

Robust Predictions in Dynamic Screening*

Daniel Garrett[†]

Alessandro Pavan[‡]

Juuso Toikka[§]

April, 2018

Abstract

We characterize properties of optimal dynamic mechanisms using a variational approach that permits us to tackle directly the full program. This allows us to make predictions for a considerably broader class of stochastic processes than can be handled by the “first-order, Myersonian, approach,” which focuses on local incentive compatibility constraints and has become standard in the literature. Among other things, we characterize the dynamics of optimal allocations when the agent’s type evolves according to a stationary Markov processes, and show that, provided the players are sufficiently patient, optimal allocations converge to the efficient ones in the long run.

JEL classification: D82

Keywords: asymmetric information, dynamic mechanism design, stochastic processes, convergence to efficiency, variational approach

*For useful comments and suggestions we thank participants at various conferences and workshops where the paper has been presented, and in particular Nicola Pavoni. Pavan also thanks Bocconi University for its hospitality during the 2017-2018 academic year and the National Science Foundation for financial support. The usual disclaimer applies. Email addresses: dfgarrett@gmail.com [Garrett], alepavan@northwestern.edu [Pavan], toikka@MIT.EDU [Toikka].

[†]Toulouse School of Economics

[‡]Northwestern University and CEPR

[§]Massachusetts Institute of Technology

1 Introduction

The ideas and tools developed in the mechanism design literature have found applications in a variety of contexts, including auctions, regulation, taxation, employment, political economy, matching, and many others. While much of the literature focuses on environments where the agents learn information only at a single point in time, and the mechanism makes one-time decisions, many environments are inherently dynamic. One class of problems that has been of particular interest involves an agent, or multiple agents, whose private information (as described by their “type”) changes over time and is serially correlated. Since a sequence of decisions needs to be made, the mechanism must elicit this information progressively over time.

Solving these dynamic mechanism design problems is often a complicated task, but a common approach has become popular in the literature on profit-maximizing mechanisms, building on ideas from static mechanism design (such as Myerson, 1981). This approach focuses on solving a “relaxed program” which accounts only for certain necessary incentive compatibility conditions, which can be derived from the requirement that agents do not have an incentive to misreport their types locally (e.g., by claiming to have a type adjacent to their true type). Of course, the solution to such a relaxed program need not correspond to an optimal mechanism in the problem of interest; in particular, some of the ignored incentive compatibility constraints may be violated. Nonetheless, the standard approach is to choose conditions on the environment (primitives that include, in particular, the evolution of the agents’ types) that guarantee global incentive compatibility. Unfortunately, the conditions imposed, typically, have little to do with the economic environment as it is naturally conceived. One is then left to wonder to what extent qualitative properties of the optimal mechanism are a consequence of the restrictions that guarantee the validity of the “relaxed approach.”

The present paper takes an alternative route to the characterization of the qualitative features of optimal dynamic mechanisms. This route yields insights for settings well beyond those for which the above approach based on the relaxed program applies. The property of optimal mechanisms that has received perhaps the most attention in the existing literature on dynamic mechanism design is the property of *vanishing distortions*; i.e., optimal mechanisms become progressively more efficient and distortions from efficient allocations eventually vanish. Examples of such work include Besanko (1985), Battaglini (2005), Pavan, Segal and Toikka (2014), and Bergemann and Strack (2015), among others. We investigate whether this and other related properties continue to hold in a broad class of environments for which the familiar “relaxed approach” need not apply.

Our approach is based on identifying “admissible perturbations” to any optimal mechanism. For any optimal, and hence incentive-compatible and individually-rational, mechanism, we obtain nearby mechanisms which continue to satisfy all the relevant incentive-compatibility and individual-rationality constraints. Of course, for the original mechanism to be optimal, the perturbed mechanism must not increase the principal’s expected payoff, which yields necessary conditions for optimality. These necessary conditions can in turn be translated into the qualitative properties of

interest.

For concreteness, we focus on a canonical procurement model in which the principal (a procurer) seeks to obtain an input in each period from an agent (the supplier). The (commonly-observed) quantity of the input is controlled by the agent whose cost for different quantity levels (the agent’s “type”) is his private information and evolves stochastically over time. An optimal mechanism must ensure both participation in the mechanism at the initial date (individual rationality), as well as the truthful revelation of the agent’s information on costs as it evolves (incentive compatibility). We focus on settings in which the agent’s types are drawn from a finite set in each period, and then comment (in Section 4) on ways in which the results can be extended to settings with a continuum of types.

Our main results are along two distinct lines. Our first result pertains to the dynamics of the ex-ante expectation of the “wedge” between the marginal benefit of additional quantity to the principal and its marginal cost to the agent. First we show (Proposition 1) that when the process governing the evolution of the agent’s private information satisfies the property of “Long-run Independence”, the expected wedge vanishes in the long run. The property of “Long-run Independence” requires that the agent’s type eventually becomes independent of its realization at the time of contracting, and is satisfied in most cases of interest considered in the literature. Next, we show that, under additional assumptions (namely, that types are stochastically ordered and follow a stationary Markov process), convergence is monotone and from above, i.e., the expected wedges are positive, they decrease over time and vanish in the long run (Proposition 2). These results hold across a broad range of processes and preferences, and, in particular, for any discount factor for the players. They are obtained by considering a particularly simple class of perturbations whereby quantity in a given period is either increased or decreased by a uniform constant amount that does not depend on the history of the agent’s types, and then adjusting the payments appropriately to maintain incentive compatibility and individual rationality.

The above results may be seen as offering guidance on the long-run properties of optimal mechanisms when the agent’s initial type is uninformative about the types that will be realized far into the relationship (mechanisms where expected wedges fail to vanish cannot be optimal). However, because distortions may in principle be *both* upwards along some histories of types and downwards along others, the results leave open the possibility that distortions away from the efficient surplus persist in the long run (even though the expected wedge between the marginal benefit and marginal cost of higher quantity vanishes). Our third result (Proposition 3) then provides a sufficient condition that guarantees distortions (i.e., discrepancies in the per-period surplus relative to the first best) vanish in the long run, and that the supplied quantities converge in probability to their first-best levels as the relationship progresses. The condition requires the process governing the evolution of the agent’s type to be a stationary Markov processes, and the players to be sufficiently patient. We provide a lower bound on the players’ discount factor, in terms of the other primitives of the model, for which distortions vanish and quantities converge in probability to the first-best levels. The result

leverages on the fact that, in a discrete-type model, the agent’s loss from misreporting his type in an efficient mechanism (be it static or dynamic) is bounded away from zero by a constant that does not depend on the reported type. We also extend these results (Proposition 4) to settings with an arbitrary discount factor but where the type process is not too persistent. Such stronger results (about convergence of surplus to the efficient level, in probability) are established through more complex perturbations whereby the putative optimal policies are replaced by a convex combination of such policies with the efficient ones, with the weights on the efficient policies growing gradually over time at a rate that guarantees the new policies satisfy all the incentive compatibility constraints. Importantly, it is worth emphasizing that simpler perturbations that replace the original policies with the efficient ones at distant dates need not guarantee incentive compatibility at earlier dates. The proposed perturbations, instead, by introducing slack in incentive compatibility gradually over time, permit one to eventually replace the original policies with the efficient ones while preserving incentive compatibility not just in the continuation but also at all earlier dates. However, because such perturbations increase the agents’ informational rents at earlier periods, the players need to be sufficiently patient for the benefits of converging to efficiency in the long-run to compensate the increased rents at earlier dates, which explains why such stronger results require the discount factor to be above a certain bound (whose value, naturally, is strictly less than one).

The intuition for the above results is somewhat related to the one proposed in the existing literature based on the “relaxed approach”. The benefit of distorting quantities (or allocations) away from the efficient levels comes from the fact that such distortions permit the principal to reduce the information rents that must be left to an agent whose initial cost of producing larger quantities is low (while also ensuring the participation of those agents whose initial cost is high). In environments where, at the time of contracting, the agent’s initial type carries little information about the costs of supplying different quantities in the distant future, such distortions at later periods have less effect on the rents expected at the time of contracting than distortions introduced earlier on. This simple logic suggests allocations should converge to the efficient levels over time. The complication with this logic relates to the fact that the allocations chosen at a given date not only affect the agent’s information rents as anticipated at the time of contracting (i.e., at the beginning of the relationship) but also the incentive compatibility of the mechanism at *all* intermediate dates. In principle, this might motivate the principal to persist with large distortions in allocations along many realizations of the agent’s type history (although possibly abandoning distortions in favor of efficiency along others). Explained differently, there is potentially a role for distortions in the mechanism at dates far in the future in order to guarantee the incentive compatibility of the mechanism at earlier dates, all the way back to the beginning of the relationship. This complication is central to the difficulty of characterizing optimal dynamic mechanisms without resorting to the “relaxed approach” described above, and it has precluded results establishing vanishing distortions for optimal mechanisms in general environments (in particular, in environments where the “relaxed approach” cannot be guaranteed to be valid). In turn, such a difficulty relates to the need to arrive at properties of optimal mechanisms without

knowing which incentive-compatibility and individual-rationality constraints bind at the optimum.

While the focus in this paper is on the long-run properties of optimal mechanisms, our analysis also yields implications for optimal mechanisms at fixed horizons. For instance, Corollary 1 provides a bound on the expected distortions in each period in terms of model parameters that, for sufficiently high discount factors, converges linearly to zero (i.e., the bound on distortions is a geometric sequence). Thus our approach is applicable also to relationships that are not expected to last indefinitely and our results also provide a conservative bound on the rate at which the allocations under optimal mechanisms converge to the efficient levels, at arbitrary discount factors.

Understanding the dynamics of distortions under optimal mechanisms may be useful for a variety of reasons. First, it helps guiding policy interventions in many markets where long-term contracting is expected to play a major role. Second, such an understanding provides guidance for the actual design of optimal long-term contracts. In this respect, our results may also be useful in settings where the choice of the mechanism is restricted. For example, the principal may be required to restrict attention to mechanisms in which the outcome at any date depends only on a limited number of past reports, as is often assumed in the optimal taxation literature (see, for example, Farhi and Werning (2013), Golosov et al (2016), and Makris and Pavan (2017), for a discussion of such restrictions). Provided that the proposed perturbations to mechanisms within the restricted class preserve the properties defining the class (e.g., respect the relevant measurability constraints), the approach developed in the present paper can yield predictions also about the dynamics of distortions for such restricted mechanisms.¹

Outline. The rest of the paper is organized as follows. Below we wrap up the Introduction with a brief discussion of the most pertinent literature. Section 2 describes the model. Section 3 contains the results about the long-run dynamics of distortions under optimal contracts. Section 4 discusses the case with a continuum of types. Section 5 offers a few concluding remarks. All formal proofs are in the Appendix at the end of the document.

1.1 Related Literature

The literature on dynamic contracts and mechanism design is too broad to be described concisely here. We refer the reader to Bergemann and Pavan (2015), Pavan (2017), and Bergemann and Välimäki (2017) for overviews. Here, we focus on the most closely related work.

As mentioned above, the approach followed in the dynamic mechanism design literature to arrive at a characterization of properties of optimal contracts in environments with evolving private information is the so-called “relaxed,” or “first-order” approach, whereby global incentive-compatibility constraints are replaced by certain local incentive-compatibility constraints. In quasilinear environ-

¹For example, when the relevant measurability constraints require the policies to depend only on current reports, the perturbations considered in the present paper guarantee that the perturbed policies continue to satisfy the relevant measurability constraints.

ments, this approach yields a convenient representation of the principal’s objective as “dynamic virtual surplus”. The latter combines the true intertemporal total surplus with time-evolving handicaps that capture the costs to the principal of leaving information rents to the agents. Such handicaps in turn combine properties of the agents’ payoffs with properties of the process controlling the evolution of the agents’ private information. Under the relaxed approach, optimal contracts are then identified by first maximizing dynamic virtual surplus over all allocation rules, including those that need not be incentive compatible, and then finding primitive conditions (on payoffs and type processes) guaranteeing that the policies that solve the relaxed program satisfy all the omitted incentive-compatibility and participation constraints. Establishing the validity of the relaxed approach involves verifying that the policies that solve the relaxed program are sufficiently monotone, in a sense that accounts for the time-varying nature of the agents’ private information and the multi-dimensionality of the decisions taken under the mechanism. Earlier contributions using the relaxed approach include Baron and Besanko (1984), Besanko (1985), and Riordan and Sappington (1987). For more recent contributions, see, among others, Courty and Li (2000), Battaglini (2005), Esó and Szentes (2007), Board (2007), and Kakade et al. (2013). Pavan, Segal, and Toikka (2014) summarize most of these contributions and extend them to a general dynamic contracting setting with a continuum of types, multiple agents, and arbitrary time horizon.

The cornerstone of the “relaxed approach” is a dynamic envelope formula that describes the response of each agent’s equilibrium payoff to the arrival of new private information. The formula combines the familiar direct effect of the agent’s type on the agent’s utility with novel effects that originate from the fact that the marginal information the agent receives in each period is also informative of the information the agent expects to receive in the future. Such novel effects can be summarized in impulse response functions describing how a change in the current type propagates throughout the entire type process. In Markov environments, the aforementioned dynamic envelope formula, when paired with appropriate monotonicity conditions on the allocation rule, provides a complete characterization of incentive compatibility (see Section 4 for a brief overview of these conditions in settings where the process governing the agents’ type can be described by a collection of continuous conditional distributions).

Two recent papers that go beyond the “relaxed approach” are Garrett and Pavan (2015) and Battaglini and Lamba (2017). The first paper uses variational arguments to identify certain properties of optimal contracts in a two-period managerial compensation model. That paper focuses on the interaction between risk aversion and the persistence of the agent’s private information for the dynamics of wedges under profit-maximizing contracts. Relative to that paper, the contribution of the present work is the identification of key properties that are responsible for the *long-run dynamics* of allocations under profit-maximizing contracts. Apart from permitting longer horizons, we study here a much broader class of stochastic processes for types.² More importantly, none of the results

²The earlier paper considers a two-period setting with continuous types, where the second-period type is determined by a linear function of the initial type and a random “shock” that is independent of the initial type.

in the present paper about the convergence of allocations (either in expectation or in probability) to the first-best levels has any counterpart in Garrett and Pavan (2015). The key methodological advance that permits these convergence results is the identification of a novel class of perturbations to proposed optimal allocations that preserve incentive compatibility. As mentioned above, such perturbations involve linear combinations of the putative optimal and efficient policies, with the weight on the efficient policies increasing gradually over time from the beginning of the relationship at a rate that guarantees incentive compatibility.

In a model with finitely many types, Battaglini and Lamba (2017) show that, with more than two types, the “relaxed” or “first-order” approach typically yields policies that fail to satisfy the intertemporal monotonicity conditions necessary for global incentive compatibility. In particular, one of the key insights of that paper is in showing that monotonicity is violated when the process governing the evolution of the agents’ private information is highly persistent. They consider a setting where the agent’s private information is drawn from a continuous-time but finite Markov process, and where the principal and the agent meet at discrete intervals. For generic transitions, as the length of the intervals vanishes, the policies that solve the relaxed program violate at least one of the ignored incentive-compatibility constraints. In a fully-solved two-period-three-type example, they show that the optimal dynamic mechanism can exhibit bunching.

Battaglini and Lamba (2017) also seek results on convergence to efficiency. They focus on mechanisms whose allocations are restricted to be “strongly monotone.” By this, it is meant that an agent who experiences a history of higher types receives a (weakly) larger allocation in each period (in their monopolistic screening model, higher types have a higher preference for additional quantities). They show that optimal “strongly monotone mechanisms” involve allocations that are always (weakly) downward distorted and converge in probability to the efficient ones with time. The key observation is that, when allocations are restricted to be strongly monotone, the optimal such allocations must be efficient at and after any date at which the agent’s type assumes its highest value. The result then follows because the probability that the agent’s type has not yet assumed its highest value vanishes (under their full-support assumption) with time. Unfortunately, optimal dynamic allocations need not be strongly monotone. One possible justification for considering strongly monotone mechanisms, as offered by Battaglini and Lamba (2017), is that strongly monotone mechanisms approximate the discounted average payoffs under optimal mechanisms as the players become infinitely patient. Note, however, that this does not imply distortions vanish with time under fully optimal mechanisms. The distortions in the distant future may remain large and serve the purpose of guaranteeing incentive compatibility at earlier dates, despite having a negligible (but not zero) effect on the expected welfare in the relationship. Understanding how allocations (not only payoffs) behave under fully optimal mechanisms can be important for an empiricist interested in testing the implications of dynamic contracting from a long time series.

Our results thus differ from those in Battaglini and Lamba (2017) in various important dimensions. First, we focus on the dynamics of distortions under *fully optimal* contracts, as opposed to

restricted ones. Second, our results are provided for *fixed* discount factors, and do not require considering the limit of infinite patience. Third, some implications of our analysis do not depend on the discount factor. In particular, our predictions in Section 3.2 for the dynamics of the expected “wedges” do not depend on the discount factor. Likewise, the bounds on distortions we identify in Section 3.3 (Corollary 1) hold for all discount factors and our results in Proposition 4 about vanishing distortions in the long run are valid for *all* discount factors.

2 The Model

Consider the following procurement problem. The principal is a procurer of an input (say a manufacturer), while the agent is a supplier. Their relationship lasts for $T \in \mathbb{N} \cup \{+\infty\}$ periods. Time is discrete and indexed by $t = 1, 2, \dots, T$.

The principal needs to procure a strictly positive quantity of the input in every period. Failure to do so results in the worst possible payoff for the principal (for instance, one may assume the principal’s payoff from this event is equal to $-\infty$). This assumption, along with other Inada conditions described above, guarantees the solution to the principal’s problem is interior at all histories, thus avoiding complications stemming from corner solutions.

At each period, the agent can produce the good in variable quantity $q_t \in (0, \bar{q})$, with $\bar{q} \in \mathbb{R}_{++}$. The principal’s payoff is quasi-linear in transfers. Her gross per-period benefit from procuring q_t units of the good is given by $B(q_t)$, where the function $B : (0, \bar{q}) \rightarrow \mathbb{R}$ is twice continuously differentiable, strictly increasing, strictly concave, and satisfies the Inada condition $\lim_{q \rightarrow 0} B(q) = -\infty$.

The agent’s per-period payoff is also quasi-linear in transfers, with the per-period cost of producing the input in quantity q_t given by the function $C(q_t, h_t)$, where h_t is the agent’s period- t “type”, and where

$$C(q_t, h_t) = h_t q_t + c(q_t), \tag{1}$$

with $c(\cdot)$ twice continuously differentiable, strictly increasing, strictly convex, and satisfying the Inada condition $\lim_{q \rightarrow \bar{q}} c(q) = +\infty$. That the agent’s cost is linear in types facilitates the exposition. In particular, it guarantees that the optimal mechanism is deterministic. In Proposition 5, however, we extend the results about vanishing distortions to more general cost functions.

The agent’s types h_t are drawn from a finite set $\Theta = \{\theta_1, \dots, \theta_N\}$, with $0 < \theta_1 < \dots < \theta_N < +\infty$, $N \geq 2$, and $\Delta\theta \equiv \theta_N - \theta_1$ (in Section 4, we discuss the case where the agent’s types are drawn from absolutely continuous distributions with compact support $\Theta = [\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} > 0$).

Both the principal and the agent have expected-utility preferences over lotteries over streams of quantities and payments and their Bernoulli utility functions take the familiar time-additively-separable form

$$U^P = \sum_t \delta^{t-1} (B(q_t) - p_t) \quad \text{and} \quad U^A = \sum_t \delta^{t-1} (p_t - C(q_t, h_t)),$$

where p_t is the total payment from the principal to the agent in period t , and $\delta \in (0, 1]$ is the common discount factor (with $\delta < 1$ in case $T = +\infty$).

The process governing the evolution of the agent's type is described by the collection of kernels (aka conditional probabilities) $F \equiv (F_t)$. Let $h_s^t \equiv (h_s, \dots, h_t)$, $h^t \equiv (h_1, \dots, h_t)$ and $h_{-s}^t \equiv (h_1, \dots, h_{s-1}, h_{s+1}, \dots, h_t)$. The function F_1 denotes the cumulative distribution function from which h_1 is drawn while, for all $t \geq 2$, $F_t(\cdot | h^{t-1})$ denotes the cumulative distribution function from which h_t is drawn, given h^{t-1} . In particular, for each $n \in \{1, \dots, N\}$, each $h^{t-1} \in \Theta^{t-1}$,

$$F_1(\theta_n) = \sum_{i=1}^n f_1(\theta_i) \quad \text{and} \quad F(\theta_n | h^{t-1}) = \sum_{i=1}^n f_t(\theta_i | h^{t-1}),$$

where, for $i \in \{1, \dots, N\}$, $f_1(\theta_i)$ denotes the probability the agent's initial type is θ_i , while $f_t(\theta_i | h^{t-1})$ denotes the probability his period- t type is θ_i , following history h^{t-1} . We assume the process corresponding to the kernels F has *full support*, meaning that, for all i, t , and h^{t-1} , $f_1(\theta_i) > 0$ and $f_t(\theta_i | h^{t-1}) > 0$.

