_i, h^0; \mathbf{q}) \right\}, \tag{22}$$

[31] In terms of the continuous-time Markov chain model of the previous section, we set τ and α at arbitrary values and choose λ accordingly.

[32] Note that in all the cases presented in Section 4 in which the FO-optimal contract coincides with the optimal contract, the optimal monotonic contract is exactly optimal.

where $U^*(\theta_i, h^0; \mathbf{q})$ is given by the envelope formula (3). Problem (22), moreover, is sufficiently tractable to allow a partial characterization of the properties of its solution.

Proposition 5. *In the optimal monotonic allocation,* $q(\theta^t|h^{t-1}) \leq \theta^t$ *for any* θ^t *and* h^{t-1}*. Moreover, for any arbitrarily small* $\varepsilon_1, \varepsilon_2 > 0$*, we have* $\Pr(|q(\theta^t|h^{t-1}) - \theta^t| > \varepsilon_1) \leq \varepsilon_2$ *for t and T sufficiently large.*

The first part of the proposition establishes that, analogous to the static model, the optimal monotonic contract is uniformly downward distorted. The second part states that the contract converges in probability to an efficient contract.

How good is the optimal monotonic contract as an approximation of the optimal contract? In the Supplemental Material, we study this question by numerically solving a model with a uniform prior and the same Markov process used in Section 5. We calculate the loss in expected profit from using (i) the optimal monotonic contract and (ii) the repeated optimal static contract, where the loss is measured as a percentage of the profit in the optimal contract. This exercise shows that the approximation by the optimal monotonic contract is quite good over the entire parametric range, with a loss of profit that is never higher than 0.06. In fact, the loss from using a monotonic contract is an inverse-U-shaped function of persistence. In both limit cases with i.i.d. and with perfectly persistent types, the optimal contract is monotonic; hence, the loss is zero. For intermediate levels, the loss first rises and then falls with types' persistence.

The following result formalizes the observation that the optimal monotonic contract is a good approximation for the optimal contract when the discount rate $r$is sufficiently small. Define $\pi_m(\Lambda, r, T)$ and $\pi^*(\Lambda, r, T)$ to be the expected average discounted profits as functions of the rate of transition Λ, discount rate r, and the time horizon T in the optimal monotonic contract and the optimal contract, respectively. Define $\pi_m(\Lambda, r) = \lim_{T\to\infty} \pi_m(\Lambda, r, T)$ and $\pi^*(\Lambda, r) = \lim_{T\to\infty} \pi^*(\Lambda, r, T)$.

Proposition 6. *Let* $\lambda = \max \Lambda$*. The profit of the optimal monotonic contract converges to the profit of the optimal contract independent of the order of limits:*

$$\lim_{r\to 0}\lim_{\lambda\to 0}\left[\pi_m(\Lambda, r) - \pi^*(\Lambda, r)\right] = \lim_{\lambda\to 0}\lim_{r\to 0}\left[\pi_m(\Lambda, r) - \pi^*(\Lambda, r)\right] = 0.$$

The fact that the convergence holds independently of the order of limits emphasizes the robustness of the result and the fact that the optimal monotonic contract converges quickly to the optimal contract in comparison to other incentive compatible contracts such as the repetition of the static optimum.

The results of this section may be useful in applied work. As mentioned in the Introduction, many works in the applied literature postulate that the first-order approach works; however, the risk is that the contracts thus characterized are not incentive compatible. Furthermore, in the most natural environments, this risk cannot be fully resolved by numerical methods. To the extent that it is not possible to check all the incentive compatibility constraints, studying optimal monotonic contracts may be a more robust option, since it guarantees implementability and it is approximately optimal when players are sufficiently forward looking.

7. Related literature

Our paper is related to four main literatures. First, we have the traditional literature that studies dynamic principal–agent models when the agent's type follows a stochastic process and the allocation is chosen in every period. The first paper to use the first-order approach to study dynamic models and to state an associated "envelope formula" is Baron and Besanko (1984).[33] Their paper states the formula in general terms and shows it to be sufficient in two benchmark cases: when types are constant over time, in which case the optimal dynamic contract corresponds to a repetition of the static optimum, and when type realizations are independently distributed over time, in which case the optimal contract is efficient starting from period 2.[34] Extensions of this approach to environments with imperfect correlation of types are presented by Besanko (1985), Laffont and Tirole (1990), and Battaglini (2005). Besanko (1985) extends the analysis to an infinite horizon with continuous types following an AR(1) process; Laffont and Tirole (1990) focus on a two-period environment with two types. Battaglini (2005) extends the two-type model to an infinite horizon. The main contributions of all these papers are in showing that the first-order approach is sufficient in their respective environments. Laffont and Tirole (1996) and, more recently, Pavan et al. (2014) have derived "envelope formulas" for continuous types applicable to more complex environments. Contrary to the previous literature, these papers are not focused on finding specific environments in which these envelope formulas are sufficient for incentive compatibility, leaving open the question of the general applicability of the first-order approach.[35]

The second literature to which our paper is related is on sequential screening that was started by Courty and Li (2000). This literature studies environments in which the agent receives information gradually over time, but the allocation is determined only in the last period. The models in this literature have two stages. At the beginning of period one, the agent receives an informative signal and the contract is signed at the end of this period, but no allocation is made; in the second period, the type is revealed to the agent and the allocation takes place. Courty and Li's paper is one of the first papers to clearly discuss the limitations of the first-order approach in dynamic environments;

[33]See Section 3 for a discussion of the first-order approach and envelope formula. See Stole (2001), Laffont and Martimort (2002), Milgrom (2004), and Bolton and Dewatripont (2005) for general discussions of the envelope formula in the static case.

[34]See also Townsend (1982) and Clementi and Hopenhayn (2006), among others, for dynamic principal–agent models in which types are serially uncorrelated.

[35]Other important contributions in the dynamic contracting literature are Dewatripont (1989), Hart and Tirole (1988), Rey and Salanie (1990), Rustichini and Wolinsky (1995), Battaglini (2007), Williams (2011), Bergemann and Välimäki (2010), Strulovici (2011), Garrett and Pavan (2012), Athey and Segal (2013), and Maestri (2017). These papers, however, focus on different aspects of the problem and limit the analysis to environments that are quite different from ours. Hart and Tirole (1988) assume that supply can have two values, 0 or 1. Rustichini and Wolinsky (1995) assume consumers are not strategic and ignore that future prices depend on their current actions. Dewatripont (1989), Rey and Salanie (1990), Battaglini (2007), Maestri (2017), and Strulovici (2011) focus on renegotiation. Garrett and Pavan (2012) look at managerial compensation when allocations are ex post monotone. Bergemann and Välimäki (2010) and Athey and Segal (2013) study implementation of efficient allocations, extending the pivot mechanism to dynamic environments. See also Bergemann and Said (2011) for a short survey.

one of their main achievements is to identify environments in which the first-order approach can be applied in the class of problems that they study. More recently, Courty and Li's work has been extended in many directions. Esö and Szentes (2007) consider the case in which the seller can choose to voluntarily disclose information in the first period. They show that the agent does not receive private rents for the disclosure of information. Li and Shi (2017) show that discriminatory disclosure of information can be optimal when the amount of additional private information that the buyer can learn depends on her type. Krähmer and Strausz (2015) argue that in this class of models the benefit of sequential screening is due to the joint relaxation of incentive and participation constraints. To solve their model, the authors propose an original approach to deal with global constraints that works in their environment with N types. In all these papers the key question is whether the contract must depend on the interim informative signal or if it can depend only on the type revealed in the last stage. In our model, because the allocation is chosen in all periods, information must be disclosed in all periods.

Third, our paper is related to a recent literature devoted to the study of approximately optimal mechanisms in environments in which fully optimal mechanisms are hard to characterize (see Madarász and Prat (2017) and Chassang (2013) for recent contributions, and Hartline (2012) for a summary of the computer science approach). While parts of this literature deal with more general environments than ours, the approach we adopt in Section 6 takes full advantage of the dynamic structure of the framework we study; this allows us to obtain an approximately optimal contract that guarantees incentive compatibility for all types at all histories.

Finally, there is a large and growing applied literature using the first-order approach to solve dynamic contracts in complex environments via numerical methods. Understanding the conditions for the applicability of the first-order approach with discrete types seems particularly important in these exercises. Even when using models with continuous types, these papers typically compute the equilibrium policies and verify incentive compatibility using discretized approximations.[36] The envelope formula presented in our paper provides an exact formula for discrete types that can be used to compute the first-order optimal contract and to verify incentive compatibility directly without approximations.

8. Conclusion

In this paper, we studied a simple principal–agent model in which the agent's type is private information and follows a Markov process. Following the standard approach in the literature, we first studied the optimal contract when only local incentive constraints are considered. We have shown that the agent's equilibrium rents can be represented

[36]This is the case, for example, in Kapička (2013), Farhi and Werning (2013), and Guvenen et al. (2016), who study models of intertemporal consumption smoothing using numerical methods. Exceptions are Zhang (2009) and Williams (2011), who use continuous-time methods. They verify that the conditions for the first-order approach are satisfied in their model. Zhang (2009), however, limits the analysis to a two-type model, and Williams (2011) limits the set of possible deviations available to the agent (who can report only incomes lower or equal to the true income).

purely as a function of the allocation through a dynamic version of the so-called envelope formula. Moreover, as in the static model, the envelope formula and a natural monotonicity condition on the allocation guarantee that the contract is implementable. Although this condition is only sufficient and quite strong, it is verified for virtually all the natural environments that have been used to justify the use of the FO approach.

The main result of our analysis is to show that the environments for which the envelope formula is sufficient to characterize the optimal dynamic contract are special. Except for nongeneric choices of the stochastic process governing the evolution of the agent's type, monotonicity and, more generally, incentive compatibility are necessarily violated by the FO-optimal contract if the frequency of interactions is sufficiently high (or, equivalently, if the discount factor, time horizon, and type persistence are sufficiently large). Numerical examples, moreover, show that we require only moderate levels of persistence to violate the first-order approach. These findings suggest that the applicability of the FO approach is problematic in environments in which types are persistent and expected continuation values are important relative to per period payoffs.

To gain insight into what the optimal contract looks like when the first-order approach does not work, we have characterized it in the simple case of three types and two periods. We have also characterized a class of easily solvable contracts—monotonic contracts—that are approximately optimal when the players are patient.

The analysis suggests a number of open research questions. The characterization of the optimal contract with three types and two periods suggests that state-dependent pooling of types plays an important role in dynamic screening. Future research should explore the extent to which these features extend to more general environments. The analysis in Section 7, moreover, suggests that even when it is not possible to fully characterize the optimal contract, useful insights can be gained by studying contracts that are approximately optimal. It is plausible to assume that more complex—but still solvable—classes of incentive compatible contracts can be found that improve upon the approximation. We leave further development of these ideas for future research.

Appendix

A.1 *Proofs of Lemma 1 and Corollary 1*

We first show that all the constraints in the relaxed problem can be assumed to hold with equality.

Lemma A1. *In a FO-relaxed problem, $IR_N(h^{t-1})$ can be assumed to hold with equality for all $h^{t-1} \in H^{t-1}$, and $IC_{i,i+1}(h^{t-1})$ can be assumed to hold with equality for all $h^{t-1} \in H^{t-1}$ and $i = 0, 1, \ldots, N-1$.*

The proofs of Lemmas A1–A3 and A9 are provided in the Supplemental Material. The appendix is available in a supplementary file on the journal website, http://econtheory.org/supp/2355/supplement.pdf.

