Robust Predictions in Dynamic Screening*

Daniel Garrett[†] **Alessandro Pavan**[‡] **Juuso Toikka**[§]

November, 2017
[PRELIMINARY AND INCOMPLETE]

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. 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 agent learns 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 these 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 (see 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 wedge decreases over time and vanishes in the long run (Proposition 2). These results hold across a broad range of preferences, and, in particular, for any discount factor for the players. These results 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., reductions in the per-period surplus relative to the first best) vanish in the long run, and that the supplied quantity converges in probability to the 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 quantity converges in probability to its first-best level. The result leverages 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 his true type. We also extend these results (in Proposition 4) to settings with an arbitrary discount factor but where the type process is not too persistent.

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 quantity 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 ones 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 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 a 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 in general (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 applied).

While our focus in this paper is on the long-run properties of optimal mechanisms, our analysis also generates 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, in general, we can provide conditions ensuring that convergence to efficiency is not merely something that must be expected in the remotely distant future.

Finally, note that 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 on 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.

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 environments, 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, properties of 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), Eső 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 how these conditions appear when 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.[1] More importantly, none of the results 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 the appropriate perturbations to proposed optimal allocations. Relative to the variational arguments in the earlier paper, the perturbations needed for our present argument are complex: the perturbed allocations in each period are linear combinations of the proposed optimal and efficient allocations, with the weight on the efficient allocation gradually increasing over time from the beginning of the relationship.

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 lies in showing that monotonicity is violated when the process governing the evolution of the agent's 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 also seek results on convergence to efficiency. They focus on mechanisms

[1] The 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 plus a random "shock" that is independent of the initial type.

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 monopoly 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, is that strongly monotone mechanisms approximate the optimal discounted average payoffs as the players become infinitely patient. Note, however, that this does not imply distortions vanish with time in the truly optimal mechanism (since distortions could, in principle, remain large in the sufficiently distant future while having a negligible effect on expected welfare in the relationship). Understanding how allocations (not only payoffs) behave in the truly optimal mechanism would be of interest, say, for an empiricist interested in testing the implications of dynamic contracting from a long time series.

Our results then 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. Our predictions for the behavior of expected "wedges" (see Section 3.2) do not depend on the discount factor, and we also obtain (see Section 3.3) sufficient conditions on primitives that ensure vanishing distortions under the optimal mechanism *irrespective* of the discount factor.

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 $t = 1, 2, ..., T$.

The agent can produce a good in variable quantities $q_t \in (0, \bar{q})$ at each date t (or he can walk away from the relationship and produce a quantity zero). The principal is *required* to procure a (strictly) positive quantity of the input in every period. The principal has payoffs quasi-linear in transfers, and enjoys a per-period benefit from the input which is an increasing function of quantity. This benefit is determined according to a function $B : (0, \bar{q}) \to \mathbb{R}$ that is twice continuously differentiable, strictly increasing, strictly concave, and satisfies an Inada condition $\lim_{q \to 0} B(q) = -\infty$.

The agent's per-period payoff is also quasi-linear in transfers, and the agent suffers a cost of producing quantity q equal to $C(q,h)$, where h is his "type" in that period. For now, we focus on the case where agent types h are drawn from a finite set $\Theta = \{\theta_1, \ldots, \theta_N\}$, with $0 < \theta_1 < \cdots < \theta_N < +\infty$ and $N \geq 2$ (later, we discuss the case where the agent's types are drawn from an absolutely continuous distribution with compact support $\Theta = [\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} > 0$). Let $\Delta\theta \equiv \theta_N - \theta_1$. We assume that

$$C(q,h) = hq + c(q),$$

where the function $c(q)$ is twice-continuously differentiable, strictly increasing, strictly convex, and satisfies the Inada condition $\lim_{q \to \bar{q}} c(q) = +\infty$ (Proposition 5, however, extends our results for vanishing distortions to more general cost functions).

Both the principal and the agent have time-additively-separable preferences given by

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

where $h_t \in \Theta$ is the agent's period-t type, p_t denotes the total payment from the principal to the agent in period t, and $\delta \in (0,1]$ denotes the common discount factor (with $\delta < 1$ in case $T = +\infty$).

We now specify the process governing the evolution of the agent's type. We use the following notation for sequences of types: $h_s^t \equiv (h_s, \ldots, h_t)$, $h^t \equiv (h_1, \ldots, h_t)$ and $h_{-s}^t \equiv (h_1, \ldots, h_{s-1}, h_{s+1}, \ldots, h_t)$. Let then $F \equiv (F_t)$ be the collection of cdfs describing the evolution of the agent's private types, with F_1 denoting the initial distribution of h_1 over Θ and, for all $t \geq 2$, $F_t(\cdot \mid h^{t-1})$ denoting the cdf of h_t given h^{t-1}. In particular, for each $n \in \{1, \ldots, N\}$, each $h^{t-1} \in \Theta^{t-1}$,

$$F_1(\theta_n) = \sum_{i=1}^{n} f_1(\theta_i), \text{ and}$$

$$F_t(\theta_n \mid h^{t-1}) = \sum_{i=1}^{n} f_t\left(\theta_i | h^{t-1}\right)$$

where, for $i \in \{1, \ldots, N\}$, $f_1(\theta_i)$ denotes the probability of type θ_i at date 1, while $f_t\left(\theta_i | h^{t-1}\right)$ denotes the probability of type θ_i at date t, following history h^{t-1}. For now, the only restriction on the stochastic process for types F is that these distributions have full support on a fixed finite set of types; i.e., for all i, t, and h^{t-1}, $f_1(\theta_i) > 0$ and $f_t\left(\theta_i | h^{t-1}\right) > 0$.

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 the collection of message spaces for each period, and $\phi \equiv (\phi_t)_{t=1}^T$ specifies mappings from messages to allocations. In particular, for each t,

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

  specifies a payment-quantity pair for each possible profile of messages $m^t \equiv (m_1, ..., m_t) \in \mathcal{M}_1 \times \cdots \times \mathcal{M}_t$ (agreement to participate in the mechanism will thus be taken as a promise to produce a strictly positive quantity in every period). A mechanism is thus equivalent to a menu of long-term contracts.

If the agent refuses to participate in φ, the game ends. This is taken to be the worst possible outcome for the principal, since it fails to secure the input (for instance, we may take the principal's payoff from this event to be $-\infty$), while the agent earns a payoff equal to zero. If the agent accepts to participate in φ, he chooses a message $m_1 \in \mathcal{M}_1$, supplies the principal with $q_1(m_1)$, receives a transfer $p_1(m_1)$, and the game then moves to period 2.

- At the beginning of period $t \geq 2$, the agent privately learns h_t. He then sends a new message $m_t \in \mathcal{M}_t$, supplies a quantity $q_t(m^t)$, receives a transfer $p_t(m^t)$, and the game moves to period $t+1$.
- $\cdots$
- At $t = T+1$ (in case T is finite), the game is over.

Remark. The game described above assumes that the principal perfectly commits to the mechanism φ. It also assumes that, at any period $t \geq 2$, the buyer is constrained to stay in the relationship if he signs the contract in period one. When the agent has deep pockets, there are simple ways to distribute the payments over time so that it is in the interest of the buyer to remain in the relationship at all periods, irrespective of what he did in the past.

The principal's problem involves designing a mechanism that disciplines the provision of quantity and the payments over time. Because the principal can commit, the revelation principle applies.[2] Hence, without loss of optimality, one can restrict attention to direct mechanisms in which $\mathcal{M}_t = \Theta$ for all t and are such that the agent finds it optimal to report truthfully at all periods. Given our focus on direct mechanisms, we hereafter omit the message spaces and identify such a mechanism directly with the policies $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$ that it induces, where $\mathbf{q} = (q_t)_{t=1}^T$ and $\mathbf{p} = (p_t)_{t=1}^T$, and where $q_t : \Theta^t \to (0, \bar{q})$ and $p_t : \Theta^t \to \mathbb{R}$. We let σ denote an arbitrary reporting strategy for the agent and $\mathbf{q}^\sigma$ and $\mathbf{p}^\sigma$ the quantity and transfer policy induced by the strategy σ.

For any mechanism ψ, let

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

denote the agent's expected continuation payoff from date t onwards, when the realized sequence of types up to period t is h^t and when the agent reported the sequence of types $\hat{h}^{t-1}$ in previous periods, given that the agent intends to report truthfully from date t onwards. (Throughout, hatted variables represent reports, while random variables are denoted with tildes.) We then require that

[2] See, among others, Myerson (1981).

the principal selects the mechanism ψ from a set Ψ of mechanisms satisfying both the *individual rationality* condition

$$V_1^{\psi}(h_1) \geq 0 \quad \text{for all } h_1 \in \Theta \tag{1}$$

and the *incentive compatibility* condition

$$\mathbb{E}\left[V_1^{\psi}(\tilde{h}_1)\right] \geq \mathbb{E}\left[\sum_{t\geq 1} \delta^{t-1}\left(p_t^{\sigma}\left(\tilde{h}^t\right) - C\left(q_t^{\sigma}\left(\tilde{h}^t\right), \tilde{h}_t\right)\right)\right]. \tag{2}$$

for all possible reporting strategies σ. Condition (1) requires that the agent prefers to participate in period one and report truthfully thereafter, rather than not participating and receiving the payoff associated with his outside option (zero). Condition (2) requires that the agent prefers to follow a truthful reporting strategy rather than any other reporting strategy σ.[3] The principal then maximizes

$$\mathbb{E}\left[\sum_{t} \delta^{t-1}\left(B\left(q_t\left(\tilde{h}^t\right)\right) - p_t\left(\tilde{h}^t\right)\right)\right] \tag{3}$$

by choice of $\psi \in \Psi$. We refer to a mechanism ψ that maximizes (3) over Ψ as an optimal mechanism.

3 Robust predictions for discrete (finite) types

3.1 Preliminary properties of optimal mechanisms

We begin by studying the optimal mechanism for the case where the type space Θ is finite, as described above. A preliminary comment relates to our focus on deterministic mechanisms. As usual, quasi-linearity of payoffs implies that deterministic transfers are without loss of generality. Perhaps more subtle, deterministic quantity allocations are also without loss of optimality. A random mechanism 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 $\mu_t\left(h^t, q^{t-1}\right)$ on the interval of quantities $(0, \bar{q})$, with $\mu_t\left(h^t, q^{t-1}\right)(q) = \Pr\left(\tilde{q} \leq q\right)$, i.e. the probability that the realized quantity is in $(0, q)$. After reporting at date t, the mechanism would draw a quantity according to the probability distribution $\mu_t\left(h^t, q^{t-1}\right)$, which the agent would be compelled to produce. Note that, for any probability distribution $\mu\left(h^t, q^{t-1}\right)$, and any i and j,

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

Hence, if the mechanism calls for a non-degenerate distribution over quantities $\mu_t\left(h^t, q^{t-1}\right)$ at history h^t, the mechanism that instead calls for a deterministic quantity $\mathbb{E}_{\mu_t(h^t,q^{t-1})}\left[\tilde{q}\right]$ at $\left(h^t, q^{t-1}\right)$, and pays the agent less by $\mathbb{E}_{\mu_t(h^t,q^{t-1})}\left[c\left(\tilde{q}\right)\right] - c\left(\mathbb{E}_{\mu_t(h^t,q^{t-1})}\left[\tilde{q}\right]\right)$, remains incentive compatible and generates the

[3] The condition in (2) is an ex-ante (i.e., period-0) incentive-compatibility condition. However, because of the assumption of full support, (2) holds if and only if incentive compatibility holds at all period-t histories (that is, at all histories $(h^t, \hat{h}^{t-1})$ such that $\hat{h}^{t-1} = h^{t-1}$), all $t \geq 0$.

same expected payoff for the agent at all histories. Because of the concavity of $B(\cdot)$, the principal's expected payoff in Equation (3) is higher under the deterministic mechanism.

A second observation is a familiar one from the existing literature: because both the principal's and the agent's payoffs are linear in transfers (and because the players are equally patient), payoffs at any date t depend on the sum of discounted future payments, but not on the precise timing of these payments. This leads us to the following preliminary 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 such that, for all $t \geq 2$, all $h^{t-1} \in \Theta^{t-1}$,*

$$\mathbb{E}\left[V_t^{\psi'}\left(\tilde{h}^t; h^{t-1}\right) | h^{t-1}\right] = 0. \tag{4}$$

Furthermore, in any mechanism satisfying Equation (4), for all $t \geq 2$, all $h^t \in \Theta^t$,

$$\left|p_t\left(h^t\right) - C\left(q_t\left(h^t\right), h_t\right)\right| \leq \frac{\bar{q}\Delta\theta}{1-\delta}. \tag{5}$$

Finally, any optimal mechanism satisfying Equation (4), all $t \geq 2$, all $h^t \in \Theta^t$, is such that Condition (5) 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. The second part of the result can then be read as a bound on the agent's continuation payoff at any history h^t. This follows from combining the first part of the lemma with agent incentive constraints. We use these observations to establish the existence of an optimal mechanism in the next result.

