

Infinite-Horizon Mechanism Design:

The Independent-Shock Approach

Alessandro Pavan

Northwestern University

Ilya Segal

Stanford University

Juuso Toikka

Stanford University

July 16, 2010

Abstract

These notes examine the problem of how to extend envelope theorems to infinite-horizon dynamic mechanism design settings, with an application to the design of "bandit auctions."

1 The environment

To facilitate the exposition, we follow the same notation as in Pavan, Segal, and Toikka (2009) (hereafter PST. While the environment considered here features a single agent, the results easily extend to settings with multiple agents under the additional assumptions of independence of types across agents and quasilinearity in payoffs discussed in Section 4 in PST (see also the application to the design of profit-maximizing-auctions for experience goods in the next section).

Set-up. Time is discrete and indexed by $t = 1, 2, \dots, T$ with $T \in \mathbb{N} \cup \{+\infty\}$. In each period t there is a contractible *decision* $y_t \in Y_t$, whose outcome is observed by the agent. (In the application considered in the next section we apply the model to a more general setup with multiple agents where y_t is the part of the decision taken in period t that is observed by the agent.) Each Y_t is assumed to be a measurable space with the sigma-algebra left implicit. The set of all possible histories of feasible decisions is denoted by $Y \subset \prod_{\tau=1}^T Y_\tau$, with y denoting a generic element of Y .¹ Likewise, the set of all feasible period- t histories of decisions is denoted by $Y^t \subset \prod_{\tau=1}^t Y_\tau$.

Before the period- t decision is taken, the agent receives some private information $\theta_t \in \Theta_t \subset \mathbb{R}$. We implicitly endow the set Θ_t with the Borel sigma-algebra. We refer to θ_t as the agent's *current type*. The set of all possible type histories at period t is then denoted by $\Theta^t \equiv \prod_{\tau=1}^t \Theta_\tau$. An element θ of $\Theta \equiv \prod_{\tau=1}^T \Theta_\tau$ is referred to as the agent's *type*. For any sequence θ and a fixed δ , we then let

$$\|\theta\|_\delta \equiv \sup_t \delta^t |\theta_t|.$$

where $\delta \in \mathbb{R}_{++}$ in case $T < \infty$ while $\delta \in (0, 1)$ for $T = \infty$.

The distribution of the current type θ_t may depend on the entire history of types and decisions $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$. In particular, we assume that the distribution of θ_t is governed by a history-dependent probability measure (“kernel”) $F_t(\cdot | \theta^{t-1}, y^{t-1})$ on Θ_t such that $F_t(A | \cdot) : \Theta^{t-1} \times Y^{t-1} \rightarrow \mathbb{R}$ is measurable for all measurable sets $A \subset \Theta_t$.² We denote the collection of kernels by

$$F \equiv \langle F_t : \Theta^{t-1} \times Y^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T,$$

where for any measurable set A , $\Delta(A)$ denotes the set of probability measures on A . We abuse notation by using $F_t(\cdot | \theta^{t-1}, y^{t-1})$ to also denote the cumulative distribution function (c.d.f.) corresponding to the measure $F_t(\theta^{t-1}, y^{t-1})$. Throughout, we assume that for all t , $\Theta_t = (\underline{\theta}_t, \bar{\theta}_t) \subset \mathbb{R}$ for some $-\infty \leq \underline{\theta}_t \leq \bar{\theta}_t \leq +\infty$.

Finally, let $\mathcal{B}(\Theta) \equiv \{\theta \in \Theta : \|\theta\|_\delta < +\infty\}$ the set of types whose norm is finite. We will assume

¹By convention, all products of measurable spaces encountered in the text are endowed with the product sigma-algebra.

²Throughout, we adopt the convention that for any set A , $A^0 \equiv \{\emptyset\}$.

that, given the stochastic process that corresponds to the kernels F , for any mechanism and any strategy of the agent in the mechanism, $\theta \in \mathcal{B}(\Theta)$ with probability one. With an abuse of notation, we then simply let $\mathcal{B}(\Theta) = \Theta$.

The agent is a von Neumann-Morgenstern decision maker whose preferences over lotteries over $\Theta \times Y$ are represented by the expectation of a (measurable) Bernoulli utility function

$$U : \Theta \times Y \rightarrow \mathbb{R}.$$

Although some form of time separability of U is typically assumed in applications, this is not needed for our results. What is essential is only that the agent's preferences be time consistent, which is captured by the assumption that the agent is an expected-utility maximizer, with a Bernoulli utility function that is constant over time.