The sequence of events is the following:

- At $t = 0$, the agent privately learns h_1 .
- At $t = 1$, the principal offers a mechanism $\varphi = (\mathcal{M}, \phi)$, where $\mathcal{M} \equiv (\mathcal{M}_t)_{t=1}^T$ is a collection of message spaces, one for each period, and $\phi \equiv (\phi_t(\cdot))_{t=1}^T$ is a collection of mappings from such spaces to payments and output levels. In particular, for each t , the mapping

$$\phi_t : \mathcal{M}_1 \times \dots \times \mathcal{M}_t \rightarrow \mathbb{R} \times (0, \bar{q}),$$

specifies a payment-quantity pair for each possible profile of messages $m^t \equiv (m_1, \dots, m_t) \in \mathcal{M}_1 \times \dots \times \mathcal{M}_t$ sent by the agent up to period t included. A mechanism is thus equivalent to a menu of long-term contracts. If the agent refuses to participate in φ , the game ends. As explained above, this is taken to be the worst possible outcome for the principal. In this case, the agent earns a payoff equal to zero. If, instead, the agent accepts to participate in φ , he is then committed to produce a strictly positive quantity in every period, the level of which depends on the agent's reports. In particular, in period one, the choice of the message $m_1 \in \mathcal{M}_1$ translates into the obligation to supply a quantity $q_1(m_1)$ in exchange of a payment $p_1(m_1)$.

- At the beginning of each period $t \geq 2$, the agent privately learns his period- t type, h_t . Provided the agent accepted to participate at $t = 1$, he then sends a new message $m_t \in \mathcal{M}_t$, is asked to supply a quantity $q_t(m^t)$, receives a payment $p_t(m^t)$, and the game moves to period $t + 1$.
- ...
- At $t = T + 1$ (in case T is finite), the game ends.

Remark. As standard in the literature on dynamic mechanism design, the game described above assumes the principal perfectly commits to the mechanism φ . It also assumes that, at any period $t \geq 2$, the agent is constrained to stay in the relationship if he signed the contract in period one. When the agent has deep pockets, there are simple ways to distribute the payments over time that guarantee that it is in the agent's interest to remain in the relationship at all periods, irrespective of what he did in the past.

The principal's problem consists in designing a mechanism that maximizes her ex-ante expected payoff. Because the principal can commit, the Revelation Principle applies.³ Without loss of optimality, one can restrict attention to direct mechanisms in which $\mathcal{M}_t = \Theta$ for all t and that induce the agent to report truthfully at all periods, conditional on having reported truthfully in the past. Because the message space is fixed across all such mechanisms, hereafter we economize on notation and identify a direct mechanism with the associated policies $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$, where $\mathbf{q} = (q_t(\cdot))_{t=1}^T$ and $\mathbf{p} = (p_t(\cdot))_{t=1}^T$, with $q_t : \Theta^t \rightarrow (0, \bar{q})$ and $p_t : \Theta^t \rightarrow \mathbb{R}$, $t \geq 1$.

Let σ denote an arbitrary reporting strategy for the agent in ψ and \mathbf{q}^σ and \mathbf{p}^σ the quantity and transfer policy induced by the strategy σ in ψ .

For any ψ , let

$$V_t^\psi(h^t; \hat{h}^{t-1}) \equiv \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(p_s(\hat{h}^{t-1}, \tilde{h}_t^s) - C(q_s(\hat{h}^{t-1}, \tilde{h}_t^s), \tilde{h}_s) \right) \mid h^t \right]$$

denote the agent's expected continuation payoff from date t onwards, when the realized sequence of types up to period t is h^t , the agent reported the sequence of types \hat{h}^{t-1} in previous periods, and the agent reports truthfully from date t onwards. Then let $V_t^\psi(h^t) \equiv V_t^\psi(h^t; h^{t-1})$ denote the agent's continuation payoff at a generic *truthful history*, i.e., when past reports coincide with the true types (throughout, hatted variables represent reports, while random variables are denoted with tildes.)

The principal selects the mechanism ψ from the set Ψ of *individually-rational* and *incentive-compatible* mechanisms. Formally, a mechanism belongs in Ψ if and only if it satisfies the *individual-rationality* constraints (in short, IR)

$$V_1^\psi(h_1) \geq 0 \quad \text{for all } h_1 \in \Theta, \tag{2}$$

and the *incentive-compatibility* constraints (in short, IC)

$$\mathbb{E} \left[V_1^\psi(\tilde{h}_1) \right] \geq \mathbb{E} \left[\sum_{t \geq 1} \delta^{t-1} \left(p_t^\sigma(\tilde{h}^t) - C(q_t^\sigma(\tilde{h}^t), \tilde{h}_t) \right) \right], \tag{3}$$

for all possible reporting strategies σ . Condition (2) requires that the agent prefers to participate in period one and report truthfully in each period, rather than not participating and receiving the

³See, among others, Myerson (1981).

payoff associated with his outside option (zero). Condition (3) requires that the agent prefers to follow a truthful reporting strategy rather than any other reporting strategy σ .⁴

The principal's problem thus consists of maximizing

$$\mathbb{E} \left[\sum_t \delta^{t-1} \left(B \left(q_t \left(\tilde{h}^t \right) \right) - p_t \left(\tilde{h}^t \right) \right) \right] \quad (4)$$

by choice of $\psi \in \Psi$. We refer to a mechanism ψ that maximizes (4) over Ψ as an *optimal mechanism*.

3 Robust predictions

3.1 Preliminary properties of optimal mechanisms

We start by establishing a few preliminary properties of optimal mechanisms, starting with the optimality of deterministic policies. As usual, quasi-linearity of payoffs in transfers implies that deterministic payments are without loss of optimality. Perhaps less obvious is the property that optimal mechanisms involve a deterministic provision of quantity at all histories. Note that, because the optimal mechanism need not coincide with the solution to the relaxed program (i.e., the various dynamic monotonicity constraints on the output schedules may bind at the optimum), the optimality of deterministic mechanisms does not follow from the arguments in Strausz (2006). To see that deterministic policies are optimal, note that, in this environment, a stochastic output policy would stipulate, for any sequence of reports $h^t \in \Theta^t$, and past quantity realizations $q^{t-1} \in (0, \bar{q})^{t-1}$, a probability distribution $\mu_t(h^t, q^{t-1})$ on the interval of quantities $(0, \bar{q})$, with $\mu_t(h^t, q^{t-1})(q) = \Pr(\tilde{q} \leq q)$ denoting the probability that the realized quantity is in $(0, q)$. After receiving the agent's period- t report, the mechanism would draw a quantity according to the probability distribution $\mu_t(h^t, q^{t-1})$, which the agent would then be compelled to produce. Note that the linearity of the cost function in the agent's type implies that, for any probability distribution $\mu_t(h^t, q^{t-1})$, any i and j ,

$$\mathbb{E}_{\mu_t(h^t, q^{t-1})} [C(\tilde{q}, \theta_i) - C(\tilde{q}, \theta_j)] = C(\mathbb{E}_{\mu_t(h^t, q^{t-1})}[\tilde{q}], \theta_i) - C(\mathbb{E}_{\mu_t(h^t, q^{t-1})}[\tilde{q}], \theta_j).$$

Hence, if the proposed stochastic mechanism is IR and IC and calls for a non-degenerate distribution over quantities $\mu_t(h^t, q^{t-1})$ at history h^t , the mechanism that coincides with the proposed mechanism at all histories other than (h^t, q^{t-1}) and that, at history (h^t, q^{t-1}) , calls for a deterministic quantity $\mathbb{E}_{\mu_t(h^t, q^{t-1})}[\tilde{q}]$ and pays the agent

$$p_t(h^t, q^{t-1}) - \left\{ \mathbb{E}_{\mu_t(h^t, q^{t-1})} [c(\tilde{q})] - c(\mathbb{E}_{\mu_t(h^t, q^{t-1})}[\tilde{q}]) \right\}$$

is also IR and IC and gives the agent the same expected payoff (at all histories) as the original mechanism. Because of the concavity of $B(\cdot)$, the principal's expected payoff is higher under the new mechanism, thus establishing the optimality of deterministic mechanisms.

⁴The condition in (3) is an ex-ante (i.e., period-0) incentive-compatibility condition. However, because of the assumption of full support, (3) holds if and only if incentive compatibility holds at all period- t truthful histories (that is, at all histories (h^t, \hat{h}^{t-1}) such that $\hat{h}^{t-1} = h^{t-1}$), all $t \geq 0$.

The second property pertains to the distribution of payments over time. Because both the principal's and the agent's payoffs are linear in transfers (and because the players are equally patient), continuation payoffs at any date t depend on the sum of discounted future payments, but not on the precise timing of these payments. This observation leads to the following result:

Lemma 1. *For any mechanism $\psi \in \Psi$, there exists a mechanism $\psi' \in \Psi$ yielding the principal a payoff at least as large as ψ and such that, for all $t \geq 2$, all $h^{t-1} \in \Theta^{t-1}$,*

$$\mathbb{E} \left[V_t^{\psi'}(\tilde{h}^t) | h^{t-1} \right] = 0. \quad (5)$$

Furthermore, in any mechanism satisfying Condition (5), for all $t \geq 2$, all $h^t \in \Theta^t$,

$$|p_t(h^t) - C(q_t(h^t), h_t)| \leq \frac{\bar{q}\Delta\theta}{1-\delta}. \quad (6)$$

Finally, any optimal mechanism satisfying Condition (5), all $t \geq 2$, all $h^t \in \Theta^t$, is such that Condition (6) holds also for $t = 1$.

The first part of the result states that there is no loss for the principal in restricting attention to mechanisms in which, when the agent follows a truthful reporting strategy, his expected continuation payoff in the subsequent period is always equal to zero. This property follows directly from the possibility of shifting the payments over time in a way that preserves the agent's incentives.

The second part of the result can then be read as a bound on the agent's flow payoff at any history h^t . This follows from combining the first part of the lemma with the agent's incentive compatibility constraints. Note that the right-hand side of (6) is the maximal difference in the expected net present value of the costs of producing any stream of quantities across any two types. Hence, when continuation payoffs satisfy Condition (5), flow payoffs must satisfy Condition (6) (otherwise, there is some date t and type h'_t such that an agent with period- t type equal to h'_t has a profitable deviation of claiming to be of type h_t at period t , and then continuing making reports at later dates that mimic the distribution of type h_t 's reports under truthful reporting).

Below, we use these properties to establish several results, including the existence of an optimal mechanism as stated in the next lemma.

Lemma 2. *An optimal mechanism $\psi^* = \langle \mathbf{q}^*, \mathbf{p}^* \rangle$ exists. Furthermore, any optimal mechanism has the same allocation rule \mathbf{q}^* .*

Given that payoffs are quasi-linear in transfers, the indeterminacy of the optimal payment rules \mathbf{p}^* is to be expected. That, instead, the optimal allocation rule \mathbf{q}^* is unique follows from the convexity of the agent's cost function $C(\cdot, h)$ for each type $h \in \Theta$, along with the concavity of the principal's gross payoff $B(\cdot)$.

3.2 Convergence of wedges in expectation

We are now ready to state our first main result.

Condition 1. [Long-run Independence] Suppose $T = +\infty$. The dependence of the date- t distribution on the date-1 types vanishes in the long run, in the following sense:

$$\lim_{t \rightarrow \infty} \max_{h_1, h'_1, h_t \in \Theta} \left| \Pr(\tilde{h}_t \leq h_t | h_1) - \Pr(\tilde{h}_t \leq h_t | h'_1) \right| = 0.$$

The property in Condition “Long-run Independence” thus captures the idea that, eventually, the agent’s initial type becomes uninformative about the agent’s later types. Note that this condition trivially holds when the agent’s type evolves according to an irreducible aperiodic Markov chain, as typically assumed in the literature.

The next result identifies robust predictions for the dynamics of the ex-ante expected distortions. To understand the result, note that efficiency requires that, at any history $h^t \in \Theta^t$, output be equal to the level $q^E(h_t)$ given by the unique solution to

$$B'(q^E(h_t)) = C_q(q^E(h_t), h_t), \quad (7)$$

where $C_q(q_t, h_t)$ denotes the partial derivative of the agent’s total cost with respect to period- t output when his period- t type is h_t , and is equal to $C_q(q_t, h_t) = h_t + c'(q_t)$ when the agent’s cost takes the linear form in (1).

Proposition 1. *Suppose $T = +\infty$ and Condition “Long-run Independence” holds.*

1. *The optimal mechanism satisfies*

$$\mathbb{E} \left[B'(q_t^*(\tilde{h}^t)) - C_q(q_t^*(\tilde{h}^t), \tilde{h}_t) \right] \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (8)$$

2. *If distortions are always downwards (i.e., $B'(q_t^*(h^t)) - C_q(q_t^*(h^t), h_t) \geq 0$ for all t , all h^t), or if distortions are always upwards (i.e., the above inequality is reversed at all t , all h^t), then*

$$\mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right] \rightarrow \mathbb{E} \left[B(q^E(\tilde{h}_t)) - C(q^E(\tilde{h}_t), \tilde{h}_t) \right] \text{ as } t \rightarrow \infty, \quad (9)$$

and hence $q_t^(h^t)$ converges in probability to $q^E(h_t)$. That is, for any $\eta > 0$,*

$$\lim_{t \rightarrow \infty} \Pr \left(\left| q_t^*(\tilde{h}^t) - q^E(\tilde{h}_t) \right| > \eta \right) = 0.$$

Part 1 of the proposition states that the ex-ante expected “wedge” between the marginal benefit to the principal and the marginal cost to the agent of higher output (to borrow a term from the dynamic public finance literature) vanishes as the relationship progresses. As explained in the Introduction, the result is established by considering a certain class of perturbations to any putative optimal mechanism ψ^* in which the ex-ante expected wedge fails to vanish. In particular, the perturbations

consist in increasing (alternatively, decreasing) by a fixed amount the quantity supplied at dates t long after the relationship has commenced and at which the ex-ante expected wedge remains away from zero. Because the perturbation in quantity is by a uniform constant (that is, it is uniform over all reported histories h^t), appropriate adjustments to the agent’s payments can be found that ensure that the new mechanism is in Ψ (i.e., it is individually rational and incentive compatible). Ensuring participation, i.e. the satisfaction of the constraints in (2), may require leaving the agent with additional expected rents. However, Condition “Long-run Independence” guarantees that such additional rents are small if the period t at which the perturbation takes place is far in the future, since the agent has little relevant private information about his period- t type at the time of contracting. Thus, when t is sufficiently large, irrespective of the discount factor, the proposed perturbations increase ex-ante expected surplus more than they increase the agent’s expected rents, and thus increase the principal’s payoff. This would contradict the optimality of the proposed mechanism implying that, under optimal mechanisms, ex-ante expected wedges must vanish in the long run.

Importantly, note that the convergence of ex-ante expected wedges to zero, by itself, need *not* imply that the surplus from the relationship converges to the efficient level over time. In particular, a pattern of allocations such that quantities are upward distorted along some histories but downward distorted along others is not inconsistent with Part 1 of Proposition 1. While downward distortions have been shown to be a fairly robust feature of policies solving the “relaxed approach,” they need not be a robust feature of optimal mechanisms in general settings where the validity of the relaxed approach cannot be taken for granted. The main difficulty in establishing whether distortions are always downwards lies in the difficulty of discerning which incentive-compatibility and individual-rationality constraints bind (see, for example, the discussion in Section 6 of Battaglini and Lamba (2017) of the challenges in establishing which incentive-compatibility constraints bind in a simple two-period setting with three possible types⁵). Part 2 of the proposition then establishes that, if the direction of distortions was known to always be either upwards or downwards (this knowledge may be a consequence of additional structure assumed in the model), then one could conclude that necessarily distortions would vanish in the long-run, not just in expectations, but in probability. This is because, when distortions are always of the same sign, failure of total surplus to converge to the efficient level cannot be consistent with convergence of the expected wedges, thus implying the result.

The result in the proposition does not require any assumption on the process governing the evolution of the agent’s type other than “Long-run Independence.” In environments where there is additional structure on the evolution of the agent’s private information, sharper predictions on the dynamics of expected wedges are possible, as we show next.

⁵That paper assumes that types are dynamically ordered by FOSD, in which case the individual-rationality constraints always bind at the extreme of the type support.

Condition 2. [FOSD] Kernels are stochastically ordered by previous realizations: For any t , any $\check{h}^{t-1}, \bar{h}^{t-1} \in \Theta^{t-1}$ such that $\check{h}^{t-1} \geq \bar{h}^{t-1}$, any $h_t \in \Theta$, $F_t(h_t|\check{h}^{t-1}) \leq F_t(h_t|\bar{h}^{t-1})$.

Condition ‘‘FOSD’’ requires that the history of the agent’s types up to any date t stochastically orders the period- t distributions. Such assumption is commonly made in the literature and implies that a type’s static comparative advantage vis-a-vis another type also brings a dynamic advantage. Another assumption often encountered in the literature is the following:

Condition 3. [Markov] The process F is a time-homogeneous first-order Markov chain: In particular for any $s, t \geq 2$, any h^t and h^s such that $(h_{t-1}, h_t) = (h_{s-1}, h_s)$, $F_t(h_t|h^{t-1}) = F_s(h_s|h^{s-1})$.

Under the full support assumption (a transition to each of the N states has strictly positive probability after all histories), Condition ‘‘Markov’’ implies that the process is aperiodic and irreducible and hence has a unique steady state (or ‘‘ergodic’’) distribution. One can then associate the Markov transitions with an $N \times N$ matrix A , whose generic element A_{ij} denotes the probability of moving from type i to type j in one period. A further possible restriction on the process is then that the initial distribution coincides with the ergodic:

Condition 4. [Stationary Markov] The process F is a first-order Markov process whose initial distribution F_1 coincides with the ergodic distribution for the time-homogeneous Markov transitions associated with F .

The next result identifies additional robust predictions that are implications of the above assumptions.

Proposition 2. *1. Suppose Condition ‘‘FOSD’’ holds. Then, irrespective of whether T is finite or infinite, the expected wedge*

$$\mathbb{E} \left[B' \left(q_t^* \left(\tilde{h}^t \right) \right) - C_q \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \quad (10)$$

is non-negative for any $t \geq 1$.

2. If, in addition to Condition ‘‘FOSD’’, Condition ‘‘Stationary Markov’’ also holds, the expected wedge (10) is decreasing in t .

3. Hence, if $T = +\infty$ and Conditions ‘‘FOSD’’ and ‘‘Stationary Markov’’ hold, the convergence of the expected wedges to zero is from above and monotone in time.

For processes satisfying Condition ‘‘FOSD’’, it is fairly straightforward to see that the participation constraint in (2) binds uniquely at the least efficient period-1 type, θ_N , (a similar observation has been made in other works, although typically by envelope-type arguments, as, e.g., in Pavan, Segal and Toikka (2014)). To appreciate the result in part 1 in the proposition, note that if the expected wedge (10) were strictly negative at some date t , one could reduce the quantity at date

t by a uniform constant, while adjusting payments to ensure incentive compatibility and that the participation constraint (2) continues to bind at θ_N . The perturbation would increase ex-ante expected surplus. Furthermore, because the initial type θ_N expects the greatest cost reduction from the change in the quantity schedule and the new mechanism leaves the initial type θ_N with the same expected rent (zero) as the original mechanism, under the new mechanism, the expected rent for any initial type below θ_N is lower than under the original mechanism. Hence, profits are unambiguously higher under the new mechanism, which would contradict the optimality of the proposed policies.

Next, consider part 2 in the proposition. That, for stationary Markov processes, the expected wedge (10) decreases monotonically with time follows from the fact that, for such processes, the initial types become progressively less informative about the distribution of the later types. Increases in quantity at later dates thus have progressively smaller effect on the agent's expected rents. When coupled with the result that expected wedges are positive at all histories, this latter property implies that expected output must increase, on average, with time under optimal mechanisms, for, otherwise, the principal could transfer output from earlier to later dates in a way that preserves surplus while reducing the agent's rents, thus contradicting the optimality of the putative policies.

The result in part 3 then follows directly from the above properties along with the result in Proposition 1.

3.3 Convergence of policies in probability

The results above establish that distortions (or, more precisely, “wedges”) vanish in expectation for a large class of processes (those satisfying Condition “Long-run Independence”). As discussed above, these results, however, draw no conclusions regarding the long-run efficiency of the relationship, except when coupled with additional (exogenous) information on the direction of the distortions. Below, we identify primitive conditions that guarantee convergence (in probability) of the allocations under optimal contracts to their efficient counterparts. Throughout, we assume Condition “Markov” holds, although we expect that arguments similar to the ones below can be used to establish that the same results extend to non-Markov finite-state processes provided that such processes are not too persistent, in the sense of Proposition 4. Notice that while Condition “Long-run Independence” is implied by Condition “Markov,” the results below do not impose either Condition “FOSD,” or Condition “Stationary Markov.” Rather, the important economic restriction is that the players are sufficiently patient.