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

Given that payoffs are quasi-linear in transfers, multiplicity of optimal payment rules $\mathbf{p}^*$ is to be expected. Lemma **2** states, nonetheless, that the optimal allocation $\mathbf{q}^*$ is unique. This follows in turn using the convexity of the agent's cost function $C(\cdot, h)$ for each $h \in \Theta$, the concavity of the principal's benefit function $B(\cdot)$.

3.2 Convergence in expectation

We are now ready to give our first prediction for long-run behavior. We make the following restriction on the stochastic process F, formalizing the idea that the agent's initial type eventually becomes uninformative about the distribution of the agent's later types.

Condition 1. [Long-run Independence] Suppose that $T=+\infty$. The dependence of the date-t distribution on the date-1 distribution vanishes in the following sense:[4]

$$\lim_{t\to\infty}\max_{h_1,h_1',h_t\in\Theta}\left|F_t\left(h_t|h_1\right)-F_t\left(h_t|h_1'\right)\right|=0.$$

The following result then characterizes ex-ante expected distortions at dates long after the relationship has commenced. To understand these distortions, it is helpful to define the efficient policy $q^E\left(h_t\right)$ at any history of types $h^t\in\Theta^t$ to be the unique value in $(0,\bar{q})$ such that

$$B'\left(q^E\left(h_t\right)\right)=h_t+c'\left(q^E\left(h_t\right)\right). \tag{6}$$

Distortions can thus be defined relative to the efficient quantity at any history of types h^t.

Proposition 1. *Suppose $T=+\infty$ and Condition "Long-run independence" holds. Then:*

1. *The optimal mechanism satisfies*

$$\mathbb{E}\left[B'\left(q_t^*\left(\tilde{h}^t\right)\right)-\tilde{h}_t-c'\left(q_t^*\left(\tilde{h}^t\right)\right)\right]\to 0 \tag{7}$$

as $t\to\infty$.

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

$$\begin{aligned}&\mathbb{E}\left[B\left(q_t^*\left(\tilde{h}^t\right)\right)-\tilde{h}_tq_t^*\left(\tilde{h}^t\right)-c\left(q_t^*\left(\tilde{h}^t\right)\right)\right]\\ &\to\mathbb{E}\left[B\left(q_t^E\left(\tilde{h}_t\right)\right)-\tilde{h}_tq_t^E\left(\tilde{h}_t\right)-c\left(q_t^E\left(\tilde{h}_t\right)\right)\right]\end{aligned} \tag{8}$$

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

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

Part 1 of the proposition states that the ex-ante expected "wedge" between the marginal benefit of quantity to the principal and marginal cost to the agent (to borrow a term from the dynamic taxation literature) vanishes as the relationship progresses. As explained in the Introduction, the result is established by considering perturbations to any putative optimal mechanism ψ^* in which the ex-ante expected wedge fails to vanish. In this case, the claimed optimal mechanism ψ^* can be perturbed by increasing or decreasing the quantity supplied at dates t long after the relationship has commenced but at which the ex-ante wedge remains away from zero. This adjustment in quantity is by a uniform constant; that is, it is uniform over reported types h^t . Appropriate adjustments to the agent's payments can be found to ensure that the mechanism remains in Ψ (i.e., that it remains incentive compatible and individually rational). Ensuring date-1 participation, i.e. the satisfaction

[4] Here, we abuse notation by letting $F_t\left(\cdot|h_1\right)$ denote the cdf of date-t types conditional on the date-1 type $h_1\in\Theta$.

of (1), may require leaving the agent with additional expected rents. However, these additional rents are naturally small if the date of adjustment t is far in the future, since Condition "Long-run independence" implies that the agent has little relevant private information about his date-t type at the time of contracting. Thus, the perturbations to quantity at date t can increase ex-ante expected surplus by an amount that dominates the increase in rents, at least when t is sufficiently large. This contradicts the optimality of ψ^*, implying that ex-ante expected wedges must vanish as claimed.

Importantly, note that the convergence of ex-ante expected wedges to zero, by itself, need *not* imply that the surplus from the relationship is expected to converge to the efficient one with time. In particular, a pattern of allocations such that quantity levels are upward distorted along some histories h^t but downward distorted along others is not inconsistent with Part 1 of Proposition 1. While downward distortions have commonly been observed in successful applications of the "relaxed approach" (as discussed in the Introduction), this does not seem a necessary or inherent feature of dynamic screening. The main reason it is difficult to address the question of whether distortions are always downwards lies in the difficulty of discerning which incentive constraints bind (see Section 6 of Battaglini and Lamba (2015) for an example of the challenges involved in establishing which incentive constraints bind in a simple two-period setting where there are $N = 3$ possible types). Part 2 of the proposition then confirms that if the direction of distortions is known to always be either upwards or downwards, then the relationship becomes efficient (in expectation) with time.

In cases where there is additional structure on the evolution of types, stronger predictions on wedges are possible. We now introduce this additional structure.

Condition 2. **[FOSD]** Date-t types are stochastically ordered by previous realizations: For any t, any $\breve{h}^{t-1}$, $\bar{h}^{t-1} \in \Theta^{t-1}$ satisfying $\breve{h}^{t-1} \geq \bar{h}^{t-1}$, and any $h_t \in \Theta$, $F_t\left(h_t|\breve{h}^{t-1}\right) \leq F_t\left(h_t|\bar{h}^{t-1}\right)$.

Condition "FOSD" requires that the history of agent types up to any date t stochastically orders the date-t realization. Another common restriction in the literature is the following:

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

Given our full support assumption (a transition to each of the N states has positive probability at all histories), Condition "Markov" implies that the process is aperiodic and irreducible and hence has a unique steady state (or "ergodic") distribution. Given our restriction to N types, we can associate the Markov transitions with an $N \times N$ transition matrix, denoted A, with element A_{ij} denoting 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 of types is ergodic:

Condition 4. **[Stationary Markov]** The process F is a first-order Markov process whose initial distribution F_1 is the ergodic distribution for the time-homogeneous Markov transitions described by $F_t(\cdot|\cdot)$.

The following result presents implications of these additional restrictions on the process.

Proposition 2. *1. Suppose that the condition "FOSD" holds. Then*

$$\mathbb{E}\left[B'\left(q_t^*\left(\tilde{h}^t\right)\right)-\tilde{h}_t-c'\left(q_t^*\left(\tilde{h}^t\right)\right)\right] \tag{9}$$

is non-negative for every date t (irrespective of whether T is finite or infinite). Hence, if $T=+\infty$ and, in addition, Condition "Long-run independence" holds, the convergence of (9) to zero is from above.

2. If, further, Condition "Stationary Markov" holds, then convergence is monotone. In particular, the expression (9) is non-increasing in t.

For processes satisfying the condition "FOSD", it is fairly straightforward to see that the participation constraint (1) binds uniquely at the worst initial type θ_N (an observation that has been made elsewhere in the literature for stochastic processes that are similarly ordered, although which, unlike our proof, usually follows after applying an envelope-type argument as in Pavan, Segal and Toikka (2014)). If the expected wedge (9) were strictly negative at some date t, we could reduce the quantity at date t by a uniform constant, while adjusting payments to ensure incentive compatibility and that the participation constraint (1) continues to bind at θ_N. This would increase ex-ante expected surplus. Given that this initial type expects the greatest cost reduction from the change in quantity schedule, the new mechanism leaves the initial type θ_N with the same expected rents (zero, i.e. equal to the outside option), but reduces the expected rents for all lower initial types. Hence, profits unambiguously increase.

For stationary Markov processes, it turns out that the expected wedge (9) decreases monotonically with time. Heuristically, this might be expected since initial types become progressively less informative about the distribution of later types. Increases in quantity at later dates thus have progressively less effect on the expected rents that must be left to the agent to ensure date-1 participation.

3.3 Convergence (in probability) to efficient policies

As discussed above, while we are able to establish that distortions (or more precisely, "wedges") vanish in expectation for a large class of processes (those satisfying the condition "Long-run independence"), we have so far drawn no conclusions regarding the long-run efficiency of the relationship, except when relying on additional information on the direction of distortions. We now provide sufficient conditions for long-run efficiency. We simplify our arguments by restricting attention to processes F satisfying the condition "Markov", although we expect that our arguments can be extended quite directly to (at least some) non-Markov finite-state processes provided these processes are not too persistent. Notice that while the condition "Long-run independence" is implied, we do not impose either Condition "FOSD" nor Condition "Stationary Markov". Rather, the important economic restriction will be that the players are sufficiently patient.

A sufficient condition on the degree of patience for our result can be stated in terms of primitives of the model. Recalling that the Markov transitions can be described by a matrix A whose ij^{th} element is A_{ij}, let $\alpha \equiv \min_{i,j\in\{1,\ldots,N\}} A_{ij}$. Since each A_{ij} is strictly positive, we have $\alpha > 0$. This value 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 below by α. In addition, let $b \equiv \sum_{i=1}^{N} \theta_i$, and let

$$\kappa \equiv \min_{i,j\in\{1,\ldots,N\}\text{ s.t. } i\neq j} \left\{ B\left(q^E(\theta_i)\right) - \theta_i q^E(\theta_i) - c\left(q^E(\theta_i)\right) - \left(B\left(q^E(\theta_j)\right) - \theta_i q^E(\theta_j) - c\left(q^E(\theta_j)\right)\right)\right\}. \tag{10}$$

Here, $\kappa > 0$ is the smallest loss in surplus due to some type θ_i supplying the efficient quantity for some other type θ_j, $j \neq i$; that is, due to supplying quantity $q^E(\theta_j)$ rather than the quantity that is efficient for θ_i, namely $q^E(\theta_i)$. Our result on long-run efficiency will then apply whenever $\delta > \bar{\delta}$, with $\bar{\delta}$ given by

$$\bar{\delta} \equiv \begin{cases} \frac{2\bar{q}b-\kappa}{2\bar{q}b-\kappa+2\kappa\alpha} \text{ if } \kappa < 2\bar{q}b \\ 0 \text{ otherwise} \end{cases}. \tag{11}$$

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

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

$$\begin{aligned} &\mathbb{E}\left[B\left(q_t^*\left(\tilde{h}^t\right)\right) - \tilde{h}_t q_t^*\left(\tilde{h}^t\right) - c\left(q_t^*\left(\tilde{h}^t\right)\right)\right] \\ &\to \mathbb{E}\left[B\left(q^E\left(\tilde{h}_t\right)\right) - \tilde{h}_t q^E\left(\tilde{h}_t\right) - c\left(q^E\left(\tilde{h}_t\right)\right)\right] \end{aligned} \tag{12}$$

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

We now summarize the key ideas in the proof of Proposition 3. This proof proceeds by considering any supposed optimal mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ for which the convergence in (12) fails, and constructing a "perturbed" version of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that achieves a strictly higher payoff for the principal. 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 Equation (6)). Intuitively, one might hope that such a mechanism would not only increase the expected surplus from the relationship, but also have a negligible effect on the information rents expected by the agent at the time of contracting (i.e., at date 1); after all, at this point in time the agent is poorly informed about his types far in the future.

While an efficient continuation policy always exists that is also incentive compatible for the agent, the difficulty with this naive approach is that it does not properly account for the effect on the agent's incentive constraints before the switch to efficiency. In particular, it is possible that the optimal mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ induces truthful reporting at early dates only because of distortions away from the efficient allocation in the distant future.

Our proof therefore overcomes this difficulty by considering a perturbed version of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ for which the allocation approaches the efficient one gradually. We show that such a perturbed mechanism can be chosen that continues to reside in Ψ. A first step is to note that checking incentive compatibility in our setting with discounted payoffs reduces to verifying the sub-optimality of one-shot deviations from truth-telling at all histories. This observation has often been made elsewhere, including in the emerging literature on dynamic mechanism design. More importantly, our second observation relates to a property of many efficient mechanisms, for instance the mechanism that prescribes an allocation $q^E(h_t)$ at date t to type $h_t \in \Theta$ and payment $p(h_t) = B\left(q^E(h_t)\right)$. For such a mechanism, the aforementioned incentive constraints are always satisfied as strict inequalities, with slack that is bounded away from zero by the amount κ identified above. One can then combine an efficient mechanism with the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ to yield an incentive-compatible mechanism. If the new output policy is a convex combination of the optimal and efficient policies, i.e. $q_t\left(h^t\right) = \gamma q^E(h_t) + (1-\gamma)\, q_t^*\left(h^t\right)$ for all t and all h^t, and for some $\gamma \in (0,1)$, then all incentive constraints can be guaranteed to hold strictly, with the agent's loss from deviating from truth-telling at least $\gamma\kappa$. In turn, this means it must be possible to choose also incentive-compatible mechanisms that place increasing weight on the efficient policy, i.e. where $q_t\left(h^t\right) = \gamma_t q^E(h_t) + (1-\gamma_t)\, q_t^*\left(h^t\right)$ for all t and all h^t, and for some increasing sequence of scalars $(\gamma_t)_{t=1}^{\infty}$.