An often encountered special case in applications is one where private information evolves in a Markovian fashion, and where the agent's payoff is Markovian in the following sense.

Definition 1 *The environment is Markov if*

M1 for all t , and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, $F_t(\cdot | \theta^{t-1}, y^{t-1})$ does not depend on θ^{t-2} , and

M2 there exists functions $\langle A_t : \Theta^t \times Y^t \rightarrow \mathbb{R}_{++} \rangle_{t=1}^{T-1}$ and $\langle B_t : \Theta_t \times Y^t \rightarrow \mathbb{R} \rangle_{t=1}^T$ such that for all $(\theta, y) \in \Theta \times Y$,

$$U(\theta, y) = \sum_{t=1}^T \left(\prod_{\tau=1}^{t-1} A_\tau(\theta_\tau, y^\tau) \right) B_t(\theta_t, y^t).$$

Condition (1) guarantees that the stochastic process governing the evolution of the agent's type is Markov, while Condition (2) ensures that in any given period t , after observing history (θ^t, y^{t-1}) , the agent's von Neumann-Morgenstern preferences over lotteries over future types and decisions depend on the type history θ^t only through the current type θ_t . In particular, it encompasses the case of *additive-separable* preferences ($A_t(\theta_t, y^t) = 1$ for all t) as well as the case of *multiplicative-separable* preferences ($B_t(\theta_t, y^t) = 0$ for all $t < T$).

Mechanisms. A mechanism in the above environment assigns a set of possible messages to the agent in each period. The agent sends a message from this set and the mechanism responds with a (possibly randomized) decision that may depend on the entire history of messages sent up to period t , and on past decisions. By the Revelation Principle (Myerson, 86), for any standard solution concept, any distribution on $\Theta \times Y$ that can be induced as an equilibrium outcome in any "indirect mechanism" can also be induced as an equilibrium outcome of a "direct mechanism" in which the agent is asked to report the current type in each period, and where, in equilibrium, he

finds it optimal to report truthfully.³

Let $m_t \in \Theta_t$ denote the agent's period- t message, and let $m^t \equiv (m_1, \dots, m_t)$.

A *direct mechanism* is a collection

$$\Omega \equiv \langle \Omega_t : \Theta^t \times Y^{t-1} \rightarrow \Delta(Y_t) \rangle_{t=1}^T$$

such that for all t , and all measurable $A \subset Y_t$, $\Omega_t(A|\cdot) : \Theta^t \times Y^{t-1} \rightarrow [0, 1]$ is measurable. The notation $\Omega_t(A|m^t, y^{t-1})$ stands for the probability that the mechanism generates $y_t \in A \subset Y_t$ given history $(m^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$.

Given a direct mechanism Ω , and a history $(\theta^{t-1}, m^{t-1}, y^{t-1}) \in \Theta^{t-1} \times \Theta^{t-1} \times Y^{t-1}$, the following sequence of events takes place in each period t :

1. The agent privately observes his current type $\theta_t \in \Theta_t$ drawn according to $F_t(\cdot|\theta^{t-1}, y^{t-1})$.
2. The agent sends a message $m_t \in \Theta_t$.
3. The mechanism selects a decision $y_t \in Y_t$ according to $\Omega_t(\cdot|m^t, y^{t-1})$.

A (*pure*) *strategy* for the agent in a direct mechanism is a collection of measurable functions

$$\sigma \equiv \langle \sigma_t : \Theta^t \times \Theta^{t-1} \times Y^{t-1} \rightarrow \Theta_t \rangle_{t=1}^T.$$

A strategy σ is *truthful* if for all t and all $((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) \in \Theta^t \times \Theta^{t-1} \times Y^{t-1}$,

$$\sigma_t((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) = \theta_t.$$

This definition identifies a unique strategy; such a strategy has the property that the agent reports his current type truthfully after any history, including non-truthful ones. Note that we are not claiming here that it is without loss of generality to restrict attention to mechanisms with the property that the truthful strategy (as defined above) is optimal at all histories. As explained above, what the Revelation Principle guarantees is only that it is without loss of generality to restrict attention to mechanisms where the agent finds it optimal to report truthfully conditional on having reported truthfully in the past; this is equivalent to requiring that the truthful strategy be optimal at all *truthful* histories.