We can now prove Lemma 1 and Corollary 1 together. We proceed by (backward) induction on t. Note that at $t = T$, Lemma A1 implies

$$U(\theta_N|h^{T-1}) = 0 \quad \text{and} \quad U(\theta_i|h^{T-1}) = \sum_{l=1}^{N-i} \Delta u(\theta_{i+l}|h^{T-1}; \mathbf{q}) \quad \forall i \leq N-1, \tag{23}$$

where $\Delta u(\theta_{i+1}|h^{t-1}; \mathbf{q}) = u(\theta_i, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1}))$. Similarly, for $t = T - 1$, we have for $i \leq N-1$,

$$\begin{aligned}
&U(\theta_i|h^{T-2}) \\
&\quad = \Delta u(\theta_{i+1}|h^{T-2}; \mathbf{q}) + U(\theta_{i+1}|h^{T-2}) + \delta \sum_{k=0}^{N} (f_{ik} - f_{(i+1)k}) U(\theta_k|h^{T-2}, \theta_{i+1}) \\
&\quad = \sum_{n=1}^{N-i} \left[\Delta u(\theta_{i+n}|h^{T-2}; \mathbf{q}) + \delta \sum_{k=0}^{N} (f_{(i+n-1)k} - f_{(i+n)k}) U(\theta_k|h^{T-2}, \theta_{i+n}) \right] \\
&\quad = \sum_{n=1}^{N-i} \left[\Delta u(\theta_{i+n}|h^{T-2}; \mathbf{q}) + \delta \sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} \Delta u(\theta_{k+l}|h^{T-2}, \theta_{i+n}; \mathbf{q})) \right].
\end{aligned}$$

Now let

$$\begin{aligned}
&\sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} \Delta u(\theta_{k+l}|h^{T-2}, \theta_{i+n}; \mathbf{q})) \\
&\qquad = \sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} V_{k+l},
\end{aligned} \tag{24}$$

where $V_j = \Delta u(\theta_j|h^{T-2}, \theta_{i+n}; \mathbf{q}))$ for any type θ_j. The right-hand side of (24) can be written *as*

$$\sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} V_{k+l} = \begin{bmatrix} (f_{(i+n-1)0} - f_{(i+n)0})(V_1 + \cdots + V_N) \\ + (f_{(i+n-1)1} - f_{(i+n)1})(V_2 + \cdots + V_N) \\ + \cdots + (f_{(i+n-1)(N-1)} - f_{(i+n)(N-1)}) V_N \end{bmatrix}.$$

Rearranging the terms yields

$$\begin{aligned}
&\sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} V_{k+l} \\
&\quad = \begin{bmatrix} (f_{(i+n-1)0} - f_{(i+n)0}) V_1 \\ + \big((f_{(i+n-1)0} + f_{(i+n-1)1}) - (f_{(i+n)0} + f_{(i+n)1})\big) V_2 \\ + \cdots + \big((f_{(i+n-1)0} + \cdots + f_{(i+n-1)(N-1)}) - (f_{(i+n)0} + \cdots + f_{(i+n)(N-1)})\big) V_N \end{bmatrix} \\
&\quad = \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n}) V_k,
\end{aligned}$$

where we recall that $\Delta F(\theta_j|\theta_i) = F(\theta_j|\theta_i) - F(\theta_j|\theta_{i-1})$. This implies

$$\sum_{k=0}^{N-1}(f_{(i+n-1)k} - f_{(i+n)k}) \sum_{l=1}^{N-k} \Delta u\big(\theta_{k+l}|h^{T-2}, \theta_{i+n}; \mathbf{q}\big))$$
$$= \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n})\Delta u\big(\theta_k|h^{T-2}, \theta_{i+n}; \mathbf{q}\big)).$$

It follows that

$$\begin{aligned} &U\big(\theta_i|h^{T-2}\big) \\ &= \sum_{n=1}^{N-i}\left[\Delta u\big(\theta_{i+n}|h^{T-2}; \mathbf{q}\big) + \delta \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n})\Delta u\big(\theta_k|h^{T-2}, \theta_{i+n}; \mathbf{q}\big))\right] \\ &= \sum_{n=1}^{N-i}\left[\begin{array}{c} \Delta u\big(\theta_{i+n}|h^{T-2}; \mathbf{q}\big) \\ + \displaystyle\sum_{\hat{h}\in\widehat{H}(h^{T-2}, \theta_{i+n})} \sum_{\tau>T-1} \delta^{\tau-T-1} \prod_{k=T}^{\tau} \Delta F(\hat{h}_k|\hat{h}_{k-1})\Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big)) \end{array}\right], \end{aligned} \tag{25}$$

where we recall that $\widehat{H}(h^t)$ is the set of histories following h^t in which all realizations after t are lower than θ_0.

It is easy to see that (23) and (25) prove the statement in Corollary 1 and in Lemma 1 for $t = T$ and $t = T - 1$, respectively. We therefore conclude that our hypothesis holds for $t \geq T - 1$. Next suppose it holds for $t + 1$, where $t \geq T - 2$. We want to show that it holds for t. We have

$$\begin{aligned} &U\big(\theta_i|h^{t-1}\big) \\ &= \Delta u\big(\theta_{i+1}|h^{t-1}; \mathbf{q}\big) + U\big(\theta_{i+1}|h^{t-1}\big) + \delta \sum_{k=0}^{N}(f_{ik} - f_{(i+1)k})U\big(\theta_k|h^{t-1}, \theta_{i+1}\big) \\ &= \sum_{n=1}^{N-i}\left[\Delta u\big(\theta_{i+n}|h^{t-1}; \mathbf{q}\big) + \delta \sum_{k=0}^{N}(f_{(i+n-1)k} - f_{(i+n)k})U\big(\theta_k|h^{t-1}, \theta_{i+n}\big)\right] \\ &= \sum_{n=1}^{N-i}\Bigg[\Delta u\big(\theta_{i+n}|h^{t-1}; \mathbf{q}\big) + \delta \sum_{k=0}^{N-1}(f_{(i+n-1)k} - f_{(i+n)k}) \sum_{m=1}^{N-k}\Bigg(\Delta u\big(\theta_{k+m}|h^{t-1}, \theta_{i+n}; \mathbf{q}\big) \\ &\quad + \sum_{\hat{h}\in\widehat{H}(h^{t-1}, \theta_{i+n}, \theta_{k+m})} \sum_{\tau>t+1} \delta^{\tau-(t+1)} \prod_{\iota=t+2}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1})\Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big)\Bigg)\Bigg], \end{aligned} \tag{26}$$

where the third equality follows from the induction hypothesis. Now

$$\sum_{k=0}^{N-1}(f_{(i+n-1)k} - f_{(i+n)k}) \sum_{m=1}^{N-k} \Delta u\big(\theta_{k+m}|h^{t-1}, \theta_{i+n}; \mathbf{q}\big)$$

$$= \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n}) \Delta u\big(\theta_k|h^{t-1}, \theta_{i+n}; \mathbf{q}\big) \tag{27}$$

and

$$\begin{aligned}
&\delta \sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{m=1}^{N-k} \sum_{\hat{h} \in \widehat{H}(h^{t-1}, \theta_{i+n}, \theta_{k+m})} \sum_{\tau > t+1} \delta^{\tau-(t+1)} \\
&\quad \times \prod_{\iota=t+2}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1}) \Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big) \\
&= \delta \sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{m=1}^{N-k} Q_{k+m},
\end{aligned}$$

where $V_l = \sum_{\hat{h} \in \widehat{H}(h^{t-1}, \theta_{i+n}, \theta_l)} \sum_{\tau > t+1} \delta^{\tau-(t+1)} \prod_{\iota=t+2}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1}) \Delta u(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q})$. As before, after some algebraic manipulation, this becomes

$$\delta \sum_{k=0}^{N-1} (f_{(i+n-1)k} - f_{(i+n)k}) \sum_{m=1}^{N-k} V_{k+m} = \delta \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n}) V_k. \tag{28}$$

Combining (27) and (28), we obtain

$$\begin{aligned}
&\delta \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n}) \big[\Delta u\big(\theta_k|h^{t-1}, \theta_{i+n}; \mathbf{q}\big) + V_k\big] \\
&\quad = \delta \sum_{k=1}^{N} \Delta F(\theta_k|\theta_{i+n}) \left[\begin{array}{c} \Delta u\big(\theta_k|h^{t-1}, \theta_{i+n}; \mathbf{q}\big) \\ + \displaystyle\sum_{\hat{h} \in \widehat{H}(h^{t-1}, \theta_{i+n}, \theta_k)} \sum_{\tau > t+1} \delta^{\tau-(t+1)} \prod_{\iota=t+2}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1}) \Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big) \end{array} \right] \\
&\quad = \sum_{\hat{h} \in \widehat{H}(h^{t-1}, \theta_{i+n})} \sum_{\tau > t} \delta^{\tau-t} \prod_{\iota=t+1}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1}) \Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big).
\end{aligned} \tag{29}$$

Combining (26) and (29), we obtain

$$\begin{aligned}
&U\big(\theta_i|h^{t-1}\big) \\
&\quad = \sum_{n=1}^{N-i} \left[\Delta u\big(\theta_{i+n}|h^{t-1}; \mathbf{q}\big) + \sum_{\hat{h} \in \widehat{H}(h^{t-1}, \theta_{i+n})} \sum_{\tau > t} \delta^{\tau-t} \prod_{\iota=t+1}^{\tau} \Delta F(\hat{h}_\iota|\hat{h}_{\iota-1}) \Delta u\big(\hat{h}_\tau|\hat{h}^{\tau-1}; \mathbf{q}\big) \right].
\end{aligned}$$

Note that

$$\begin{aligned}
\Delta u\big(\theta_{i+1}|h^{t-1}; \mathbf{q}\big) &= u\big(\theta_i, q\big(\theta_{i+1}|h^{t-1}\big)\big) - u\big(\theta_{i+1}, q\big(\theta_{i+1}|h^{t-1}\big)\big) \\
&= \int_{\theta_{i+1}}^{\theta_i} u_\theta\big(x, q\big(\theta_{i+1}|h^{t-1}\big)\big)\, dx.
\end{aligned}$$

It follows that

$$
\begin{aligned}
&U\big(\theta_i|h^{t-1}\big)\\
&=\sum_{n=1}^{N-i}\left[\begin{array}{c}\displaystyle\int_{\theta_{i+n}}^{\theta_{i+n-1}} u_\theta\big(x,q\big(\theta_{i+n}|h^{t-1}\big)\big)\,dx\\
\displaystyle +\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_{i+n})}\sum_{\tau>t}\delta^{\tau-t}\prod_{\iota=t+1}^{\tau}\Delta F(\hat{h}_\iota|\hat{h}_{\iota-1})\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta\big(x,q\big(\hat{h}_\tau|\hat{h}^{\tau-1}\big)\big)\,dx\end{array}\right].
\end{aligned}
$$

This proves Corollary 1. Subtracting $U(\theta_{i+1}|h^{t-1})$ and dividing by $\Delta\theta$, the above expression gives us Lemma 1.