Let A denote the transition matrix associated with the time-autonomous Markov process and define

$$\alpha \equiv \min_{i,j \in \{1, \dots, N\}} A_{ij}.$$

Since each A_{ij} is strictly positive, $\alpha > 0$. Note that the value of α is related to the persistence of the process: for larger values of α , the process cannot be too persistent, since the probability of

transiting from any state i to any other state j is bounded from below by α . In addition, let

$$b \equiv \sum_{i=1}^N \theta_i$$

and

$$\kappa \equiv \min_{i,j \in \{1, \dots, N\} \text{ s.t. } i \neq j} \{B(q^E(\theta_i)) - C(q^E(\theta_i), \theta_i) - (B(q^E(\theta_j)) - C(q^E(\theta_j), \theta_i))\}. \quad (11)$$

Note that $\kappa > 0$ is the smallest loss in total surplus due to some type θ_i supplying the efficient quantity for some other type θ_j , $j \neq i$. Next, let

$$\bar{\delta} \equiv \begin{cases} \frac{2\bar{q}b - \kappa}{2\bar{q}b - \kappa + 2\kappa\alpha} & \text{if } \kappa < 2\bar{q}b \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

Note that $\bar{\delta}$ is decreasing in κ and α and increasing in b and \bar{q} (a brief explanation is provided below).

Proposition 3. *Suppose $T = +\infty$ and the process F satisfies Condition “Markov.” Then, for all $\delta \in (\bar{\delta}, 1)$,*

$$\mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right] \rightarrow \mathbb{E} \left[B(q^E(\tilde{h}_t)) - C(q^E(\tilde{h}_t), \tilde{h}_t) \right] \text{ as } t \rightarrow \infty. \quad (13)$$

Hence, the quantity supplied at date t under an optimal mechanism, $q_t^(h^t)$, converges in probability to the efficient quantity $q^E(h_t)$.*

The proof of Proposition 3 in the Appendix shows that if a mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ in Ψ violates the convergence property in Condition (13), then one can identify a “perturbation” of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that yields the principal a strictly higher payoff. A naive attempt at such a construction might involve a switch at some date long after the relationship has commenced to a fully efficient policy (with allocation $q^E(\cdot)$ as given in (7)). Intuitively, one may conjecture that such a mechanism increases the expected surplus from the relationship at a negligible cost in terms of additional rents to the agent at the time of contracting (i.e., at date 1); after all, if the perturbation occurs far in the future, the agent is poorly informed about his types at the time of the perturbation implying that the effect of such perturbation on his informational rent should be negligible.

The problem with such a naive approach is that it does not account for the effect of the perturbation on the agent’s incentive constraints prior to the switch to the efficient policy. In particular, it does not account for the fact that the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ might have induced truthful reporting at early dates precisely because of distortions away from the efficient allocation in the distant future.

As anticipated in the Introduction, the proof in the Appendix overcomes such difficulty by considering a more complex class of perturbations whereby the efficient allocations are approached gradually over time. We show that, provided the rate of convergence is chosen appropriately, such perturbations can be constructed to guarantee that the perturbed mechanism continues to reside in Ψ (that is, is IR and IC).

A first observation is that, in a setting with time-additively separable payoffs and Markov processes, the sub-optimality of one-stage deviations from truth-telling at all histories implies incentive compatibility. A second observation is that a static mechanism that prescribes an efficient allocation $q^E(h_t)$ along with a payment

$$p^E(h_t) = B(q^E(h_t)) \quad (14)$$

guarantees that the agent's incentive-compatibility constraints are satisfied as strict inequalities, with slack that is bounded away from zero by the amount κ defined in (11). In turn, such observation implies that, if the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is IC, then there exists a mechanism $\langle \mathbf{q}, \mathbf{p} \rangle$ in which the output policy is given by

$$q_t(h^t) = \gamma q^E(h_t) + (1 - \gamma) q_t^*(h^t)$$

for all t , all h^t , in which the agent's incentive-compatibility constraints hold strictly at all histories, and where the agent's loss from deviating from truth-telling is at least $\gamma\kappa$.

The above observations also imply existence of incentive-compatible mechanisms that place increasing weight on the efficient policy, i.e. such that $q_t(h^t) = \gamma_t q^E(h_t) + (1 - \gamma_t) q_t^*(h^t)$ for all t , all h^t , for some increasing sequence of scalars $(\gamma_t)_{t=1}^\infty$. Furthermore, the sequence $(\gamma_t)_{t=1}^\infty$ can be chosen so that the perturbed mechanism becomes fully efficient, i.e. $\gamma_t = 1$, after any period t where the allocation in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is distorted. Note, however, that, for the perturbed mechanism to be incentive compatible, it is essential that the allocations be perturbed also in earlier periods, i.e., at $s < t$. In particular, it may be necessary to start the adjustments directly in period one by setting $\gamma_1 > 0$. Since, from the perspective of the time of contracting, the agent is relatively well informed about his type early in the relationship, such adjustments may necessitate leaving the agent with additional rents to guarantee participation, and these additional rents are larger, the larger are γ_t for small t . Whether the gains in surplus can be guaranteed to exceed the increase in additional rents then depends on the primitives of the problem, the dates and size of the distortions in the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, and the weights assigned by the perturbed mechanism to the efficient output policy early in the relationship. The key difficulty in establishing the proposition comes precisely from the need to relate the *lower* bounds on the weights γ_s to the efficient policies at periods $s < t$ that guarantee incentive compatibility of the perturbed mechanism to the *upper* bounds on the same weights that guarantee the profitability of the proposed perturbation.

When the players are sufficiently patient, namely when $\delta > \bar{\delta}$, with $\bar{\delta}$ as given in (12), the size of distortions in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that are necessary for one to be able to find a perturbation along the lines indicated above that improves upon $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ shrink with the date at which such distortions occur,

which is precisely what allows us to establish the convergence result in Proposition 3. Intuitively, because the perturbations may involve an increase in the agent's rents from the very early periods and an increase in total surplus only in the distant future, players need to be sufficiently patient for such perturbations to be profitable. We will come back to the dependence of the threshold $\bar{\delta}$ on the primitive parameters momentarily. Before we do so, we highlight that the same arguments that establish Proposition 3 above also imply that distortions under optimal contracts must be bounded *at any date*, for else perturbations along the lines indicated above exist yielding a strictly higher payoff to the principal. Furthermore, such bounds can be established for any degree of patience. The following Corollary summarizes the above observations and is a direct implication of the proof of Proposition 3 in the Appendix:

Corollary 1. *Suppose the process F satisfies Condition “Markov.” (The horizon length T may be finite or infinite.) Let $\lambda = \frac{\bar{q}^b}{1-\delta(1-2\alpha)}$. Then, for any date t , and any $\delta \in [0, 1]$:*

$$\mathbb{E} \left[B \left(q^E \left(\tilde{h}_t \right) \right) - C \left(q^E \left(\tilde{h}_t \right), \tilde{h}_t \right) \right] - \mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) \right) - C \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \leq \frac{2\lambda}{\left(\delta + \frac{\kappa}{2\lambda} \right)^{t-1}},$$

When $\delta \in (\bar{\delta}, 1)$, $\delta + \frac{\kappa}{2\lambda} > 1$, and hence the above bounds decrease monotonically with t and vanish in the long run.

Note that, in certain environments, the bounds in the above corollary need not be sharp. However, when coupled with the assumption that players are sufficiently patient (namely, that $\delta \in (\bar{\delta}, 1)$), the above bounds imply that convergence of the output policies to the efficient level cannot occur too slowly: distortions must decline with t at least at a geometric rate, whose value $(\delta + \frac{\kappa}{2\lambda})^{t-1}$ can be easily related to the various parameters of the model.

Next, consider the dependence of the patience threshold $\bar{\delta}$ on the primitive parameters. First, note that the threshold $\bar{\delta}$ can be reduced if the sequence $(\gamma_t)_{t=1}^{\infty}$ can be taken to increase more quickly (while still rendering an incentive-compatible and individually-rational perturbed mechanism). We can then observe that, as κ increases, the amount of slack in the incentive constraints of an (appropriately chosen) efficient mechanism increases, which permits $(\gamma_t)_{t=1}^{\infty}$ to increase more quickly (that is, less weight γ_t must be placed on the efficient mechanism at early periods). In turn, this reduces the rents granted to the agent to ensure participation and hence the critical level of patience $\bar{\delta}$ guaranteeing the profitability of the above perturbations. The same is true as α increases, or as b or \bar{q} are reduced.

Related to the dependence of $\bar{\delta}$ on the persistence parameter α , one might naturally expect the convergence in (13) to occur for any *fixed* discount factor δ when the persistence of the process is sufficiently small. We formalize this idea in the following result.

Proposition 4. *Suppose $T = +\infty$ and F satisfies Condition “Markov”. For any $\delta \in (0, 1)$, there exists $\varepsilon(\delta) > 0$ such that the following is true. If, for all $t \geq 2$, all $(h_{t-1}, h'_{t-1}, h_t) \in \Theta^3$,*

$|f_t(h_t|h_{t-1}) - f_t(h_t|h'_{t-1})| < \varepsilon(\delta)$, then

$$\mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) \right) - C \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \rightarrow \mathbb{E} \left[B \left(q^E \left(\tilde{h}_t \right) \right) - C \left(q^E \left(\tilde{h}_t \right), \tilde{h}_t \right) \right] \text{ as } t \rightarrow \infty.$$

The proof of Proposition 4 is similar to that of Proposition 3 and involves many of the same steps. The key difference is that the lack of persistence of the agent's types (i.e., the condition that, for all $(h_{t-1}, h'_{t-1}, h_t) \in \Theta^3$, $|f_t(h_t|h_{t-1}) - f_t(h_t|h'_{t-1})| < \varepsilon(\delta)$ where $\varepsilon(\delta)$ is small) is used to guarantee the existence of a perturbation of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that is incentive compatible and that prescribes the efficient allocation early in the relationship, while prescribing outcomes close to those prescribed by the original mechanism (i.e., $(q_1^*(h_1), p_1^*(h_1))$) at the very first period. The condition in the proposition then implies that, because the persistence of types is limited, the additional rents expected by the agent under the perturbed mechanism are small relative to the increase in the discounted surplus, thus making such perturbations profitable to the principal.

Finally, note that the result in Proposition 3 can be extended to more general cost functions $C(\cdot, \cdot) : (0, \bar{q}) \times \Theta \rightarrow \mathbb{R}_+$ (subject to weak restrictions to be introduced momentarily) provided we allow for stochastic mechanisms. In particular, we now consider mechanisms $\langle \mu, \mathbf{p} \rangle$ with the mechanism specifying distributions over quantities according to $\mu = (\mu_t(\cdot))_{t=1}^T$, and with payments given (as before) by $\mathbf{p} = (p_t(\cdot))_{t=1}^T$. As noted above, deterministic payment rules are without loss of optimality in light of the quasi-linearity of payoffs in transfers. Given that payoffs are separable across time, the fact that payments and the distribution of quantities at each date are independent of past quantity realizations q^{t-1} is also without loss of optimality.

There are two important reasons for considering stochastic mechanisms. First, for reasons similar to those explained in Strausz (2006) for static environments, a restriction to deterministic mechanisms can in general imply a reduction in the principal's payoff (as explained above, this is not the case when $C(q, h) = hq + c(q)$). Strausz (2006) shows that, in static environments, the optimality of deterministic mechanisms is implied when the "relaxed approach" that focuses on *local* (in particular, *upwards* in the present setting) incentive-compatibility constraints succeeds, but also suggests that this conclusion need not hold more generally. While Strausz's result for the optimality of deterministic mechanisms applies also to dynamic environments where the "relaxed approach" succeeds, the main objective of our work is precisely the identification of robust predictions that also apply to environments where this approach is not guaranteed to be valid.

The second reason relates to our approach of extending Proposition 3 to more general cost functions. Our approach relies on perturbations to the proposed mechanisms that fail to satisfy the property of interest, i.e., the convergence of the expected surplus to the efficient levels. The perturbed mechanisms we construct involve randomizations between the policies of the mechanism of interest and the policies of efficient mechanisms. That is, the perturbed mechanisms are stochastic.

Motivated by the aforementioned considerations, hereafter we thus extend the analysis to stochastic mechanisms. We start by explaining how the incentive compatibility and the participation con-

straints must be adapted to accommodate for stochastic mechanisms $\langle \mu, \mathbf{p} \rangle$. Hereafter, we confine attention to processes F that satisfy Condition ‘‘Markov.’’ In this case, the agent’s expected (continuation) payoff from reporting truthfully in a mechanism $\langle \mu, \mathbf{p} \rangle$, given past reports h^{t-1} and period- t type h_t , is given by

$$V_t^{\langle \mu, \mathbf{p} \rangle}(h^t) \equiv \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(p_s(h^{t-1}, \tilde{h}_t^s) - \int C(\tilde{q}, \tilde{h}_s) d\mu_t(h^{t-1}, \tilde{h}_t^s) \right) | h_t \right].$$

The participation constraints then amount to the requirement that

$$V_1^{\langle \mu, \mathbf{p} \rangle}(h_1) \geq 0 \quad \text{for all } h_1 \in \Theta, \quad (15)$$

while the incentive-compatibility constraints amount to the requirement that, for all reporting strategies σ ,

$$\mathbb{E} \left[V_1^{\langle \mu, \mathbf{p} \rangle}(\tilde{h}_1) \right] \geq \mathbb{E} \left[\sum_{t \geq 1} \delta^{t-1} \left(p_t^\sigma(\tilde{h}^t) - \int C(\tilde{q}, \tilde{h}_t) d\mu_t^\sigma(\tilde{h}^t) \right) \right], \quad (16)$$

where $(\mu_t^\sigma)_{t=1}^T$ and $(p_t^\sigma)_{t=1}^T$ are, respectively, the probability measures over the quantities supplied and the payments induced by the reporting strategy σ .⁶ Let Ψ^S denote the set of (possibly stochastic) mechanisms satisfying these above constraints. The principal’s problem thus consists in maximizing

$$\Pi(\mu, \mathbf{p}) = \mathbb{E} \left[\sum_{t \geq 1} \delta^{t-1} \left(\int B(\tilde{q}) d\mu_t(\tilde{h}^t) - p_t(\tilde{h}^t) \right) \right],$$

over the set Ψ^S .

Our result below establishes convergence of the output policies to the efficient ones when the players are sufficiently patient, with the degree of patience given by a value $\bar{\delta}^S$ in terms of the model parameters. We restrict attention to continuously differentiable and non-negative cost functions $C(q, h)$ such that the following condition holds.

Condition 5. [Cost restriction] The following two properties are satisfied:

1. There exists a unique function $q^E : \Theta \rightarrow [0, \bar{q}]$ such that, for each $h_t \in \Theta$, $q^E(h_t)$ maximizes $B(q) - C(q, h_t)$; moreover $q^E(\cdot)$ is one-to-one.
2. There exists $u \in \mathbb{R}_+$ such that, for all $q \in (0, \bar{q})$ and $h, h' \in \Theta$, $|C(q, h) - C(q, h')| \leq u$.

Recall that $\alpha \equiv \min_{i,j \in \{1, \dots, N\}} A_{ij}$, and let κ be given as before by (11). Note that $\kappa > 0$ by the first requirement in Condition ‘‘Cost restriction.’’ The result below on long-run efficiency then applies to any environment in which $\delta > \bar{\delta}^S$, with $\bar{\delta}^S$ given by

$$\bar{\delta}^S \equiv \begin{cases} \frac{2Nu - \kappa}{2Nu - \kappa + 2\kappa\alpha} & \text{if } \kappa < 2Nu \\ 0 & \text{otherwise} \end{cases}. \quad (17)$$

⁶Note that the strategy σ is allowed to condition the reports on past output realizations.

Note that, for this more general environment, we lack an argument guaranteeing existence of an optimal mechanism. The result in the next proposition is thus stated in a way that permits us to deliver interesting properties even in the event an optimal mechanism does not exist.

Proposition 5. *Suppose $T = +\infty$ and F satisfies Condition “Markov”. Suppose, in addition, that the cost function $C(\cdot, \cdot)$ satisfies Condition “Cost restriction”. If the mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ is optimal, then, for all $\delta \in (\bar{\delta}^S, 1)$,*

$$\mathbb{E} \left[\int \left(B(\tilde{q}) - C(\tilde{q}, \tilde{h}_t) \right) d\mu_t^* \left(\tilde{h}^t \right) \right] \rightarrow \mathbb{E} \left[B \left(q^E \left(\tilde{h}_t \right) \right) - C \left(q^E \left(\tilde{h}_t \right), \tilde{h}_t \right) \right] \text{ as } t \rightarrow \infty. \quad (18)$$

More generally, take any $\delta \in (\bar{\delta}^S, 1)$ and consider any sequence of stochastic mechanisms $(\langle \mu^k, \mathbf{p}^k \rangle)$ with $\lim_{k \rightarrow \infty} \Pi(\mu^k, \mathbf{p}^k) = \sup_{\langle \mu, \mathbf{p} \rangle \in \Psi^S} \Pi(\mu, \mathbf{p})$. For any $\varepsilon > 0$, there exists a $\bar{t} \in \mathbb{N}$ and a sequence (s_k) , $s_k \rightarrow \infty$, such that, for all $\bar{t} \leq t \leq \bar{t} + s_k$,

$$\mathbb{E} \left[B \left(q^E \left(\tilde{h}_t \right) \right) - C \left(q^E \left(\tilde{h}_t \right), \tilde{h}_t \right) \right] - \mathbb{E} \left[\int \left(B(\tilde{q}) - C(\tilde{q}, \tilde{h}_t) \right) d\mu_t^k \left(\tilde{h}^t \right) \right] < \varepsilon. \quad (19)$$

The first part of Proposition 5 confirms that, when an optimal mechanism exists, the convergence of expected surplus under the optimal mechanism to the efficient level occurs also for the more general cost functions covered by Condition “Cost restriction”. In environments where existence of an optimal mechanism cannot be guaranteed, the second part of the result provides a weaker sense in which distortions vanish that applies to mechanisms for which the principal earns a payoff close to the supremum.

4 Continuum of types

We now turn to the case where the agent’s types are drawn from continuous distributions with compact support $\Theta = [\underline{\theta}, \bar{\theta}]$. As for most of the analysis above, we return to the case where the agent’s cost function takes the form $C(q, h) = hq + c(q)$, and where mechanisms are deterministic.

The main economic trade-offs involved in the design of optimal dynamic mechanisms are closely related to those in the discrete-type case. However, there are a few differences from a methodological perspective. Most notably, the expected net present value of the payments to the agent is uniquely pinned down by the quantity schedule, up to a scalar that depends on the surplus expected by the least productive period-1 type. Such a property, which is referred to in the literature as “payoff equivalence” (see, e.g., Pavan, Segal, and Toikka, 2014) implies a tighter relationship between the quantity schedule in any incentive-compatible mechanism and the agent’s expected rents (and hence the principal’s expected profits). As we show below, such a property in turn may permit a direct calculation of the expected “wedge” between the marginal benefit of higher output to the principal and its marginal cost to the agent under any optimal mechanism. Below, we first revisit the steps that permit one to leverage on the above “payoff equivalence” result to express the principal’s payoff

as “dynamic virtual surplus” (i.e., as total surplus, net of handicaps that depend entirely on the quantity schedule and of the surplus expected by the least efficient period-1 type). We then show how variational arguments similar to those in the previous section permit us to identify robust predictions for the dynamics of wedges under optimal mechanisms akin to those in the discrete case.

Consider an environment identical to that in Section 2 except for the stochastic process $F \equiv (F_t)$ governing the evolution of the agent’s types, which, in addition to Conditions “Markov” and “FOSD,” is assumed to satisfy the following properties. For any $t \geq 2$, any $h_{t-1} \in \Theta$, $F(\cdot|h_{t-1})$ is absolutely continuous over the entire real line with density $f(\cdot|h_{t-1})$ strictly positive over Θ (full support). The same absolute continuity property applies to the period-1 distribution F_1 (with density f_1). Given the process F , we say that Z is a *state representation* of F , if, for any $s \geq 1$, any $t \geq s$, h_t can be written as $h_t = Z_{(s),t}(h_s, \varepsilon)$, where $\varepsilon \equiv (\varepsilon_\tau) \in \mathcal{E} \subset \mathbb{R}^\infty$ is a random vector drawn independently from F , and where $Z_{(s),t}(\cdot)$ are functions satisfying the following properties.⁷ Let $\|\cdot\|$ denote the discounted L1 norm on \mathbb{R}^∞ , defined by $\|y\| \equiv \sum_{t=0}^\infty \delta^t |y_t|$.