In order to describe expected payoffs, it is convenient to develop some more notation. First we define histories. For all $t = 0, 1, \dots$, let

$$H_t \equiv (\Theta^t \times \Theta^{t-1} \times Y^{t-1}) \cup (\Theta^t \times \Theta^t \times Y^{t-1}) \cup (\Theta^t \times \Theta^t \times Y^t),$$

³With a single agent, “in equilibrium” means conditional on having reported truthfully in the past.

where by convention $H_0 = \{\emptyset\}$, and $H_1 = \Theta_1 \cup (\Theta_1 \times \Theta_1) \cup (\Theta_1 \times \Theta_1 \times Y_1)$. Then H_t is the set of all histories terminating within period t , and, accordingly, any $h \in H_t$ is referred to as a *period- t history*. We let

$$H \equiv \bigcup_{t=0}^T H_t$$

denote the set of all histories. A history $(\theta^s, m^t, y^u) \in H$ is a *successor* to history $(\hat{\theta}^j, \hat{m}^k, \hat{y}^l) \in H$ if (1) $(s, t, u) \geq (j, k, l)$, and (2) $(\theta^j, m^k, y^l) = (\hat{\theta}^j, \hat{m}^k, \hat{y}^l)$. A history $h = (\theta^s, m^t, y^u) \in H$ is a *truthful history* if $\theta^t = m^t$.

Fix a direct mechanism Ω , a strategy σ , and a history $h \in H$. Let $\mu[\Omega, \sigma] | h$ denote the (unique) probability measure on $\Theta \times \Theta \times Y$ —the product space of types, messages, and decisions—induced by assuming that following history h in mechanism Ω , the agent follows strategy σ in the future. More precisely, let $h = (\theta^s, m^t, y^u)$. Then $\mu[\Omega, \sigma] | h$ assigns probability one to $(\tilde{\theta}, \tilde{m}, \tilde{y})$ such that $(\tilde{\theta}^s, \tilde{m}^t, \tilde{y}^u) = (\theta^s, m^t, y^u)$. Its behavior on $\Theta \times \Theta \times Y$ is otherwise induced by the stochastic process that starts in period s with history h , and whose transitions are determined by the strategy σ , the mechanism Ω , and the kernels F . If h is the null history we then simply write $\mu[\Omega, \sigma]$. We also adopt the convention of omitting σ from the arguments of μ when σ is the truthful strategy. Thus $\mu[\Omega]$ is the ex-ante measure induced by truth-telling while $\mu[\Omega] | h$ is the measure induced by the truthful strategy following history h .

Given this notation, we write the agent’s expected payoff when following history h he plays according to strategy σ in the future as $\mathbb{E}^{\mu[\Omega, \sigma] | h} [U(\tilde{\theta}, \tilde{y})]$.⁴ Now, given a direct mechanism Ω , let the agent’s *value function* $V^\Omega : H \rightarrow \mathbb{R}$ be a mapping such that for all histories $h \in H$,

$$V^\Omega(h) = \sup_{\sigma} \mathbb{E}^{\mu[\Omega, \sigma] | h} [U(\tilde{\theta}, \tilde{y})].$$

Incentive compatibility at a generic history h is then defined as follows.

Definition 2 *Let $h \in H$. A direct mechanism Ω is incentive compatible at history h (IC at h) if*

$$\mathbb{E}^{\mu[\Omega] | h} [U(\tilde{\theta}, \tilde{y})] = V^\Omega(h).$$

In words, Ω is IC at h if truthful reporting in the future maximizes the agent’s expected payoff following history h . This definition is flexible in that it allows us to generate different notions of IC as special cases by requiring IC at all histories in a particular subset. For example, ex-ante IC is equivalent to requiring IC only at the null history. Or in a static model (i.e., if $T = 1$), the standard definition of interim incentive compatibility obtains by requiring Ω to be IC at all histories where

⁴Throughout we use “tildes” to denote random variables with the same symbol without the tilde corresponding to a particular realization.

the agent knows only his type. In a dynamic model a natural alternative is to require that if the agent has been truthful in the past, he finds it optimal to continue to report truthfully. This is obtained by requiring Ω to be IC at all truthful histories.