The sequence $(\gamma_t)_{t=1}^{\infty}$ can be chosen so that the perturbed mechanism becomes fully efficient, i.e. $\gamma_t = 1$, in a period t when the allocation in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is distorted. However, for the perturbed mechanism to be guaranteed incentive compatible, it is necessary to choose $\gamma_1 > 0$; in particular, it is necessary to make adjustments to the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ also early in the relationship. Since, from the perspective of the time of contracting, the agent is relatively well informed about his type early in the relationship, this necessitates leaving the agent with additional rents to guarantee participation, and these rents need to be larger when γ_t is larger for small t. Whether the gains in surplus can be guaranteed to exceed these additional rents depends on both the primitives of the problem, the date and size of the distortions in the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, and the amount of weight assigned to the efficient mechanism early in the relationship.

When $\delta > \bar{\delta}$, with $\bar{\delta}$ given in Equation (11), the size of distortions needed for us to find a more profitable perturbed mechanism shrink with the date of these distortions, which is precisely what allows us to establish convergence to efficiency in Proposition 3. A direct implication of our proof is thus an explicit bound the size of distortions in each period.

Corollary 1. *Suppose that the process F satisfies Condition "Markov" and let $\delta >\in \left(\bar{\delta}, 1\right)$. (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,*

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

where $\delta + \frac{\kappa}{2\lambda} > 1$.

According to Corollary 1, when $T = +\infty$, a bound exists for (ex-ante expected) distortions that vanishes linearly with rate $\frac{1}{\delta+\frac{\kappa}{2\lambda}}$. That $\delta > \bar{\delta}$ is what ensures $\delta + \frac{\kappa}{2\lambda} > 1$.

The dependence of $\bar{\delta}$ on parameters can be explained as follows. 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). Less weight must be placed on the efficient mechanism at early periods (γ_t can be chosen smaller for small t), and this reduces the rents granted to the agent to ensure participation. 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 that $(\gamma_t)_{t=1}^{\infty}$ increases more quickly. 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 (12) to occur instead 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 $\left(h_{t-1}, h'_{t-1}, h_t\right) \in \Theta^3$, $\left|f_t\left(h_t|h_{t-1}\right) - f_t\left(h_{t-1}|h'_{t-1}\right)\right| < \varepsilon(\delta)$, then*

$$\begin{aligned} &\mathbb{E}\left[B\left(q_t^*\left(\tilde{h}^t\right)\right) - \tilde{h}_t q_t^*\left(\tilde{h}^t\right) - c\left(q_t^*\left(\tilde{h}^t\right)\right)\right] \\ &\quad \to \mathbb{E}\left[B\left(q^E\left(\tilde{h}_t\right)\right) - \tilde{h}_t q^E\left(\tilde{h}_t\right) - c\left(q^E\left(\tilde{h}_t\right)\right)\right] \end{aligned}$$

as $t \to \infty$.

The proof of Proposition 4 is close to that of Proposition 3 and involves many of the same steps. The key difference is that the lack of persistence of agent types (i.e., the condition that, for all $\left(h_{t-1}, h'_{t-1}, h_t\right) \in \Theta^3$, $\left|f_t\left(h_t|h_{t-1}\right) - f_t\left(h_t|h'_{t-1}\right)\right| < \varepsilon(\delta)$ where $\varepsilon(\delta)$ is small) is used to guarantee the existence of a perturbed version of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that is incentive compatible and prescribes the efficient allocation early in the relationship, while still prescribing close to the outcome of the original mechanism (i.e., $\left(q_1^*\left(h_1\right), p_1^*\left(h_1\right)\right)$) at the very first date. Given that the persistence of types is not strong, any additional rents expected by the agent in the perturbed mechanism are then small, yet the increase in discounted surplus is relatively large (and this holds for the fixed value of the discount factor $\delta \in (0,1)$ under consideration).

Finally, note that the result in Proposition 3 can be extended to more general cost functions $C(\cdot,\cdot) : (0,\bar{q}) \times \Theta \to \mathbb{R}_+$ (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 $\mu = \left(\mu_t\left(h^t\right)\right)_{t=1}^T$, and with payments given (as before) by $\mathbf{p} = \left(p_t\left(h^t\right)\right)_{t=1}^T$. 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 will be two important reasons to consider stochastic mechanisms. First, for reasons explained by Strausz (2006) for static environments, a restriction to deterministic mechanisms can in general imply a reduction in the principal's payoff (as demonstrated above, this is not the case when $C(q,h) = hq + c(q)$). Strausz 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 constraints succeeds, but also suggests that this conclusion need not hold more generally. While Strausz's result for the optimality of deterministic mechanisms should be expected to apply also to dynamic environments where the "relaxed approach" succeeds, the main objective of our work is precisely the treatment of environments where this approach does not apply.

The second reason relates to our approach to extending Proposition 3 to more general cost functions. In particular, we again will not be able to exhibit an optimal mechanism, but rely on perturbations to mechanisms that fail the property of interest; i.e., convergence of the expected surplus to the efficient one. The perturbed mechanisms we construct will involve randomizations between the policies of the mechanism of interest and the policies of another, efficient, mechanism. That is, we anticipate the perturbed mechanism will be stochastic.

Let us now state the incentive and participation constraints for any stochastic mechanism $\langle \mu, \mathbf{p} \rangle$. Assuming that F satisfies the condition "Markov", the agent's expected (continuation) payoff from reporting truthfully in any mechanism $\langle \mu, \mathbf{p} \rangle$, given past reports h^{t-1} and a date-t type h_t, is

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

The participation constraint is then

$$V_1^{\langle \mu, \mathbf{p} \rangle}(h_1) \geq 0 \quad \text{for all } h_1 \in \Theta, \tag{13}$$

while the incentive compatibility constraint is that, for all reporting strategies σ,

$$\mathbb{E}\left[V_1^{\langle \mu, \mathbf{p} \rangle}(\tilde{h}_1)\right] \geq \mathbb{E}\left[\sum_{t \geq 1} \delta^{t-1} \left(p_t^\sigma\left(\tilde{h}^t\right) - \int C\left(\tilde{q}, \tilde{h}_t\right) d\mu_t^\sigma\left(\tilde{h}^t\right)\right)\right], \tag{14}$$

where $(\mu_t^\sigma)_{t=1}^T$ and $(p_t^\sigma)_{t=1}^T$ are probability measures over quantities and payments induced by the reporting strategy σ. Let Ψ^S denote the set of (possibly stochastic) mechanisms satisfying these constraints. The principal's expected payoff can be written

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

and the design problem is to maximize $\Pi(\mu, \mathbf{p})$ over the set of mechanisms Ψ^S.

As before, we will provide a result establishing convergence to the efficient quantities when the players are sufficiently patient, with the degree of patience given by a value $\bar{\delta}^S$ in terms of 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 $q^E : \Theta \to [0, \bar{q}]$ such that $q^E(h_t)$ maximizes $B(q) - C(q, h_t)$ for each $h_t \in \Theta$; 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 $h, h' \in \Theta$,

$$\left|C(q, h) - C\left(q, h'\right)\right| \leq u.$$

Recall that $\alpha \equiv \min_{i,j\in\{1,\ldots,N\}} A_{ij}$, and let κ be given as before by Equation (10), which is strictly positive by the first requirement in Condition "Cost restriction". Our result on long-run efficiency will then apply whenever $\delta > \bar{\delta}$, with $\bar{\delta}$ given by

$$\bar{\delta} \equiv \begin{cases} \frac{2Nu-\kappa}{2Nu-\kappa+2\kappa\alpha} \text{ if } \kappa < 2Nu \\ 0 \text{ otherwise} \end{cases}. \tag{15}$$

Note that, for the statement of the following result, we lack an argument to guarantee existence of an optimal mechanism. The result is thus stated in a way that permits possible non-existence.

Proposition 5. *Suppose $T = +\infty$ and F satisfies Condition "Markov". Suppose, in addition, that the agent's cost function $C(\cdot, \cdot)$ is of the more general form specified above, that is satisfying Condition "Cost restriction". If a mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ that is optimal in Ψ^S exists, then, for all $\delta \in \left(\bar{\delta}^S, 1\right)$,*

$$\begin{aligned} &\mathbb{E}\left[\int \left(B(\tilde{q}) - C\left(\tilde{q}, \tilde{h}_t\right)\right) d\mu_t^*\left(\tilde{h}^t\right)\right] \\ &\to \mathbb{E}\left[B\left(q^E\left(\tilde{h}_t\right)\right) - C\left(q^E\left(\tilde{h}_t\right), \tilde{h}_t\right)\right] \end{aligned} \tag{16}$$

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

$$\left|\begin{array}{c} \mathbb{E}\left[\int \left(B(\tilde{q}) - C\left(\tilde{q}, \tilde{h}_t\right)\right) d\mu_t^k\left(\tilde{h}^t\right)\right] \\ -\mathbb{E}\left[B\left(q^E\left(\tilde{h}_t\right)\right) - C\left(q^E\left(\tilde{h}_t\right), \tilde{h}_t\right)\right] \end{array}\right| < \varepsilon. \tag{17}$$

The first part of Proposition 5 confirms that the convergence in (12) occurs also for the more general cost functions whenever an optimal mechanism exists. In cases where existence of an optimal mechanism cannot be guaranteed, the second part of our result provides a weaker sense in which distortions vanish that applies to mechanisms for which the principal earns a payoff close to the supremum. Since allocations in the very distant future have only a negligible effect on payoffs, our arguments do not yield predictions about such allocations for any fixed near-optimal mechanism.

4 Continuous types: The Myersonian Approach

We now turn to the case where the agent's types are drawn from a continuous distribution with compact support $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$. As for most of the analysis above, we will assume that agent costs take the separable form $C\left(q, h\right) = hq + c\left(q\right)$, where $q \in [0, \bar{q}]$ and $h \in \Theta$.

While the economic trade-offs involved in the design of optimal dynamic mechanisms will be closely related to what we have seen with discrete types, there are a few differences from a methodological perspective. On the one hand, as in papers featuring continuous types dating back to Baron and Besanko (1984), one can exploit a payoff equivalence property of incentive-compatible mechanisms. This permits, in particular, a better understanding of the relationship between the quantity allocation prescribed by an incentive-compatible mechanism and the agent's expected rents (and hence principal's expected profits). As we show below, this potentially permits direct calculation of the expected "wedge" between the marginal benefit of quantity to the principal and its marginal cost to the agent in any optimal mechanism.

The only differences in the environment relative to the model set-up of Section 2 relate to the stochastic process for types, and are as follows. As noted, we take $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$. We impose the following restrictions on the stochastic process $F \equiv (F_t)$. To simplify, we assume the process satisfies the conditions "Markov" and "FOSD". We then abuse notation by letting, for all $t \geq 2$, $F\left(\cdot|h_{t-1}\right) = F_t\left(\cdot|h_{t-1}\right)$ denote the distribution of the date-t type h_t conditional on the previous realization h_{t-1}. Each distribution function $F\left(\cdot|h_{t-1}\right)$ is taken to be absolutely continuous with density $f\left(\cdot|h_{t-1}\right)$, and to have full support on Θ (while the full support assumption is convenient and consistent with the analysis in the previous section, it is not essential to our arguments). We then denote the distribution of the date-1 type, h_1, by F_1, again absolutely continuous with a density $f_1\left(\theta_1\right) > 0$ for almost all $\theta_1 \in \Theta$.

To identify necessary and sufficient conditions for incentive compatibility, Pavan, Segal and Toikka (2014) use a state representation of the evolution of the agents' private information similar to the one in Eso and Szentes (2007) — see also Eso and Szentes (2015). In our environment, where F is a first-order Markov process, this can be described as follows. For any $s > 1$ and $t > s$, let $h_t = Z_{(s),t}(h_s, \varepsilon)$ denote the representation of h_t, where ε is a vector $(\varepsilon_\tau) \in \mathcal{E} \subset \mathbb{R}^\infty$ independent of the types $h^s = (h_1, \ldots, h_s)$ realized up to date s. We let $Z^t_{(s)}(h_s, \varepsilon) \equiv \left(Z_{(s),s}(h_s, \varepsilon), \ldots, Z_{(s),t}(h_s, \varepsilon)\right)$. Any first-order Markov stochastic process F can be described this way. For example, given the kernels

F, for any $t \geq 2$, one can let $F^{-1}(\varepsilon_t|h_{t-1}) \equiv \inf\{h_t : F(h_t|h_{t-1}) \geq \varepsilon_t\}$. If ε_t is drawn from a Uniform distribution over $(0,1)$, then the random variable $F_t^{-1}(\varepsilon_t|h_{t-1})$ is distributed according to the c.d.f. $F_t(\cdot|h_{t-1})$. We can therefore let the evolution of the process following the realization of θ_s at date s be defined recursively by $Z_{(s),t}(\theta_s, \varepsilon) = F_t^{-1}(\varepsilon_t|Z_{t-1}(\theta_s, \varepsilon))$. The above representation is referred to as the *canonical representation* of F in Pavan, Segal and Toikka (which generalizes to the case of non-Markov processes).