Condition 6. [Regularity] The process F is “regular” if it admits a state representation Z satisfying the following conditions:

1. There exist functions $K_{(s)} : \mathcal{E} \rightarrow \mathbb{R}^\infty$, $s \geq 1$, with $\mathbb{E}[\|K_{(s)}(\tilde{\varepsilon})\|] \leq B$ for some constant B independent of s , such that for all $t \geq s$, $h_s \in \Theta_s$, and $\varepsilon \in \mathcal{E}$, $Z_{(s),t}(h_s, \varepsilon)$ is a differentiable function of h_s with $|\partial Z_{(s),t}(h_s, \varepsilon)/\partial h_s| \leq K_{(s),t-s}(\varepsilon)$.⁸
2. For each s , $\log [\partial Z_{(s),t}(h_s, \varepsilon)/\partial h_s]$ is continuous in h_s uniformly over $t \geq s$ and $(h_s, \varepsilon) \in \Theta_s \times \mathcal{E}$.

Given the above notation, for any $t > s$, any h^t , then let

$$I_{(s),t}(h_s^t) = \left. \frac{\partial Z_{(s),t}(h_s, \varepsilon)}{\partial h_s} \right|_{\varepsilon: Z_{(s),t}^t(h_s, \varepsilon) = h_s^t}$$

denote the impulse response of h_t to h_s . The impulse response of h_t to h_s captures the effects of a marginal variation in h_s on h_t across all histories of shocks that, starting from h_s lead to h_s^t . When $s = 1$, we simplify the notation by dropping (s) from the subscripts and define such functions by $I_t(h^t)$. Finally, we let $I_{(s),s}(h_s) = 1$, all $s \geq 1$, all $h_s \in \Theta$.⁹

⁷For example, the following is a state representation of F . For any $t \geq 2$, any $\varepsilon_t \in (0, 1)$, any $h_{t-1} \in \Theta$, let $F^{-1}(\varepsilon_t|h_{t-1}) \equiv \inf\{h_t : F(h_t|h_{t-1}) \geq \varepsilon_t\}$. Let each ε_t be drawn from a Uniform distribution over $(0, 1)$, independently across t . For any $s \geq 1$, then let $Z_{(s),s}(h_s, \varepsilon) = h_s$, while for any $t > s$, let $Z_{(s),t}(h_s, \varepsilon) = F_t^{-1}(\varepsilon_t|Z_{(s),t-1}(h_s, \varepsilon))$. The above representation is referred to as the *canonical representation* of F in Pavan, Segal and Toikka (2014).

⁸For any ε , the term $K_{(s),t-s}(\varepsilon)$ is the $(t-s)$ -component of the sequence $K_{(s)}(\varepsilon)$.

⁹As an example, note that, when h_t follows an AR(1) process, then $I_t(h^t) = \gamma^{t-1}$. More generally, such impulse response functions are themselves stochastic processes. Also note that, when the process is Markov, as assumed here, and the kernels $F(h_t|h_{t-1})$ are continuously differentiable in (h_t, h_{t-1}) — a property not assumed here, for inconsequential to the results — the canonical representation introduced above yields the following expression for the impulse responses:

$$I_{(s),t}(h_s^t) = \prod_{\tau=s+1}^t \left(-\frac{\partial F(h_\tau|h_{\tau-1})/\partial h_{\tau-1}}{f(h_\tau|h_{\tau-1})} \right).$$

Next, for any t , any $h^{t-1} \in \Theta^{t-1}$, any $h_t, \hat{h}_t \in \Theta_t$, let

$$D_t(h^t; \hat{h}_t) \equiv -\mathbb{E} \left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s(\tilde{h}_{-t}^s, \hat{h}_t) \mid h^t \right].$$

Under Conditions “Markov”, “FOSD”, and “Regularity”, a mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$ is individually rational and incentive compatible if and only if the following conditions jointly hold for all $t \geq 1$, all h^{t-1} , all $h_t, \hat{h}_t \in \Theta$ (see Pavan, Segal and Toikka (2014)): (a) $V_1^{(\mathbf{q}, \mathbf{p})}(\bar{\theta}) \geq 0$, (b) $V_t^{(\mathbf{q}, \mathbf{p})}(h^t)$ is equi-Lipschitz continuous in h_t with

$$\frac{\partial V_t^{(\mathbf{q}, \mathbf{p})}(h^t)}{\partial h_t} = -\mathbb{E} \left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s(\tilde{h}^s) \mid h^t \right] \text{ a.e. } h_t \in \Theta, \quad (20)$$

and (c)

$$\int_{\hat{h}_t}^{h_t} [D_t((h^{t-1}, x); x) - D_t((h^{t-1}, x); \hat{h}_t)] dx \geq 0. \quad (21)$$

Condition (20) is the dynamic analog of the usual envelope formula for static environments. In a Markov environment, such condition, when paired with the integral monotonicity condition (21), yields a complete characterization of incentive compatibility.^{10, 11}

That $V_1^{(\mathbf{q}, \mathbf{p})}(h_1)$ must satisfy Condition (20) for almost all $h_1 \in \Theta$ in turn implies that the principal’s payoff in any mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$ that is individually rational and incentive compatible is equal to “*Dynamic Virtual Surplus*”, defined as¹²

$$\mathbb{E} \left[\sum_t \delta^{t-1} \left(B(q_t(\tilde{h}^t)) - C(q_t(\tilde{h}^t), \tilde{h}_t) - \frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} I_t(\tilde{h}^t) q_t(\tilde{h}^t) \right) \right] - V_1^{(\mathbf{q}, \mathbf{p})}(\bar{\theta}). \quad (22)$$

¹⁰The role played by the Markov assumption is that it implies that an agent’s incentives in any period depend only on his current true type and his past reports, but not on his past true types. In turn this implies that, when a single departure from truthful reporting is suboptimal, then truthful reporting (at all histories) yields a higher payoff than any other strategy.

¹¹The integral monotonicity condition (21) generalizes the more familiar monotonicity condition for static settings requiring the allocation rule to be nondecreasing. As the above condition reveals, what is required by incentive compatibility in dynamic settings is that the derivative of the agent’s payoff with respect to his true type be sufficiently monotone in the reported type. In particular, note that (21) holds in the dynamic environment under examination here if the NPV of expected future output, *discounted by impulse responses*,

$$\mathbb{E} \left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s(\tilde{h}_{-t}^s, \hat{h}_t) \mid h^t \right]$$

is *nonincreasing* in the current report \hat{h}_t . Output need not be monotone in each period. It suffices it is sufficiently monotone, on average, where the average is both over states and time.

¹²The expression in (22) can be obtained from (20) integrating by parts.

The terms $F_1(\tilde{h}_1) I_t(\tilde{h}^t) q_t(\tilde{h}^t) / f_1(\tilde{h}_1)$ in (22) are handicaps that control for the cost to the principal of asking for higher output at various histories. The assumption that the process satisfies Condition “FOSD” implies that impulse responses are non-negative. Condition (20) then implies that the agent’s period-1 expected payoff, $V_1^{(\mathbf{q}, \mathbf{p})}(\cdot)$, reaches its minimum at $h_1 = \bar{\theta}$. Hence, under any optimal mechanism, the participation constraint (2) binds for the least efficient period-1 type, i.e., $V_1^{(\mathbf{q}, \mathbf{p})}(\bar{\theta}) = 0$, whereas all other participation constraints are slack. The principal’s problem thus reduces to the maximization of (22) over all quantity schedules \mathbf{q} satisfying the integral monotonicity conditions in (21).

As discussed above, solving such a problem is tedious. The approach typically followed in the literature is to solve a relaxed program where the integral monotonicity conditions are dropped and verified ex-post. In the context of the model under consideration here, given the concavity of $B(\cdot)$ and the convexity of $C(\cdot)$ in output, the policies that solve such relaxed program are those identified by the following first-order conditions

$$B'(q_t(h^t)) = C_q(q_t(h^t), h_t) + \frac{F_1(h_1)}{f_1(h_1)} I_t(h^t)$$

for all $t \geq 1$, all $h^t \in \Theta^t$.

In general, however, there is no guarantee that the allocation rule that solves the relaxed program satisfies all the “integral monotonicity” conditions in (21), unless additional conditions are imposed. As explained above, such conditions are typically identified through “reverse engineering,” i.e., by using monotone comparative statics results to guarantee that the quantity schedule that solves the above first-order conditions is sufficiently monotone. Unfortunately, such conditions can be stringent and may lack compelling economic foundations.

Consistently with the analysis in the previous section, the approach followed here is therefore different. We look for perturbations to the putative optimal quantity rule \mathbf{q}^* that preserve the integral monotonicity conditions in (21). Note that, given any such perturbed rule, there always exists a payment rule that (i) yields an individually rational and incentive compatible mechanism and (ii) gives the least efficient period-1 type a payoff $V_1^{(\mathbf{q}, \mathbf{p})}(\bar{\theta}) = 0$. One such rule that satisfies the additional property of Lemma 1 that, for all $t \geq 2$, all $h^{t-1} \in \Theta^{t-1}$, $\mathbb{E}[V_t^{(\mathbf{q}^*, \mathbf{p}^*)}(\tilde{h}^t) | h^{t-1}] = 0$, is the following:

$$p_1(h_1) = C(q_1(h_1), h_1) + \int_{h_1}^{\bar{\theta}} \mathbb{E} \left[\sum_{s=1}^{\infty} \delta^{s-1} I_s(\tilde{h}^s) q_s(\tilde{h}^s) \mid x \right] dx, \quad (23)$$

for all $h_1 \in \Theta$, while, for all $t \geq 2$, all h^t ,

$$p_t(h^t) = C(q_t(h^t), h_t) + \int_{h_t}^{\bar{\theta}} \mathbb{E} \left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s^*(\tilde{h}^s) \mid (h^{t-1}, x) \right] dx \\ - \mathbb{E} \left[\frac{F(\tilde{h}_t | h_{t-1})}{f(\tilde{h}_t | h_{t-1})} \sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s(\tilde{h}^s) \mid h^{t-1} \right]. \quad (24)$$

Clearly, for the putative mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ to be optimal it must be that no perturbation to the quantity schedule q^* preserving the integral monotonicity conditions in (21) increases the principal’s

expected payoff, as given by (22). In particular, suppose that \mathbf{q}^* is such that the period- t schedule $q_t^*(\cdot)$ is bounded away from the boundaries of the set of feasible quantity levels; namely, suppose that there exist $\underline{L}_t, \bar{L}_t \in \mathbb{R}_{++}$, with $0 < \underline{L}_t \leq \bar{L}_t < \bar{q}$, such that, for all h^t , $q_t^*(h^t) \in (\underline{L}_t, \bar{L}_t)$. Then adding a constant $\nu \in \mathbb{R}$ to the schedule $q_t^*(\cdot)$ yields a new policy that continues to satisfy the integral monotonicity conditions in (21) (the arguments are similar to those yielding to Proposition 1). The variation in the principal's expected payoff relative to her payoff under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is then equal to

$$\begin{aligned} \delta^{t-1} \mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) + \nu \right) - C \left(q_t^* \left(\tilde{h}^t \right) + \nu, \tilde{h}_t \right) - \left(\frac{F_1 \left(\tilde{h}_1 \right)}{f_1 \left(\tilde{h}_1 \right)} I_t \left(\tilde{h}^t \right) \right) \left(q_t^* \left(\tilde{h}^t \right) + \nu \right) \right] \\ - \delta^{t-1} \mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) \right) - C \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) - \left(\frac{F_1 \left(\tilde{h}_1 \right)}{f_1 \left(\tilde{h}_1 \right)} I_t \left(\tilde{h}^t \right) \right) q_t^* \left(\tilde{h}^t \right) \right]. \end{aligned} \quad (25)$$

Dividing the expression in (25) by ν and taking the limit for ν going to zero yields the following expression for the derivative of the principal's expected payoff with respect to ν evaluated at $\nu = 0$:

$$\delta^{t-1} \mathbb{E} \left[\left(B' \left(q_t^* \left(\tilde{h}^t \right) \right) - C_q \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) - \left(\frac{F_1 \left(\tilde{h}_1 \right)}{f_1 \left(\tilde{h}_1 \right)} I_t \left(\tilde{h}^t \right) \right) \right) \right].$$

A necessary condition for the optimality of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is that such derivative vanishes. We then have the following result:

Definition 1. The mechanism $\langle \mathbf{q}, \mathbf{p} \rangle$ is “eventually interior” if there exists T and sequences of scalars (\underline{L}_t) and (\bar{L}_t) , with $0 < \underline{L}_t \leq \bar{L}_t < \bar{q}$, all t , such that, for any $t > T$, any h^t , $q_t(h^t) \in [\underline{L}_t, \bar{L}_t]$.

Proposition 6. Assume the process F satisfies Conditions “Markov,” “FOSD,” and “Regularity”. If $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is “eventually interior,” then, for all t large enough,

$$\mathbb{E} \left[B' \left(q_t^* \left(\tilde{h}^t \right) \right) - C_q \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] = \mathbb{E} \left[\frac{F_1 \left(\tilde{h}_1 \right)}{f_1 \left(\tilde{h}_1 \right)} I_t \left(\tilde{h}^t \right) \right]. \quad (26)$$

That the quantity schedule under the optimal mechanism is eventually interior appears difficult to guarantee in general. Naturally, given the Inada conditions ($\lim_{q \rightarrow \bar{q}} c(q) = +\infty$ and $\lim_{q \rightarrow 0} B(q) = -\infty$), such interiority can be guaranteed if one imposes continuity restrictions on the optimal schedules, such as requiring allocations to be Lipschitz continuous with a fixed Lipschitz constant. Under such a restriction, the perturbations to the optimal policies described above are then feasible, which in turn guarantees that the expected wedge under the optimal mechanism is given by the right hand side of (26).

Note that the result in Proposition 6 is related to the result in the literature that focuses on environments for which the relaxed approach is valid, but with important differences. When the optimal policies are those that solve the relaxed approach, and such policies are interior, the wedge,

at any history h^t , is given by the *marginal handicap* $F_1(h_1)I_t(h^t)/f_1(h_1)$. This need not be true more generally, i.e., for environment where the relaxed approach is not guaranteed to be valid. What remains true, though, is that the *expected wedge* is equal to the *expected marginal handicap*. By studying the dynamics of the expected marginal handicaps, one can then identify useful properties for the dynamics of the expected wedges, as we show below. Let $\mathcal{B}(\Theta)$ be the Borel sigma algebra associated with the set Θ . For any $A \in \mathcal{B}(\Theta)$, $h \in \Theta$, and $t \geq 1$, then let

$$P^t(h, A) \equiv \Pr\left(\tilde{h}_t \in A \mid \tilde{h}_1 = h\right).$$

Condition 7. [Ergodicity] The process F is *ergodic* if there exists a unique (invariant) probability measure π on $\mathcal{B}(\Theta)$ whose support has a nonempty interior such that, for all $h \in \Theta$,

$$\sup_{A \in \mathcal{B}(\Theta)} |P^t(h, A) - \pi(A)| \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (27)$$

Under the additional Condition ‘‘Ergodicity,’’ we are able to establish a result analogous to the one in (8) for discrete types.

Proposition 7. *Assume F satisfies Conditions ‘‘Regularity’’ and ‘‘Ergodicity.’’ Then*

$$\mathbb{E} \left[\frac{F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t) \right] \rightarrow 0 \text{ as } t \rightarrow \infty.$$

If, in addition, F satisfies Condition ‘‘FOSD,’’ then convergence is from above, i.e., $\mathbb{E} \left[\frac{F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t) \right] \geq 0$ for all t . Furthermore, if, in addition to the above conditions, the period-1 distribution is the stationary (ergodic) distribution of F , then convergence is monotone in time, i.e., $\mathbb{E} \left[\frac{F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t) \right]$ is decreasing in t .

Together, Propositions 6 and 7 thus provide implications for the dynamics of expected wedges analogous to those in the discrete case. We summarize such implications in the following Corollary:

Corollary 2. *Suppose F satisfies Conditions ‘‘Markov,’’ ‘‘Regularity,’’ ‘‘FOSD,’’ and ‘‘Ergodicity,’’ and that the optimal policies are ‘‘eventually interior.’’ Then the expected wedges are eventually positive and vanish in the long run. If, in addition, the period-1 distribution of F coincides with the ergodic distribution, then, eventually, convergence becomes monotone in time.*

As in the discrete case, the result in Corollary 2 does not imply convergence of the quantity schedules to the efficient ones in probability. Nonetheless, we expect that the same restrictions on the quantity schedules that guarantee the eventual interiority of the policies, such as the requirement that $q_t(\cdot)$ be Lipschitz with known Lipschitz constants, also rule out pathological behavior, permitting arguments analogous to those in the discrete case to establish convergence of the allocations to the efficient levels (in probability) when players are sufficiently patient.

5 Conclusions

We develop a novel variational approach that permits us to study the long-run dynamics of allocations under fully optimal contracts. The approach permits us to bypass many of the technical restrictions required to validate the “relaxed” approach typically followed in the literature. In particular, the analysis identifies primitive conditions guaranteeing convergence of the allocations under fully optimal contracts to the first-best levels.

In future work, it would be desirable to extend the analysis to a richer class of dynamic contracting problems, for example by accommodating for endogenous private information and for the competition between multiple privately informed agents. It would also be interesting to apply a similar methodology to study the dynamics of allocations under restricted contracts, where the restrictions could originate in a quest for “simplicity” such as the requirement that allocations be invariant in past reported types, as in the new dynamic public finance literature (e.g., Farhi and Werning (2013), Golosov et al. (2016), and Makris and Pavan (2017)), or the quest for “robustness to model mis-specification,” as in the recent literature on robustly optimal contracts (e.g., Carroll (2015)). In particular, the variational approach developed in the present paper could also be useful to identify certain properties of optimal contracts in settings in which the principal lacks detailed information about the process governing the evolution of the agents’ private information.

References

- [1] Baron, D., Besanko, D., 1984. Regulation and information in a continuing relationship. *Inf. Econ. Policy* 1, 267–302.
- [2] Battaglini, M., 2005. Long-term contracting with Markovian consumers. *Am. Econ. Rev.* 95, 637–658.
- [3] Battaglini, M., Lamba, R., 2017. Optimal dynamic contracting: the first-order approach and beyond. Discussion paper, Cornell University and Penn State University.
- [4] Bergemann, D., Pavan, A., 2015. Introduction to JET Symposium on Dynamic Contracts and Mechanism Design. *J. Econ. Theory* 159, 679-701.
- [5] Bergemann, D., Strack, P., 2015. Dynamic revenue maximization: a continuous time approach. *J. Econ. Theory* 159 (Part B), 819–853.
- [6] Bergemann, D., Välimäki, J., 2017. Dynamic Mechanism Design: an Introduction. Discussion paper, Yale University.

- [7] Besanko, D., 1985. Multi-period contracts between principal and agent with adverse selection. *Econ. Lett.* 17, 33–37. Bester, H., Strausz, R., 2000. Imperfect commitment and the revelation principle: the multi-agent case. *Econ. Lett.* 69, 165–171.
- [8] Board, S., 2007. Selling options. *J. Econ. Theory* 136, 324–340.
- [9] Carroll, G., 2015. Robustness and Linear Contracts, *American Economic Review* 105 (2), 536–563.
- [10] Courty, P., Li, H., 2000. Sequential screening. *Rev. Econ. Stud.* 67, 697–717.
- [11] Eső, P., Szentes B., 2007. Optimal information disclosure in auctions. *Rev. Econ. Stud.* 74, 705–731.
- [12] Garrett, D., Pavan, A., 2012. Managerial turnover in a changing world. *J. Polit. Econ.* 120, 879–925.
- [13] Farhi, E., Werning, I., 2013. Insurance and taxation over the life cycle. *Rev. Econ. Stud.* 80, 596–635.
- [14] Garrett, D., Pavan, A., 2015. Dynamic managerial compensation: a variational approach. *J. Econ. Theory* 159 (Part B), 775–818.
- [15] Golosov, M., Troschkin, M., Tsyvinski, A., 2016. Optimal Dynamic Taxes. *Am. Econ. Rev.*, 106(2), 359–386.
- [16] Kakade, S., Lobel, I., Nazerzadeh, H., 2013. Optimal dynamic mechanism design and the virtual pivot mechanism. *Oper. Res.* 61, 837–854.
- [17] Makris, M., Pavan, A., 2017. Taxation under learning by doing. Discussion paper. Northwestern University.
- [18] Myerson, R., 1981. Optimal auction design. *Math. Oper. Res.* 6, 58–73.
- [19] Pavan, A., Segal, I., Toikka, J., 2014. Dynamic mechanism design: a Myersonian approach. *Econometrica* 82, 601–653.
- [20] Riordan, M., Sappington, D., 1987. Information, incentives, and organizational mode. *Q. J. Econ.* 102, 243–264.
- [21] Strausz, R., 2006, Deterministic versus Stochastic Mechanisms in Principal-Agent Models, *Journal of Economic Theory*, 127(1), 306–314.