A.2 *Proof of Proposition 2*

Recall that $\Delta U(\theta_k|h^{t-1},\theta_i)=U(\theta_k|h^{t-1},\theta_i)-U(\theta_k|h^{t-1},\theta_{i+1})$. We start with some useful lemmas.

Lemma A2. *If $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k|h^{t-1})$ are non-increasing in i and k for any h^{t-1}, respectively, then (3) implies that local upward incentive compatibility constraints are satisfied.*

Lemma A3. *If $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k|h^{t-1})$ are non-increasing in i and k for any h^{t-1}, respectively, and (3) holds, then the local incentive compatibility constraints imply the global incentive compatibility constraints.*

Given the lemmas presented above, Proposition 2 is proven if we establish that when the allocation is monotonic as defined in Definition 2, then $q(\theta_i|h^{t-1})$ and $\Delta U(\theta_k|h^{t-1},\theta_i)$ are non-increasing in i for any h^{t-1}. The fact that $q(\theta_i|h^{t-1})$ is non-increasing in i for any h^{t-1} is an immediate consequence of the monotonicity. The fact that $\Delta U(\theta_k|h^{t-1},\theta_i)$ is non-increasing in k for any h^{t-1} is established by the following result.

Lemma A4. *If the allocation is monotonic as in Definition 2, then $\Delta U(\theta_k|h^{t-1})$ is non-increasing in k $\forall h^{t-1}$.*

Proof. Note first that $U(\theta_N|h^{t-1},\theta_i)=U(\theta_N|h^{t-1},\theta_{i+1})=0$, so $\Delta U(\theta_N|h^{t-1},\theta_i)=0$. By Lemma 1, we have

$$
\begin{aligned}
&U\big(\theta_{N-1}|h^{t-1},\theta_i\big)\\
&\quad=\int_{\theta_N}^{\theta_{N-1}} u_\theta\big(x,q\big(\theta_N|h^{t-1},\theta_i\big)\big)\,dx\\
&\qquad+\sum_{\hat{h}\in\widehat{H}(h^{t-1},\theta_i,\theta_{N-1})}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\displaystyle\prod_{\iota=t+2}^{\tau}\Delta F(\hat{h}_\iota|\hat{h}_{\iota-1})\\
\displaystyle\cdot\int_{\hat{h}_\tau}^{\hat{h}_\tau+\Delta\theta} u_\theta\big(x,q\big(\hat{h}_\tau|\hat{h}^{\tau-1}\big)\big)\,dx\end{array}\right].
\end{aligned}
$$

It is useful to write this expression with different notation. Let $\widehat{H}_t(i)$ be the set of realizations of length $T-t$ that start with the first element equal to θ_i (we denote by ${}_th$ a typical element of $\widehat{H}_t(i)$, so ${}_th_1=\theta_i$). A history $h^\tau \in \widehat{H}(h^t)$ with $(t+1)$th element equal to θ_i ($h^\tau_{t+1}=\theta_i$) is then $h^\tau=\{h^t,{}_t h^{\tau-t}\}$ for ${}_th\in\widehat{H}_t(i)$ (by convention we write $h^t=\{h^t,{}_t h^0\}$). We can then write

$$\begin{aligned} &U\big(\theta_{N-1}|h^{t-1},\theta_i\big)\\ &\quad=\int_{\theta_N}^{\theta_{N-1}} u_\theta\big(x,q\big(\theta_N|h^{t-1},\theta_i\big)\big)\,dx\\ &\qquad+\sum_{{}_th\in\widehat{H}_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\displaystyle\prod_{l=t+2}^{\tau}\Delta F({}_th_l|{}_th_{l-1})\\ \displaystyle\cdot\int_{{}_th_\tau}^{{}_th_\tau+\Delta\theta} u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_i,{}_t h^{\tau-t-1}\big)\big)\end{array}\right].\end{aligned}$$

Similarly, we can write

$$\begin{aligned} &U\big(\theta_{N-1}|h^{t-1},\theta_{i+1}\big)\\ &\quad=\int_{\theta_N}^{\theta_{N-1}} u_\theta\big(x,q\big(\theta_N|h^{t-1},\theta_{i+1}\big)\big)\,dx\\ &\qquad+\sum_{{}_th\in\widehat{H}_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\displaystyle\prod_{l=t+2}^{\tau}\Delta F({}_th_l|{}_th_{l-1})\\ \displaystyle\cdot\int_{{}_th_\tau}^{{}_th_\tau+\Delta\theta} u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1}\big)\big)\,dx\end{array}\right].\end{aligned}$$

Therefore,

$$\begin{aligned} &\Delta U\big(\theta_{N-1}|h^{t-1},\theta_i\big)\\ &\quad=\int_{\theta_N}^{\theta_{N-1}} \big[u_\theta\big(x,q\big(\theta_N|h^{t-1},\theta_i\big)\big)-u_\theta\big(x,q\big(\theta_N|h^{t-1},\theta_{i+1}\big)\big)\big]\,dx\\ &\qquad+\sum_{{}_th\in H_t(N-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\displaystyle\prod_{l=t+2}^{\tau}\Delta F({}_th_l|{}_th_{l-1})\\ \displaystyle\cdot\int_{{}_th_\tau}^{{}_th_\tau-\Delta\theta}\left[\begin{array}{c}u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_i,{}_t h^{\tau-t-1}\big)\big)\\ -\,u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1}\big)\big)\end{array}\right]dx\end{array}\right].\end{aligned}$$

Note that by monotonicity, we must have $q(\theta_N|h^{t-1},\theta_i)-q(\theta_N|h^{t-1},\theta_{i+1})\geq 0$ and

$$q\big({}_th_\tau|h^{t-1},\theta_i,{}_t h^{\tau-t-1}\big)-q\big({}_th_\tau|h^{t-1},\theta_{i+1},{}_t h^{\tau-t-1}\big)\geq 0.$$

The above condition plus the single-crossing condition (Assumption 1) imply that $\Delta U(\theta_{N-1}|h^{t-1},\theta_i)\geq\Delta U(\theta_N|h^{t-1},\theta_i)$. Assume now that $\Delta U(\theta_j|h^{t-1},\theta_i)$ is monotonic in j for $j\geq m$. We show below that $\Delta U(\theta_{m-1}|h^{t-1},\theta_i)\geq\Delta U(\theta_m|h^{t-1},\theta_i)$ and the result then follows by induction. Applying Lemma 1 and using the notation developed above,

$$\begin{aligned}
&\Delta U\big(\theta_{m-1}|h^{t-1},\theta_i\big)\\
&\quad=\Delta U\big(\theta_m|h^{t-1},\theta_i\big)+\int_{\theta_N}^{\theta_{N-1}}\big[u_\theta\big(x,q\big(\theta_m|h^{t-1},\theta_i\big)\big)-u_\theta\big(x,q\big(\theta_m|h^{t-1},\theta_{i+1}\big)\big)\big]\,dx\\
&\qquad+\sum_{{}_th\in H_t(m-1)}\sum_{\tau>t+1}\delta^{\tau-t-1}\left[\begin{array}{c}\displaystyle\prod_{l=t+2}^{\tau}\Delta F({}_th_l|{}_th_{l-1})\\ \displaystyle\cdot\int_{{}_th_\tau}^{{}_th_\tau-\Delta\theta}\left[\begin{array}{c}u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_i,{}_th^{\tau-t-1}\big)\big)\\-u_\theta\big(x,q\big({}_th_\tau|h^{t-1},\theta_{i+1},{}_th^{\tau-t-1}\big)\big)\end{array}\right]dx\end{array}\right].
\end{aligned}$$

Thus, the single-crossing condition and monotonicity of the allocation imply $\Delta U(\theta_{m-1}|h^{t-1},\theta_i)\geq\Delta U(\theta_m|h^{t-1},\theta_i)$. □

A.3 *Preliminary results for Propositions 3 and 4*

Define $f_\tau(\theta_j|\theta_i)$ to be the probability that type i will transition to type j when the length of time between two consecutive interactions is given by τ, with the shorthand f_{ij}^τ. In the standard model, $\tau=1$. The underlying continuous Markov chain is uniquely defined by an exponential transition probability with parameters described by an $N+1$ dimensional vector $\Lambda=(\lambda_1,\dots,\lambda_n)$ and an $N+1$ dimensional matrix P, where the following definitions hold:

- The rate of switching out of the current state i is λ_i.
- The conditional transition rate of switching to state j from i, with $P_{ii}=0$, $P_{ij}>0$, and $\sum_j P_{ij}=1$ for all i is $P_{i,j}$.

We refer to the transition function as (Λ,P) and denote the set of all such functions by $\hat{\mathcal{Q}}$. Note that $\hat{\mathcal{Q}}$ is independent of τ.

Define $Q_{ij}=\lambda_iP_{ij}$ for all $i\neq j$. From Lemma 6.2 in Ross (2014), we have the limit properties

$$1-f_{ii}^\tau=\lambda_i\tau+o(\tau)\quad\text{and}\quad\lim_{\tau\to0}\frac{1-f_{ii}^\tau}{\tau}=\lambda_i,\tag{30}$$

$$f_{ij}^\tau=\tau p_{ij}+o(\tau)\quad\text{and}\quad\lim_{\tau\to0}\frac{f_{ij}^\tau}{\tau}=Q_{ij}\quad\forall i\neq j,\tag{31}$$

where $o(\tau)$ is such that $o(\tau)\to0$ and $\frac{o(\tau)}{\tau}\to0$ as $\tau\to0$.

Let $\mathcal{Q}\subset\hat{\mathcal{Q}}$ be the set of all transition functions that satisfy first-order stochastic dominance. It can be shown that if (Λ,P) satisfies first-order stochastic dominance for $\tau=1$ (the benchmark model), then it satisfies first-order stochastic dominance for all $\tau\leq1$.

Next define the dynamic distortion associated with the history $h_i'=\{\theta_i,\theta_{i+1},\theta_i\}$ as

$$\Psi_i(f_\tau)=\frac{\Delta F_\tau(\theta_{i+1}|\theta_i)}{f_\tau(\theta_{i+1}|\theta_i)}\cdot\frac{\Delta F_\tau(\theta_i|\theta_{i+1})}{f_\tau(\theta_i|\theta_{i+1})}$$

$$= \frac{\sum_{k=i+1}^{N} [f_{ik}^{\tau} - f_{i-1,k}^{\tau}]}{f_{i,i+1}^{\tau}} \cdot \frac{\sum_{k=i}^{N} [f_{i+1,k}^{\tau} - f_{ik}^{\tau}]}{f_{i+1,i}^{\tau}}$$

$$= \frac{\sum_{k=i+1}^{N} \left[\frac{f_{\tau}^{ik}}{\tau} - \frac{f_{\tau}^{i-1,k}}{\tau}\right]}{\frac{f_{i,i+1}^{\tau}}{\tau}} \cdot \frac{\sum_{k=i}^{N} \left[\frac{f_{\tau}^{i+1,k}}{\tau} - \frac{f_{\tau}^{i,k}}{\tau}\right]}{\frac{f_{i+1,i}^{\tau}}{\tau}}.$$

Taking limits on both sides, and using (30) and (31), we get

$$\lim_{\tau \to 0} \Psi_i(f_\tau) = \frac{\sum_{k=i+1}^{N} [Q_{ik} - Q_{i-1,k}]}{p_{i,i+1}} \cdot \frac{\begin{bmatrix} (Q_{i+1,i} + \lambda_i - \lambda_{i+1} - Q_{i,i+1}) \\ + \sum_{k=i+2}^{N} (Q_{ik} - Q_{i+1,k}) \end{bmatrix}}{p_{i,i+1}} = \zeta(\Lambda, P). \tag{32}$$

Note that for each $(\Lambda, P) \in \mathcal{Q}$ and τ, a unique transition probability matrix, f_τ, is defined. Define Υ_i to be the subset of $\mathcal{Q}$ where the above defined limit is not equal to 1:

$$\Upsilon_i = \left\{(\Lambda, P) \in \mathcal{Q} \mid \lim_{\tau \to 0} \Psi_i(f_\tau) \neq 1\right\}.$$

We show that this set is generic in $\mathcal{Q}$. A necessary and sufficient condition is that Υ_i must be open and dense in $\mathcal{Q}$.