The Principle of Optimality then implies that if Ω is IC at h , then for $\mu[\Omega]|h$ -almost all successors h' to h , Ω is IC at h' . In particular, if Ω is ex-ante IC, then truthtelling is also sequentially optimal at truthful future histories h with probability one, and the agent's equilibrium payoff at such histories is given by $V^\Omega(h)$ with probability one. We will sometimes find it convenient to focus on such histories, and they are the only ones that matter for ex-ante expectations.

2 Independent-Shock Representations

We now propose a way of characterizing the agent's payoff in an incentive-compatible mechanism based on the idea that the information the agent receives over time can be conveniently described as a function of "shocks" that are serially independent. This approach complements the one illustrated in the Pavan, Segal and Toikka (2009) in two ways: first, it permits us to accommodate the case where $T = +\infty$; second, even when restricted to the case $T < +\infty$, it permits us to identify a different set of assumptions on the primitive environment that guarantee that the agent's payoff in any incentive-compatible mechanism satisfies a certain envelope condition.

We start by defining what we mean when we say that a process admits an independent-shock representation. Next, we define in what sense this representation is "strategically equivalent" to the original one and hence can be used to characterize incentive-compatible mechanisms. We then proceed by showing how the formula for the (derivative of the) agent's payoff function simplifies when the agent is asked to report the shocks instead of his types and identify conditions on the agent's reduced-form payoff (i.e., his payoff expressed as a function of the shocks) that validate this formula. Finally, we conclude by showing that *any* stochastic process admits a particular independent-shock representation which we use to identify conditions for the primitive environment that guarantee that in the corresponding independent-shock representation the agent's reduced-form payoff is "well-behaved" in the sense that it satisfies an envelope formula analogous to the one derived in static settings. While these conditions differ from the ones identified in PST using a backward-induction approach, the formula for the derivative of the agent's payoff reduces to the one in PST when expressed in terms of the primitive representation.

Definition 3 Fix $T \in \mathbb{N} \cup \{+\infty\}$ and let $\tilde{\varepsilon} \equiv (\tilde{\varepsilon}_t)_{t=1}^T$ denote a collection of random variables with support $\mathcal{E} \equiv \times_{t=1}^T \mathcal{E}_t \subset \mathbb{R}^T$ and distribution $G(\cdot; y)$ and $z \equiv \langle z_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Theta_t \rangle_{t=1}^T$ denote a collection of measurable functions of these variables and of the decisions y . We say that (G, z) , where $G \equiv \langle G(\cdot; y) : y \in Y \rangle$, is an independent-shock (IS) representation for the stochastic process that

corresponds to the kernels $F \equiv \langle F_t : \Theta^{t-1} \times Y^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T$ if

(i) for each t , each $y^{t-1} \in Y^{t-1}$, there exists a probability measure $G_t(\cdot; y^{t-1})$ on \mathcal{E}_t such that, for any y $G(\cdot; y) = \times_{t=1}^T G_t(\cdot; y^{t-1})$; and

(ii) for any t , $\varepsilon^{t-1} \in \mathcal{E}^{t-1}$ and $y^{t-1} \in Y^{t-1}$, the distribution of $z_t(\tilde{\varepsilon}^t; y^{t-1})$ given y^{t-1} and $\tilde{\varepsilon}^{t-1} = \varepsilon^{t-1}$ is the same as the distribution of θ_t given y^{t-1} and $\theta^{t-1} = z^{t-1}(\varepsilon^{t-1}; y^{t-2}) \equiv (z_\tau(\varepsilon^\tau; y^{\tau-1}))_{\tau=1}^{t-1}$.

Together, conditions (i) and (ii) say that, for any y , one can think of the agent's primitive information θ as being generated by the independent "shocks" $\tilde{\varepsilon}$.