Appendix: Proofs

Proof of Lemma 1. Consider any mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$. Let ψ' be the mechanism constructed from ψ as follows. The payment rule $\mathbf{p}' \equiv (p'_t(\cdot))_{t \geq 1}$ is such that, for $t = 1$,

$$p'_1(h_1) = p_1(h_1) + \delta \mathbb{E} \left[V_2^\psi(\tilde{h}^2) | h_1 \right]$$

for all $h_1 \in \Theta$, whereas for $t > 1$,

$$p'_t(h^t) = p_t(h^t) - \mathbb{E} \left[V_t^\psi(\tilde{h}^t) | h^{t-1} \right] + \delta \mathbb{E} \left[V_{t+1}^\psi(\tilde{h}^{t+1}) | h^t \right]$$

all $h^t \in \Theta^t$. The allocation rule $\mathbf{q}' \equiv (q'_t(\cdot))_{t \geq 1}$ is the same as in ψ , i.e., $\mathbf{q}' = \mathbf{q}$. The mechanism $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ satisfies (5) for all $t \geq 2$, all $h^t \in \Theta^t$. Moreover, if $\psi \in \Psi$, then $\psi' \in \Psi$. Clearly, ψ and ψ' yield the principal the same expected payoff. This establishes the first claim.

Next, notice that, in any mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$ satisfying (5), for all $t \geq 2$, all $h^t \in \Theta^t$, necessarily

$$V_t^\psi(h^t) \leq \frac{\bar{q}\Delta\theta}{1-\delta}.$$

To see this, suppose that, instead, $V_t^\psi(h^t) > \frac{\bar{q}\Delta\theta}{1-\delta}$ for some h^t . That the mechanism ψ satisfies the property in the first part of the lemma implies that there exists a type $h'_t \in \Theta$ such that $V_t^\psi(h^{t-1}, h'_t) \leq 0$. An agent whose history of private types is (h^{t-1}, h'_t) , who reported truthfully up to period $t-1$, could then replicate the distribution of reports from period t onwards of an agent whose period- t history of private types is (h^{t-1}, h_t) . By misreporting in this way, the former agent (the one with history of private types (h^{t-1}, h'_t)) earns an expected payoff (from t onwards) at least equal to

$$V_t^\psi(h^t) - \frac{\bar{q}\Delta\theta}{1-\delta}.$$

This is because, in each period, the additional costs of the former type (relative to the one being mimicked) is bounded by $\bar{q}\Delta\theta$. That $V_t^\psi(h^t) > \frac{\bar{q}\Delta\theta}{1-\delta}$ thus implies that the former agent has a profitable deviation. Since (h^{t-1}, h'_t) has positive probability by the full-support assumption, there exists a reporting strategy σ that induces a strictly higher expected payoff for the agent than the truthful one, thus violating Condition (3).

A similar argument implies that $V_t^\psi(h^t) \geq -\frac{\bar{q}\Delta\theta}{1-\delta}$ for all t , all h^t . Again, suppose this is not the case. That the mechanism satisfies the property in the first part of the lemma implies that there exists a h'_t such that $V_t^\psi(h^{t-1}, h'_t) \geq 0$. An agent whose period- t continuation payoff satisfies $V_t^\psi(h^t) < -\frac{\bar{q}\Delta\theta}{1-\delta}$ could then replicate the distribution of reports of an agent whose period- t type is h'_t (who truthfully reported h^{t-1} in the past) and secure herself a payoff at least equal to $V_t^\psi(h^{t-1}, h'_t) - \frac{\bar{q}\Delta\theta}{1-\delta}$ which is larger than $V_t^\psi(h^t)$. The former agent would thus have a profitable deviation, contradicting that $\psi \in \Psi$.

That in any mechanism satisfying (5), all $t \geq 2$, all $h^t \in \Theta^t$, necessarily

$$|p_t(h^t) - C(q_t(h^t), h_t)| \leq \frac{\bar{q}\Delta\theta}{1-\delta}$$

then follows from the above properties along with the fact that, in any such mechanism,

$$V_t^\psi(h^t) = p_t(h^t) - C(q_t(h^t), h_t),$$

all $t \geq 2$, all h^t .

Finally, to see that, in any *optimal* mechanism satisfying (5), all $t \geq 2$, all $h^t \in \Theta^t$, Condition (6) holds also for $t = 1$, note that, in any such mechanism, the date-1 participation constraint (2) necessarily binds for some type, for otherwise the principal could reduce the payments of all types uniformly by $\varepsilon > 0$ increasing her payoff. That the mechanism satisfies (5), all $t \geq 2$, all $h^t \in \Theta^t$, then implies that condition (6) holds also for $t = 1$. The arguments are the same as for the case $t \geq 2$ discussed above. Q.E.D.

Proof of Lemma 2. It is easy to see that there exist sequences of scalars (\underline{b}_t) and (\bar{b}_t) , with $0 < \underline{b}_t < \bar{b}_t < \bar{q}$, such that the following is true: For any mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle \in \Psi$ with $q_t(h^t) \notin [\underline{b}_t, \bar{b}_t]$ for some t and h^t , there exists another mechanism $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ with $\psi' \in \Psi$ and $q'_t(h^t) \in [\underline{b}_t, \bar{b}_t]$ for all t and all h^t , that yields the principal a higher payoff. The existence of such bounds follows from the combination of the Inada conditions with the discreteness of the process. Let $\bar{\Psi}$ denote those mechanisms in Ψ that satisfy the above bounds on allocations in each period as well as both (5) and (6) for all $t \geq 2$, all h^{t-1} , with the latter also holding at $t = 1$. When T is finite, the design problem amounts to maximizing the principal's continuous objective on the compact set $\bar{\Psi}$, so existence of a solution follows from standard results.

Next, consider the case where $T = +\infty$. Let $(\psi^k) = (\langle \mathbf{q}^k, \mathbf{p}^k \rangle)$ denote a sequence of mechanisms in $\bar{\Psi}$ such that

$$\sup_{\psi \in \bar{\Psi}} \left\{ \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(B(q_t(\tilde{h}^t)) - p_t(\tilde{h}^t) \right) \right] \right\} - \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(B(q_t^k(\tilde{h}^t)) - p_t^k(\tilde{h}^t) \right) \right] < 1/k. \quad (28)$$

We can then construct an optimal policy as follows. Let $((q_1^k(\cdot), p_1^k(\cdot)))$ denote the sequence of period-1 policies defined by the above sequence of mechanisms. From the Bolzano-Weierstrass Theorem, there exists a subsequence of $((q_1^k(\cdot), p_1^k(\cdot)))$ converging to some $(q_1^*(\cdot), p_1^*(\cdot))$. Letting (k_l) index this subsequence, there exists a further subsequence, indexed by $(k_{l(m)})$, of the same original sequence of mechanisms such that $((q_1^{k_{l(m)}}(\cdot), p_1^{k_{l(m)}}(\cdot)))$ converges to some $(q_1^*(\cdot), p_1^*(\cdot))$. Clearly, this also implies that, along the subsequence indexed by $(k_{l(m)})$, $((q_1^{k_{l(m)}}(\cdot), p_1^{k_{l(m)}}(\cdot)))$ converges to $(q_1^*(\cdot), p_1^*(\cdot))$. Proceeding this way, we obtain a mechanism $\psi^* = \langle \mathbf{q}^*, \mathbf{p}^* \rangle$.

We now show that ψ^* is incentive compatible. To do so, we make use of the fact that each ψ^k satisfies Condition (6) (since $\psi^k \in \bar{\Psi}$) and that the maximal difference across types is $\Delta\theta$. This implies that the (absolute value of the) agent's per-period payoff in a mechanism ψ^k is bounded by a constant $M > 0$, uniformly over k , dates $t \geq 2$, histories $h^t \in \Theta^t$ and strategies σ . Furthermore, by construction of ψ^* , the same property (and the same bound M) applies to ψ^* .

Now, suppose that ψ^* is *not* incentive compatible. Then, there exists a reporting strategy σ such that the agent's ex-ante expected payoff under σ is higher than under the truthful reporting strategy by some amount $\eta > 0$. Let $\varepsilon = \eta/3$. The above property (i.e., the boundedness of the agent's per-period payoff) implies that there exists k such that the agent's ex-ante payoffs in ψ^* , under σ and the truthful reporting strategy, respectively, are within an ε -ball of the respective payoffs in ψ^k . Abusing notation, let $\psi^{*,\sigma}$ and $\psi^{k,\sigma}$ denote the outcomes induced under σ in the mechanisms ψ^* and ψ^k , respectively. Then,

$$\begin{aligned} \mathbb{E} \left[V_1^{\psi^k}(\tilde{h}_1) \right] - \mathbb{E} \left[V_1^{\psi^{k,\sigma}}(\tilde{h}_1) \right] &\leq \mathbb{E} \left[V_1^{\psi^*}(\tilde{h}_1) \right] - \mathbb{E} \left[V_1^{\psi^{*,\sigma}}(\tilde{h}_1) \right] + 2\varepsilon \\ &= -\varepsilon \\ &< 0, \end{aligned}$$

where the equality follows from the assumption that $\mathbb{E} \left[V_1^{\psi^*}(\tilde{h}_1) \right] - \mathbb{E} \left[V_1^{\psi^{*,\sigma}}(\tilde{h}_1) \right] = \eta = 3\varepsilon$. This implies that the mechanism ψ^k is not incentive compatible, a contradiction. A similar logic implies that ψ^* is individually rational, and hence $\psi^* \in \Psi$.

We now show that ψ^* is optimal in Ψ . Using the fact that the agent's per-period payoffs in ψ^* are bounded in absolute value by $M > 0$ uniformly over histories h^t , $t \geq 2$, it is easy to see that ψ^* satisfies (5) for all $t \geq 2$ (a property which is inherited from the mechanisms ψ^k). It follows that, for all $t \geq 2$,

$$\mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - p_t^*(\tilde{h}^t) \right] = \mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right]$$

(where the expectations are ex-ante). Moreover,

$$B(q_t^*(h^t)) - C(q_t^*(h^t), h_t) \leq \bar{S} \equiv \max_{q \in (0, \bar{q})} \{B(q) - C(q, \theta_1)\}$$

uniformly over t and h^t . Hence,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(B(q_t^*(\tilde{h}^t)) - p_t^*(\tilde{h}^t) \right) \right] &= \mathbb{E} \left[B(q_1^*(\tilde{h}_1)) - p_1^*(\tilde{h}_1) \right] \\ &\quad + \mathbb{E} \left[\sum_{t=2}^{\infty} \delta^{t-1} \left(B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right) \right] \\ &= \mathbb{E} \left[B(q_1^*(\tilde{h}_1)) - p_1^*(\tilde{h}_1) \right] \\ &\quad + \lim_{T \rightarrow \infty} \sum_{t=2}^T \delta^{t-1} \mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right]. \quad (29) \end{aligned}$$

The first equality uses the fact that, by construction of ψ^* , $\mathbb{E} \left[B(q_1^*(\tilde{h}_1)) - p_1^*(\tilde{h}_1) \right]$ is finite. The second equality, and the existence of the limit in the last line of (29), follow because per-period surplus is uniformly bounded from above by \bar{S} . Hence, expected profits under ψ^* are well defined and equal to some value $\pi^* \in \mathbb{R} \cup \{-\infty\}$.

Now, let $\pi^{\text{sup}} \equiv \sup_{\psi \in \bar{\Psi}} \left\{ \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(B \left(q_t \left(\tilde{h}^t \right) \right) - p_t \left(\tilde{h}^t \right) \right) \right] \right\} \in \mathbb{R}$. Suppose then that ψ^* is not optimal. This means that there exists a (finite) $\eta > 0$ such that $\pi^{\text{sup}} - \pi^* \geq \eta$. Because the series in the last line of (29) is convergent, we have that, for any $\varepsilon > 0$, there exists $T^*(\varepsilon)$, such that, for any $T \geq T^*(\varepsilon)$,

$$\mathbb{E} \left[B \left(q_1^* \left(\tilde{h}_1 \right) \right) - p_1^* \left(\tilde{h}_1 \right) \right] + \sum_{t=2}^T \delta^{t-1} \mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) \right) - C \left(q_t^* \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] < \pi^{\text{sup}} - \eta + \varepsilon.$$

Now fix $\varepsilon > 0$ and let $T^{**}(\varepsilon)$ be the smallest positive integer such that $\frac{\delta^{T^{**}(\varepsilon)-1}}{1-\delta} < \varepsilon$. Then let $\bar{T}(\varepsilon) = \max \{ T^*(\varepsilon), T^{**}(\varepsilon) \}$. Considering mechanisms ψ^k in the original sequence, for any $k \geq \frac{1}{\varepsilon}$ and any $T \geq \bar{T}(\varepsilon)$, we have

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \left(B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right) \right] &= \sum_{t=1}^T \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] \\ &= \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] \\ &\quad - \sum_{t=T+1}^{\infty} \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] \\ &> \pi^{\text{sup}} - 2\varepsilon. \end{aligned}$$

The inequality follows from the fact that (a) the original sequence was constructed so that

$$\begin{aligned} \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] &= \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right) \right] \\ &> \pi^{\text{sup}} - \frac{1}{k} \\ &\geq \pi^{\text{sup}} - \varepsilon \end{aligned}$$

and (b)

$$\sum_{t=T+1}^{\infty} \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] < \frac{\delta^{T^{**}(\varepsilon)-1}}{1-\delta} \bar{S} < \varepsilon,$$

which in turn follows because, for all k and all $t \geq 2$,

$$\mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] = \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - C \left(q_t^k \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \leq \bar{S}.$$

Now, take $\varepsilon = \eta/6$, where η is the constant defined above. We then have, for any $T \geq \bar{T}(\varepsilon)$,

$$\sum_{t=1}^T \delta^{t-1} \mathbb{E} \left[B \left(q_t^k \left(\tilde{h}^t \right) \right) - p_t^k \left(\tilde{h}^t \right) \right] - \sum_{t=1}^T \delta^{t-1} \mathbb{E} \left[B \left(q_t^* \left(\tilde{h}^t \right) \right) - p_t^* \left(\tilde{h}^t \right) \right] > \eta/2. \quad (30)$$

Then fix $T \geq \bar{T}(\varepsilon)$. By construction of ψ^* , for any $\nu > 0$, we can find a $k > 1/\varepsilon$ such that $|q_t^k(h^t) - q_t^*(h^t)|, |p_t^k(h^t) - p_t^*(h^t)| \leq \nu$ for all $t = 1, \dots, T$, all $h^t \in \Theta^t$. Taking ν sufficiently small,

and using the fact that $B(q) - p$ is continuous in (q, p) , we can then find a $k > 1/\varepsilon$ such that (30) fails, a contradiction. This proves that the principal's expected payoff under ψ^* is equal to π^{sup} and hence that ψ^* is optimal.

Finally, to see that the allocation rule in any optimal mechanism is unique, suppose, on the contrary, that $\psi^A = \langle \mathbf{q}^A, \mathbf{p}^A \rangle$ and $\psi^B = \langle \mathbf{q}^B, \mathbf{p}^B \rangle$ are optimal mechanisms with $\mathbf{q}^A \neq \mathbf{q}^B$. Let $\gamma \in (0, 1)$ and consider the mechanism $\psi^\gamma = \langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ defined by

$$q_t^\gamma(h^t) = \gamma q_t^A(h^t) + (1 - \gamma) q_t^B(h^t)$$

and

$$p_t^\gamma(h^t) = \gamma p^A(h^t) + (1 - \gamma) p^B(h^t) + c(q_t^\gamma(h^t)) - \gamma c(q_t^A(h^t)) - (1 - \gamma) c(q_t^B(h^t)),$$

all t , all h^t . For each initial type h_1 , the mechanism ψ^γ so constructed yields the agent an expected payoff

$$V_1^{\psi^\gamma}(h_1) = \gamma V_1^{\psi^A}(h_1) + (1 - \gamma) V_1^{\psi^B}(h_1)$$

under a truthful reporting strategy, and hence it is individually rational, since $\psi^A, \psi^B \in \Psi$. Similarly, the agent's payoff in the mechanism ψ^γ under any reporting strategy σ is a convex combination (with weights γ and $1 - \gamma$) of the agent's payoffs under the same strategy σ in ψ^A and ψ^B . Hence, ψ^γ is incentive compatible; i.e. $\psi^\gamma \in \Psi$. The result then follows from the fact that the principal's payoff under any mechanism is equal to expected total surplus net of the agent's expected rent, along with the fact that the strict concavity of $B(q) - C(q, h)$ in q , for all h , implies that, for all t ,

$$\begin{aligned} \mathbb{E} \left[B \left(q_t^\gamma \left(\tilde{h}^t \right) \right) - C \left(q_t^\gamma \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] &\geq \gamma \mathbb{E} \left[B \left(q_t^A \left(\tilde{h}^t \right) \right) - C \left(q_t^A \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \\ &\quad + (1 - \gamma) \mathbb{E} \left[B \left(q_t^B \left(\tilde{h}^t \right) \right) - C \left(q_t^B \left(\tilde{h}^t \right), \tilde{h}_t \right) \right] \end{aligned}$$

with the inequality strict whenever $q_t^A(h^t) \neq q_t^B(h^t)$ for some h^t . Q.E.D.

Proof of Proposition 1. [Part 1]. Suppose the result is not true and let $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ denote an optimal mechanism (recall that the quantity schedule \mathbf{q}^* is unique across all optimal mechanisms). Then there exists $\kappa > 0$ and a strictly increasing sequence of dates (t_k) such that either

$$\mathbb{E} \left[B' \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) - \tilde{h}_{t_k} - c' \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) \right] > \kappa$$

for all k , or

$$\mathbb{E} \left[B' \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) - \tilde{h}_{t_k} - c' \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) \right] < -\kappa$$

for all k . Consider the first case. Then, for any k , consider the mechanism $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ whose allocation rule is given by $q'_t(\cdot) = q_t^*(\cdot)$ for all $t \neq t_k$, and by $q'_{t_k}(\cdot) = q_{t_k}^*(\cdot) + \nu_k$ for $t = t_k$, for some constant $\nu_k \in (0, \bar{q} - \max_{h^{t_k} \in \Theta^{t_k}} \{q(h^{t_k})\})$. As for the payments, the new mechanism is defined by $p'_1(\cdot) = p_1^*(\cdot) + \delta^{t_k-1} \nu_k \max_{h_1 \in \Theta} \left\{ \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right\}$, $p'_t(\cdot) = p_t^*(\cdot)$ for all $t \geq 2$, $t \neq t_k$, and by $p'_{t_k}(h^{t_k}) = p_{t_k}^*(h^{t_k}) + c(q'_{t_k}(h^{t_k})) - c(q_{t_k}^*(h^{t_k}))$ for all $h^{t_k} \in \Theta^{t_k}$. Note that the new mechanism ψ'

is incentive compatible. This follows from the fact that the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is incentive compatible, along with the fact that, for any reporting strategy σ , any h_1 ,¹³

$$V_1^{\langle \mathbf{q}', \mathbf{p}' \rangle}(h_1) - V_1^{\langle \mathbf{q}', \sigma, \mathbf{p}', \sigma \rangle}(h_1) = V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(h_1) - V_1^{\langle \mathbf{q}^*, \sigma, \mathbf{p}^*, \sigma \rangle}(h_1).$$

Moreover, the adjustment in payments ensures that ψ' is individually rational, and hence $\psi' \in \Psi$.

For ν_k sufficiently small, the increase in total surplus exceeds $\delta^{t_k-1}\nu_k\kappa$, while the expected rent of the agent conditional on his initial type being h_1 increases by

$$\delta^{t_k-1}\nu_k \left(\max_{h_1 \in \Theta} \left\{ \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right\} - \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right).$$

Since $\max_{h_1 \in \Theta} \left\{ \mathbb{E} \left[\tilde{h}_t | h_1 \right] \right\} - \mathbb{E} \left[\tilde{h}_t | h_1 \right]$ vanishes with t by Condition ‘‘Long-run independence,’’ it follows that the increase in the principal’s payoff (i.e., in total surplus, net of the agent’s rents) is positive for k sufficiently large. The new mechanism thus improves over the original one, contradicting the optimality of the latter.