We now impose two further regularity conditions on F. To do so, let $\|\cdot\|$ denote the discounted L1 norm on $\mathbb{R}^\infty$ defined by $\|y\| \equiv \sum_{t=0}^{\infty} \delta^t |y_t|$.

Condition 6. [**Regularity**] The process is "regular" if the following conditions hold.

1. There exist functions $K_{(s)} : \mathcal{E} \to \mathbb{R}^\infty$, $s \geq 0$, 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 $\left|\partial Z_{(s),t}(h_s, \varepsilon)/\partial h_s\right| \leq K_{(s),t-s}(\varepsilon)$.[5]

2. For each s, $\log\left[\partial Z_{(s),t}(h_s, \varepsilon)/\partial h_s\right]$ 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, let

$$I_{(s),t}(h_s^t) = \frac{\partial Z_{(s),t}(h_s, \varepsilon)}{\partial h_s}$$

denote the impulse response of h_t to h_s, where ε is such that $Z_{(s)}^t(h_s, \varepsilon) = h_s^t$. The impulse response 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 all h_s.

Note that, when h_t follows the AR(1) process specified above, $I_t(h^t) = \gamma^{t-1}$. More generally, these functions are themselves stochastic processes. Also note that, when the process is Markov and the kernels $F(h_t|h_{t-1})$ are continuously differentiable in (h_t, h_{t-1}), 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 all t, all $h^{t-1} \in \Theta^{t-1}$, all $h_t, \hat{h}_t \in \Theta_t$, let

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

The following result then follows from Pavan, Segal and Toikka (2014).

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

Theorem 1 (Pavan, Segal and Toikka, 2014)**.** *Assume the stochastic process F, taking values in $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$, satisfies the conditions "Markov", "FOSD", and "Regularity". A mechanism $\psi = \langle \mathbf{q}, \mathbf{p} \rangle$ is incentive compatible if and only if the following conditions jointly hold for all $t \geq 1$, all h^{t-1}, all $h_t, \hat{h}_t \in \Theta$: (a) $V_t^{\langle \mathbf{q},\mathbf{p} \rangle}(h^t)$ is equi-Lipschitz continuous in h_t with*

$$\frac{\partial V_t^{\langle \mathbf{q},\mathbf{p} \rangle}(h^t)}{\partial h_t} = -\mathbb{E}\left[\sum_{s \geq t} \delta^{s-t} I_{(t),s}(\tilde{h}_t^s) q_s(\tilde{h}^s) \mid h^t\right] \text{ a.e.} \tag{18}$$

and (b)

$$\int_{\hat{h}_t}^{h_t} [D_t((h^{t-1}, x); x) - D_t((h^{t-1}, x); \hat{h}_t)] dx \geq 0. \tag{19}$$

Condition (18) is the dynamic analog of the usual envelope formula for static environments. In Markov environments such as the one under consideration here, the above necessary condition, when paired with the integral monotonicity condition (19) yields a complete characterization of incentive compatibility (for a general treatment in richer environments, see Pavan, Segal, and Toikka (2014)).[6]

The integral monotonicity condition (19) 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 (19) holds in the dynamic environment under examination here if the NPV of expected future output, *discounted by impulse responses*

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

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

We now show how it may be possible to explicitly calculate the expectation of the optimal "wedges" between the marginal benefit and marginal cost of output in an optimal mechanism, and then show that they vanish as the relationship progresses. Using that $V_1^{\langle \mathbf{q},\mathbf{p} \rangle}(h_1)$ must satisfy (18) for almost all $h_1 \in \Theta$, and integration by parts, we are able to write the principal's objective as *"Dynamic Virtual Surplus"*; that is,

$$\mathbb{E}\left[\sum_t \delta^{t-1}\left(B\left(q_t\left(\tilde{h}^t\right)\right) - q_t\left(\tilde{h}^t\right)\left(\tilde{h}_t + \frac{F_1\left(\tilde{h}_1\right)}{f_1\left(\tilde{h}_1\right)} I_t(\tilde{h}^t)\right) - c\left(q_t\left(\tilde{h}^t\right)\right)\right)\right]$$
$$- V_1^{\langle \mathbf{q},\mathbf{p} \rangle}(\bar{\theta}). \tag{20}$$

[6]The 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) dominates any other strategy.

Given condition "FOSD", (18) implies that $V_1^{\langle \mathbf{q},\mathbf{p}\rangle}(\cdot)$ obtains its minimum at $h_1 = \bar{\theta}$. Hence, for an optimal mechanism, where the participation constraint (1) binds, we have $V_1^{\langle \mathbf{q},\mathbf{p}\rangle}(\bar{\theta}) = 0$. Then, the principal's problem can be understood as maximizing (20) over allocations $\mathbf{q}$ such that, for some payment rule $\mathbf{p}$, the conditions of Theorem 1 are satisfied. In turn, this amounts to maximizing (20) by choice of $\mathbf{q}$ subject to the integral monotonicity conditions in (19).

It is worth mentioning that the approach to dynamic mechanism design problems of this nature, common in the literature, is to simply maximize pointwise under the expectation in (20), while ignoring the incentive constraints (analogously, ignoring the integral monotonicity constraints). Given the concavity of $B(\cdot)$ and convexity of $c(\cdot)$, this leads to the first-order condition

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

also being sufficient for a maximum, for all $h^t \in \Theta^t$. This quantity allocation coincides with the solution to the principal's (full) problem provided it satisfies all the integral monotonicity conditions in (19). A sufficient condition for this to be the case is that

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

is nonincreasing in $\hat{h}_t$ for all t and all h^t. With an optimal allocation $\mathbf{q}^*$ in hand, an appropriate payment rule can be specified. For instance, we can put, for all $h_1 \in \Theta$,

$$p_1^*(h_1) = h_1 q_1^*(h_1) + c\left(q_1^*(h_1)\right) + \int_{h_1}^{\bar{\theta}} \mathbb{E}\left[\sum_{s=1}^{\infty} \delta^{s-1} I_s(\tilde{h}^s) q_s^*(\tilde{h}^s) \mid \tilde{h}_1 = x\right] dx,$$

and, for all $t \geq 2$, and all h^t,

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

The mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is then incentive compatible and individually rational, with the participation constraint (1) binding for $h_1 = \bar{\theta}$. Moreover, for all $t \geq 2$, and all $h^{t-1} \in \Theta^{t-1}$, $\mathbb{E}\left[V_t^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(\tilde{h}^t\right) |h^{t-1}\right] = 0$.

In general, however, there is no guarantee that the allocation rule that solves the relaxed program (that is, that maximizes the virtual surplus pointwise over all possible allocation rules) satisfies all the "integral monotonicity" conditions in Part (b) of Theorem 1 (Battaglini and Lamba, 2015, explain, in the context of discrete-type models, that in fact such an integral monotonicity condition will typically fail; i.e., they argue the conditions on primitives to ensure integral monotonicity are rather special). Our approach, analogous to that for discrete types above, is therefore to look for perturbations to the putative optimal allocation rule $\mathbf{q}^*$ that preserve the integral monotonicity conditions in (19).

Such a perturbed rule should not increase the principal's expected payoff in (20), which can lead to necessary conditions that any optimal allocation $\mathbf{q}^*$ must satisfy.

In particular, suppose that $\mathbf{q}^*$ is an optimal allocation, and that $\mathbf{p}^*$ is a payment rule such that $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is incentive compatible and $V_1^{\langle \mathbf{q}, \mathbf{p} \rangle}(\bar{\theta}) = 0$; without loss of generality, we can take the payment rule specified above. One adjustment to this mechanism, if feasible, is to add a constant $\nu \in \mathbb{R}$ to the allocation $q_t^*(\cdot)$ at date t, adjusting the payments accordingly. The principal's expected payoff (20) then changes by

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

A derivative with respect to ν evaluated at $\nu = 0$ yields

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

A necessary condition for the optimality of $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is that this equal zero.

The simple perturbation described above is feasible provided the date-t allocation $q_t^*(\cdot)$ is bounded away from the boundaries of the possible quantities, i.e. 0 and $\bar{q}$, and in this case it renders an incentive-compatible and individually-rational mechanism. We can therefore conclude the following:

Proposition 6. *Assume the stochastic process F, taking values in $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$, satisfies the conditions "Markov", "FOSD", and "Regularity". If $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is an optimal mechanism whose allocation remains bounded away from the boundaries zero and $\bar{q}$, then*

$$\mathbb{E}\left[B'\left(q_t^*\left(\tilde{h}^t\right)\right)-\tilde{h}_t-c'\left(q_t^*\left(\tilde{h}^t\right)\right)\right]=\mathbb{E}\left[\frac{F_1\left(\tilde{h}_1\right)}{f_1\left(\tilde{h}_1\right)}I_t(\tilde{h}^t)\right]. \tag{21}$$

That optimal allocations satisfy the stipulated interiority condition appears difficult to guarantee in general. Naturally, given the Inada conditions ($\lim_{q\to\bar{q}} c(q) = +\infty$ and $\lim_{q\to 0} B(q) = -\infty$), such interiority can be guaranteed if one imposes continuity restrictions on the optimal mechanism, such as requiring allocations to be Lipschitz continuous with a fixed Lipschitz constant. Under such an assumption, the aforementioned perturbations to an optimal mechanism are then feasible, so that the result in Proposition 6 holds. We can then study the dynamics of the expected wedges, by considering the evolution of the right-hand side of Equation (21) when the stochastic process F satisfies the following condition.

Condition 7. [Ergodic] Let $P^t(h, A)$ be defined, for any period s, by $\Pr\left(\tilde{h}_{s+t} \in A | \tilde{h}_s = h\right)$. 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 $h \in \Theta$,

$$\sup_{A \in \mathcal{B}(\Theta)} \left| P^t(h, A) - \pi(A) \right| \to 0 \text{ as } t \to \infty. \tag{22}$$

Under the additional condition "Ergodic", we are able to provide an analogue of the convergence in (7) for discrete types.

Proposition 7. *Assume F satisfies conditions "Markov", "Regularity", "Ergodic" and "FOSD". Then*

$$\mathbb{E}\left[\frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] \to 0 \text{ as } t \to \infty.$$

Given that F satisfies "FOSD", convergence is from above ($\mathbb{E}\left[\frac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] \geq 0$ for all t). If, further, the distribution of the initial type F_1 is the stationary (ergodic) distribution of F, then convergence is monotone in time ($\mathbb{E}\left[\frac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right]$ is decreasing in t).

Proposition 7, combined with Proposition 6, is thus informative about the long-run behavior of the expected wedges $\mathbb{E}\left[B'\left(q_t^*\left(\tilde{h}^t\right)\right) - \tilde{h}_t - c'\left(q_t^*\left(\tilde{h}^t\right)\right)\right]$. As before, this does not allow us to conclude anything regarding the expected efficiency of the relationship at long horizons, without further information. We nonetheless expect that restrictions on the continuity of allocations, such as requiring the quantities to be Lipschitz in reports, with known Lipschitz constant, would allow us to rule out pathological behavior, permitting arguments analogous to the ones made for discrete types that guarantee (ex-ante expected) distortions vanish 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 "first-order 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 discussed in the new dynamic public finance literature (e.g., Farhi and Werning (2013) and Golosov et al. (2016)), 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 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., Troshkin, 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., 2016. 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.

Appendix: Omitted 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 $\mathbf{p}' \equiv (p'_t)_{t \geq 1}$ is such that, for $t = 1$,

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

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

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

all $h^t \in \Theta^t$. The allocation rule $\mathbf{q}' \equiv (q'_t)_{t \geq 1}$ is the same as in ψ, i.e., $\mathbf{q}' = \mathbf{q}$. The mechanism $\psi' = \langle \mathbf{q}', \mathbf{p}' \rangle$ satisfies (4) for all $t \geq 2$ and $h^t \in \Theta^t$. 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 (4), for all $t \geq 2$, all $h^t \in \Theta^t$, necessarily

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

for all $t \geq 2$, all h^t. To see this, suppose that, instead, $V_t^{\psi} \left(h^t; h^{t-1} \right) > \frac{\bar{q} \Delta \theta}{1 - \delta}$ 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 $V_t^{\psi} \left(\left(h^{t-1}, h'_t \right); h^{t-1} \right) \leq 0$. An agent whose history of private types is $\left(h^{t-1}, h'_t \right)$, 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 $\left(h^{t-1}, h_t \right)$. By misreporting in this way, the former agent (the one with history of private types $\left(h^{t-1},h'_t\right)$) earns an expected payoff (from t onwards) at least equal to

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

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 $V_t^{\psi}\left(h^t;h^{t-1}\right)>\frac{\bar{q}\Delta\theta}{1-\delta}$ thus implies that the former agent has a profitable deviation. Since $\left(h^{t-1},h'_t\right)$ has positive probability by the full-support assumption, there exists a misreporting strategy σ that induces a higher expected payoff for the agent that the truthful one, thus violating Condition (2).