Example 1 Assume that θ_t evolves according to an $AR(k)$ process whose kernels are independent of past decisions:

$$\theta_t = \sum_{j=1}^k \phi_j \theta_{t-j} + \varepsilon_t,$$

where $\theta_t = 0$ for any $t \leq 0$, $\phi_j \in \mathbb{R}$ for any $j = 1, \dots, k$, and ε_t is the realization of the random variable $\tilde{\varepsilon}_t$ distributed according to some c.d.f. G_t with strictly positive density over \mathbb{R} , independent from all $\tilde{\varepsilon}_s$, $s \neq t$. In this example, the functions z_t do not depend on y and are given by

$$\begin{aligned} z_1(\varepsilon_1) &= \varepsilon_1 \\ z_2(\varepsilon^2) &= \phi_1 \varepsilon_1 + \varepsilon_2 \\ z_3(\varepsilon^3) &= \phi_1(\phi_1 \varepsilon_1 + \varepsilon_2) + \phi_2 \varepsilon_1 + \varepsilon_3 = (\phi_1^2 + \phi_2) \varepsilon_1 + \phi_1 \varepsilon_2 + \varepsilon_3 \\ &\dots \\ z_t(\varepsilon^t) &= \sum_{j=1}^t \left[\sum_{M \in \mathbb{N}, l \in \mathbb{N}^{M+1}; j=l_0 < \dots < l_M=t} \prod_{m=1}^M \phi_{l_m - l_{m-1}} \right] \varepsilon_j. \end{aligned}$$

Suppose now that the agent's information θ is generated by the independent shocks ε and let $z : \mathcal{E} \times Y \rightarrow \Theta$ denote the function defined by

$$z(\varepsilon; y) \equiv (z_\tau(\varepsilon^\tau; y^{\tau-1}))_{\tau=1}^T.$$

Assume further that the agent observes not only θ but also the shocks ε . The agent's payoff can then be expressed in terms of the shocks ε and the decisions y by the function $\hat{U} : \mathcal{E} \times Y \rightarrow \mathbb{R}$ defined by

$$\hat{U}(\varepsilon, y) \equiv U(z(\varepsilon; y), y). \quad (1)$$

Next, consider a (randomized direct) mechanism

$$\hat{\Omega} \equiv \langle \hat{\Omega}_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Delta(Y_t) \rangle_{t=1}^T,$$

in which the agent reports the shocks ε instead of his primitive payoff-relevant information θ . For any t any $y^{t-1} \in Y^{t-1}$, then let $\hat{G}^t(\cdot|z^t(\tilde{\varepsilon}^t; y^{t-1}))$ denote any regular conditional probability distribution for the random vector $\tilde{\varepsilon}^t$ given the sigma-algebra $\Sigma(z^t(\tilde{\varepsilon}^t; y^{t-1}))$ generated by the random vector $z^t(\tilde{\varepsilon}^t; y^{t-1})$.⁵

The primitive representation (U, F) is equivalent to the representation (\hat{U}, G, Z) in the following sense.

Lemma 1 (a) *Given any ex-ante IC mechanism Ω for the primitive representation (U, F) , there exists an ex-ante IC mechanism $\hat{\Omega}$ for the corresponding independent-shock representation (\hat{U}, G, z) such that, for any t , any measurable set $A \subseteq Y_t$, and any (θ^t, y^{t-1}) ,*

$$\int \hat{\Omega}_t(A|\varepsilon^t, y^{t-1}) d\hat{G}^t(\varepsilon^t|z^t(\varepsilon^t; y^{t-1}) = \theta^t) = \Omega_t(A|\theta^t, y^{t-1}). \quad (2)$$

(b) *Given any ex-ante IC mechanism $\hat{\Omega}$ for the independent-shock representation (\hat{U}, G, z) , there exists an ex-ante IC mechanism Ω for the primitive representation (U, F) such that, for any t , any measurable set $A \subseteq Y_t$, and any (θ^t, y^{t-1}) , (2) holds.*

Hence any outcome (i.e., any joint distribution over $\Theta \times Y$) that can be sustained by having the agent report the payoff-relevant information θ can also be sustained by having him report the shocks ε , and vice versa. Note that Part (a) follows directly from the fact that if the mechanism Ω is ex-ante IC, then the mechanism $\hat{\Omega}$ defined by

$$\hat{\Omega}_t(\cdot|\varepsilon^t, y^{t-1}) = \Omega_t(\cdot|z^t(\varepsilon^t; y^{t-1}), y^{t-1}) \quad \forall (\varepsilon^t, y^{t-1}) \quad (3)$$

is also ex-ante IC. This mechanism de facto uses the same information as Ω , in the sense that it depends on ε only through $z(\varepsilon; y)$. Part (b) is also trivially satisfied. It suffices to construct Ω from $\hat{\Omega}$ using the transformation defined in (2). To see that if $\hat{\Omega}$ is ex-ante IC, so is Ω , it suffices to note that (i) payoffs depend on the shocks ε only through $z(\varepsilon; y)$, (ii) Ω induces the same distribution over $\Theta \times Y$ as $\hat{\Omega}$, and (iii) any distribution over $\Theta \times Y$ that the agent can induce given Ω could also have been induced given $\hat{\Omega}$.