Next, consider the second case. The proof is analogous to the one above, except that the mechanism ψ' used to establish the improvement over ψ^* is such that $\nu_k \in (-\min_{h^{t_k} \in \Theta^{t_k}} \{q_{t_k}(h^{t_k})\}, 0)$, and $p'_1(\cdot) = p_1^*(\cdot) + \delta^{t_k-1}\nu_k \min_{h_1 \in \Theta} \left\{ \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right\}$. Now expected surplus increases by at least $-\delta^{t_k-1}\nu_k\kappa$, while the rent expected by an agent whose initial type is h_1 increases by

$$\delta^{t_k-1}\nu_k \left(\min_{h_1 \in \Theta} \left\{ \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right\} - \mathbb{E} \left[\tilde{h}_{t_k} | h_1 \right] \right).$$

Again, the quantity in brackets vanishes as $k \rightarrow \infty$, by virtue of Condition ‘‘Long-run independence,’’ thus establishing the result.

[Part 2]. Suppose that distortions are always downwards, but that the convergence in (9) does not occur. Note that, because total surplus is concave in output, for any history h^t ,

$$\begin{aligned} & B(q_t^E(h^t)) - C(q_t^E(h^t), h^t) - (B(q_t^*(h^t)) - C(q_t^*(h^t), h^t)) \\ & \leq (B'(q_t^*(h^t)) - C_q(q_t^*(h^t), h^t)) (q_t^E(h^t) - q_t^*(h^t)). \end{aligned}$$

Hence, if the convergence in (9) fails, there exists $\varepsilon > 0$ and a sequence of dates (t_k) , such that, for any t_k ,

$$\mathbb{E} \left[\left(B'(q_{t_k}^*(\tilde{h}^{t_k})) - C_q(q_{t_k}^*(\tilde{h}^{t_k}), \tilde{h}_{t_k}) \right) \left(q_{t_k}^E(\tilde{h}_{t_k}) - q_{t_k}^*(\tilde{h}^{t_k}) \right) \right] \geq \varepsilon.$$

This implies that

$$\mathbb{E} \left[B'(q_{t_k}^*(\tilde{h}^{t_k})) - C_q(q_{t_k}^*(\tilde{h}^{t_k}), \tilde{h}_{t_k}) \right] \geq \frac{\varepsilon}{\bar{q}}.$$

¹³The equality follows because the quantity adjustments used to obtain the new mechanism $\langle \mathbf{q}', \mathbf{p}' \rangle$ are uniform over histories of types h^{t_k} , and hence the effect of the adjustment on the linear part of the agent’s cost is the same under truth-telling and σ . Furthermore, any effect on the convex part of the cost is undone by the adjustment in the payment (i.e., by the adjustment $c(q_{t_k}^*(h^{t_k})) - c(q_{t_k}^*(h^{t_k}))$, for each h^{t_k}).

But this contradicts the convergence of the wedges established in Part 1. The argument for upward distortions is analogous.

Finally, to see that convergence in expected surplus implies convergence in probability of the allocation rule, suppose this is not the case. Then there exists a sequence of dates (t_k) , together with a constant $\eta > 0$, such that

$$\Pr \left(\left| q_{t_k}^* \left(\tilde{h}^{t_k} \right) - q^E \left(\tilde{h}_{t_k} \right) \right| > \eta \right) > \eta$$

along the sequence (t_k) . It is then easy to check, given the strict concavity of $B(\cdot)$ and the convexity of $C(\cdot, h)$, that

$$\mathbb{E} \left[B \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) - C \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right), \tilde{h}_{t_k} \right) \right]$$

remains strictly below

$$\mathbb{E} \left[B \left(q_{t_k}^E \left(\tilde{h}_{t_k} \right) \right) - C \left(q_{t_k}^E \left(\tilde{h}_{t_k} \right), \tilde{h}_{t_k} \right) \right]$$

by a constant uniform along the sequence (t_k) , contradicting (9). Q.E.D.

Proof of Proposition 2 [Part 1]. When Condition ‘‘FOSD’’ holds, the agent’s period-1 rent $V_1^{\psi^*}(\cdot)$ is strictly decreasing in type, under any optimal mechanism $\psi^* = \langle \mathbf{q}^*, \mathbf{p}^* \rangle$. To see this, we introduce the following ‘‘canonical’’ representation of the evolution of the agent’s types, in which the agent is viewed as receiving independently distributed ‘‘shocks’’ to his cost in each period. In particular, for any $t \geq 2$, let $\tilde{\varepsilon}_t$ be drawn from a uniform distribution over $[0, 1]$, independently of \tilde{h}_1 and of any $\tilde{\varepsilon}_s$, $s \neq t$. The evolution of the agent’s types can then be thought of as determined by the initial draw of h_1 (according to F_1) and by the subsequent ‘‘shocks’’ ε_t (drawn independently from a uniform distribution over $[0, 1]$) via the equation

$$h_t = F_t^{-1}(\varepsilon_t | h_1, F_2^{-1}(\varepsilon_2 | h_1), F_3^{-1}(\varepsilon_3 | h_1, F_2^{-1}(\varepsilon_2 | h_1)), \dots),$$

where, for any $t \geq 2$, any $\varepsilon_t \in [0, 1]$, any $h^{t-1} \in \Theta^{t-1}$, $F_t^{-1}(\varepsilon_t | h^{t-1}) \equiv \inf \{ h_t : F(h_t | h^{t-1}) \geq \varepsilon_t \}$. By the Integral Transform Probability Theorem, the above representation is ‘‘equivalent’’ to the original representation, in the sense that, for any $t \geq 2$, any $h^{t-1} \in \Theta^{t-1}$, the distribution of \tilde{h}_t given h^{t-1} , as given by the kernel $F_t(\cdot | h^{t-1})$, is the same as the distribution of

$$F_t^{-1}(\tilde{\varepsilon}_t | h_1, F_2^{-1}(\varepsilon_2 | h_1), F_3^{-1}(\varepsilon_3 | h_1, F_2^{-1}(\varepsilon_2 | h_1)), \dots)$$

for any vector of shocks $(\varepsilon_2, \dots, \varepsilon_{t-1})$ such that $F_2^{-1}(\varepsilon_2 | h_1) = h_2$, $F_3^{-1}(\varepsilon_3 | h_1, F_2^{-1}(\varepsilon_2 | h_1)) = h_3, \dots$,

$$F_{t-1}^{-1}(\varepsilon_{t-1} | h_1, F_2^{-1}(\varepsilon_2 | h_1), F_3^{-1}(\varepsilon_3 | h_1, F_2^{-1}(\varepsilon_2 | h_1)), \dots) = h_{t-1}.$$

Now, consider a period-1 type θ_i who reports θ_{i+1} in period 1, and then goes on to report in subsequent periods in such a way that the process for reports is indistinguishable from that for an agent whose true initial type is θ_{i+1} and who reports truthfully in each period. In particular, in an arbitrary period $t \geq 2$ in which the agent has already reported \hat{h}^{t-1} and his true type history is h^{t-1} ,

with $h_t = \theta_j$, the agent draws $\tilde{\varepsilon}_t$ from the uniform distribution on $[F_t(\theta_{j-1}|h^{t-1}), F_t(\theta_j|h^{t-1})]$ if $j > 1$, and from the uniform distribution over $[0, F_t(\theta_1|h^{t-1})]$ otherwise, and then reports

$$\hat{h}_t = F_t^{-1}(\varepsilon_t|\hat{h}^{t-1}).$$

Then note that, given F satisfies Condition ‘‘FOSD’’, for each $t \geq 1$, we must have $h_t \leq \hat{h}_t$ (with a strict inequality occurring with positive probability in at least one period, in particular, for $t = 1$). Hence, given that output is strictly positive at each period, by following the above reporting strategy the agent with initial type θ_i must expect the same payoff $V_1^{\psi^*}(\theta_{i+1})$ as the type θ_{i+1} mimicked in period 1, augmented by the strictly positive cost savings that occur due to having a lower true cost than the reported one in each period (with the cost saving strictly positive at least in period one). This implies that, if the mechanism ψ^* is incentive compatible, then $V_1^{\psi^*}(\cdot)$ must be strictly decreasing, as claimed.

Now, suppose, towards a contradiction, that (10) is strictly negative for some date t' under the putative optimal mechanism ψ^* . Consider decreasing $q_{t'}^*(\cdot)$ uniformly by an arbitrary constant $\nu \in (0, \min_{h^{t'} \in \Theta^{t'}} \{q_{t'}^*(h^{t'})\})$. Formally, consider a new mechanism whose allocation rule \mathbf{q}' is such that $q'_t(\cdot) = q_t^*(\cdot)$ for $t \neq t'$, while, for $t = t'$, $q'_{t'}(\cdot) = q_{t'}^*(\cdot) - \nu$, and whose payment rule \mathbf{p}' is such that $p'_t(\cdot) = p_t^*(\cdot)$ for all $t \neq t'$, while, for $t = t'$, $p'_{t'}(h^{t'}) = p_{t'}^*(h^{t'}) + c(q'_{t'}(h^{t'})) - c(q_{t'}^*(h^{t'})) - \nu \mathbb{E}[\tilde{h}_{t'}|\theta_N]$ for all $h^{t'} \in \Theta^{t'}$. Note that the incentive constraints (3) are unaffected by these changes, i.e., the mechanism $\langle \mathbf{q}', \mathbf{p}' \rangle$ so constructed is incentive compatible. Furthermore, the payoff expected by the least efficient period-1 type θ_N under the new mechanism $\langle \mathbf{q}', \mathbf{p}' \rangle$ is the same as under the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. That the participation constraints of all other types in the new mechanism are also satisfied follows from the fact that the new mechanism is incentive compatible and that, under ‘‘FOSD’’, the participation of type θ_N implies the participation of any other type (this property follows from the the arguments above). Hence the new mechanism is in Ψ .

Next observe that, under the new mechanism, for each period-1 type $h'_1 \neq \theta_N$, the expected payoff under the new mechanism is equal to

$$V_1^{\langle \mathbf{q}', \mathbf{p}' \rangle}(h'_1) = V_1^{\psi^*}(h'_1) - \delta^{t'-1} \nu \left(\mathbb{E}[\tilde{h}_{t'}|\theta_N] - \mathbb{E}[\tilde{h}_{t'}|h'_1] \right) < V_1^{\psi^*}(h'_1).$$

Furthermore, because (10) is strictly negative at t' , provided ν is small, the reduction in output at t' increases expected total surplus. The new mechanism, by reducing rents and increasing total surplus, thus increases the principal’s profits, contradicting the optimality of ψ^* .

[Part 2]. Suppose, towards a contradiction, that the result is not true. Then there are adjacent periods s and $s' = s + 1$ such that the expected wedge in (10) is larger at $t = s'$ than at $t = s$ under the optimal mechanism ψ^* . Let $\mathbf{x}_t \in \mathbb{R}^N$ represent the period- t marginal probability distribution over Θ (each element x_{it} denotes the probability of type i). That the process satisfies ‘‘Stationary Markov’’ implies that $\mathbf{x}_t = \mathbf{x}$ for all t , where \mathbf{x} is the unique ergodic distribution for the Markov process with transition matrix A .

Now let $\mathbf{e}_N = (0, \dots, 0, 1)$ represent the vector whose elements are all 0, except the N^{th} element which is equal to 1. The distribution of types at date t , conditional on the initial type being θ_N is then $\mathbf{e}_N A^{t-1}$. Clearly, \mathbf{e}_N first-order stochastically dominates $\mathbf{e}_N A$. Furthermore, the assumption that F satisfies ‘‘FOSD’’ implies that A is a ‘‘stochastically monotone matrix’’ in the sense of Daley (1968). Given that $\mathbf{e}_N A^{s'-1} = (\mathbf{e}_N A) A^{s-1}$, Corollary 1a of Daley implies that $\mathbf{e}_N A^{s-1}$ stochastically dominates $\mathbf{e}_N A^{s'-1}$. This implies that

$$\mathbb{E} [\tilde{h}_{s'} | \theta_N] \leq \mathbb{E} [\tilde{h}_s | \theta_N]. \quad (31)$$

Now, consider the mechanism $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ defined as follows. Let

$$\nu \in \left(0, \min \left\{ \min_{h^s \in \Theta^s} \{q_s^*(h^s)\}, \delta \left(\bar{q} - \max_{h^{s'} \in \Theta^{s'}} \{q_{s'}^*(h^{s'})\} \right) \right\} \right).$$

The allocation rule \mathbf{q}' is such that

$$\begin{aligned} q'_s(\cdot) &= q_s^*(\cdot) - \nu \\ q'_{s'}(\cdot) &= q_{s'}^*(\cdot) + \frac{\nu}{\delta} \end{aligned}$$

whereas, for all $t \notin \{s, s'\}$, $q'_t(\cdot) = q_t^*(\cdot)$. In turn, the payment rule \mathbf{p}' is such that $p'_t(\cdot) = p_t^*(\cdot)$ for $t \notin \{s, s'\}$, whereas, for $t \in \{s, s'\}$,

$$p'_t(h^t) = p_t^*(h^t) + c(q'_t(h^t)) - c(q_t^*(h^t))$$

all h^t .

Because the stochastic process F satisfies Condition ‘‘Stationary Markov,’’ the marginal distribution of types is the same at each date. Hence, the ex-ante expected (discounted) payoff of the agent is the same under the original and the new mechanisms.

That the expected wedge under the original mechanism is larger at $t = s'$ than at $t = s$ also implies that, when ν is sufficiently small, the new mechanism improves over the original one in terms of ex-ante expected (discounted) surplus. Because the new mechanism improves over the original one both in terms of rents and expected surplus, it yields the principal a higher ex-ante expected payoff.

That the new mechanism is incentive compatible follows for the same reason as in the proof of Proposition 1; i.e., because, for any reporting strategy σ , any h_1 ,

$$V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(h_1) - V_1^{\langle \mathbf{q}^*, \sigma, \mathbf{p}^*, \sigma \rangle}(h_1) = V_1^{\langle \mathbf{q}', \mathbf{p}' \rangle}(h_1) - V_1^{\langle \mathbf{q}', \sigma, \mathbf{p}', \sigma \rangle}(h_1).$$

Finally, to see that the new mechanism is also individually rational, consider type θ_N . By (31), under $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$, the reduction in the linear part of the expected costs at period s is larger (in present value terms) than the increase in the linear part of the expected costs at s' . Furthermore, by the way the payments \mathbf{p}' are constructed, the variation in the convex part of the agent’s costs both

at $t = s$ and at $t = s'$ is neutralized by the adjustment in the payments at these two dates. As a result, type θ_N 's period-1 expected payoff under ψ' is higher than under ψ^* . That the participation of all other types is also guaranteed in ψ' follows from the fact that $V_1^{\psi'}(\theta_N) > 0$, along with the fact that the mechanism ψ' is incentive compatible and F satisfies Condition ‘‘FOSD.’’

We conclude that the new mechanism ψ' is in Ψ and achieves strictly higher ex-ante profits than ψ^* , contradicting the optimality of the latter. Q.E.D.

Proof of Proposition 3. When the process satisfies Condition ‘‘Markov’’, the agent’s continuation payoff at any date t depends on past reports, \hat{h}^{t-1} , and current true period- t type h_t , but not on past true types h^{t-1} . As a result, for any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle$, hereafter we abuse notation by denoting the agent’s expected continuation payoff at any period- t history (h^t, \hat{h}^{t-1}) by $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle}(\hat{h}^{t-1}, h_t)$.

The idea of the proof is the following. Let $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ be a putative optimal mechanism. From Lemmas 1 and 2, without loss of generality, we can assume that $\langle \mathbf{q}^*, \mathbf{p}^* \rangle \in \bar{\Psi}$, where $\bar{\Psi}$ is the class of mechanisms defined in the proof of Lemma 2 (recall that mechanisms in this class satisfy Conditions (5) and (6) for all $t \geq 2$, all h^{t-1} , with (6) holding also at $t = 1$, and that the quantity schedules $q_t^*(\cdot)$ are bounded by time-varying constants, uniformly across histories). Now suppose that the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ fails to satisfy the convergence property in the first part of the proposition. We then show that there exists another individually-rational and incentive-compatible mechanism that yields the principal a strictly higher expected payoff. The new mechanism is obtained by taking linear convex combinations of the policies of the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ with the policies of a certain efficient mechanism that also belongs in $\bar{\Psi}$. Importantly the weights that the new mechanism assigns to the efficient policies increase gradually over time at a rate that preserves incentive compatibility while also guaranteeing the profitability of the proposed perturbation.

The proof unfolds through several lemmas. The first one illustrates the incentive properties of the specific efficient mechanism we use below to construct the perturbation.

Lemma 3. *Let \mathbf{q}^E denote the efficient quantity rule, as defined in (7). Let $\hat{\mathbf{p}}^E$ denote the payment scheme defined, for all t , all h^t , by $\hat{p}_t^E(h^t) = B(q^E(h_t))$. Then let $\bar{\mathbf{p}}^E$ be the payment rule obtained from $\hat{\mathbf{p}}^E$ using the transformation in the proof of Lemma 1 to guarantee that the mechanism $\langle \mathbf{q}^E, \bar{\mathbf{p}}^E \rangle$ satisfies Equation (5) for all $t \geq 2$, all $h^{t-1} \in \Theta^{t-1}$. Finally, let \mathbf{p}^E be the payment rule obtained from $\bar{\mathbf{p}}^E$ by adding a (possibly negative) constant M to the period-1 payment rule \bar{p}_1^E (independent across the period-1 reports h_1) so as to guarantee that the agent’s participation constraints (2) are satisfied in $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ for all types, i.e., $V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}(h_1) \geq 0$ for all $h_1 \in \Theta$, with the inequality holding as equality for at least one value of h_1 . The mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ so constructed is individually rational, incentive compatible, and implements the efficient allocation rule. Since, for $t \geq 1$, the bounds \underline{b}_t and \bar{b}_t on the quantity schedules defining the set $\bar{\Psi}$ can be taken to be arbitrarily close to zero and \bar{q} , respectively, the set $\bar{\Psi}$ can be chosen to include the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$. Furthermore,*

the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ so constructed is such that, for any t , any $(h^{t-1}, h_t) \in \Theta^t$, and any $h'_t \in \Theta$, with $h'_t \neq h_t$,

$$V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^t) \geq p_t^E (h^{t-1}, h'_t) - C(q^E(h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] + \kappa, \quad (32)$$

where κ is the constant defined in (11) in the main text.

Proof of Lemma 3. That the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ is individually rational, incentive compatible, and implements the efficient allocation rule is immediate, given the way the payment rule is constructed (note that the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ gives the agent the entire surplus, net of a constant that guarantees participation; incentive compatibility then trivially holds). Moreover, using the same arguments as in Lemma 1, the construction of $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ in the lemma guarantees that $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ belongs to $\bar{\Psi}$ (with the bounds on output (\underline{b}_t) and (\bar{b}_t) specified appropriately, as explained in the lemma).

Thus consider the final claim of the lemma, i.e., the property that continuation payoffs in $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ satisfy Condition (7). Because F is Markov, it is readily verified that, for any t , any h^{t-1} , the difference in discounted expected payoffs between truthful reporting from t onwards and reporting $h'_t \neq h_t$ in period t and then reporting truthfully thereafter is equal to

$$\begin{aligned} & V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^t) - \left(p_t^E (h^{t-1}, h'_t) - C(q^E(h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ & = B(q^E(h_t)) - C(q^E(h_t), h_t) - (B(q^E(h'_t)) - C(q^E(h'_t), h_t)). \end{aligned} \quad (33)$$

Note that, if $t = 1$, the equality in (33) follows by noting that, given the way the payments in $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ are constructed, an agent with initial type h_1 who reports truthfully at all periods expects a payoff equal to

$$\begin{aligned} V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h_1) & = M + B(q^E(h_1)) - C(q^E(h_1), h_1) \\ & \quad + \mathbb{E} \left[\sum_{t=2}^{\infty} \delta^{t-1} (B(q^E(\tilde{h}_t)) - C(q^E(\tilde{h}_t), \tilde{h}_t)) | h_1 \right]. \end{aligned}$$

On the other hand, an agent whose initial type is h_1 , who reports \hat{h}_1 at $t = 1$ and then reports truthfully from $t = 2$ onwards, expects a payoff equal to

$$\begin{aligned} & M + B(q^E(\hat{h}_1)) - C(q^E(\hat{h}_1), h_1) \\ & \quad + \mathbb{E} \left[\sum_{t=2}^{\infty} \delta^{t-1} (B(q^E(\tilde{h}_t)) - C(q^E(\tilde{h}_t), \tilde{h}_t)) | h_1 \right]. \end{aligned}$$

If, instead, $t \geq 2$, the equality in (33) follows from the observation that an agent who truthfully

reports from period t onwards expects a payoff equal to

$$\begin{aligned} V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^t) = & - \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(B \left(q^E \left(\tilde{h}_s \right) \right) - C \left(q^E \left(\tilde{h}_s \right), \tilde{h}_s \right) \right) | h_{t-1} \right] \\ & + \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(B \left(q^E \left(\tilde{h}_s \right) \right) - C \left(q^E \left(\tilde{h}_s \right), \tilde{h}_s \right) \right) | h_t \right], \end{aligned}$$

whereas an agent who lies in period t by reporting h'_t and then reports truthfully from $t+1$ onwards expects a payoff equal to

$$\begin{aligned} & - \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \left(B \left(q^E \left(\tilde{h}_s \right) \right) - C \left(q^E \left(\tilde{h}_s \right), \tilde{h}_s \right) \right) | h_{t-1} \right] \\ & + B \left(q^E \left(h'_t \right) \right) - C \left(q^E \left(h'_t \right), h_t \right) \\ & + \mathbb{E} \left[\sum_{s=t+1}^{\infty} \delta^{s-t} \left(B \left(q^E \left(\tilde{h}_s \right) \right) - C \left(q^E \left(\tilde{h}_s \right), \tilde{h}_s \right) \right) | h_t \right]. \end{aligned}$$

That the equality in (33) holds for all histories then follows from the fact that, by the definition of κ , for any $h_t, h'_t \in \Theta$ with $h_t \neq h'_t$,

$$B \left(q^E \left(h_t \right) \right) - C \left(q^E \left(h_t \right), h_t \right) - \left[B \left(q^E \left(h'_t \right) \right) - C \left(q^E \left(h'_t \right), h_t \right) \right] \geq \kappa.$$

Q.E.D.