A similar argument implies that $V_t^{\psi}\left(h^t;h^{t-1}\right)\geq-\frac{\bar{q}\Delta\theta}{1-\delta}$ 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 $V_t^{\psi}\left(\left(h^{t-1},h'_t\right);h^{t-1}\right)\geq 0$. An agent whose period-$t$ continuation payoff satisfies $V_t^{\psi}\left(h^t;h^{t-1}\right)<-\frac{\bar{q}\Delta\theta}{1-\delta}$ 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 $V_t^{\psi}\left(\left(h^{t-1},h'_t\right);h^{t-1}\right)-\frac{\bar{q}\Delta\theta}{1-\delta}$ which is larger than $V_t^{\psi}\left(h^t;h^{t-1}\right)$. The former agent would thus have a profitable deviation, contradicting that $\psi\in\Psi$.

That in any mechanism satisfying (4), all $t\geq 2$, all $h^t\in\Theta^t$, necessarily

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

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

$$V_t^{\psi}\left(h^t;h^{t-1}\right)=p_t\left(h^t\right)-C\left(q_t\left(h^t\right),h_t\right),$$

all $t\geq 2$, all h^{t-1}.

Finally, to see that, in any *optimal* mechanism satisfying (4), all $t\geq 2$, all $h^t\in\Theta^t$, Condition (5) holds also for $t=1$, note that, in any such mechanism, the date-1 participation constraint (1) 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 (4), all $t\geq 2$, all $h^t\in\Theta^t$, then implies that condition (5) 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 $\left(\bar{b}_t\right)$, 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 $q_t\left(h^t\right)\notin[\underline{b}_t,\bar{b}_t]$ for some t and h^t, there exists another mechanism $\psi'=\langle\mathbf{q},\mathbf{p}\rangle$ with $\psi'\in\Psi$ and $q'_t\left(h^t\right)\in[\underline{b}_t,\bar{b}_t]$ 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 (4) and (5) for all $t\geq 2$, and all h^{t-1}, with the latter also holding also 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

$$\sup_{\psi\in\bar{\Psi}}\left\{\mathbb{E}\left[\sum_{t=1}^{\infty}\delta^{t-1}\left(B\left(q_t\left(\tilde{h}^t\right)\right)-p_t\left(\tilde{h}^t\right)\right)\right]\right\}-\mathbb{E}\left[\sum_{t=1}^{\infty}\delta^{t-1}\left(B\left(q_t^k\left(\tilde{h}^t\right)\right)-p_t^k\left(\tilde{h}^t\right)\right)\right]<1/k. \tag{23}$$

We can then construct an optimal policy as follows. Let $\left(\left(q_1^k(\cdot),p_1^k(\cdot)\right)\right)$ denote the sequence of period-1 policies defined by the above sequence of mechanisms. From Bolzano-Weierstrass, there exists a subsequence of $\left(\left(q_1^k(\cdot),p_1^k(\cdot)\right)\right)$ 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 $\left(\left(q_2^{k_{l(m)}}(\cdot),p_2^{k_{l(m)}}(\cdot)\right)\right)$ converges to some $(q_2^*(\cdot),p_2^*(\cdot))$. Clearly, this also implies that, along the subsequence indexed $(k_{l(m)})$, $\left(\left(q_1^{k_{l(m)}}(\cdot),p_1^{k_{l(m)}}(\cdot)\right)\right)$ 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 (5) (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,

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

where the equality follows from the assumption that $\mathbb{E}\left[V_1^{\psi^*}\left(\widetilde{h}_1\right)\right]-\mathbb{E}\left[V_1^{\psi^{*,\sigma}}\left(\widetilde{h}_1\right)\right]=\eta=3\varepsilon$. 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 aforementioned bound on agent per-period payoffs, it is easy to see that ψ^* satisfies (4) for all $t\geq 2$ (a property which is inherited from the mechanisms ψ^k). It follows that, for all $t\geq 2$,

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

(where note that these expectations are ex-ante). Moreover,

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

uniformly over t and h^t. Hence,

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{\infty}\delta^{t-1}\left(B\left(q_t^*\left(\tilde{h}^t\right)\right)-p_t^*\left(\tilde{h}^t\right)\right)\right] &= \mathbb{E}\left[B\left(q_1^*\left(\tilde{h}_1\right)\right)-p_1^*\left(\tilde{h}_1\right)\right] \\
&\quad +\mathbb{E}\left[\sum_{t=2}^{\infty}\delta^{t-1}\left(\begin{array}{c} B\left(q_t^*\left(\tilde{h}^t\right)\right) \\ -\tilde{h}_t q_t^*\left(\tilde{h}^t\right)-c\left(q_t^*\left(\tilde{h}^t\right)\right)\end{array}\right)\right] \\
&= \mathbb{E}\left[B\left(q_1^*\left(\tilde{h}_1\right)\right)-p_1^*\left(\tilde{h}_1\right)\right] \\
&\quad +\lim_{T\to\infty}\sum_{t=2}^{T}\delta^{t-1}\mathbb{E}\left[\begin{array}{c} B\left(q_t^*\left(\tilde{h}^t\right)\right) \\ -\tilde{h}_t q_t^*\left(\tilde{h}^t\right)-c\left(q_t^*\left(\tilde{h}^t\right)\right)\end{array}\right]. \qquad (24)
\end{aligned}$$

The first equality uses that, by construction of ψ^*, $\mathbb{E}\left[B\left(q_1^*\left(\tilde{h}_1\right)\right)-p_1^*\left(\tilde{h}_1\right)\right]$ is finite. The second equality and the existence of the limit in the last line of (24) follows because per-period surplus is uniformly bounded above by $\overline{S}$. Hence, expected profits under ψ^* are well-defined and equal to some value $\pi^* \in \mathbb{R}\cup\{-\infty\}$.

Now, let $\pi^{\sup} = \sup_{\psi\in\bar{\Psi}}\left\{\mathbb{E}\left[\sum_t \delta^{t-1}\left(B\left(q_t\left(\tilde{h}^t\right)\right)-p_t\left(\tilde{h}^t\right)\right)\right]\right\} \in \mathbb{R}$. Suppose then that ψ^* is not optimal. This means that there exists a (finite) $\eta > 0$ such that $\pi^{\sup}-\pi^* \geq \eta$. Because the series in the last line of (24) is convergent, we have that, for any $\varepsilon > 0$, there exists $T^*(\varepsilon)$, such that, for any $T \geq T^*(\varepsilon)$,

$$\mathbb{E}\left[B\left(q_1^*\left(\tilde{h_1}\right)\right)-p_1^*\left(\tilde{h}_1\right)\right]+\sum_{t=2}^{T}\delta^{t-1}\mathbb{E}\left[B\left(q_t^*\left(\tilde{h}^t\right)\right)-\tilde{h}_t q_t^*\left(\tilde{h}^t\right)-c\left(q_t^*\left(\tilde{h}^t\right)\right)\right] < \pi^{\sup}-\eta+\varepsilon.$$

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

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

The inequality follows from the fact (a) that 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) that

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

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

$$\begin{aligned}\mathbb{E}\left[B\left(q_t^k\left(\tilde{h}^t\right)\right)-p_t^k\left(\tilde{h}^t\right)\right] &= \mathbb{E}\left[B\left(q_t^k\left(\tilde{h}^t\right)\right)-\tilde{h}_t q_t^k\left(\tilde{h}^t\right)-c\left(q_t^k\left(\tilde{h}^t\right)\right)\right] \\ &\leq \bar{S}.\end{aligned}$$

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

$$\sum_{t=1}^{T}\delta^{t-1}\mathbb{E}\left[B\left(q_t^k\left(\tilde{h}^t\right)\right)-p_t^k\left(\tilde{h}^t\right)\right]-\sum_{t=1}^{T}\delta^{t-1}\mathbb{E}\left[B\left(q_t^*\left(\tilde{h}^t\right)\right)-p_t^*\left(\tilde{h}^t\right)\right] \quad > \quad \eta/2. \tag{25}$$

Then fix $T \geq \bar{T}(\varepsilon)$. By construction of ψ^*, for any $\nu > 0$, we can find a $k > 1/\varepsilon$ such that $|q_t^*\left(h^t\right)-q_t^k\left(h^t\right)|, |p_t^*\left(h^t\right)-p_t^k\left(h^t\right)| \leq \nu$ for all $t = 1, \ldots, T$, 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 (25) 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 $\gamma \in (0,1)$. For each t, each h^t, let $\psi^\gamma = \langle \mathbf{q}^\gamma, \mathbf{p}^\gamma\rangle$

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

and

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

The mechanism ψ^γ so constructed yields the agent an expected payoff

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

conditional on initial type h_1 under a truthful reporting strategy, and hence is individually rational, since $\psi^A, \psi^B \in \Psi$. Similarly, the agent's payoff in the mechanism ψ^γ for any reporting strategy σ is a convex combination of the payoffs for the strategy σ in $\psi^A$and ψ^B, with weight γ on the former. Hence, ψ^γ is incentive compatible; i.e. $\psi^\gamma \in \Psi$. The result then follows using that, for all t,

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

with strict inequality whenever $q_t^A\left(h^t\right) \neq q_t^B\left(h^t\right)$ for some h^t.

Proof of Proposition 1. [Part 1]. Suppose the result is not true and let $\langle (q_t^*)_{t=1}^\infty, (p_t^*)_{t=1}^\infty \rangle$ denote an optimal mechanism. Then there is $\kappa > 0$ 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 $\psi' = \left\langle \left(q_t^{'}\right)_{t=1}^\infty, \left(p_t^{'}\right)_{t=1}^\infty \right\rangle$ whose allocation rule is given by $q_t'(\cdot) = q_t^*(\cdot)$ for $t \neq t_k$, and by $q_{t_k}'(\cdot) = q_{t_k}^*(\cdot) + \nu_k$ for $t = t_k$, for some constant $\nu_k \in \left(0, \bar{q} - \max_{h^{t_k} \in \Theta^{t_k}}\left\{q\left(h^{t_k}\right)\right\}\right)$. 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 $t \geq 2$ such that $t \neq t_k$, and by $p_{t_k}'\left(h^{t_k}\right) = p_{t_k}^*\left(h^{t_k}\right) + c\left(q_{t_k}'\left(h^{t_k}\right)\right) - c\left(q_{t_k}^*\left(h^{t_k}\right)\right)$ for all $h^{t_k} \in \Theta^{t_k}$. Note that the new mechanism ψ' is incentive compatible, which follows because, for any reporting strategy σ, and any h_1,[7]

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

Moreover, the adjustment in date-1 payoffs 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 $\nu_k \in \left(-\min_{h^{t_k}\in\Theta^{t_k}}\left\{q_{t_k}\left(h^{t_k}\right)\right\}, 0\right)$,

[7] The equality follows because the quantity adjustments used to obtain the 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 cost is the same under truthtelling and σ. Furthermore, any effect on the convex part of the cost is undone by the adjustment in the payment (i.e., the adjustment $c\left(q_{t_k}'\left(h^{t_k}\right)\right) - c\left(q_{t_k}^*\left(h^{t_k}\right)\right)$ for each h^{t_k}).

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 $k\to\infty$, by virtue of Condition "Long-run independence", thus establishing the result.

[Part 2]. Suppose that distortions are always downwards, but that the convergence in (8) does not occur. Note that, because total surplus is concave in output, for any history h^t,

$$\begin{aligned}&B\left(q_t^E(h_t)\right) - h_t q_t^E(h_t) - c\left(q_t^E(h_t)\right) - \left(B\left(q_t^*\left(h^t\right)\right) - h_t q_t^*\left(h^t\right) - c\left(q_t^*\left(h^t\right)\right)\right)\\ &\leq \left(B'\left(q_t^*\left(h^t\right)\right) - h_t - c'\left(q_t^*\left(h^t\right)\right)\right)\left(q_t^E(h_t) - q_t^*\left(h^t\right)\right).\end{aligned}$$

Hence, if the convergence in (8) fails, there exists $\varepsilon > 0$ and a sequence of dates (t_k) such that, for any t_k

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

This implies that

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

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 strict concavity of $B(\cdot)$ and convexity of $c(\cdot)$, that

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

remains strictly below

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

by a constant uniform along the sequence (t_k), contradicting (8). Q.E.D.