2.1 Necessary conditions for Incentive Compatibility

Suppose now that the primitive environment (U, F) admits an independent-shock representation (\hat{U}, G, z) — we will show below that this is always the case. One can then use this representation as an instrument to characterize the properties of incentive-compatible mechanisms. In particular, as mentioned above, one can treat the shocks as the agent's private information and then express

⁵Such a regular conditional probability distribution here exists since $\varepsilon^t \in \mathbb{R}^t$.

the (dynamics of the) agent's equilibrium payoff in terms of the (derivative of the) value function with respect to the shocks. To this aim, let

$$\hat{H} \equiv \{(\varepsilon^s, m^t, y^u) \in \mathcal{E}^s \times \mathcal{E}^t \times Y^u \quad \text{with } T \geq s \geq t \geq u \geq s-1\}$$

denote the set of all possible histories in the extensive form corresponding to the mechanism $\hat{\Omega}$ for the IS representation. For any $\hat{h} \in \hat{H}$, then let $\hat{\mu}[\hat{\Omega}|\hat{h}]$ denote the (unique) probability measure over $\mathcal{E} \times \mathcal{E} \times Y$ induced by assuming that, following the history \hat{h} in the mechanism $\hat{\Omega}$, the agent reports truthfully at any subsequent information set. Finally, let $\hat{V}^{\hat{\Omega}}(\hat{h})$ denote the agent's value function in $\hat{\Omega}$ evaluated at history \hat{h} . We then have the following result.

Proposition 1 *Fix t and suppose that $\mathcal{E}_t = (\underline{\varepsilon}_t, \bar{\varepsilon}_t) \subset \mathbb{R}$ for some $-\infty \leq \underline{\varepsilon}_t \leq \bar{\varepsilon}_t \leq +\infty$. In addition, suppose that there exists $A_t \in \mathbb{R}_+$ such that, for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$,⁶ the function $\hat{U}((\cdot, \varepsilon_{-t}), y) : \mathcal{E}_t \rightarrow \mathbb{R}$ is A_t -Lipschitz continuous and differentiable in ε_t . Then after any history $\hat{h}^{t-1} = (\varepsilon^{t-1}, \hat{\varepsilon}^{t-1}, y^{t-1})$, the value function $\hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})$ is Lipschitz continuous in ε_t .*

Furthermore, at any period- t history $(\varepsilon_t, \hat{h}^{t-1})$ at which the mechanism $\hat{\Omega}$ is IC and the value function is differentiable in ε_t ,

$$\frac{\partial \hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})}{\partial \varepsilon_t} = \mathbb{E}^{\hat{\mu}[\hat{\Omega}|\varepsilon_t, \hat{h}^{t-1}]} \left[\frac{\partial \hat{U}(\tilde{\varepsilon}, \tilde{y})}{\partial \varepsilon_t} \right].$$

The proof of this result is quite simple and follows from arguments similar to those that establish the envelope theorem in a static setting.

Condition (1) provides a convenient representation of how the agent's payoff must vary with the agent's private information in an IC mechanism. In certain applications (e.g. the AR(k) example described above), working directly with the reduced-form payoff \hat{U} may facilitate the characterization of the properties of optimal mechanisms. For the result in Proposition 1 to be useful, it is however important to understand what properties of the payoff function U and of the functions z corresponding to the kernels F of the primitive representation guarantee that the agent's reduced-form payoff \hat{U} is equi-Lipschitz continuous and differentiable in ε_t . This is what we address next.⁷ We start with an example which we believe is prominent in applications.

⁶Throughout, the notation \mathcal{E}_{-t} stands for $\mathcal{E}_{-t} \equiv \times_{\tau \neq t} \mathcal{E}_\tau$.

⁷When $T = +\infty$, the properties of Fréchet differentiability and equi-Lipschitz continuity are always meant to apply with respect to the $\|\cdot\|_\delta$ norm. When, instead, T is finite, the specification of the norm is irrelevant, for all norms are equivalent.