Lemma A5. *The subset Υ_i is open and dense in Q.*

Proof. Define $\overline{\Upsilon}_i = \mathcal{Q} \setminus \Upsilon_i$ to be the complement of Υ_i.

Step 1. We first prove that Υ_i is open in $\mathcal{Q}$. Suppose not. Then there exists $(\Lambda, P) \in \Upsilon_i$ and a sequence $(\Lambda^n, P^n) \in \overline{\Upsilon}_i$ such that$(\Lambda^n, P^n) \to (\Lambda, P)$. For a fixed τ, let f_τ be the stochastic matrix corresponding to (Λ, P) and let f_τ^n be the stochastic matrix corresponding to (Λ^n, P^n). From (32), we have $\lim_{\tau \to 0} \Psi(f_\tau) = \zeta(\Lambda, P) = \lim_{n \to \infty} \zeta(\Lambda^n, P^n)$, where $\zeta(\Lambda, P)$ and $\zeta(\Lambda^n, P^n)$ are independent of τ, and the latter equality follows from the fact that ζ is continuous in its arguments. By definition, along the sequence we have $\lim_{\tau \to 0} \Psi(f_\tau^n) = 1$ for any n; thus, we have that $\zeta(\Lambda^n, P^n) = \lim_{\tau \to 0} \Psi_i(f_\tau^n) = 1$ for any n and, therefore, $\lim_{n \to \infty} \zeta(\Lambda^n, P^n) = 1$. We conclude that

$$\lim_{\tau \to 0} \Psi(f_\tau) = \zeta(\Lambda, P) = \lim_{n \to \infty} \zeta\big(\Lambda^n, P^n\big) = 1,$$

proving that $(\Lambda, P) \in \overline{\Upsilon}_i$, which is a contradiction.

Step 2. The subset Υ_i is dense in Q. We want to show that $\forall(\Lambda, P) \in \mathcal{Q}$ and $\varepsilon > 0$, $\exists(\Lambda', P')$ such that $d[(\Lambda, P), (\Lambda', P')] < \varepsilon$ and $(\Lambda', P') \in \Upsilon_i$. If $(\Lambda, P) \in \Upsilon_i$, the result is immediate. Therefore, let $(\Lambda, P) \in \overline{\Upsilon}_i$ and fix $\varepsilon > 0$. Consider the following change to (Λ, P): increase $P_{i+1,i-1}$ by a very small amount and decrease $P_{0,i-1}$ by the same amount. We can do this because we assume that P has full nondiagonal support and that first-order stochastic dominance holds strictly.[37] Call this modification (Λ', P'). Then it is easy to show that (Λ', P') satisfies the first-order stochastic dominance condition. Moreover, $(\Lambda', P') \in \Upsilon_i$ and $d[(\Lambda, P), (\Lambda', P')] < \varepsilon$. □

A.4 *Proof of Proposition 3*

Fix $(\Lambda, P) \in \mathcal{Q}$, with the corresponding stochastic matrix f_τ and let $\lim_{\tau\to 0} \Psi_i(f_\tau) = D$. Then by Lemma A5, we have $D < 1$ or $D > 1$ generically. First consider the case $D < 1$. We shall show that $q(\theta_i|\theta_i, \theta_{i+1}) > q(\theta_i|\theta_i, \theta_i)$ for small enough τ, thereby violating monotonicity. Note that

$$\begin{aligned} & s_q\big(\theta_i, q(\theta_i|\theta_i, \theta_i)\big) \\ & \quad \le \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \left[\frac{\Delta F_\tau(\theta_i|\theta_i)}{f_\tau(\theta_i|\theta_i)} \cdot \frac{\Delta F_\tau(\theta_i|\theta_i)}{f_\tau(\theta_i|\theta_i)}\right] \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q(\theta_i|\theta_i, \theta_i)\, dx \end{aligned}$$

and

$$\begin{aligned} & s_q\big(\theta_i, q(\theta_i|\theta_i, \theta_{i+1})\big) \\ & \quad \le \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \left[\frac{\Delta F_\tau(\theta_i|\theta_{i+1})}{f_\tau(\theta_i|\theta_{i+1})} \cdot \frac{\Delta F_\tau(\theta_{i+1}|\theta_i)}{f_\tau(\theta_{i+1}|\theta_i)}\right] \\ & \qquad \times \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q(\theta_i|\theta_i, \theta_{i+1})\, dx. \end{aligned}$$

Let $q_1 = \lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_i)$. Since $\lim_{\tau\to 0} \frac{\Delta F_\tau(\theta_i|\theta_i)}{f_\tau(\theta_i|\theta_i)} = 1$, distortions converge to 1 along constant histories, and so $q_1 = q(\theta_i|h^0)$. By Assumption 3, since the static optimum is an interior solution, we have

$$s_q(\theta_i, q_1) = \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_1)\, dx.$$

Also, letting $q_2 = \lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_{i+1})$, we have

$$s_q(\theta_i, q_2) \le \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot D \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)\, dx.$$

[37]If first-order stochastic dominance was weak, that is, $\Delta F_\tau(\theta_{i+1}|\theta_i) = 0$ or $\Delta F_\tau(\theta_i|\theta_{i+1}) = 0$, then $\Psi_i(f_\tau) = 0$.

It follows that

$$\begin{aligned}\Phi_q(\theta_i, q_1) &= s_q(\theta_i, q_1) - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\cdot\int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_1)\,dx\\ &= 0 \geq s_q(\theta_i, q_2) - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\cdot D\cdot\int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)\,dx\\ &> s_q(\theta_i, q_2) - \frac{1-\sum_{k=i}^{N}\mu_k}{\mu_i}\cdot\int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)\,dx = \Phi_q(\theta_i, q_2),\end{aligned}$$

where the strict inequality follows from $D < 1$. Since Φ is concave, we have $q_2 > q_1$.

Next suppose $\lim_{\tau\to 0}\Psi_i(f_\tau) = D > 1$. Then, analogous to the steps above, we show that $\lim_{\tau\to 0} q(\theta_{i+1}|\theta_{i+1}, \theta_{i+1}) > \lim_{\tau\to 0} q(\theta_{i+1}|\theta_{i+1}, \theta_i)$. Letting $q_3 = \lim_{\tau\to 0} q(\theta_{i+1}|\theta_{i+1}, \theta_i)$ and $q_4 = \lim_{\tau\to 0} q(\theta_{i+1}|\theta_{i+1}, \theta_{i+1})$, we get

$$s_q(\theta_{i+1}, q_4) = \frac{1-\sum_{k=i+1}^{N}\mu_k}{\mu_{i+1}}\cdot\int_{\theta_{i+1}}^{\theta_i} u_{\theta,q}(x, q_4)\,dx$$

and

$$s_q(\theta_{i+1}, q_3) \leq \frac{1-\sum_{k=i+1}^{N}\mu_k}{\mu_{i+1}}\cdot D\cdot\int_{\theta_{i+1}}^{\theta_i} u_{\theta,q}(x, q_3)\,dx.$$

Thus, using $D > 1$, we obtain $\Phi_q(\theta_{i+1}, q_3) < \Phi_q(\theta_{i+1}, q_4)$, implying $q_4 > q_3$.

A.5 *Proof of Proposition 4*

Fix $(\Lambda, P) \in \mathcal{Q}$, with the corresponding stochastic matrix f_τ and let $\lim_{\tau\to 0}\Psi_i(f_\tau) = D$. We consider the two cases, $D < 1$ and $D > 1$, separately.

Case 1: $D < 1$. We prove the result by showing that the FO-optimal contract violates the second period global incentive constraint $IC_{i-1,i+1}(\theta_i)$. To this end, we first make a useful observation.

Lemma A6. *The constraint $IC_{i-1,i+1}(h^{t-1})$ holds if and only if*

$$\begin{aligned}&\int_{\theta_i}^{\theta_{i-1}}\begin{bmatrix} u_\theta\big(x, q(\theta_i|h^{t-1})\big)\\ -\,u_\theta\big(x, q(\theta_{i+1}|h^{t-1})\big)\end{bmatrix}dx + \delta\sum_{k=0}^{N}\begin{bmatrix} (f_{i-1,k}^{\tau} - f_{ik}^{\tau})\\ \cdot\big(U\big(\theta_k|h^{t-1}, \theta_i\big) - U\big(\theta_k|h^{t-1}, \theta_{i+1}\big)\big)\end{bmatrix}\\ &\quad\geq 0,\end{aligned} \tag{33}$$

where $U(\theta_k|h^{t-1}, \theta_i) = U^*(\theta_k|h^{t-1}, \theta_i; \mathbf{q})$, as defined in (3).

Proof. The global incentive compatibility constraint $IC_{i-1,i+1}(h^{t-1})$ can be written *as*

$$\begin{aligned} &U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) \\ &\quad \geq u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1})) \\ &\qquad + \delta \sum_{k=0}^{N} (f^{\tau}_{i-1,k} - f^{\tau}_{i+1,k}) U(\theta_k|h^{t-1}, \theta_{i+1}). \end{aligned} \tag{34}$$

Note that

$$U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) = (U(\theta_{i-1}|h^{t-1}) - U(\theta_i|h^{t-1})) + (U(\theta_i|h^{t-1}) - U(\theta_{i+1}|h^{t-1})).$$

So using $IC_{i-1,i}(h^{t-1})$ and $IC_{i,i+1}(h^{t-1})$, we have

$$\begin{aligned} &U(\theta_{i-1}|h^{t-1}) - U(\theta_{i+1}|h^{t-1}) \\ &\quad - \begin{bmatrix} u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i+1}, q(\theta_{i+1}|h^{t-1})) \\ + \delta \sum_{k=0}^{N} (f^{\tau}_{i-1,k} - f^{\tau}_{i+1,k}) U(\theta_k|h^{t}, \theta_{i+1}) \end{bmatrix} \\ &= \begin{bmatrix} u(\theta_{i-1}, q(\theta_i|h^{t-1})) - u(\theta_i, q(\theta_i|h^{t-1})) \\ + u(\theta_i, q(\theta_{i+1}|h^{t-1})) - u(\theta_{i-1}, q(\theta_{i+1}|h^{t-1})) \end{bmatrix} \\ &\quad + \delta \sum_{k=0}^{N} \left(\begin{array}{c} (f^{\tau}_{i-1,k} - f^{\tau}_{ik}) \\ \cdot \begin{pmatrix} U(\theta_k|h^{t-1}, \theta_i) \\ - U(\theta_k|h^{t-1}, \theta_{i+1}) \end{pmatrix} \end{array} \right). \end{aligned} \tag{35}$$