The result in Lemma 3 implies that, under the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$, incentive constraints do not bind at any history and the difference in continuation payoffs between telling the truth from period t onwards and lying in period t and then reverting to truthtelling from $t+1$ onwards is bounded by a constant κ , uniformly over histories.

Equipped with the result in Lemma 3, we now specify the proposed perturbation to the putative optimal mechanism. Let $\gamma \equiv (\gamma_t)_{t=1}^{\infty}$ be any collection of constants with $\gamma_t \in [0, 1]$ all t . Let $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle = \langle (q_t^\gamma)_{t=1}^{\infty}, (p_t^\gamma)_{t=1}^{\infty} \rangle$ be the mechanism constructed from taking a linear convex combination of the putative optimal mechanism the efficient $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and the efficient mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ of Lemma 3 as follows. For each t , each h^t ,

$$q_t^\gamma(h^t) = \gamma_t q_t^E(h^t) + (1 - \gamma_t) q_t^*(h^t) \tag{34}$$

and

$$p_t^\gamma(h^t) = \gamma_t p_t^E(h^t) + (1 - \gamma_t) p_t^*(h^t) + c(q_t^\gamma(h^t)) - \gamma_t c(q_t^E(h^t)) - (1 - \gamma_t) c(q_t^*(h^t)).$$

Note that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ so constructed ensures that, for all $t \geq 1$, all h^t ,

$$\begin{aligned} V_t^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle} (h^t) = & \gamma_t V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^t) + (1 - \gamma_t) V_t^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^t) \\ & + \sum_{s=t+1}^{\infty} \delta^{s-t} (\gamma_s - \gamma_{s-1}) \mathbb{E} \left[\mathbb{E} \left[V_s^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(\tilde{h}^s \right) - V_s^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} \left(\tilde{h}^s \right) | \tilde{h}^{s-1} \right] | h^t \right] \\ = & \gamma_t V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^t) + (1 - \gamma_t) V_t^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^t). \end{aligned} \tag{35}$$

The first equality follows from the definition of the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ along with the law of iterated expectations, while the second equality follows from the fact that $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ satisfy Condition (5).

We now establish that, if the sequence γ is appropriately chosen, the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is individually rational and incentive compatible (that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ belongs in $\bar{\Psi}$ then follows directly from (35)).

To see that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is individual rational note that the payoff that each period-1 type expects in $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ under truthful reporting is a linear convex combination of the payoffs that the same type expects in $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ and $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ (with linear weight equal to γ_1 ; see Equation (35)). The individual rationality of $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ then follows directly from the individual rationality of $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ and $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$.

Next, consider incentive compatibility. Note that, for arbitrary sequences of weights γ , the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is not guaranteed to be incentive compatible; that is, incentive compatibility of $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ need not be inherited from the incentive compatibility of $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ and $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. In particular, the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ may fail to be incentive compatible if the weights $(\gamma_t)_{t=1}^\infty$ grow too fast, e.g., if, for some $t > 1$, $\gamma_s = 0$ for all $s \in \{1, \dots, t\}$, while $\gamma_s = 1$ for all $s \in \{t+1, \dots, +\infty\}$, in which case incentive compatibility may fail at date t . This may happen because, in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, truth-telling from period t onwards may be optimal for the agent only because of the dependence of the allocations $(q_s^*(\hat{h}^s), p_s^*(\hat{h}^s))$ at dates $s > t$ on the period- t report \hat{h}_t .

Building on the above observations, below we show how, despite the problems discussed above, the incentive compatibility of $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ can be guaranteed by appropriate choice of the sequence of weights γ , whereby the weights to the efficient policy increase gradually over time at an appropriate rate. The results rely on two preliminary observations. The first makes use of the assumption that every element of the transition matrix A is positive to provide a sense in which, after enough time, an agent's expected type becomes independent of the initial conditions, in a sense made precise in the following lemma:

Lemma 4. *For all $j \in \{1, \dots, N\}$, and all $s \geq 1$,*

$$\left(\max_k A_{kj}^s - \min_k A_{kj}^s \right) \leq (1 - 2\alpha)^s.$$

Furthermore, for any $t \geq 2$, any $s \geq 0$,

$$\max_{k,l} \left| \mathbb{E} \left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_k \right] - \mathbb{E} \left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_l \right] \right| \leq b(1 - 2\alpha)^s. \quad (36)$$

Proof of Lemma 4. The first result is a well-known property of positive transition matrices; see Lemma 4.3.2 of Gallager (2013). For the second claim, it is enough to note that

$$\left| \mathbb{E} \left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_k \right] - \mathbb{E} \left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_l \right] \right| \leq (1 - 2\alpha)^s \sum_{i=1}^N \theta_i.$$

Q.E.D.

The second observation is that the agent's continuation payoff in any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$ satisfying the participation constraints in (2) as an equality for at least one type $h_1 \in \Theta$ is bounded uniformly across histories. The result is presented in the next lemma (Note that the value of the bound in Lemma 5 is an alternative to the one in Condition (6), which is a defining feature of the set $\bar{\Psi}$).

Lemma 5. *Consider any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$ satisfying the participation constraints in (2) as an equality for at least one type $h_1 \in \Theta$. Then for any $t \geq 1$, any $h^t \in \Theta^t$*

$$\left| V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^t) \right| \leq \frac{\bar{q}b}{1 - \delta(1 - 2\alpha)} \equiv \lambda.$$

Proof of Lemma 5. The proof follows from arguments similar to those establishing Lemma 1. Suppose the inequality fails to hold at some history h^t , and assume $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^t) > \frac{\bar{q}b}{1 - \delta(1 - 2\alpha)}$. The assumption that the participation constraint in (2) holds as an equality for at least one type $h_1 \in \Theta$, along with the assumption that $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$ (which implies that Condition (5) holds for all $t \geq 2$) jointly imply that, irrespective of whether $t = 1$ or $t > 1$, there exists a type θ_j for whom $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^{t-1}, \theta_j) \leq 0$ (with $h^{t-1} = \emptyset$ if $t = 1$). Suppose such a type uses the “canonical” representation of the process F described in the proof of Proposition 2 to mimic the distribution of reports of type h_t in the continuation starting with period t (see the proof of Proposition 2 for the details). Lemma 4, along with the fact that output in each period is bounded from above by \bar{q} implies that, in any period $t + s$, for $s \geq 0$, the difference in expected per-period payoffs across the two types (i.e., type θ_j using the canonical representation to mimic type h_t and type h_t reporting truthfully) is no more than

$$\bar{q}b(1 - 2\alpha)^s.$$

This implies that, by mimicking type h_t from period t onwards, an agent whose true period- t type is θ_j can guarantee himself a continuation payoff at least equal to

$$V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^t) - \frac{\bar{q}b}{1 - \delta(1 - 2\alpha)}.$$

Hence, if $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^t) > \frac{\bar{q}b}{1 - \delta(1 - 2\alpha)}$, the mechanism is not incentive compatible. A similar argument implies that a necessary condition for incentive compatibility is that $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle} (h^t) \geq -\frac{\bar{q}b}{1 - \delta(1 - 2\alpha)}$ all t , all h^t . Combining the two properties leads to the result in the lemma. Q.E.D.

We now use the properties in the previous two lemmas to specify a sequence $\gamma \in [0, 1]^\infty$ for which the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible and hence belongs to $\bar{\Psi}$.

Lemma 6. Let $\gamma \in [0, 1]^\infty$ be any non-decreasing sequence of scalars such that $\gamma_1 \in (0, 1]$ and, for all $t \geq 2$,

$$\gamma_t = \min \left\{ \gamma_1 \left(1 + \frac{\kappa}{2\delta\lambda} \right)^{t-1}, 1 \right\}, \quad (37)$$

where λ is the constant defined in Lemma 5. The mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is in $\bar{\Psi}$.

Proof of Lemma 6. First, observe that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (5) for all $t \geq 2$, all h^{t-1} , and Condition (6) for all $t \geq 1$, all h^t . These properties follow directly from the fact that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (35) along with the fact that the mechanisms $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ satisfy the same properties.

Second, observe that, because $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ belong in $\bar{\Psi}$, there exist sequences of scalars (\underline{b}_t) and (\bar{b}_t) , as specified in the definition of $\bar{\Psi}$, with $0 < \underline{b}_t < \bar{b}_t < \bar{q}$, such that $q_t^*(h^t), q_t^E(h^t) \in [\underline{b}_t, \bar{b}_t]$, all t and all h^t . That the output schedule \mathbf{q}^γ satisfies the same property (i.e., that $q_t^\gamma(h^t) \in [\underline{b}_t, \bar{b}_t]$ for all t , all h^t) then follows directly from (34).

Third, observe that, in $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, each period-1 type expects a non-negative payoff under truthtelling. This follows again from the fact that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (35) along with the fact that the mechanisms $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ are individually rational. Hence, $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is also individually rational.

It remains to show that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible. That $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (6) for all t implies that the agent's flow payoffs under $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ are bounded for all t , all (h_t, \hat{h}^t) . This means that payoffs under $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ are *continuous at infinity*, for any reporting strategy σ , implying that the one-shot-deviation principle holds. Together with Condition "Markov" and the fact that every element of A is strictly positive, this implies that, if, in $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, no one-shot deviation from truthtelling is profitable at any truthful history (that is, at any history (h^t, \hat{h}^{t-1}) such that $\hat{h}^{t-1} = h^{t-1}$), then no deviation from truthtelling is profitable at any history and hence $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible.

To establish incentive compatibility of the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ it thus suffices to show that, for all t , all $h^t \in \Theta^t$, all $h'_t \in \Theta$,

$$V_t^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle}(h^t) \geq p_t^\gamma(h^{t-1}, h'_t) - C(q_t^\gamma(h^{t-1}, h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) \mid h_t \right]. \quad (38)$$

To see that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (38) note that, for all t , all $h^t \in \Theta^t$, all $h'_t \in \Theta$,

$$\begin{aligned} V_t^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle}(h^t) &= \gamma_t V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}(h^t) + (1 - \gamma_t) V_t^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(h^t) \\ &\geq \gamma_t \left(p^E(h'_t) - C(q^E(h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) \mid h_t \right] \right) \\ &\quad + (1 - \gamma_t) \left(p_t^*(h^{t-1}, h'_t) - C(q_t^*(h^{t-1}, h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) \mid h_t \right] \right) \\ &\quad + \gamma_t \kappa. \end{aligned} \quad (39)$$

The equality is (39) simply follows from the fact that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (35), whereas the inequality follows from the fact that the efficient mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ satisfies Condition (32),

together with the fact that the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is incentive compatible. Next note that, by virtue of (37),

$$\gamma_t \kappa \geq 2\lambda\delta(\gamma_{t+1} - \gamma_t).$$

Hence, using the fact that the mechanisms $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ both satisfy the condition in Lemma 5 and the triangle inequality, we have that

$$\gamma_t \kappa \geq \delta(\gamma_{t+1} - \gamma_t) \left(\begin{array}{c} \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \\ - \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \end{array} \right).$$

Combining the last inequality with (39), we then have that

$$\begin{aligned} V_t^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle} (h^t) &\geq \gamma_t \left(p_t^E (h^{t-1}, h'_t) - C(q^E(h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + (1 - \gamma_t) \left(p_t^* (h^{t-1}, h'_t) - C(q_t^* (h^{t-1}, h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + \delta(\gamma_{t+1} - \gamma_t) \left(\begin{array}{c} \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \\ - \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \end{array} \right). \end{aligned} \tag{40}$$

Now note that the right-hand-side of (40) is equal to the right-hand side of (38). To see this it suffices to note that, for $(h^{t-1}, h'_t, h_{t+1}) \in \Theta^{t+1}$,

$$(1 - \gamma_{t+1})V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} (h^{t-1}, h'_t, h_{t+1}) + \gamma_{t+1}V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, h_{t+1}) = V_{t+1}^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle} (h^{t-1}, h'_t, h_{t+1}).$$

The last equality holds by Condition (35) (note that, while this was established for truthful histories, Condition ‘‘Markov’’ ensures the same property holds also for date t misreports h'_t , provided the agent reports truthfully after date t). The inequality in (40) thus implies that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible. We conclude that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle \in \bar{\Psi}$, as claimed. Q.E.D.

When the sequence γ satisfies Condition (37), the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ thus places progressively more weight on the outcomes of the efficient mechanism of Lemma 3 as time passes. Eventually, i.e. as soon as $\gamma_t = 1$, the mechanism is fully efficient. The less weight is initially placed on the efficient mechanism, i.e. the smaller γ_1 , the longer it takes for the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ to induce an efficient output.

Building on the above results, the final lemma in the proof (Lemma 7 below) shows that, when players are sufficiently patient, i.e., when $\delta > \bar{\delta}$, with $\bar{\delta}$ as defined in (12), if expected surplus under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ fails to converge to the efficient level, then there exists a mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, obtained by taking γ_1 sufficiently close to zero, which is not only individually rational and incentive compatible, but generates an expected payoff for the principal strictly higher than $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$.

Lemma 7. *Suppose $\delta > \bar{\delta}$, with $\bar{\delta}$ as defined in (12). Then any optimal mechanism must satisfy the convergence results in the proposition.*

Proof of Lemma 7. Consider first convergence of expected surplus. Suppose $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is an optimal mechanism and, under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, convergence of expected surplus to the efficient level fails. Recall that this means that exists $\varepsilon > 0$ and a strictly increasing sequence of dates (t_k) such that, for all k ,

$$\mathbb{E} \left[B \left(q^E \left(\tilde{h}_{t_k} \right) \right) - C \left(q^E \left(\tilde{h}_{t_k} \right), \tilde{h}_{t_k} \right) \right] - \mathbb{E} \left[B \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right) \right) - C \left(q_{t_k}^* \left(\tilde{h}^{t_k} \right), \tilde{h}_{t_k} \right) \right] > \varepsilon.$$

Now, for any $k \in \mathbb{N}$, let

$$\gamma_1(k) \equiv \left(1 + \frac{\kappa}{2\delta\lambda} \right)^{1-t_k}.$$

Then consider the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ whose weights are given by Condition (37), with $\gamma_1 = \gamma_1(k)$. Note that, when $\gamma_1 = \gamma_1(k)$, $\gamma_{t_k} = 1$. Now suppose the principal replaces the original mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ with the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$. Clearly, the increase in ex-ante expected surplus is more than $\delta^{t_k-1}\varepsilon$ (this follows from (a) the fact that the new mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ delivers at least ε more than the original mechanism in expected surplus in period t_k along with (b) the fact that the new mechanism implements the efficient output policies at all $t > t_k$ and (c) the fact that, by strict concavity of the total surplus function $B(q) - C(q, h)$, the new mechanism, by implementing policies that are a convex combination of those in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ with the efficient ones at all $t < t_k$, yields more expected surplus than the original mechanism also in all periods $t < t_k$).

Then use Condition (35) to observe that the increase in the agent's ex-ante expected rent is equal to

$$\left(1 + \frac{\kappa}{2\delta\lambda} \right)^{1-t_k} \left(\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right).$$

Combining the two effects, we have that the total change in the principal's ex-ante expected payoff is such that

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^{+\infty} \delta^{t-1} \left(B \left(q_t^\gamma \left(\tilde{h}^t \right) \right) - p_t^\gamma \left(\tilde{h}^t \right) \right) \right] - \mathbb{E} \left[\sum_{t=1}^{+\infty} \delta^{t-1} \left(B \left(q_t^* \left(\tilde{h}^t \right) \right) - p_t^* \left(\tilde{h}^t \right) \right) \right] \\ & \geq \delta^{t_k-1}\varepsilon - \left(1 + \frac{\kappa}{2\delta\lambda} \right)^{1-t_k} \left(\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right) \\ & = \delta^{t_k-1} \left[\varepsilon - \left(\delta + \frac{\kappa}{2\lambda} \right)^{1-t_k} \left(\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right) \right]. \end{aligned} \quad (41)$$

Next observe that, when $\delta > \bar{\delta}$,

$$\begin{aligned}
\delta + \frac{\kappa}{2\lambda} &= \delta + \frac{\kappa(1 - \delta(1 - 2\alpha))}{2\bar{q}b} \\
&= \delta \left(1 - \frac{\kappa(1 - 2\alpha)}{2\bar{q}b}\right) + \frac{\kappa}{2\bar{q}b} \\
&> \frac{2\bar{q}b - \kappa}{2\bar{q}b - \kappa + 2\kappa\alpha} \left(1 - \frac{\kappa(1 - 2\alpha)}{2\bar{q}b}\right) + \frac{\kappa}{2\bar{q}b} \\
&= \frac{1 - \frac{\kappa}{2\bar{q}b}}{1 - \frac{\kappa(1 - 2\alpha)}{2\bar{q}b}} \left(1 - \frac{\kappa(1 - 2\alpha)}{2\bar{q}b}\right) + \frac{\kappa}{2\bar{q}b} \\
&= 1.
\end{aligned}$$

Hence, by taking k large enough, the right-hand side of the inequality (41) is strictly positive. This means that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is more profitable than $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, contradicting the optimality of the latter.

Given that expected surplus under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ converges to expected surplus under the efficient policy, we then have that $q_t^*(h^t)$ must converge, in probability, to $q^E(h_t)$. The arguments are the same as those establishing Part 2 of Proposition 1. Q.E.D.

This completes the proof of the proposition. Q.E.D.

Proof of Corollary 1. The result follows from essentially the same arguments that establish the convergence results in Proposition 3. Suppose the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ fails to satisfy the property in the corollary for some t . Then let $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ be the mechanism constructed in the proof of Proposition 3 with the weights γ satisfying Condition (37) in Lemma 6 and with $\gamma_1 = (1 + \frac{\kappa}{2\delta\lambda})^{1-t}$. Recall that such a mechanism is individually rational and incentive compatible and implements the efficient policies from period t onwards. Furthermore, from the same arguments that lead to Condition (41) in the proof of Lemma 7, we have that the differential in the principal's ex-ante expected payoff under $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ and $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, respectively, is such that

$$\begin{aligned}
&\mathbb{E} \left[\sum_{t=1}^{+\infty} \delta^{t-1} \left(B(q_t^\gamma(\tilde{h}^t)) - p_t^\gamma(\tilde{h}^t) \right) \right] - \mathbb{E} \left[\sum_{t=1}^{+\infty} \delta^{t-1} \left(B(q_t^*(\tilde{h}^t)) - p_t^*(\tilde{h}^t) \right) \right] \\
&\geq \delta^{t-1} \left[\Delta - \left(\delta + \frac{\kappa}{2\lambda} \right)^{1-t} \left(\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}(\tilde{h}_1) \right] - \mathbb{E} \left[V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(\tilde{h}_1) \right] \right) \right], \quad (42)
\end{aligned}$$

where

$$\Delta = \mathbb{E} \left[B(q^E(\tilde{h}_t)) - C(q^E(\tilde{h}_t), \tilde{h}_t) \right] - \mathbb{E} \left[B(q_t^*(\tilde{h}^t)) - C(q_t^*(\tilde{h}^t), \tilde{h}_t) \right]$$

is the loss in period- t expected surplus under the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. Now use the result in Lemma 5 to observe that

$$\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}(\tilde{h}_1) \right] - \mathbb{E} \left[V_1^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}(\tilde{h}_1) \right] \leq 2\lambda.$$

Irrespective of whether the horizon is finite or infinite, we thus have that, whenever $\Delta > 2\lambda \left(\delta + \frac{\kappa}{2\lambda}\right)^{1-t}$, the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ strictly improves upon $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. We conclude that, under any optimal mechanism, the loss in period- t expected surplus relative to the efficient level must be bounded from above by $2\lambda \left(\delta + \frac{\kappa}{2\lambda}\right)^{-(t-1)}$, as claimed. Q.E.D.