Proof of Proposition 2 [Part 1]. When Condition "FOSD" holds, it is easy to show that $V_1^{\psi^*}(\cdot)$ is strictly decreasing under any optimal mechanism $\psi^* = \langle\mathbf{q}^*, \mathbf{p}^*\rangle$. To see this, we introduce the following "canonical" representation of the evolution of agent types, in which the agent is viewed as receiving independently distributed "shocks" to his value ε_t in each period t. In particular, for any $t \geq 2$ and any $h^t \in \Theta^t$ with $h_t = \theta_j$, $j \in \{1, \ldots, N\}$, let $\tilde{\varepsilon}_t|h^t$ be distributed uniformly on

$\left[F_t\left(\theta_{j-1}|h^{t-1}\right), F_t\left(\theta_j|h^{t-1}\right)\right]$ if $j>1$ and on $\left[0, F_t\left(\theta_1|h^{t-1}\right)\right]$ otherwise. Thus, $\tilde{\varepsilon}_t|h^{t-1}$ is distributed uniformly on $[0,1]$. We can then view the evolution of the agent's types as determined by the initial draw of h_1 (according to F_1) and subsequently by "shocks" ε_t (drawn uniformly from $[0,1]$) via the equation

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

where $F_t^{-1}(\varepsilon_t|h^{t-1}) \equiv \inf\{h_t : F(h_t|h^{t-1}) \geq \varepsilon_t\}$.

Now, consider a period-1 type θ_i who reports θ_{i+1} in period 1, but goes on to report in such a way that the process for reports is indistinguishable from that for an agent whose true initial type is θ_{i+1}. In particular, in an arbitrary period $t \geq 2$ in which the agent has already reported $\hat{h}^{t-1}$ and he observes true type $h_t = \theta_j$, he draws ε_t from the uniform distribution on $\left[F_t\left(\theta_{j-1}|h^{t-1}\right), F_t\left(\theta_j|h^{t-1}\right)\right]$ if $j>1$ and on $\left[0, F_t\left(\theta_1|h^{t-1}\right)\right]$ otherwise (where h^{t-1} are the agent's true past types), and then makes a period-t report equal to

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

Then note that, given F satisfies "FOSD", for each $t \geq 1$, we must have $h_t \leq \hat{h}_t$ (with a strict inequality occurring with positive probability on at least one date; in particular, for $t=1$). Hence, given that output is strictly positive at all dates, the agent with initial type θ_i expects the same payoff $V_1^{\psi^*}(\theta_{i+1})$ as an agent with initial type θ_{i+1}, less the strictly positive cost savings that occur due to having a lower true cost than reported (which clearly occurs with probability one at the initial date, and may also occur at later periods). This implies $V_1^{\psi^*}(\cdot)$ is strictly decreasing as claimed.

Now, suppose, towards a contradiction, that (9) is strictly negative for some date t' under the optimal mechanism ψ^*. Consider decreasing $q_{t'}^*(\cdot)$ uniformly by an arbitrary constant $\nu \in \left(0, \min_{h^{t'} \in \Theta^{t'}}\left\{q_{t'}^*\left(h^{t'}\right)\right\}\right)$. Formally, consider a new mechanism whose allocation rule is $q_t'(\cdot) = q_t^*(\cdot)$ for $t \neq t'$, and for which $q_{t'}'(\cdot) = q_{t'}^*(\cdot) - \nu$. Then let $p_t'(\cdot) = p_t^*(\cdot)$ for all $t \neq t'$, and put $p_{t'}'\left(h^{t'}\right) = p_{t'}^*\left(h^{t'}\right) + c\left(q_{t'}'\left(h^{t'}\right)\right) - c\left(q_{t'}^*\left(h^{t'}\right)\right) - \nu\mathbb{E}\left[\tilde{h}_{t'}|\tilde{h}_1 = \theta_N\right]$ for all $h^{t'} \in \Theta^{t'}$. Note that the incentive constraints (2) are unaffected by these changes; i.e., the mechanism $\langle (q_t')_{t=1}^{\infty}, (p_t')_{t=1}^{\infty}\rangle$ remains incentive compatible. Furthermore, the expected payoff conditional on the agent's initial type being θ_N is the same as under the original mechanism. 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, under "FOSD", the participation of type θ_N implies the participation of any other type (see the arguments above).

Next observe that, under the new mechanism, for each initial type $h_1' \neq \theta_N$, the period-1 expected payoff is

$$\delta^{t'-1}\nu\left(\mathbb{E}\left[\tilde{h}_{t'}|\tilde{h}_1 = \theta_N\right] - \mathbb{E}\left[\tilde{h}_{t'}|\tilde{h}_1 = h_1'\right]\right)$$

less than in the original mechanism. Furthermore, because (9) 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^*.

The rest of the claim in Part 1 follows from the result in Proposition 1.

[Part 2]. Suppose, towards a contradiction, that the result is not true. Then then there are adjacent periods s and $s' = s + 1$ such that the expected wedge in (9) 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 $\mathbf{e}_N = (0, ..., 0, 1)$ 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). Give 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

$$\mathbb{E}\left[\tilde{h}_{s'}|\tilde{h}_1 = \theta_N\right] \leq \mathbb{E}\left[\tilde{h}_s|\tilde{h}_1 = \theta_N\right]. \tag{26}$$

Now, consider the mechanism ψ' defined as follows. Let

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

Then let the allocation $\mathbf{q}'$ be such that

$$q_s'(\cdot) = q_s^*(\cdot) - \nu$$
$$q_{s'}'(\cdot) = q_{s'}^*(\cdot) + \frac{\nu}{\delta}$$

whereas, for all $t \notin \{s, s'\}$, $q_t'(\cdot) = q_t^*(\cdot)$. Then, let the payment rule $\mathbf{p}'$ be such that $p_t'(\cdot) = p_t^*(\cdot)$ for $t \notin \{s, s'\}$, whereas, for $t \in \{s, s'\}$,

$$p_t'\left(h^t\right) = p_t^*\left(h^t\right) + c\left(q_t'\left(h^t\right)\right) - c\left(q_t^*\left(h^t\right)\right)$$

all h^t.

Because the stochastic process F satisfies "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 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 (26), 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 "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. Given the Markov assumption, the agent's payoff from any date t onwards depends on past reports $\hat{h}^{t-1}$, and his true date-t type h_t, but not on past true types h^{t-1}. As a result, for any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle$, we may abuse notation by denoting the agent's expected payoff from date t onwards, having reported $\hat{h}^{t-1}$ at past dates and experienced true types h^t, by $V_t^{\langle \mathbf{q}, \mathbf{p} \rangle}(\hat{h}^{t-1}, h_t) = V_t^{\langle \mathbf{q}, \mathbf{p} \rangle}(h^t; \hat{h}^{t-1})$.

The central idea of the proof is as follows. For any claimed optimal mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ that fails to satisfy the convergence property in the first part of Proposition 3, we can obtain an adjusted mechanism that is incentive compatible, individually rational and increases the principal's expected payoff. Without loss of generality, we can take $\langle \mathbf{q}^*, \mathbf{p}^* \rangle \in \bar{\Psi}$, where $\bar{\Psi}$ is the set defined in the proof of Lemma **2** (in particular, this means that $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ satisfies both Equations (4) and (5) for all $t \geq 2$, all h^{t-1}, with the latter also holding at $t = 1$). The adjusted mechanism is then obtained by combining $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ with an efficient mechanism, chosen again to belong to $\bar{\Psi}$ (for an appropriate choice of bounds on the quantities $\underline{b}_t$ and $\bar{b}_t$, for $t \geq 1$). To define this efficient mechanism, and elucidate its incentive properties (which will prove crucial in the argument below), we begin with the following lemma.

Lemma 3. *Let $\mathbf{q}^E$ denote the efficient allocation rule defined according to Equation (6). Let $\hat{\mathbf{p}}^E$ denote the payment scheme defined, for all t and 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 (4) 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 in Equation (1) 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 we may choose, for $t \geq 1$, the bounds on the quantities $\underline{b}_t$ and $\bar{b}_t$ close enough to zero and $\bar{q}$, respectively, we may then take $\bar{\Psi}$ to include the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$. Furthermore, for any $(h^{t-1}, h_t) \in \Theta^t$, and any $h_t' \in \Theta$, with $h_t' \neq h_t$,*

$$V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^t\right) \geq p_t^E\left(h^{t-1}, h_t'\right) - h_t q^E\left(h_t'\right) - c\left(q^E\left(h_t'\right)\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right] + \kappa. \quad (27)$$

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 has been 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 it belongs to $\bar{\Psi}$ (with the bounds on quantities, i.e. $\underline{b}_t$ and $\bar{b}_t$ for $t \geq 1$, specified appropriately, as noted in the lemma).

Thus consider the final claim of the lemma, i.e. the one made in Equation (6). Because F is Markov, it is readily verified that, for any t and any h^{t-1}, the difference in discounted expected payoffs between truthful reporting from t onwards and instead reporting $h_t' \neq h_t$ in period t and reporting truthfully thereafter (as is optimal, conditional on the lie at date t) is equal to

$$\begin{aligned} & V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^t\right) - \left(p_t^E\left(h^{t-1}, h_t'\right) - h_t q^E\left(h_t'\right) - c\left(q^E\left(h_t'\right)\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right]\right) \quad (28) \\ & = B\left(q^E\left(h_t\right)\right) - h_t q^E\left(h_t\right) - c\left(q^E\left(h_t\right)\right) - \left(B\left(q^E\left(h_t'\right)\right) - h_t q^E\left(h_t'\right) - c\left(q^E\left(h_t'\right)\right)\right). \end{aligned}$$

The equality in Equation (28) follows for date $t = 1$ by noting that an agent with initial type h_1 who plans to report truthfully in the mechanism $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$ expects a payoff

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

On the other hand, by construction of $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$, an agent whose initial type is h_1, who misreports at date 1 sending $\hat{h}_1$, but who plans to report truthfully thereafter, expects a payoff

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

For dates $t \geq 2$, it follows because an agent who truthfully reports h_t after making reports h^{t-1} expects payoff

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

whereas an agent who instead lies by reporting h'_t at date t and reports truthfully thereafter expects a payoff

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

The result in the lemma then follows from the fact that, by the definition of κ, for any $h_t, h'_t \in \Theta$ with $h_t \neq h'_t$,

$$B\left(q^E\left(h_t\right)\right)-h_t q^E\left(h_t\right)-c\left(q^E\left(h_t\right)\right)-\left(B\left(q^E\left(h'_t\right)\right)-h_t q^E\left(h'_t\right)-c\left(q^E\left(h'_t\right)\right)\right)\geq\kappa.$$

Q.E.D.

The result in the previous lemma means that, under the mechanism $\left\langle \mathbf{q}^E, \mathbf{p}^E\right\rangle$, incentive constraints do not bind 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.

We now specify the form of our adjusted mechanism. We let $\boldsymbol{\gamma}$ be any collection of constants $(\gamma_t)_{t=1}^{\infty}$, with $\gamma_t \in [0,1]$ for all t. We then define $\langle \mathbf{q}^{\boldsymbol{\gamma}}, \mathbf{p}^{\boldsymbol{\gamma}}\rangle = \left\langle \left(q_t^{\boldsymbol{\gamma}}\right)_{t=1}^{\infty}, \left(p_t^{\boldsymbol{\gamma}}\right)_{t=1}^{\infty}\right\rangle$ to be the mechanism constructed from the efficient and assumed optimal mechanisms as follows. For each t, each h^t, let

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

and

$$p_t^{\boldsymbol{\gamma}}(h^t) = \gamma_t p_t^E(h^t) + (1-\gamma_t)p_t^*(h^t) + c(q_t^{\boldsymbol{\gamma}}(h^t)) - \gamma_t c\left(q_t^E(h^t)\right) - (1-\gamma_t)c\left(q_t^*(h^t)\right).$$

This construction ensures that, for all $t \geq 1$, all h^t,

$$\begin{aligned}V_t^{\langle \mathbf{q}^{\boldsymbol{\gamma}},\mathbf{p}^{\boldsymbol{\gamma}}\rangle}\left(h^t\right) =&\gamma_t V_t^{\langle \mathbf{q}^E,\mathbf{p}^E\rangle}\left(h^t\right) + (1-\gamma_t)V_t^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(h^t\right)\\ &+\sum_{s=t+1}^{\infty}\delta^{s-t}\left(\gamma_s-\gamma_{s-1}\right)\mathbb{E}\left[\mathbb{E}\left[V_s^{\langle \mathbf{q}^E,\mathbf{p}^E\rangle}\left(\tilde{h}^s\right)-V_s^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(\tilde{h}^s\right)|\tilde{h}^{s-1}\right]|h^t\right]\\ =&\gamma_t V_t^{\langle \mathbf{q}^E,\mathbf{p}^E\rangle}\left(h^t\right) + (1-\gamma_t)V_t^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(h^t\right)\;.\end{aligned} \tag{29}$$

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 $\left\langle \mathbf{q}^{E}, \mathbf{p}^{E} \right\rangle$ satisfy the condition in Equation (4).