Using (34) and (35), it follows that that $IC_{i-1,i+1}(h^{t-1})$ holds if and only if (33) holds. $\square$

We conclude that $IC_{i-1,i+1}(\theta_i)$ holds if and only if

$$\int_{\theta_i}^{\theta_{i-1}} \begin{bmatrix} u_\theta(x, q(\theta_i|\theta_i)) \\ - u_\theta(x, q(\theta_{i+1}|\theta_i)) \end{bmatrix} dx + \delta \sum_{k=0}^{N} \begin{bmatrix} (f^{\tau}_{i-1,k} - f^{\tau}_{ik}) \\ \cdot (U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1})) \end{bmatrix} \geq 0. \tag{36}$$

In the formulation described in Section 4, we have $\delta = e^{-r\tau}$. Moreover, the per period payoff is bounded above by some $\overline{u}$. Thus,

$$U(\theta_k|h^{t-1}) \leq \frac{1-\delta^{\frac{T-\tau t}{\tau}}}{1-\delta}\overline{u} = \frac{1-e^{-r(T-\tau t)}}{1-e^{-r\tau}}\overline{u} \quad \forall k, \forall t.$$ [38]

Next note that

$$f^{\tau}_{ij} \cdot U(\theta_k|h^{t-1}) \leq [\tau Q_{ij} + o(\tau)] \frac{1-e^{-r(T-\tau t)}}{1-e^{-r\tau}}\overline{u}$$

[38]As stated in the text, we indulge in a slight abuse of notation here in that each 'time' t now is divided into intervals of length $\frac{1}{\tau}$. So, at time t, the principal and agent have interacted $\frac{t}{\tau}$ times. The expected utility should technically be written as $U(\theta_k|h^{\frac{t-1}{\tau}})$.

$$= \left(1 - e^{-r(T-\tau t)}\right) \frac{\tau}{1 - e^{-r\tau}} \left[Q_{ij} + \frac{o(\tau)}{\tau}\right] \overline{u}.$$

Therefore,

$$\begin{aligned} \lim_{\tau \to 0} \left[f_{ij}^{\tau} \cdot U\left(\theta_k | h^{t-1}\right)\right] &\leq \left(1 - e^{-rT)}\right) Q_{ij} \lim_{\tau \to 0} \frac{1}{re^{-r\tau}} \overline{u} \\ &= \frac{\left(1 - e^{-rT}\right) Q_{ij} \overline{u}}{r}, \end{aligned}$$

where we have used l'Hospital's rule to get $\lim_{\tau \to 0} \frac{\tau}{1-e^{-r\tau}} = \lim_{\tau \to 0} \frac{1}{re^{-r\tau}} = \frac{1}{r}$. It follows that for any ε, $\exists \tau_0$ such that for all $\tau \leq \tau_0$,

$$\begin{aligned} &\sum_{k=0}^{N} \left(f_{i-1,k}^{\tau} - f_{ik}^{\tau}\right)\left(U(\theta_k | \theta_i, \theta_i) - U(\theta_k | \theta_i, \theta_{i+1})\right) \\ &\quad \leq f_{i-1,i-1}^{\tau}\left[\left(U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_{i-1} | \theta_i, \theta_{i+1})\right)\right] \\ &\qquad - f_{ii}^{\tau}\left[\left(U(\theta_i | \theta_i, \theta_i) - U(\theta_i | \theta_i, \theta_{i+1})\right)\right] + \kappa_1 + \varepsilon, \end{aligned}$$

where $\kappa_1 = \frac{1-e^{-rT}}{r} 4N\overline{p} \cdot \overline{u}$ and $\overline{p}$ is the maximal $p_{i,j}$. We can write the right-hand side of the above expression (modulo κ_1 and ε) as

$$\begin{aligned} &f_{i-1,i-1}^{\tau}\left[U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_{i-1} | \theta_i, \theta_{i+1})\right] - f_{ii}^{\tau}\left[U(\theta_i | \theta_i, \theta_i) - U(\theta_i | \theta_i, \theta_{i+1})\right] \\ &\quad = f_{ii}^{\tau}\left[\left(U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_i | \theta_i, \theta_i)\right) - \left(U(\theta_{i-1} | \theta_i, \theta_{i+1}) - U(\theta_i | \theta_i, \theta_{i+1})\right)\right] \\ &\qquad + \left[f_{i-1,i-1}^{\tau} - f_{ii}^{\tau}\right]\left[U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_{i-1} | \theta_i, \theta_{i+1})\right]. \end{aligned}$$

Note that

$$\begin{aligned} &\left[f_{i-1,i-1}^{\tau} - f_{ii}^{\tau}\right]\left[U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_{i-1} | \theta_i, \theta_{i+1})\right] \\ &\quad \leq \left|f_{i-1,i-1}^{\tau} - f_{ii}^{\tau}\right| \cdot \left[\frac{1 - e^{-r(T-2\tau)}}{1 - e^{-r\tau}} 2\overline{u}\right] = \left|(\lambda_i - \lambda_{i-1})\tau + o(\tau)\right| \cdot \left[\frac{1 - e^{-r(T-2\tau)}}{1 - e^{-r\tau}} 2\overline{u}\right] \\ &\quad = \left|\lambda_i - \lambda_{i-1} + \frac{o(\tau)}{\tau}\right| \cdot \left[\frac{\tau}{1 - e^{-r\tau}}\right]\left(1 - e^{-r(T-2\tau)}\right) 2\overline{u}. \end{aligned}$$

Taking limits on both sides yields

$$\lim_{\tau \to 0} \left[f_{i-1,i-1}^{\tau} - f_{ii}^{\tau}\right]\left[U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_{i-1} | \theta_i, \theta_{i+1})\right] \leq \frac{|\lambda_i - \lambda_{i-1}|\left(1 - e^{-rT}\right) 2\overline{u}}{r}.$$

Thus, we have that for any $\varepsilon > 0$, there exists a τ_0 such that for all $\tau \leq \tau_0$,

$$\begin{aligned} &\sum_{k=1}^{N} \left(f_{i-1,k}^{\tau} - f_{ik}^{\tau}\right)\left[U(\theta_k | \theta_i, \theta_i) - U(\theta_k | \theta_i, \theta_{i+1})\right] \\ &\quad \leq f_{ii}^{\tau}\left[\left(U(\theta_{i-1} | \theta_i, \theta_i) - U(\theta_i | \theta_i, \theta_i)\right) - \left(U(\theta_{i-1} | \theta_i, \theta_{i+1}) - U(\theta_i | \theta_i, \theta_{i+1})\right)\right] \\ &\qquad + \kappa_2 + \varepsilon, \end{aligned} \tag{37}$$

where $\kappa_2 = \kappa_1 + \frac{(1-e^{-r(T-2)})|\lambda_i - \lambda_{i-1}|2\overline{u}}{r}$. We are interested in bounding the right-hand side of the above inequality from above. Define $D(\theta_i|\theta_i, \theta_i)$ and $D(\theta_i|\theta_i, \theta_{i+1})$ as the terms in the square brackets in (13) and (14), respectively.

We need two additional auxiliary results.

Lemma A7. *Suppose $D < 1$. There exists a $\tau_0 > 0$ such that for $\tau \leq \tau_0$, $q(\theta_i|\theta_i, \theta_i) \leq q(\theta_i|\theta_i, \theta_{i+1})$ and $q(\theta_j|h_i, h^k) \leq q(\theta_j|h_i', h^k)$ for any history h^k, where $h_i = \{\theta_i, \theta_i, \theta_i\}$ and $h_i' = \{\theta_i, \theta_{i+1}, \theta_i\}$.*

Proof. Note that $q(\theta_i|\theta_i, \theta_i) \leq q(\theta_i|\theta_i, \theta_{i+1})$ if and only if $\frac{D(\theta_i|\theta_i, \theta_{i+1})}{D(\theta_i|\theta_i, \theta_i)} \leq 1$. As $\tau \to 0$, the limit of the ratio is well defined; the denominator converges to 1, and the numerator converges to $D < 1$. Similarly, $q(\theta_j|h_i, h^k) \leq q(\theta_j|h_i', h^k)$ if and only if $\frac{D(\theta_j|h_i, h^k)}{D(\theta_j|\theta_i^3, \theta_i, h^k)} \leq 1$. This ratio is exactly equal to $\frac{D(\theta_i|\theta_i, \theta_{i+1})}{D(\theta_i|\theta_i, \theta_i)}$, which we just proved is less than 1 as $\tau \to 0$. □

Lemma A8. *Suppose $D < 1$. There exists a $\tau_0 > 0$ such that for $\tau \leq \tau_0$, $\Delta u(\theta_i|\theta_i, \theta_i) \leq \Delta u(\theta_i|\theta_i, \theta_{i+1})$ and $\Delta u(\theta_j|h_i, h^k) \leq \Delta u(\theta_j|h_i', h^k)$ for any h^k, where $h_i = \{\theta_i, \theta_i, \theta_i\}$ and $h_i' = \{\theta_i, \theta_{i+1}, \theta_i\}$.*

Proof. We have

$$\Delta u(\theta_i|\theta_i, \theta_i) = \int_{\theta_{i+1}}^{\theta_i} u_\theta\big(x, q(\theta_i|\theta_i, \theta_i)\big)\, dx \leq \int_{\theta_{i+1}}^{\theta_i} u_\theta(x, q(\theta_i|\theta_i, \theta_{i+1})\, dx = \Delta u(\theta_i|\theta_i, \theta_{i+1})$$

and

$$\begin{aligned}\Delta u\big(\theta_j|h_i, h^k\big) &= \int_{\theta_{j+1}}^{\theta_j} u_\theta(x, q\big(\theta_j|h_i, h^k\big)\, dx \\ &\leq \int_{\theta_{j+1}}^{\theta_j} u_\theta(x, q\big(\theta_j|h_i', h^k\big)\, dx = \Delta u\big(\theta_j|h_i', h^k\big),\end{aligned}$$

where the inequalities above follow from Lemma A5 and the assumption $u_{\theta q} > 0$. □

Let h_i^t be a history in which the realization is θ_i in every period for t periods. We have

$$\begin{aligned}&\big[U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i)\big] - \big[U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})\big] \\ &\quad \leq \sum_{t=3}^{K+2} \big(\delta \Delta F_\tau(\theta_i|\theta_i)\big)^{t-3} \left[\int_{\theta_i}^{\theta_{i-1}} \left[\begin{array}{c} u_\theta(x, q\big(\theta_i|\theta_i, \theta_i, h_i^{t-3}\big) \\ - u_\theta(x, q\big(\theta_i|\theta_i, \theta_{i+1}, h_i^{t-3}\big)\end{array}\right] dx\right].\end{aligned} \tag{38}$$

To obtain inequality (38), we eliminate all quantities after $K+2$ interactions and all quantities along nonconstant histories before $K+2$, where K will be chosen at a later stage. The inequality follows from the fact that we can sign the omitted terms using Lemma A8.