Assumption 1 *There exists a collection of functions $u \equiv \langle u_t : \Theta^t \times Y^t \rightarrow \mathbb{R} \rangle_{t=1}^T$ and a collection of scalars $B \equiv (B_t)_{t=1}^T$ with $B_t \in \mathbb{R}_+$ for all t and $\sum_{t=1}^T B_t < +\infty$ such that: (i) for any $(\theta, y) \in \Theta \times Y$,*

$$U(\theta, y) = \sum_{t=1}^T u_t(\theta^t, y^t) \quad (5)$$

and (ii) for any t any $y^t \in Y^t$, $u_t(\cdot, y^t)$ is B_t -Lipschitz continuous and differentiable.

With a finite horizon, part (i) is trivially satisfied and Assumption 1 is equivalent to assuming that the function $U(\theta, y)$ is equi-Lipschitz and differentiable (as a multi-variate function) in θ . With an infinite horizon, assuming that U admits the additive representation of (5) is clearly not without loss of generality. However, such a representation is quite standard in applications. Note that, in an infinite-horizon setting, the condition on the summability of the Lipschitz constants is satisfied for example when for any $t \geq 1$, $u_t(\theta^t, y^t) = \rho^{t-1}u(\theta_t, y_t)$ with $u(\cdot, y_t)$ K -Lipschitz continuous and differentiable and $\rho \in (0, 1)$. We then have the following result.

Proposition 2 *Suppose that assumption 1 holds. Fix t and suppose that $\mathcal{E}_t = (\underline{\varepsilon}_t, \bar{\varepsilon}_t) \subset \mathbb{R}$ for some $-\infty \leq \underline{\varepsilon}_t \leq \bar{\varepsilon}_t \leq +\infty$. In addition, assume that, for any $\tau \geq t$, there exists a $C_{t,\tau} \in \mathbb{R}_+$ such that (a) for all $(\varepsilon_{-\tau}^\tau, y^{\tau-1}) \in \mathcal{E}_{-\tau}^\tau \times Y^{\tau-1}$,⁸ the function $z_\tau((\cdot, \varepsilon_{-\tau}^\tau); y^{\tau-1}) : \mathcal{E}_t \rightarrow \Theta_t$ is $C_{t,\tau}$ -Lipschitz continuous and differentiable, and (b) $\sum_{\tau=t}^T C_{t,\tau} < +\infty$. Then, for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$, the function $\hat{U}((\cdot, \varepsilon_{-t}), y) : \mathcal{E}_t \rightarrow \mathbb{R}$ is A_t -Lipschitz continuous and differentiable and its derivative is given by*

$$\frac{\partial \hat{U}(\varepsilon, y)}{\partial \varepsilon_t} = \sum_{s=t}^T \sum_{\tau=s}^T \frac{\partial u_\tau(z^\tau(\varepsilon^\tau; y^{\tau-1}), y^\tau)}{\partial \theta_s} \frac{\partial z_s(\varepsilon^s; y^{s-1})}{\partial \varepsilon_t}.$$

One can verify that the conditions on the functions z_t assumed in the proposition are satisfied for example when θ_t evolves according to an AR(1) process with coefficient of linear dependence $|\phi_1| < 1$.

The result in the previous proposition can be generalized as follows.

Assumption 2 *The function $U(\cdot, y) : \Theta \rightarrow \mathbb{R}$ is K -Lipschitz continuous and (Frechet) differentiable in θ .*

Proposition 3 *Fix t and suppose that, in addition to assumption 2, $\mathcal{E}_t = (\underline{\varepsilon}_t, \bar{\varepsilon}_t) \subset \mathbb{R}$ for some $-\infty \leq \underline{\varepsilon}_t \leq \bar{\varepsilon}_t \leq +\infty$ and that there exists a scalar $Q_t \in \mathbb{R}_+$ such that, for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$, the function $z(\cdot, \varepsilon_{-t}; y) : \mathcal{E}_t \rightarrow \Theta$ is Q_t -Lipschitz continuous and (Frechet) differentiable in ε_t . Then*

⁸By $\mathcal{E}_{-\tau}^\tau$ we mean $\mathcal{E}_{-\tau}^\tau \equiv \times_{j \in \mathbb{N} \setminus \{t\}, j \leq \tau} \mathcal{E}_j$.