Our task will now be to establish that, for an appropriate choice of the sequence $\boldsymbol{\gamma}$, the mechanism $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ is incentive compatible and individually rational (it will then follow from Equation (29) that $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ belongs to $\bar{\Psi}$). Individual rationality follows because the payoffs expected in the mechanism $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ at date 1, for any initial type h_1, assuming truthful reporting, is a linear combination of the payoffs in $\left\langle \mathbf{q}^{E}, \mathbf{p}^{E} \right\rangle$ and $\langle \mathbf{q}^{*}, \mathbf{p}^{*} \rangle$ (with linear weight equal to γ_1; see Equation (29)). However, the mechanism $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ is not guaranteed to be incentive compatible; i.e., incentive compatibility need not be inherited from $\left\langle \mathbf{q}^{E}, \mathbf{p}^{E} \right\rangle$ and $\langle \mathbf{q}^{*}, \mathbf{p}^{*} \rangle$. (The mechanism $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ fails to be incentive compatible, in particular, if the weights $(\gamma_t)_{t=1}^{\infty}$ grow too fast. For instance, if, for some $t > 1$, $\gamma_s = 0$ for all $s \in \{1, \ldots, t\}$, while $\gamma_s = 1$ for all $s \in \{t+1, \ldots, +\infty\}$, then incentive compatibility may fail at date t. This occurs when truthtelling is optimal for the agent in $\langle \mathbf{q}^{*}, \mathbf{p}^{*} \rangle$ at date t *only* due to the dependence of $\left(q_s^{*}\left(\hat{h}^{s}\right), p_s^{*}\left(\hat{h}^{s}\right)\right)$ at (some) dates $s > t$ on the agent's date-t report $\hat{h}_t$.)

To establish incentive compatibility of $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ for certain choices of $\boldsymbol{\gamma}$, we will rely on two preliminary observations. The first makes use of the assumption that every element of the transition matrix A is positive in order to provide a sense in which, after enough time, an agent's expected type becomes independent of initial conditions.

Lemma 4. *For all j 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 $s \geq 0$,

$$\max_{k,l} \left| \mathbb{E}\left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_k\right] - \mathbb{E}\left[\tilde{h}_{t+s} | \tilde{h}_t = \theta_l\right] \right| \leq b\,(1 - 2\alpha)^{s}. \tag{30}$$

Proof of Lemma 4. The first result is 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

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

Q.E.D.

The second is a bound on the agent's continuation payoff for any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$ satisfying the participation constraint in Condition (1) as an equality for at least one type $h_1 \in \Theta$. Note that this bound is an alternative to the one in Condition (5) (which, recall, is a defining condition of the set $\bar{\Psi}$).

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

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

Proof of Lemma 5. The proof follows arguments similar to those establishing Lemma 1. Suppose the inequality fails to hold at some h^t, and assume $V_t^{\langle \mathbf{q},\mathbf{p}\rangle}\left(h^t\right) > \frac{\bar{q}b}{1-\delta(1-2\alpha)}$. The assumption that the participation constraint in Equation (1) 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 (4) 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 $V_t^{\langle \mathbf{q},\mathbf{p}\rangle}\left(h^{t-1},\theta_j\right) \leq 0$. Suppose such a type uses the "canonical" representation of the evolution of the process F 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

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

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

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

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

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

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

$$\gamma_t = \min\left\{\gamma_1\left(1+\frac{\kappa}{2\delta\lambda}\right)^{t-1}, 1\right\}, \tag{31}$$

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. Note that Equation (29) can be applied to determine several properties of the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma\rangle$ that, coupled with incentive compatibility, will guarantee membership of $\bar{\Psi}$. First, $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma\rangle$ satisfies Equation (4) for all $t \geq 2$ and all h^{t-1}, and it satisfies Condition (5), for all $t \geq 1$, and for all h^t (properties that are inherited from $\langle \mathbf{q}^*, \mathbf{p}^*\rangle$ and $\left\langle \mathbf{q}^E, \mathbf{p}^E\right\rangle$). Second, because $\langle \mathbf{q}^*, \mathbf{p}^*\rangle$ and $\left\langle \mathbf{q}^E, \mathbf{p}^E\right\rangle$ belong to $\bar{\Psi}$, there exist sequences of scalars (as specified in the definition of $\bar{\Psi}$) given by $(\underline{b}_t)$ and $\left(\bar{b}_t\right)$, with $0 < \underline{b}_t < \bar{b}_t < \bar{q}$, and such that $q_t^*\left(h^t\right), q_t^E\left(h^t\right) \in [\underline{b}_t, \bar{b}_t]$ for all t and all h^t. Equation (29) then implies that this property holds also for $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma\rangle$ (i.e., $q_t^\gamma\left(h^t\right) \in [\underline{b}_t, \bar{b}_t]$ for all t and all h^t). Third, that the agent earns a non-negative expected payoff in $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma\rangle$ for all realizations of the initial type h_1, i.e. the mechanism is individually rational, is similarly inherited from $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ and $\langle \mathbf{q}^E, \mathbf{p}^E \rangle$.

It therefore 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 (5) for all t implies that 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.

Thus, the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ is incentive compatible if and only if, for all t, all $h^t \in \Theta^t$, all $h'_t \in \Theta$,

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

To see that $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ satisfies Condition (32) note that, for all t, all $h^t \in \Theta^t$, all $h'_t \in \Theta$,

$$\begin{aligned} V_t^{\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle}\left(h^t\right) &= \gamma_t V_t^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^t\right) + \left(1-\gamma_t\right) V_t^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}\left(h^t\right) \\ &\geq \gamma_t \left\{p^E\left(h'_t\right) - h_t q^E\left(h'_t\right) - c\left(q^E\left(h'_t\right)\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h'_t, \tilde{h}_{t+1}\right) | h_t\right]\right\} \\ &+ \left(1-\gamma_t\right)\left\{p_t^*\left(h^{t-1}, h'_t\right) - h_t q_t^*\left(h^{t-1}, h'_t\right) - c\left(q_t^*\left(h^{t-1}, h'_t\right)\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h'_t, \tilde{h}_{t+1}\right) | h_t\right]\right\} + \gamma_t \kappa. \end{aligned} \tag{33}$$

The equality is simply Equation (29), whereas the inequality follows from (27), together with the fact that the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ is incentive compatible. Next note that, by virtue of (31),

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

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

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

Combining the last inequality with (33), we then have that

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

$$+ (1-\gamma_t)\left\{ p_t^*\left(h^{t-1}, h_t'\right) - h_t q_t^*\left(h^{t-1}, h_t'\right) - c\left(q_t^*\left(h^{t-1}, h_t'\right)\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t \right] \right\}$$

$$+\delta\left(\gamma_{t+1} - \gamma_t\right)\left\{ \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t \right] - \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q}^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t \right] \right\}. \tag{34}$$

Note that the right-hand-side of (34) is equal to the right-hand side of (32). To see this it suffices to note that, for $\left(h^{t-1}, h_t', h_{t+1}\right) \in \Theta^{t+1}$,

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

This holds by Equation (29) (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 (34) 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.

The mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, for $\boldsymbol{\gamma}$ satisfying Equation (31), places progressively more weight on the outcomes of the efficient mechanism defined in 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 full efficiency to be obtained. We now show that, when $\delta > \bar{\delta}$, with $\bar{\delta}$ defined by Equation (11), if $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$ fails to satisfy the convergence properties in the proposition, then there exists a mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, obtained by taking γ_1 sufficiently close to zero, which generates a higher expected payoff for the principal.

Lemma 7. *If $\delta > \bar{\delta}$, with $\bar{\delta}$ defined by Equation (11), then any optimal mechanism satisfies 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 that, under $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, convergence of expected surplus fails. Then consider the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$ defined in Lemma 6. Since convergence fails, there exists $\varepsilon > 0$ and a strictly increasing sequence of dates (t_k) such that

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

for all k. For any $k \in \mathbb{N}$, then let

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

Note that, when $\gamma_1 = \gamma_1(k)$, $\gamma_{t_k} = 1$. Then, the increase in ex-ante expected surplus by switching from the putative optimal mechanism to the mechanism $\langle \mathbf{q}^\gamma, \mathbf{p}^\gamma \rangle$, with $\boldsymbol{\gamma}$ determined using Equation (31) after setting $\gamma_1 = \gamma_1(k)$, is more than $\delta^{t_k - 1}\varepsilon$.

Next, use Equation (29) to observe that the increase in the agent's ex-ante expected rent is equal to

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

The associated change in the principal's expected payoff is thus equal to

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{+\infty}\delta^{t-1}\left(B\left(q_t^{\gamma}\left(\tilde{h}^t\right)\right)-p_t^{\gamma}\left(\tilde{h}^t\right)\right)\right]-\mathbb{E}\left[\sum_{t=1}^{+\infty}\delta^{t-1}\left(B\left(q_t^{*}\left(\tilde{h}^t\right)\right)-p_t^{*}\left(\tilde{h}^t\right)\right)\right]\\
&\qquad\geq \delta^{t_k-1}\varepsilon-\left(1+\frac{\kappa}{2\delta\lambda}\right)^{1-t_k}\left(\mathbb{E}\left[V_1^{\langle \mathbf{q}^E,\mathbf{p}^E\rangle}\left(\tilde{h}_1\right)\right]-\mathbb{E}\left[V_1^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(\tilde{h}_1\right)\right]\right)\\
&\qquad= \delta^{t_k-1}\left[\varepsilon-\left(\delta+\frac{\kappa}{2\lambda}\right)^{1-t_k}\left(\mathbb{E}\left[V_1^{\langle \mathbf{q}^E,\mathbf{p}^E\rangle}\left(\tilde{h}_1\right)\right]-\mathbb{E}\left[V_1^{\langle \mathbf{q}^*,\mathbf{p}^*\rangle}\left(\tilde{h}_1\right)\right]\right)\right].
\end{aligned}\tag{35}$$

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

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

Hence, by taking k large enough, the right-hand side of the inequality (35) 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 $q_t^*\left(h^t\right)$ must converge, in probability, to $q^E\left(h_t\right)$. 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 is immediate from the inequality (35), after noting that

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

by Lemma 5 and by the triangle inequality. The application to a finite horizon $T<\infty$ follows after a straightforward adaptation of the arguments in Proposition 3. Q.E.D.

Proof of Proposition 4. The proof follows an argument similar to that for Proposition 3, and relies on many of the same steps. We consider, for a contradiction, a mechanism $\langle \mathbf{q}^*,\mathbf{p}^*\rangle\in\bar{\Psi}$

for which the convergence in the proposition fails to occur. One can then define $\left\langle \mathbf{q}^E, \mathbf{p}^E \right\rangle \in \bar{\Psi}$ as in Lemma 3 (again choosing the bounds on allocations for mechanisms in $\bar{\Psi}$ to include efficient provision), and then define $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ by letting $\boldsymbol{\gamma}$ be any collection of constants $(\gamma_t)_{t=1}^{\infty}$, with $\gamma_t \in [0,1]$ for all t, and letting, for each t, each h^t,

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

and

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

As before, that $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle \in \bar{\Psi}$ follows using Equation (29) and the fact that $\langle \mathbf{q}^*, \mathbf{p}^* \rangle, \left\langle \mathbf{q}^E, \mathbf{p}^E \right\rangle \in \bar{\Psi}$, provided that $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ can also be shown to be incentive compatible.

Our departure from the proof of Proposition 3 will be that, rather than relying on the observation in Lemma 5, ensuring incentive compatibility of $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ (for an appropriate choice of the sequence $\boldsymbol{\gamma}$) will rest on the following observation. In any mechanism $\langle \mathbf{q}, \mathbf{p} \rangle \in \bar{\Psi}$, for any $t \geq 1$, $h^{t-1} \in \Theta^{t-1}$, $h_t, h_t' \in \Theta$,

$$\begin{aligned}
&\left| \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q},\mathbf{p} \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t \right] \right| \\
&= \left| \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q},\mathbf{p} \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t \right] - \mathbb{E}\left[V_{t+1}^{\langle \mathbf{q},\mathbf{p} \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t' \right] \right| \\
&\leq \frac{\bar{q}\Delta\theta}{1-\delta} N \varepsilon(\delta).
\end{aligned} \tag{36}$$

Here, the equality uses the fact that $\langle \mathbf{q}, \mathbf{p} \rangle$ satisfies the conditions in Equation (4). 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 assumption 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 $(\gamma_t)_{t=1}^{\infty} \in [0,1]^{\infty}$ be any sequence of scalars satisfying $\gamma_1 \in (0,1)$ and, for all $t \geq 2$,

$$\gamma_t = \min\left\{ \gamma_1 \left(\frac{(1-\delta)\kappa + 2\delta\bar{q}\Delta\theta N\varepsilon(\delta)}{2\delta\bar{q}\Delta\theta N\varepsilon(\delta)} \right)^{t-1}, 1 \right\}. \tag{37}$$

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 the conditions in Equation (32), all t, $h^t \in \Theta^t$, and $h_t' \in \Theta$.