Let $q_1 = \lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_i)$ and $q_2 = \lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_{i+1})$. Then

$$s_q(\theta_i, q_1) = \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta q}(x, q_1)\, dx,$$

$$s_q(\theta_i, q_2) \le \frac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \cdot D \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)\, dx.$$

Note that, as in the proof of Proposition 3, $q_2 > q_1$. Also, it is easy to see that $\lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_i, h_i^{t-3}) = q_1$ and $\lim_{\tau\to 0} q(\theta_i|\theta_i, \theta_{i+1}, h_i^{t-3}) = q_2$. So, for any $\varepsilon > 0$, we have that there exists τ_1 such that for $\tau \le \tau_1$,

$$\begin{aligned}&\sum_{t=3}^{K+2} \big(\delta \Delta F_\tau(\theta_i|\theta_i)\big)^{t-3} \bigg[\int_{\theta_i}^{\theta_{i-1}} \big[u_\theta\big(x, q\big(\theta_i|\theta_i, \theta_i, h_i^{t-3}\big) - u_\theta\big(x, q\big(\theta_i|\theta_i, \theta_{i+1}, h_i^{t-3}\big)\big]\, dx\bigg]\\ &\quad\le \frac{1 - \big[e^{-r\tau}\Delta F_\tau(\theta_i|\theta_i)\big]^K}{1 - e^{-r\tau}\Delta F_\tau(\theta_i|\theta_i)} \int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_1) - u_\theta(x, q_2)\big]\, dx + \varepsilon.\end{aligned}$$

Now $\lim_{\delta\to 1}(1 - \delta^K)/(1 - \delta) = K$. Since $e^{-r\tau}\Delta F_\tau(\theta_i|\theta_i) \to 1$ as $\tau \to 0$, we must have $\frac{1-[e^{-r\tau}\Delta F_\tau(\theta_i|\theta_i)]^K}{1-e^{-r\tau}\Delta F_\tau(\theta_i|\theta_i)} \to K$ as $\tau \to 0$. It follows that for any $\varepsilon > 0$, there must be a τ_2 such that for $\tau \le \tau_2$,

$$\begin{aligned}&U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i) - U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})\\ &\quad\le K \int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_1) - u_\theta(x, q_2)\big]\, dx + \varepsilon.\end{aligned} \tag{39}$$

Thus, by inequalities (36), (37), and (39), we get that a necessary condition for $IC_{i-1,i+1}(\theta_i)$ to hold as $\tau \to 0$ is

$$\int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_1) - u_\theta(x, q_5)\big]\, dx + K \int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_1) - u_\theta(x, q_2)\big]\, dx + \kappa_2 \ge 0, \tag{40}$$

where $q_5 = \lim_{\tau\to 0} q(\theta_{i+1}|\theta_i)$. Finally, choose K large enough so that inequality (40) is violated; for example,

$$K = \left\lceil 1 + \frac{\kappa_2 + \int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_1) - u_\theta(x, q_5)\big]\, dx}{\int_{\theta_i}^{\theta_{i-1}} \big[u_\theta(x, q_2) - u_\theta(x, q_1)\big]\, dx} \right\rceil,$$

where $\lceil x \rceil$ is the smallest integer larger than x. Therefore, $\exists \bar{\tau} > 0$ such that $IC_{i-1,i+1}(\theta_i)$ is violated for all $\tau \le \bar{\tau}$.

Case 2: $D > 1$. Following steps as in Case 1, we can analogously show that in this case, $\exists \bar{\tau} > 0$ such that the local upward constraint $IC_{i+1,i}(\theta_{i+1})$ is violated for all $\tau \le \bar{\tau}$.

A.6 *Proof of Corollary 2*

Recollect that a (finite state) Markov process in continuous time sampled at discrete intervals can be described by (Λ, P), where $\Lambda = (\lambda_0, \dots, \lambda_n)$ is the vector of rates of jump, and P is the transition probability conditional on the jump, with $P_{ii} = 0$. Fixing τ, the interval of sampling, it can be shown that the Markov process has an exponential representation given by

$$f_\tau = e^{\tau Q} = \sum_{k=0}^{\infty} \frac{\tau^k Q^k}{k!},$$

where $Q = \Lambda(P - I)$. In particular, note that

$$Q_{ii} = -\lambda_i \quad \text{and} \quad Q_{ij} = P_{ij}\lambda_i.$$

Define $\lambda = \max_i \lambda_i$. Now, Q can be rewritten as $Q = \lambda(\widehat{P} - I)$, where

$$\widehat{P}_{ij} = \begin{cases} \dfrac{Q_{ij}}{\lambda} & i \neq j, \\ 1 + \dfrac{Q_{ii}}{\lambda} & i = j. \end{cases}$$

Thus,

$$\widehat{P}_{ij} = \begin{cases} \dfrac{\lambda_i}{\lambda} P_{ij} & i \neq j, \\ 1 - \dfrac{\lambda_i}{\lambda} & i = j. \end{cases}$$

Therefore, the Markov process can be rewritten as

$$f_\tau = e^{\tau Q} = e^{-\lambda t I} e^{\tau\lambda\widehat{P}} = e^{-\lambda\tau} \sum_{k=0}^{\infty} \frac{(\lambda\tau)^k}{k!} \widehat{P}.^k$$

First, we establish a limit on the magnitudes of the persistent parameters Λ, the proof of which is provided in the Supplemental Material. Fix $\tau = 1$. Let $\eta_i = \lim_{\lambda\to 0} \frac{\lambda_i}{\lambda}$.

Lemma A9. *We have* $\lim_{\lambda\to 0} \frac{1-f_{ii}}{\lambda} \to \eta_i$ *and* $\lim_{\lambda\to 0} \frac{f_{ij}}{\lambda} = \eta_i P_{ij}$.

For the proof, see the Supplemental Material.

The dynamic distortion for history $h_i' = \{\theta_i, \theta_{i+1}, \theta_i\}$ is proportional to

$$\Psi_i(f) = \frac{\Delta F(\theta_{i+1}|\theta_i)}{f(\theta_{i+1}|\theta_i)} \cdot \frac{\Delta F(\theta_i|\theta_{i+1})}{f(\theta_i|\theta_{i+1})}.$$

Following steps as in Section A.3, and using Lemma A9, we get

$$\lim_{\lambda\to 0} \Psi_i(f) = \frac{\sum_{k=i+1}^{N} [R_{ik} - R_{i-1,k}]}{R_{i,i+1}} \cdot \frac{(R_{i+1,i} + \eta_i - \eta_{i+1} - R_{i,i+1}) + \sum_{k=i+2}^{N} [R_{ik} - R_{i+1,k}]}{R_{i,i+1}},$$

where $R_{i,j} = \eta_i P_{ij}$. Following analogous steps, we get that $D = \lim_{\lambda\to 0} \Psi_i(f)$ is generically different than 1. We focus here on the case in which $D < 1$; the other case in which $D > 1$ is analogous and omitted. Recall from (36) that $IC_{i-1,i+1}(\theta_i)$ holds if and only if

$$\int_{\theta_i}^{\theta_{i-1}} \begin{bmatrix} u_\theta\big(x, q(\theta_i|\theta_i)\big) \\ -\, u_\theta\big(x, q(\theta_{i+1}|\theta_i)\big) \end{bmatrix} dx + \delta \sum_{k=0}^{N} \begin{bmatrix} (f_{i-1,k} - f_{ik}) \\ \cdot \big(U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1})\big) \end{bmatrix} \geq 0.$$

Following the same steps as in Proposition 4, we have that for any ε, there is a λ_0 such that for $\lambda \leq \lambda_0$,

$$\begin{aligned} &\sum_{k=1}^{N} (f_{i-1,k} - f_{ik})\big[U(\theta_k|\theta_i, \theta_i) - U(\theta_k|\theta_i, \theta_{i+1})\big] \\ &\quad \leq f_{ii}\big[\big(U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i)\big) \\ &\qquad - \big(U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})\big)\big] + \kappa + \varepsilon, \end{aligned} \tag{41}$$

where κ is a finite constant. Moreover, using Lemma A8 and (2), we can now bound the right-hand side of (41). Let $\tilde{h}^t(\theta)$ be a history in which the realization is θ in every period for t periods. Then

$$\begin{aligned} &\big[U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i)\big] - \big[U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})\big] \\ &\quad \leq \sum_{t=3}^{T-2} \big(\delta \Delta F_\tau(\theta_i|\theta_i)\big)^{t-3} \left[\int_{\theta_i}^{\theta_{i-1}} \begin{bmatrix} u_\theta\big(x, q\big(\theta_i|\theta_i, \theta_i, \tilde{h}^{t-3}(\theta_i)\big) \\ -\, u_\theta\big(x, q\big(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)\big) \end{bmatrix} dx\right]. \end{aligned} \tag{42}$$

The inequality follows from the fact that all quantities that comprise the expected utility can be ranked using the λ analog of Lemma A7. We derive (42) using (2) and ignoring all histories following θ_i, θ_i and θ_i, θ_{i+1} in which the type changes. The ignored terms are all nonnegative and, thus, do not change the direction of the inequality. Note that a key difference from the proof of Proposition 4 is that now we do not need to truncate after K period since T is finite. Define q_1 as the unique solution of

$$s_q(\theta_i, q_1) = \frac{1 - \sum\limits_{k=i}^{N} \mu_k}{\mu_i} \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta q}(x, q_1)\, dx$$

and define q_2 as

$$s_q(\theta_i, q_2) \leq \frac{1 - \sum\limits_{k=i}^{N} \mu_k}{\mu_i} \cdot D \cdot \int_{\theta_i}^{\theta_{i-1}} u_{\theta,q}(x, q_2)\, dx.$$

As in the proof of Proposition 3, we have $q_2 > q_1$. As $\lambda \to 0$, we have $q(\theta_i|\theta_i, \theta_i, \tilde{h}^{t-3}(\theta_i)) \to q_1$ for all histories $\theta_i, \theta_i, \tilde{h}^{t-3}(\theta_i)$ for $t \leq T - 2$ and $q(\theta_i|\theta_i, \theta_{i+1},$

$\tilde{h}^{t-3}(\theta_i)) \to q_2$ for all histories $\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)$ for $t \le T-2$. So for any $\varepsilon > 0$, there must be a λ_1 such that for $\lambda \le \lambda_1$,

$$\sum_{t=3}^{T-2} \bigl(\delta \Delta F_\tau(\theta_i|\theta_i)\bigr)^{t-3} \biggl[\int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q\bigl(\theta_i|\theta_i, \theta_i, \tilde{h}^{t-3}(\theta_i)\bigr) - u_\theta(x, q\bigl(\theta_i|\theta_i, \theta_{i+1}, \tilde{h}^{t-3}(\theta_i)\bigr)\bigr]\,dx\biggr]$$
$$\le \frac{1-\delta^{T-2}}{1-\delta} \int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_1) - u_\theta(x, q_2)\bigr]\,dx + \varepsilon. \tag{43}$$

It follows from (42) and (43) that for any $\varepsilon > 0$, there must be a λ_2 such that for $\lambda \le \lambda_2$,

$$\delta\bigl[U(\theta_{i-1}|\theta_i, \theta_i) - U(\theta_i|\theta_i, \theta_i) - U(\theta_{i-1}|\theta_i, \theta_{i+1}) - U(\theta_i|\theta_i, \theta_{i+1})\bigr]$$
$$\le \delta \frac{1-\delta^{T-2}}{1-\delta} \int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_1) - u_\theta(x, q_2)\bigr]\,dx + \varepsilon. \tag{44}$$

Putting together (36) and (41)–(44), and since ε can be chosen to be arbitrarily small, $IC_{i-1,i+1}(\theta_i)$ holds as $\lambda \to 0$ only if

$$\int_{\theta_i}^{\theta_{i-1}} \left[\begin{array}{c} u_\theta\bigl(x, q(\theta_i|\theta_i)\bigr) \\ - u_\theta\bigl(x, q(\theta_{i+1}|\theta_i)\bigr) \end{array}\right] dx + \delta \frac{1-\delta^{T-2}}{1-\delta} \int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_1) - u_\theta(x, q_2)\bigr]\,dx \ge 0.$$

Now note that $\lim_{\lambda \to 0} q(\theta_i|\theta_i) = q_1$ and let $\lim_{\lambda \to 0} q(\theta_{i+1}|\theta_i) = q_5$. Assume that

$$\frac{\delta(1-\delta^T)}{1-\delta} > \frac{\displaystyle\int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_1) - u_\theta(x, q_5)\bigr]}{\displaystyle\int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_2) - u_\theta(x, q_1)\bigr]\,dx} = \zeta^* \quad \text{(a constant)}.$$

Then as $\lambda \to 0$, we have that $IC_{i-1,i+1}(\theta_i)$ is violated.