Proof of Proposition 4. The proof follows from arguments similar to those establishing Proposition 3. Without loss of generality, assume $\langle \mathbf{q}^*, \mathbf{p}^* \rangle \in \bar{\Psi}$ and suppose $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ violates the property in the proposition. Then let $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ be the efficient mechanism defined in Lemma 3 (again, choose the sequence of bounds on output defining the set $\bar{\Psi}$ so that $\bar{\Psi}$ includes $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$). Now let $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ be a mechanism derived from $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ using the same construction as in the proof of Proposition 3, for some sequence of weights $\gamma = (\gamma_t)_{t=1}^\infty$, with $\gamma_t \in [0, 1]$ for all t .

The point of departure from the proof of Proposition 3 is in establishing the incentive compatibility of the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, for an appropriate choice of the sequence γ . We proceed as follows. In any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$, for any $t \geq 1$, $h^{t-1} \in \Theta^{t-1}$, $h_t, h'_t \in \Theta$,

$$\begin{aligned} & \left| \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}, \mathbf{p} \rangle} \left(h^{t-1}, h'_t, \tilde{h}_{t+1} \right) \mid h_t \right] \right| \\ &= \left| \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}, \mathbf{p} \rangle} \left(h^{t-1}, h'_t, \tilde{h}_{t+1} \right) \mid h_t \right] - \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}, \mathbf{p} \rangle} \left(h^{t-1}, h'_t, \tilde{h}_{t+1} \right) \mid h'_t \right] \right| \\ &\leq \frac{\bar{q}\Delta\theta}{1-\delta} N\varepsilon(\delta). \end{aligned} \quad (43)$$

Note that the equality in (43) follows from the fact that $\langle \mathbf{q}, \mathbf{p} \rangle$ satisfies Condition (5), while the inequality follows from the fact that $\left| V_{t+1}^{\langle \mathbf{q}, \mathbf{p} \rangle} \left(h^{t+1} \right) \right|$ is bounded by $\frac{\bar{q}\Delta\theta}{1-\delta}$ uniformly over $t \geq 1$ and $h^{t+1} \in \Theta^{t+1}$, as established in Lemma 1, together with the fact that $|f_{t+1}(h_{t+1}|h_t) - f_{t+1}(h_{t+1}|h'_t)| < \varepsilon(\delta)$ for any (h_t, h'_t, h_{t+1}) .

Now, let $\gamma = (\gamma_t)_{t=1}^\infty \in [0, 1]^\infty$ be any sequence of scalars satisfying $\gamma_1 \in (0, 1)$ and, for all $t \geq 2$,

$$\gamma_t = \min \left\{ \gamma_1 \left(\frac{(1-\delta)\kappa + 2\delta\bar{q}\Delta\theta N\varepsilon(\delta)}{2\delta\bar{q}\Delta\theta N\varepsilon(\delta)} \right)^{t-1}, 1 \right\}. \quad (44)$$

From the same arguments as in the proof of Lemma 6, the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible if and only if it satisfies Condition (38) for all t , $h^t \in \Theta^t$, and $h'_t \in \Theta$.

Now observe that (44) implies that the sequence $(\gamma_t)_{t=1}^\infty$ described above satisfies

$$\gamma_t \kappa \geq 2\delta (\gamma_{t+1} - \gamma_t) \frac{\bar{q}\Delta\theta}{1-\delta} N\varepsilon(\delta)$$

for all $t \geq 1$. Furthermore, because $\langle \mathbf{q}^E, \mathbf{p}^E \rangle, \langle \mathbf{q}^*, \mathbf{p}^* \rangle \in \bar{\Psi}$, Condition (43) above (together with the triangle inequality) implies that, for all $t \geq 1$, all h^{t-1} , h_t and h'_t ,

$$2\delta (\gamma_{t+1} - \gamma_t) \frac{\bar{q}\Delta\theta}{1-\delta} N\varepsilon(\delta) \geq \delta (\gamma_{t+1} - \gamma_t) \left(\begin{array}{c} \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(h^{t-1}, h'_t, \tilde{h}_{t+1} \right) \mid h_t \right] \\ - \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle} \left(h^{t-1}, h'_t, \tilde{h}_{t+1} \right) \mid h_t \right] \end{array} \right).$$

The above two inequalities, combined with the fact that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (39) for all t , all h^t , all h'_t , then imply that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies all the inequalities in (40). As explained in the proof of Proposition 3, these inequalities coincide with those in (38). This means that the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible.

Arguments similar to those establishing Lemma 7 in the proof of Proposition 3 then imply that, if, under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, convergence in expected surplus fails, then, provided

$$\delta \left(\frac{(1 - \delta) \kappa + 2\delta\bar{q}\Delta\theta N\varepsilon(\delta)}{2\delta\bar{q}\Delta\theta N\varepsilon(\delta)} \right) > 1, \quad (45)$$

one can construct an incentive compatible and individually rational mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, with an appropriate choice of γ_1 and $(\gamma_t)_{t=2}^\infty$ satisfying (44), that yields the principal a higher expected payoff than $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. The result in the proposition then follows from the fact that the inequality in (45) is satisfied for $\varepsilon(\delta)$ sufficiently small. Q.E.D.

Proof of Proposition 5. The proof follows from arguments similar to those in the proof of Proposition 3, with the key difference being that the perturbations now involve randomizations over the schedules in the putative optimal mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ and in an efficient mechanism $\langle \mu^E, \mathbf{p}^E \rangle$, with the agent's expected production cost thus linear in the probabilities the perturbation assigns the two mechanisms $\langle \mu^*, \mathbf{p}^* \rangle$ and $\langle \mu^E, \mathbf{p}^E \rangle$, (the details are explained below).

Arguments similar to those establishing Lemma 1 imply that we can restrict attention to mechanisms $\langle \mu, \mathbf{p} \rangle \in \Psi^S$ such that (a) for all $t \geq 2$ and $h^{t-1} \in \Theta^{t-1}$,

$$\mathbb{E} \left[V_t^{\langle \mu, \mathbf{p} \rangle} (h^{t-1}, \tilde{h}_t) | h_{t-1} \right] = 0, \quad (46)$$

and (b) the participation constraint (15) is satisfied with equality for at least one type h_1 . Arguments analogous to those establishing Lemma 5 then imply that such mechanisms satisfy

$$\left| V_t^{\langle \mu, \mathbf{p} \rangle} (h^t) \right| \leq \frac{Nu}{1 - \delta(1 - 2\alpha)} \equiv \bar{\lambda}, \quad (47)$$

for all t and all $h^t \in \Theta^t$, where u is the Lipschitz constant in Property 2 of Condition ‘‘Cost restriction’’.

Next, for any $t \geq 1$, any h^t , let $\mu_t^E(h^t)$ be the degenerate distribution assigning probability mass one on the efficient quality $q^E(h^t)$. Then observe that arguments similar to those establishing Lemma 3 imply the existence of a (deterministic) efficient mechanism $\langle \mu^E, \mathbf{p}^E \rangle$ satisfying the above conditions and such that, for all $t \geq 1$, all h^t , and all $h'_t \neq h_t$,

$$V_t^{\langle \mu^E, \mathbf{p}^E \rangle} (h^t) \geq p_t^E(h^{t-1}, h'_t) - C(q^E(h'_t), h_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} (h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] + \kappa, \quad (48)$$

where $\kappa > 0$ is the constant defined in (11).

We now show how one can construct a mechanism $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle \in \Psi^S$ whose policies are obtained by taking convex combinations of the policies in $\langle \mu^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ with time-varying weights γ_t

that grow over time at an appropriate rate that guarantees incentive compatibility of the mechanism $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$. The arguments here parallel those establishing Lemma 6. Let $\gamma = (\gamma_s) \in [0, 1]^\infty$ be any non-decreasing sequence of positive scalars satisfying $\gamma_1 \in (0, 1)$ and

$$\gamma_t = \min \left\{ \gamma_1 \left(1 + \frac{\kappa}{2\delta\bar{\lambda}} \right)^{t-1}, 1 \right\} \quad (49)$$

for all $t \geq 2$. Then let $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ be the mechanism whose policies are given, for each $t \geq 1$ and each h^t , by

$$\mu_t^\gamma(h^t) = \gamma_t \mu_t^E(h^t) + (1 - \gamma_t) \mu_t^*(h^t) \quad (50)$$

and

$$p_t^\gamma(h^t) = \gamma_t p_t^E(h^t) + (1 - \gamma_t) p_t^*(h^t). \quad (51)$$

Note that, unless μ^* happens to coincide with the efficient policy, the mechanism $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ involves randomizations over quantities (with probability γ_t placed on the efficient quantity and $1 - \gamma_t$ on a draw from the distribution $\mu_t^*(h^t)$).

For any t , h^t and $h'_t \neq h_t$,

$$\begin{aligned} & V_t^{\langle \mu^\gamma, \mathbf{p}^\gamma \rangle}(h^t) \\ &= \gamma_t V_t^{\langle \mu^E, \mathbf{p}^E \rangle}(h^t) + (1 - \gamma_t) V_t^{\langle \mu^*, \mathbf{p}^* \rangle}(h^t) \\ &\geq \gamma_t \left(p_t^E(h^{t-1}, h'_t) - \int C(\tilde{q}, h_t) d\mu_t^E(h^{t-1}, h'_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + (1 - \gamma_t) \left(p_t^*(h^{t-1}, h'_t) - \int C(\tilde{q}, h_t) d\mu_t^*(h^{t-1}, h'_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + \gamma_t \kappa \\ &\geq \gamma_t \left(p_t^E(h^{t-1}, h'_t) - \int C(\tilde{q}, h_t) d\mu_t^E(h^{t-1}, h'_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + (1 - \gamma_t) \left(p_t^*(h^{t-1}, h'_t) - \int C(\tilde{q}, h_t) d\mu_t^*(h^{t-1}, h'_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \right) \\ &\quad + \delta (\gamma_{t+1} - \gamma_t) \left(\begin{array}{c} \mathbb{E} \left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \\ - \mathbb{E} \left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right] \end{array} \right) \\ &= p_t^\gamma(h^{t-1}, h'_t) - \int C(\tilde{q}, h_t) d\mu_t^\gamma(h^{t-1}, h'_t) + \delta \mathbb{E} \left[V_{t+1}^{\langle \mu^\gamma, \mathbf{p}^\gamma \rangle}(h^{t-1}, h'_t, \tilde{h}_{t+1}) | h_t \right]. \end{aligned} \quad (52)$$

The first equality follows by the fact that both $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ satisfy (46), implying that, at any truthful history, the agent's continuation payoff under $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ satisfies the same decomposition as in (35) with $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ replacing $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ and $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. The first inequality follows from the incentive compatibility of the mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ along with Condition ‘‘Markov’’ and the fact that $\langle \mu^E, \mathbf{p}^E \rangle$ satisfies (48). The second inequality follows by the choice of the sequence $\gamma = (\gamma_s)$ and the fact that the agent's continuation payoffs under both $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ satisfy the bound

in (47) (the arguments are the same as in the proof of Lemma 6). The final equality follows from the same equality as in (35). The last inequality in (52) implies that one-shot deviations from truthful reporting are unprofitable for the agent. This property, along with the fact that the process satisfies Condition “Markov” and that payoffs are continuous at infinity then implies incentive compatibility of $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ (the details of the arguments are omitted because they are essentially the same as in the proof of Proposition 3). Individual rationality of $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ in turn follows from the individual rationality of $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ along with the fact that the agent’s payoffs in $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ satisfy the same decomposition as in (35). We thus conclude that $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle \in \Psi^S$.

The first claim in the proposition then follows from arguments similar to those in the proof of Lemma 7. In particular, suppose that the assumed optimal mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ fails to satisfy the property of convergence of expected surplus to the efficient level specified in the proposition. This means that there exists $\varepsilon > 0$ and a strictly increasing sequence (t_k) such that, for all k ,

$$\mathbb{E} \left[B \left(q^E \left(\tilde{h}_{t_k} \right) \right) - C \left(q^E \left(\tilde{h}_{t_k} \right), \tilde{h}_{t_k} \right) \right] - \mathbb{E} \left[\int \left(B(\tilde{q}) - C(\tilde{q}, \tilde{h}_{t_k}) \right) d\mu_{t_k}^* \left(\tilde{h}^{t_k} \right) \right] > \varepsilon.$$

Now, for any $k \in \mathbb{N}$, let $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ be the mechanism defined above with weights γ given by $\gamma_1 = \gamma_1^s(k)$ with

$$\gamma_1^s(k) \equiv \left(1 + \frac{\kappa}{2\delta\bar{\lambda}} \right)^{1-t_k}$$

and, for any $t \geq 2$, γ_t satisfying (49). Note that the $\gamma_1 = \gamma_1^s(k)$ guarantees that, in such a mechanism, $\gamma_{t_k} = 1$. The increase in ex-ante expected surplus from switching from $\langle \mu^*, \mathbf{p}^* \rangle$ to $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ is at least $\delta^{t_k-1}\varepsilon$ (the arguments are similar to those in the proof of Proposition 3). Furthermore, because the agent’s period-1 rents satisfy the decomposition property in (35), the increase in the agent’s ex-ante expected rents is equal to

$$\gamma_1(k) \left\{ \mathbb{E} \left[V_1^{\langle \mu^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mu^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right\}.$$

It follows that the change in the principal’s ex-ante expected payoff associated with the switch from $\langle \mu^*, \mathbf{p}^* \rangle$ to $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ is equal to

$$\begin{aligned} & \mathbb{E} \left[\sum_t \delta^{t-1} \left(\int B(\tilde{q}) d\mu_t^\gamma \left(\tilde{h}^t \right) - p_t^\gamma \left(\tilde{h}^t \right) \right) \right] - \mathbb{E} \left[\sum_t \delta^{t-1} \left(\int B(\tilde{q}) d\mu_t^* \left(\tilde{h}^t \right) - p_t^* \left(\tilde{h}^t \right) \right) \right] \\ & \geq \delta^{t_k-1}\varepsilon - \left(1 + \frac{\kappa}{2\delta\bar{\lambda}} \right)^{1-t_k} \left(\mathbb{E} \left[V_1^{\langle \mu^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mu^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right) \\ & = \delta^{t_k-1} \left(\varepsilon - \left(\delta + \frac{\kappa}{2\bar{\lambda}} \right)^{1-t_k} \left(\mathbb{E} \left[V_1^{\langle \mu^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mu^*, \mathbf{p}^* \rangle} \left(\tilde{h}_1 \right) \right] \right) \right). \end{aligned} \quad (53)$$

Provided $\delta + \frac{\kappa}{2\bar{\lambda}} > 1$, the expression above is strictly positive for t_k sufficiently large. This is guaranteed for $\delta > \bar{\delta}^s$, with $\bar{\delta}^s$ as given in (17), thus establishing the claim in the first part of the proposition.

Next, consider the second claim in the proposition. Let $(\langle \mu^k, \mathbf{p}^k \rangle)$ denote a sequence of mechanisms in Ψ^S , with $\lim_{k \rightarrow \infty} \Pi(\mu^k, \mathbf{p}^k) = \sup_{\langle \mu, \mathbf{p} \rangle \in \Psi^S} \Pi(\mu, \mathbf{p})$, and satisfying the same properties as

the mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ considered above for the case where an optimal mechanism exists. Suppose, towards a contradiction, that there is $\varepsilon > 0$ for which there is no value $\bar{t} \in \mathbb{N}$ and corresponding sequence (s_k) , with $s_k \rightarrow \infty$, for which (19) is satisfied for all $\bar{t} \leq t \leq \bar{t} + s_k$. Then, for any \bar{t} , there is $\bar{s}(\bar{t})$ such that, for any $k \in \mathbb{N}$, there exists $t(\bar{t}; k) \in \{\bar{t}, \bar{t} + 1, \dots, \bar{t} + \bar{s}(\bar{t})\}$ with

$$\begin{aligned} & \mathbb{E} \left[B \left(q^E \left(\tilde{h}_{t(\bar{t}; k)} \right) \right) - C \left(q^E \left(\tilde{h}_{t(\bar{t}; k)} \right), \tilde{h}_{t(\bar{t}; k)} \right) \right] \\ & - \mathbb{E} \left[\int \left(B(\tilde{q}) - C(\tilde{q}, \tilde{h}_{t(\bar{t}; k)}) \right) d\mu_{t(\bar{t}; k)}^k \left(\tilde{h}^{t(\bar{t}; k)} \right) \right] \geq \varepsilon. \end{aligned}$$

Now, consider mechanisms $\langle \mu^{k, \gamma}, \mathbf{p}^{k, \gamma} \rangle$ whose policies are given, for each $t \geq 1$ and each h^t , by

$$\mu_t^{k, \gamma} (h^t) = \gamma_t \mu_t^E (h^t) + (1 - \gamma_t) \mu_t^k (h^t)$$

and

$$p_t^{k, \gamma} (h^t) = \gamma_t p_t^E (h^t) + (1 - \gamma_t) p_t^k (h^t),$$

with $\gamma_t = \min \left\{ \gamma_1 \left(1 + \frac{\kappa}{2\delta\lambda} \right)^{t-1}, 1 \right\}$ and $\gamma_1 = \left(1 + \frac{\kappa}{2\delta\lambda} \right)^{1-t(\bar{t}; k)}$. Arguments analogous to those used above to establish the first claim in the proposition imply that the mechanism $\langle \mu^{k, \gamma}, \mathbf{p}^{k, \gamma} \rangle$ increases the principal's expected profits by at least

$$\delta^{t(\bar{t}; k)-1} \left(\varepsilon - \left(\delta + \frac{\kappa}{2\lambda} \right)^{1-t(\bar{t}; k)} \left(\mathbb{E} \left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle} \left(\tilde{h}_1 \right) \right] - \mathbb{E} \left[V_1^{\langle \mu^k, \mathbf{p}^k \rangle} \left(\tilde{h}_1 \right) \right] \right) \right),$$

which, provided \bar{t} is taken sufficiently large, exceeds zero by at least some amount that is independent of k . For k large enough, this implies the principal obtains expected profits that are strictly higher than $\sup_{\langle \mu, \mathbf{p} \rangle \in \Psi^S} \Pi(\mu, \mathbf{p})$, a contradiction. Q.E.D.

Proof of Proposition 6. The result follows from the arguments in the main text. Q.E.D.

Proof of Proposition 7. Observe that, when the process satisfies Condition ‘‘Regularity,’’ $\mathbb{E}[I_t(\tilde{h}^t)h_1] = \frac{d}{d\tilde{h}_1} \mathbb{E}[\tilde{h}_t|h_1]$. Thus,

$$\begin{aligned} \mathbb{E} \left[\frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} I_t(\tilde{h}^t) \right] &= \mathbb{E} \left[\frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} \mathbb{E}[I_t(\tilde{h}^t) | \tilde{h}_1] \right] = \int_{\underline{\theta}}^{\bar{\theta}} F_1(\theta_1) \mathbb{E}[I_t(\tilde{h}^t) | \theta_1] d\theta_1 \\ &= F_1(\theta_1) \mathbb{E}[\tilde{h}_t | \theta_1] \Big|_{\theta_1=\underline{\theta}}^{\theta_1=\bar{\theta}} + \int_{\underline{\theta}}^{\bar{\theta}} f_1(\theta_1) \mathbb{E}[\tilde{h}_t | \theta_1] d\theta_1 \\ &= \mathbb{E}[\tilde{h}_t | \bar{\theta}] - \mathbb{E}[\tilde{h}_t]. \end{aligned}$$

When, in addition, F satisfies Condition ‘‘Ergodicity,’’ then $\mathbb{E}[\tilde{h}_t | \bar{\theta}] - \mathbb{E}[\tilde{h}_t] \rightarrow 0$, as $t \rightarrow \infty$, implying that $\mathbb{E} \left[\frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} I_t(\tilde{h}^t) \right] \rightarrow 0$, as $t \rightarrow \infty$, as claimed.

If, in addition to Condition ‘‘Ergodicity,’’ F satisfies Condition ‘‘FOSD,’’ then

$$\mathbb{E}[\tilde{h}_t | \bar{\theta}] - \mathbb{E}[\tilde{h}_t] \geq 0$$

so that the convergence is from above.

Finally, if, in addition to the conditions above, F is stationary, then

$$\mathbb{E} \left[\frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} I_t(\tilde{h}^t) \right] - \mathbb{E} \left[\frac{F_1(\tilde{h}_1)}{f_1(\tilde{h}_1)} I_s(\tilde{h}^s) \right] = \mathbb{E}[\tilde{h}_t | \bar{\theta}] - \mathbb{E}[\tilde{h}_s | \bar{\theta}] \leq 0$$

for any $t > s$, which implies that convergence is monotone in time. Q.E.D.