Now observe that (37) 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 $\left\langle \mathbf{q}^E, \mathbf{p}^E \right\rangle, \langle \mathbf{q}^*, \mathbf{p}^* \rangle \in \bar{\Psi}$, Equation (36) above (together with the triangle inequality) implies that, for all $t \geq 1$, all h^{t-1}, h_t and h_t',

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

The above two inequalities, when combined with the fact that, for the mechanism $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$, the conditions in Equation (33) are satisfied for all t, all h^t, and all h_t', then implies that the inequalities in Equation (34) are also satisfied. As before, these inequalities coincide with those in Equation (32); i.e., $\langle \mathbf{q}^{\gamma}, \mathbf{p}^{\gamma} \rangle$ is incentive compatible.

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

$$\delta \left(\frac{(1-\delta)\,\kappa + 2\delta \bar{q} \Delta\theta N \varepsilon(\delta)}{2\delta \bar{q} \Delta\theta N \varepsilon(\delta)} \right) > 1, \tag{38}$$

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 $(\gamma_t)_{t=1}^{\infty}$ satisfying (37), that yields the principal a higher expected payoff than in $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$. The result in the proposition then follows from the fact that the inequality in (38) is satisfied for $\varepsilon(\delta)$ sufficiently small. Q.E.D.

Proof of Proposition 5. The proof again mimics the arguments in the proof of Proposition 3, with the key difference now that the perturbed policy will involve randomization over the quantities in the assumed optimal mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ and an efficient mechanism $\left\langle \mu^E, \mathbf{p}^E \right\rangle$, with the agent's expected production cost thus linear in the probabilities assigned to each (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 $h^{t-1} \in \Theta^{t-1}$,

$$\mathbb{E}\left[V_t^{\langle \mu, \mathbf{p} \rangle} \left(h^{t-1}, \tilde{h}_t \right) | h_{t-1} \right] = 0, \tag{39}$$

and (b) the participation constraint (13) is satisfied with equality for at least one type h_1. Arguments analogous to those establishing Lemma 5 then imply that such mechanisms satisfy

$$\left| V_t^{\langle \mu, \mathbf{p} \rangle} \left(h^t \right) \right| \leq \frac{Nu}{1 - \delta\,(1-2\alpha)} \equiv \bar{\lambda}, \tag{40}$$

for all t and all $h^t \in \Theta^t$, with u given in Property 2 of Condition "Cost restriction".

Next, for each $t \geq 1$ and each h^t, let $\mu_t^E\left(h^t\right)$ put 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 $\left\langle \mu^E, \mathbf{p}^E \right\rangle$ satisfying the above conditions and such that, for all $t \geq 1$, all h^t, and all $h_t' \neq h_t$,

$$V_t^{\left\langle \mu^E, \mathbf{p}^E \right\rangle} \left(h^t \right) \geq p_t^E \left(h^{t-1}, h_t' \right) - C\left(q^E\left(h_t'\right), h_t \right) + \delta \mathbb{E}\left[V_{t+1}^{\left\langle \mathbf{q}^E, \mathbf{p}^E \right\rangle} \left(h^{t-1}, h_t', \tilde{h}_{t+1} \right) | h_t \right] + \kappa, \tag{41}$$

where $\kappa > 0$ is the constant identified above.

We now show how one can construct a mechanism $\langle \mu^{\gamma}, \mathbf{p}^{\gamma} \rangle \in \Psi^S$ whose policies are obtained by combining the policies in $\langle \mu^*, \mathbf{p}^* \rangle$ and $\left\langle \mathbf{q}^E, \mathbf{p}^E \right\rangle$. The arguments here parallel those establishing

Lemma 6. Let $(\gamma_s) \in [0,1]^\infty$ be any non-decreasing sequence of positive scalars satisfying $\gamma_1 \in (0,1)$ and

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

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

$$\mu_t^\gamma\left(h^t\right) = \gamma_t \mu_t^E\left(h^t\right) + (1-\gamma_t)\,\mu_t^*\left(h^t\right) \tag{42}$$

and

$$p_t^\gamma\left(h^t\right) = \gamma_t p_t^E\left(h^t\right) + (1-\gamma_t)\,p_t^*\left(h^t\right). \tag{43}$$

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 for optimal quantities, as given by $\mu_t^*\left(h^t\right)$).

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

$$\begin{aligned}
& V_t^{\langle \mu^\gamma, \mathbf{p}^\gamma \rangle}\left(h^t\right) \\
=& \gamma_t V_t^{\langle \mu^E, \mathbf{p}^E \rangle}\left(h^t\right) + (1-\gamma_t)\,V_t^{\langle \mu^*, \mathbf{p}^* \rangle}\left(h^t\right) \\
\geq& \gamma_t \left(p_t^E\left(h^{t-1}, h_t'\right) - \int C\left(\tilde{q}, h_t\right) d\mu_t^E\left(h^{t-1}, h_t'\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right]\right) \\
& + (1-\gamma_t)\left(p_t^*\left(h^{t-1}, h_t'\right) - \int C\left(\tilde{q}, h_t\right) d\mu_t^*\left(h^{t-1}, h_t'\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right]\right) \\
& + \gamma_t \kappa \\
\geq& \gamma_t \left(p_t^E\left(h^{t-1}, h_t'\right) - \int C\left(\tilde{q}, h_t\right) d\mu_t^E\left(h^{t-1}, h_t'\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right]\right) \\
& + (1-\gamma_t)\left(p_t^*\left(h^{t-1}, h_t'\right) - \int C\left(\tilde{q}, h_t\right) d\mu_t^*\left(h^{t-1}, h_t'\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right]\right) \\
& + \delta\left(\gamma_{t+1} - \gamma_t\right)\left(\begin{array}{c} \mathbb{E}\left[V_{t+1}^{\langle \mu^E, \mathbf{p}^E \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right] \\ -\mathbb{E}\left[V_{t+1}^{\langle \mu^*, \mathbf{p}^* \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right] \end{array}\right) \\
=& p_t^\gamma\left(h^{t-1}, h_t'\right) - \int C\left(\tilde{q}, h_t\right) d\mu_t^\gamma\left(h^{t-1}, h_t'\right) + \delta \mathbb{E}\left[V_{t+1}^{\langle \mu^\gamma, \mathbf{p}^\gamma \rangle}\left(h^{t-1}, h_t', \tilde{h}_{t+1}\right) | h_t\right].
\end{aligned} \tag{44}$$

The first equality follows by the fact that both $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ satisfy (39), implying that, at any truthful history, the agent's continuation payoff under $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ satisfies the same equality as in (29). 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 (41). The second inequality follows by the choice of the sequence (γ_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 (40) (the arguments are the same as in the proof of Lemma 6). The final equality follows from the same equality as in (29). Using again Condition

"Markov", as in the proof of Proposition 3, we thus establish the incentive compatibility of $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$. Individual rationality of $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ is inherited from that of $\langle \mu^E, \mathbf{p}^E \rangle$ and $\langle \mu^*, \mathbf{p}^* \rangle$ because the equality in (29) holds, and hence $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle \in \Psi^S$.

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 exhibit the convergence in expected surplus specified in the proposition. This means that there exists $\varepsilon > 0$ and a strictly increasing sequence (t_k) such that, for all k,

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

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

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

Note that the choice of γ_1 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 $\delta^{t_k - 1}\varepsilon$. Furthermore, because the equality in (29) holds, the increase in the agent's ex-ante expected rents is equal to

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

Hence, the change in the principal's expected payoff associated with the switch from $\langle \mu^*, \mathbf{p}^* \rangle$ to $\langle \mu^\gamma, \mathbf{p}^\gamma \rangle$ is equal to

$$\begin{aligned}
&\mathbb{E}\left[\sum_t \delta^{t-1}\left(\int B\left(\tilde{q}\right) d\mu^\gamma_t\left(\tilde{h}^t\right) - p^\gamma_t\left(\tilde{h}^t\right)\right)\right] - \mathbb{E}\left[\sum_t \delta^{t-1}\left(\int B\left(\tilde{q}\right) d\mu^*_t\left(\tilde{h}^t\right) - p^*_t\left(\tilde{h}^t\right)\right)\right] \\
&\geq \delta^{t_k-1}\varepsilon - \left(1 + \frac{\kappa}{2\delta\bar{\lambda}}\right)^{1-t_k}\left(\mathbb{E}\left[V_1^{\langle \mu^E, \mathbf{p}^E \rangle}\left(\tilde{h}_1\right)\right] - \mathbb{E}\left[V_1^{\langle \mu^*, \mathbf{p}^* \rangle}\left(\tilde{h}_1\right)\right]\right) \\
&= \delta^{t_k-1}\left(\varepsilon - \left(\delta + \frac{\kappa}{2\bar{\lambda}}\right)^{1-t_k}\left(\mathbb{E}\left[V_1^{\langle \mathbf{q}^E, \mathbf{p}^E \rangle}\left(\tilde{h}_1\right)\right] - \mathbb{E}\left[V_1^{\langle \mu^*, \mathbf{p}^* \rangle}\left(\tilde{h}_1\right)\right]\right)\right).
\end{aligned} \tag{45}$$

Provided $\delta + \frac{\kappa}{2\bar{\lambda}} > 1$, the expression above is strictly positive for t_k sufficiently large. This is guaranteed for $\delta > \bar{\delta}$, with $\bar{\delta}$ given in (15), thus establishing the claim in the first part of the proposition.

Finally, consider the general case where an optimal mechanism may not exist. Let $\left(\langle \mu^k, \mathbf{p}^k \rangle\right)$ now be a sequence of mechanisms in Ψ^S, with $\lim_{k\to\infty} \Pi\left(\mu^k, \mathbf{p}^k\right) = \sup_{\langle \mu, \mathbf{p} \rangle \in \Psi^S} \Pi\left(\mu, \mathbf{p}\right)$, and satisfying the same properties that we could assume to hold for the mechanism $\langle \mu^*, \mathbf{p}^* \rangle$ (which, recall, was assumed optimal in the proof above). Suppose, contrary to the proposition, that there is $\varepsilon > 0$ for which there is no value $\bar{t} \in \mathbb{N}$ and corresponding sequence (s_k), with $s_k \to \infty$, for which (17) is satisfied for all $\bar{t} \leq t \leq \bar{t} + s_k$. Then, for any $\bar{t}$, there is $\bar{s}\left(\bar{t}\right)$ such that, for any $k \in \mathbb{N}$, there exists $t\left(\bar{t}; k\right) \in \left\{\bar{t}, \bar{t}+1, \ldots, \bar{t} + \bar{s}\left(\bar{t}\right)\right\}$ with

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

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

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

and

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

with $\gamma_t = \min\left\{\gamma_1\left(1+\frac{\kappa}{2\delta\bar{\lambda}}\right)^{t-1}, 1\right\}$ and $\gamma_1 = \gamma_1(\bar{t};k) \equiv \left(1+\frac{\kappa}{2\delta\bar{\lambda}}\right)^{1-t(\bar{t};k)}$. Analogous to the argument above, the mechanism $\langle \mu^{k,\gamma}, \mathbf{p}^{k,\gamma} \rangle$ increases the principal's expected profits by at least

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

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 Theorem 1. The result follows from combining Theorems 1 and 3 of Pavan, Segal, and Toikka (2014). Q.E.D.

Proof of Proposition 6. Follows from arguments in the main text. Q.E.D.

Proof of Proposition 7. Observe that $\mathbb{E}[I_t(\theta^t) \mid \theta_1] = \frac{d}{d\theta_1}\mathbb{E}[\theta_t \mid \theta_1]$. Thus,

$$\begin{aligned}
\mathbb{E}\left[\tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] &= \mathbb{E}\left[\tfrac{1-F_1(\theta_1)}{f_1(\theta_1)}\mathbb{E}[I_t(\theta^t) \mid \theta_1]\right] = \int_{\underline{\theta}}^{\bar{\theta}} \left[1-F_1(\theta_1)\right]\mathbb{E}[I_t(\theta^t) \mid \theta_1] d\theta_1 \\
&= \left[1-F_1(\theta_1)\right]\mathbb{E}[\theta_t \mid \theta_1]|_{\theta_1=\underline{\theta}}^{\theta_1=\bar{\theta}} + \int_{\underline{\theta}}^{\bar{\theta}} f_1(\theta_1)\mathbb{E}[\theta_t \mid \theta_1] d\theta_1 \\
&= \mathbb{E}[\theta_t] - \mathbb{E}[\theta_t \mid \underline{\theta}] \to 0 \text{ as } t \to \infty
\end{aligned}$$

by ergodicity. Since F satisfies "FOSD",

$$\mathbb{E}[\theta_t] - \mathbb{E}[\theta_t \mid \underline{\theta}] \geq 0.$$

If, in addition, F is stationary, then

$$\mathbb{E}\left[\tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_t(\theta^t)\right] - \mathbb{E}\left[\tfrac{1-F_1(\theta_1)}{f_1(\theta_1)} I_s(\theta^s)\right] = \mathbb{E}[\theta_s \mid \underline{\theta}] - \mathbb{E}[\theta_t \mid \underline{\theta}] \leq 0$$

for $t > s$. Q.E.D.