It is worth observing that ζ^* depends on the primitives of the environment sans $\Lambda = (\lambda_i)_{i=1}^N$; it is, in fact, independent of Λ. For example, if we assume $u = \theta q$, then

$$\zeta^* = \frac{\displaystyle\int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_1) - u_\theta(x, q_5)\bigr]\,dx}{\displaystyle\int_{\theta_i}^{\theta_{i-1}} \bigl[u_\theta(x, q_2) - u_\theta(x, q_1)\bigr]\,dx}$$

$$= \frac{\left(\theta_i - \dfrac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \Delta\theta\right) - \left(\theta_{i+1} - \dfrac{1 - \sum_{k=i}^{N} \mu_k}{\mu_i} \tilde{D} \Delta\theta\right)}{\left(\theta_i - \dfrac{\sum_{k=i}^{N} \mu_k}{\mu_i} D \Delta\theta\right) - \left(\theta_i - \dfrac{\sum_{k=i}^{N} \mu_k}{\mu_i} \Delta\theta\right)},$$

where $D = \lim_{\lambda\to 0} \Psi_i(f)$ and $\tilde{D} = \lim_{\lambda\to 0} \frac{\Delta F(\theta_{i+1}|\theta_i)}{f(\theta_{i+1}|\theta_i)}$ are functions of (P_{ij}) and (η_i), and independent of Λ.

A.7 *Proof of Proposition 5*

For simplicity of notation, we present the proof for Mussa and Rosen (1978) preferences: $u(\theta, q) = \theta q$ and $c(q) = (1/2)q^2$. The more general case follows analogously. We proceed in two steps.

Step 1. We say that a quantity $q(\theta_i|h^{t-1})$ is distorted downward (respectively, upward) if $q(\theta_i|h^{t-1}) \le \theta_i$ (respectively, $q(\theta_i|h^{t-1}) > \theta_i$). We first show that in the optimal monotonic contract, distortions are all downward. Consider the constraint set of (22), as described by $\mathcal{M}$. Define

$$\Gamma(h^{t-1}) = \left\{ \begin{array}{c} \widehat{h}^{t-1} | \exists k \le t-1 \text{ s.t. given } h_k^{t-1} = \theta_l \text{ for some } l = 0, 1, \dots, N-1, \\ \text{we have } \widehat{h}_k^{t-1} = \theta_{l+1} \text{ and } h_j^{t-1} = \widehat{h}_j^{t-1}\ \forall j \neq k \end{array} \right\}.$$

Thus, $\Gamma(h^{t-1})$ is the set of histories that differ from h^{t-1} only once; that is, the type in period $k \le t-1$ is replaced by the contiguous lower type. It is easy to see that a contract is monotonic if and only if, for any history h^{t-1}, (i) $q(\theta_i|h^{t-1}) \ge q(\theta_{i+1}|h^{t-1})$ for all $i < N$ and (ii) $q(\theta_i|h^{t-1}) \ge q(\theta_i|\widehat{h}^{t-1})$ for all i and for all $\widehat{h}^{t-1} \in \Gamma(h^{t-1})$.

Next, we introduce the following complete order on the set of all histories at time t. For any two histories h^{t-1} and $\widehat{h}^{t-1}$, let $\tau^*(h^{t-1}, \widehat{h}^{t-1})$ be the first period in which they diverge: $\tau^*(h^{t-1}, \widehat{h}^{t-1}) = \min_j\{0 \le j \le t-1 \text{ s.t. } h_j^{t-1} \neq \widehat{h}_j^{t-1}\}$, with $\tau^*(h^{t-1}, \widehat{h}^{t-1}) = t-1$ if $h^{t-1} = \widehat{h}^{t-1}$. We say that $h^{t-1} \succeq^* \widehat{h}^{t-1}$ if $h^{t-1}_{\tau^*(h^{t-1},\widehat{h}^{t-1})} \ge \widehat{h}^{t-1}_{\tau^*(h^{t-1},\widehat{h}^{t-1})}$, i.e., if it is higher at the first point of divergence. It is easy to verify that the order $\succeq^*$ is complete, so without loss of generality we can order the histories at time t from largest ($\overline{h}^{t-1}$) to smallest ($\underline{h}^{t-1}$), where the largest (smallest) history has all realizations equal to θ_0 (θ_N). Also, note that $h^{t-1} \succeq^* \widehat{h}^{t-1}$ for all $\widehat{h}^{t-1} \in \Gamma(h^{t-1})$.

Consider period t and the smallest history of length $t-1$ (denoted $\underline{h}^{t-1}$) in which all the realizations are θ_N. It is immediate to see that $q(\theta_N|\underline{h}^{t-1})$ cannot be distorted upward. To see this, note that $q(\theta_N|\underline{h}^{t-1})$ is on the left-hand side of no constraint.[39] If it were distorted upward, then a marginal decrease in $q(\theta_N|\underline{h}^{t-1})$ would relax all constraints and increase surplus. Now consider $q(\theta_{N-1}|\underline{h}^{t-1})$. This quantity appears on the left-hand side of only one constraint, $q(\theta_{N-1}|\underline{h}^{t-1}) \ge q(\theta_N|\underline{h}^{t-1})$. If this constraint is not binding, then by the argument presented above, $q(\theta_{N-1}|\underline{h}^{t-1}) \le \theta_{N-1}$. Assume it is binding. In this case, $q(\theta_{N-1}|\underline{h}^{t-1}) = q(\theta_N|\underline{h}^{t-1}) \le \theta_N \le \theta_{N-1}$. Proceeding inductively with a similar argument, we can prove that $q(\theta_i|\underline{h}^{t-1}) \le \theta_i$ for all i.

Note that the case for first period quantities, when the history is just the empty set, is already covered by the above paragraph. Thus, now we consider $t \ge 2$. Assume, as an induction step, that there is a history $\widehat{h}^{t-1}$, where $\widehat{h}^{t-1} \succeq^* \underline{h}^{t-1}$ such that $\widehat{h}^{t-1} \succeq^* h^{t-1} \succeq^* \underline{h}^{t-1}$ implies $q(\theta_i|h^{t-1}) \le \theta_i$ for all i. Let us also introduce a useful definition. For any

[39] We say that a quantity is on the left-hand side of a given constraint if, in that constraint, it must be larger than some other quantity.

h^{t-1} with $\overline{h}^{t-1} \succeq^* h^{t-1}$, $h^{t-1} \neq \overline{h}^{t-1}$, and $t \geq 2$, define $[h^{t-1}]^+$ to be the smallest t-period history larger than h^{t-1} according to the order $\succeq^*$ in the following inductive way. If $t = 2$, then $[h^{t-1}]^+ = \{\kappa_{t-1}(h^{t-1}), h^{t-1}_{t-1} + \Delta\theta\}$; if $t > 2$, then

$$[h^{t-1}]^+ = \begin{cases} (\kappa_{t-1}(h^{t-1}), h^{t-1}_{t-1} + \Delta\theta) & \text{if } h^{t-1}_{t-1} < \theta_0, \\ ([\kappa_{t-1}(h^{t-1})]^+, \theta_N) & \text{if } h^{t-1}_{t-1} = \theta_0, \end{cases}$$

where κ_s projects the first s elements of a vector.[40] We intend to show that $q(\theta_i|[\widehat{h}^{t-1}]^+) \leq \theta_i$ for all i. Now $q(\theta_N|[\widehat{h}^{t-1}]^+)$ appears on the left-hand side of $q(\theta_N|[\widehat{h}^{t-1}]^+) \geq q(\theta_N|\widetilde{h}^{t-1})$ for all $\widetilde{h}^{t-1} \in \Gamma([\widehat{h}^{t-1}]^+)$. If none of these constraints binds, then as before, we have the desired inequality. Suppose at least one of them binds. Clearly, by the definition of $[\widehat{h}^{t-1}]^+$, we have $\widehat{h}^{t-1} \succeq^* \widetilde{h}^{t-1}$ for all $\widetilde{h}^{t-1} \in \Gamma([\widehat{h}^{t-1}]^+)$. Thus, by the induction hypothesis, $q(\theta_N|\widetilde{h}^{t-1}) \leq \theta_N$ for all $\widetilde{h}^{t-1} \in \Gamma([\widehat{h}^{t-1}]^+)$. Since the inequality constraint binds for some $\widetilde{h}^{t-1}$, we have $q(\theta_N|[\widehat{h}^{t-1}]^+) = q(\theta_N|\widetilde{h}^{t-1}) \leq \theta_N$.

Next, consider $q(\theta_{N-1}|[\widehat{h}^{t-1}]^+)$. It appears on the left-hand side of $q(\theta_{N-1}|[\widehat{h}^{t-1}]^+) \geq q(\theta_N|[\widehat{h}^{t-1}]^+)$ and $q(\theta_{N-1}|[\widehat{h}^{t-1}]^+) \geq q(\theta_{N-1}|\widetilde{h}^{t-1})$ for all $\widetilde{h}^{t-1} \in \Gamma([\widehat{h}^{t-1}]^+)$. If none of these constraints binds, then as before, we have the desired inequality. If the first one binds, then $q(\theta_{N-1}|[\widehat{h}^{t-1}]^+) \leq \theta_N < \theta_{N-1}$. If any of the latter constraints binds, then by invoking the induction hypothesis, as argued in the case above, we have the desired inequality. Proceeding inductively, we can show $q(\theta_i|h^{t-1}) \leq \theta_i$ for all i and h^{t-1}.

Step 2. We now prove that the allocation is asymptotically efficient. Consider problem (22). From this problem, eliminate the constraint $q(\theta_0|h^0) \geq q(\theta_1|h^0)$ and all the monotonicity constraints that involve quantities following a history in which the agent reports to be type θ_0. It is easy to see that in this problem the quantities offered after the agent reports (or has reported) to be θ_0 are efficient: $q(\theta_i|h^{t-1}) = \theta_i$ for $i = 0$ and/or $\forall h^{t-1} \in \overline{H}^{t-1}$, $t \geq 2$, where $\overline{H}^{t-1} = \{h^{t-1}|\exists \tau \leq t-1 \text{ s.t. } h^{t-1}_\tau = \theta_0\}$. Following the same approach as in Step 1, it can be shown that the solution of this relaxed problem is monotonic and so it coincides with the optimal monotonic contract. Since the probability of the event in which no type realization in t periods is equal to θ_0 converges to zero as $t \to \infty$, this solution is asymptotically efficient, and so is the optimal monotonic contract.

A.8 *Proof of Proposition 6*

The first part of the result, viz. $\lim_{r\to 0} \lim_{\lambda\to 0} \pi_m(\Lambda, r) = \lim_{r\to 0} \lim_{\lambda\to 0} \pi^*(\Lambda, r)$, follows from the fact that for each fixed r, the optimal monotonic contract and the optimal contract both converge to the repetition of the static optimum, that is, $\lim_{\lambda\to 0} \pi_m(\Lambda, r) = \lim_{\lambda\to 0} \pi^*(\Lambda, r)$.

We now prove the second part. For a T horizon and discount factor r, let $S^e(\Lambda, r, T)$ be the total expected surplus generated in the efficient contract, let $S^e_i(\Lambda, r, T)$ be the surplus obtained with the efficient contract conditional on being type i at $t = 1$, let

[40] Recollect that h^{t-1} is a vector of length t: $h^{t-1} = (h^{t-1}_0, h^{t-1}_1, \ldots, h^{t-1}_{t-1})$, where $h^{t-1}_0 = \varnothing$. So, $\kappa_{t-1}(h^{t-1}) = (h^{t-1}_0, \ldots, h^{t-1}_{t-2})$.

$\Pi^*(\Lambda, r, T)$ be the profit in the corresponding optimal contract, and let $\Pi_m(\Lambda, r, T)$ be the profit in the corresponding optimal monotonic contract. Define also $s^e(\Lambda, r, T) = (1 - e^{-\tau r})S^e(\Lambda, r, T)$ and, similarly, $s_i^e(\Lambda, r, T)$, $\pi_m(\Lambda, r, T)$, and $\pi^*(\Lambda, r, T)$. Finally, let $s^e(\Lambda, r) = \lim_{T\to\infty} s(\Lambda, r, T)$, $s_i^*(\Lambda, r) = \lim_{T\to\infty} s_i^*(\Lambda, r, T)$, $\pi_m(\Lambda, r) = \lim_{T\to\infty} \pi_m(\Lambda, r, T)$, and $\pi^* = \lim_{T\to\infty} \pi^*(\Lambda, r, T)$.

Now the average per period profit $\pi_m(\Lambda, r)$ must be larger than or equal to the average profit obtained by offering the efficient quantity at cost and charging a fixed per period "entry fee" equal to $s_N^*(\Lambda, r)$, since this is an incentive compatible monotonic contract. This implies

$$\pi_m(\Lambda) = \lim_{r\to 0} \pi_m(\Lambda, r) \geq \lim_{r\to 0} s_N^*(\Lambda, r) \geq s^e(\Lambda) - \varepsilon,$$

where $s^e(\Lambda) = \lim_{r\to 0} s^e(\Lambda, r)$ and the last inequality follows from the fact that the process is ergodic, so $\lim_{r\to 0} |s_i^e(\Lambda, r) - s_j^e(\Lambda, r)| \to 0$ as $r \to 0$. We conclude that for any ε, we can choose an r^* such that $\pi^*(\Lambda, r) \geq \pi_m(\Lambda, r) \geq s^e(\Lambda, r) - \varepsilon$ for $r \leq r^*$. Since $s^e(\Lambda, r) \geq \pi^*(\Lambda, r) \geq \pi_m(\Lambda, r)$, we have that for any ε, we can choose an r^* such that $|\pi^*(\Lambda, r) - \pi_m(\Lambda, r)| \leq \varepsilon$ for $r \leq r^*$. Thus, $\lim_{r\to 0} \pi_m(\Lambda, r) = \lim_{r\to 0} \pi^*(\Lambda, r)$ and the second part of the result follows.

References

Athey, Susan and Ilya Segal (2013), "An efficient dynamic mechanism." *Econometrica*, 81, 2463–2485. [1458]

Baron, David P. and David Besanko (1984), "Regulation and information in a continuing relationship." *Information Economics and Policy*, 1, 267–302. [1436, 1442, 1445, 1458]

Battaglini, Marco (2005), "Long-term contracting with Markovian consumers." *American Economic Review*, 95, 637–658. [1436, 1441, 1442, 1443, 1444, 1446, 1447, 1458]

Battaglini, Marco (2007), "Optimality and renegotiation in dynamic contracting." *Games and Economic Behavior*, 60, 213–246. [1458]

Battaglini, Marco and Stephen Coate (2008), "Pareto efficient income taxation with stochastic abilities." *Journal of Public Economics*, 92, 844–868. [1455]

Battaglini, Marco and Rohit Lamba (2015), "Optimal dynamic contracting: The first-order approach and beyond." Unpublished paper, SSRN 2697589. [1437, 1446]

Bergemann, Dirk and Maher Said (2011), "Dynamic auctions: A survey." In *Wiley Encyclopedia of Operations Research and Management Science* (James J. Cochran, Louis Anthony Cox Jr., Pinar Keskinocak, Jeffrey P. Kharoufeh, and J. Cole Smith, eds.), 1511–1522, John Wiley and Sons Ltd., Hoboken, United States. [1458]

Bergemann, Dirk and Juuso Välimäki (2010), "The dynamic pivot mechanism." *Econometrica*, 78, 771–789. [1458]

Besanko, David (1985), "Multiperiod contracts between principal and agent with adverse selection." *Economic Letters*, 17, 33–37. [1436, 1441, 1442, 1446, 1447, 1458]

Bolton, Patrick and Mathias Dewatripont (2005), *Contract Theory*. MIT press, Cambridge, Massachusetts. [1458]

Chassang, Sylvain (2013), "Calibrating incentive contracts." *Econometrica*, 81, 1935–1971. [1459]

Clementi, Gian Luca and Hugo A. Hopenhayn (2006), "A theory of financing constraints and firm dynamics." *Quarterly Journal of Economics*, 121, 229–265. [1458]

Courty, Pascal and Hao Li (2000), "Sequential screening." *Review of Economic Studies*, 67, 697–717. [1436, 1441, 1458]

Dewatripont, Mathias (1989), "Renegotiation and information revelation over time: The case of optimal labor contracts." *Quarterly Journal of Economics*, 104, 589–619. [1458]

Esö, Péter and Balázs Szentes (2007), "Optimal information disclosure in auctions and the handicap auction." *Review of Economic Studies*, 74, 705–731. [1459]

Esö, Péter and Balász Szentes (2017), "Dynamic contracting: An irrelevance result." *Theoretical Economics*, 12, 109–139. [1454]

Farhi, Emmanuel and Iván Werning (2013), "Insurance and taxation over the life cycle." *Review of Economic Studies*, 80, 596–635. [1455, 1459]

Garrett, Daniel and Alessandro Pavan (2012), "Managerial turnover in a changing world." *Journal of Political Economy*, 120, 879–925. [1458]

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski (2016), "Redistribution and social insurance." *American Economic Review*, 106, 359–386. [1455]

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2016), "What do data on millions of U.S. workers reveal about life-cycle earnings risk?" Unpublished paper, University of Minnesota, FRB of Minneapolis, and NBER. [1454, 1459]

Guvenen, Fatih, Serdar Ozkan, and Jae Song (2014), "The nature of countercyclical income risk." *Journal of Political Economy*, 122, 621–660. [1454]

Hart, Oliver and Jean Tirole (1988), "Contract renegotiation and coasian dynamics." *Review of Economic Studies*, 55, 509–540. [1458]

Hartline, Jason (2012), "Approximation in mechanism design." *American Economic Review: Papers & Proceedings*, 102, 330–336. [1459]

Kapička, Marek (2013), "Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach." *Review of Economic Studies*, 80, 1027–1054. [1455, 1459]

Krähmer, Daniel and Roland Strausz (2015), "Optimal sales contracts with withdrawal rights." *Review of Economic Studies*, 82, 762–790. [1459]

Laffont, Jean-Jacques and David Martimort (2002), *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton, New Jersey. [1458]

Laffont, Jean-Jacques and Jean Tirole (1990), "Adverse selection and renegotiation in procurement." *Review of Economic Studies*, 57, 597–625. [1458]

Laffont, Jean-Jacques and Jean Tirole (1996), "Pollution permits and compliance strategies." *Journal of Public Economics*, 62, 85–125. [1436, 1458]

Li, Hao and Xianwen Shi (2017), "Discriminatory information disclosure." *American Economic Review*, 107, 3363–3385. [1459]

Madarász, Kristóf and Andrea Prat (2017), "Sellers with misspecified models." *Review of Economic Studies*, 84, 790–815. [1459]

Maestri, Lucas (2017), "Dynamic contracting under adverse selection and renegotiation." *Journal of Economic Theory*, 171, 136–173. [1458]

Milgrom, Paul (2004), *Putting Auction Theory to Work*. Cambridge University Press. [1458]

Mussa, Michael and Sherwin Rosen (1978), "Monopoly and product quality." *Journal of Economic Theory*, 18, 301–317. [1438, 1443, 1447, 1456, 1477]

Myerson, Roger (1986), "Multistage games with communication." *Econometrica*, 54, 323–358. [1439]

Myerson, Roger B. (1981), "Optimal auction design." *Mathematics of Operations Research*, 6, 58–73. [1442, 1456]

Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014), "Dynamic mechanism design: A Myersonian approach." *Econometrica*, 82, 601–653. [1436, 1442, 1444, 1458]

Rey, Patrick and Bernard Salanie (1990), "Long-term, short-term and renegotiation: On the value of commitment in contracting." *Econometrica*, 58, 597–619. [1458]

Ross, Sheldon M. (2014), *Introduction to Probability Models*, 11th edition. Academic Press, Amsterdam, The Netherlands. [1466]

Royden, Halsey L. (1988), *Real Analysis*, third edition. Macmillan, New York. [1450]

Rustichini, Aldo and Asher Wolinsky (1995), "Learning about variable demand in the long run." *Journal of Economic Dynamics and Control*, 19, 1283–1292. [1458]

Stole, Lars (2001), "Lectures on the theory of contracts and organizations." Unpublished paper, University of Chicago, Booth School of Business. [1438, 1441, 1443, 1458]

Strulovici, Bruno (2011), "Renegotiation-proof contracts with moral hazard and persistent private information." Unpublished paper, Department of Economics, Northwestern University. [1458]

Townsend, Robert M. (1982), "Optimal multiperiod contracts and the gain from enduring relationships under private information." *Journal of Political Economy*, 90, 1166–1186. [1458]

Williams, Noah (2011), "Persistent private information." *Econometrica*, 79, 1233–1275. [1458, 1459]

Zhang, Yuzhe (2009), "Dynamic contracting with persistent shocks." *Journal of Economic Theory*, 144, 635–675. [1459]

Co-editor Johannes Hörner handled this manuscript.

Manuscript received 20 November, 2015; final version accepted 9 January, 2019; available online 6 March, 2019.