there exists an $A_t \in \mathbb{R}_+$ such that, for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$, the function $\hat{U}((\cdot, \varepsilon_{-t}), y) : \mathcal{E}_t \rightarrow \mathbb{R}$ is A_t -Lipschitz continuous and differentiable and its derivative is given by

$$\frac{\partial \hat{U}(\varepsilon, y)}{\partial \varepsilon_t} = \sum_{s=t}^T \frac{\partial U(z(\varepsilon; y), y)}{\partial \theta_s} \frac{\partial z_s(\varepsilon^s; y^{s-1})}{\partial \varepsilon_t}.$$

The proof for this proposition follows directly from the chain rule of Frechet differentiability. As mentioned above, when T is finite, then Frechet differentiability reduces to standard multivariate differentiability. In this case, a sufficient condition for $z(\cdot, \varepsilon_{-t}; y) : \mathcal{E}_t \rightarrow \Theta$ to be differentiable and equi-Lipschitz continuous is that each $z_s((\cdot, \varepsilon_{-t}^s); y^{s-1}) : \mathcal{E}_t \rightarrow \Theta_s$ is differentiable and equi-Lipschitz continuous in ε_t , $t < s$.

Comparing Proposition 3 to Proposition 2, it is immediate that the only difference between the two emerges when $T = \infty$. While Proposition 3 does not require U to take the additive form of Assumption 1, it requires to check Frechet differentiability and equi-Lipschitz continuity of $U(\cdot, y) : \Theta \rightarrow \mathbb{R}$ and $z(\cdot, \varepsilon_{-t}; y) : \mathcal{E}_t \rightarrow \Theta$. Proposition 2, on the other hand, presumes that preferences admit the additive form of Assumption 1, but then requires to check differentiability and equi-Lipschitz continuity of the single component functions $z_\tau((\cdot, \varepsilon_{-t}^\tau); y^{\tau-1}) : \mathcal{E}_t \rightarrow \Theta_t$ as opposed to Frechet differentiability and equi-Lipschitz continuity of the entire infinite-dimensional mapping $z(\cdot, \varepsilon_{-t}; y) : \mathcal{E}_t \rightarrow \Theta$. The two propositions thus complement each other.

2.2 The canonical IS representation

While the results in the previous section apply to *any* IS representation, at this point one may wonder which processes F admit an IS representation and which ones admit an IS representation for which the corresponding z function satisfies the properties of Propositions 2 and 3. We address each of these questions in turn.

First, we show that *any* process admits a particular independent-shock representation, which henceforth we refer to as the *canonical representation*. This representation is derived from the kernels F as follows. Let $\tilde{\varepsilon}$ denote a (possibly infinitely dimensional) vector of independent random variables, each uniformly distributed over $(0, 1)$, independently from any other. Next, for any t , any $\varepsilon \in (0, 1)$, any (θ^{t-1}, y^{t-1}) , let

$$F_t^{-1}(\varepsilon | \theta^{t-1}, y^{t-1}) \equiv \inf \{ \theta_t : F_t(\theta_t | \theta^{t-1}, y^{t-1}) \geq \varepsilon \}$$

denote the generalized inverse of the kernel F_t . Now let $z : \mathcal{E} \times Y \rightarrow \Theta$ be the mapping recursively defined by

$$z_t(\varepsilon^t; y^{t-1}) \equiv F_t^{-1}(\varepsilon_t | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1}) \quad \forall t \quad (6)$$

Applying the probability integral transform theorem recursively, one can then show that, given any $y^{t-1} \in Y^{t-1}$ and any $\varepsilon^{t-1} \in (0, 1)^{t-1}$, the distribution of $z_t(\tilde{\varepsilon}^t; y^{t-1})$ given y^{t-1} and $\tilde{\varepsilon}^{t-1} = \varepsilon^{t-1}$ is the same as the distribution of θ_t given y^{t-1} and $\theta^{t-1} = (F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1})$. Hence, any process admits an independent-shock representation in which, for any t and y^{t-1} , $G_t(\cdot; y^{t-1})$ is simply the uniform distribution over $(0, 1)$ and where the functions $z_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Theta_t$ are the ones defined in (6).

Using the canonical representation, one can then identify conditions on the kernels F that guarantee that the corresponding z function, as defined in (6), satisfies either the properties of Proposition 2, or those of Proposition 3. We start with the following two preliminary conditions.

Assumption 3 For any $t \geq 1$, any $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the function $F_t^{-1}(\cdot | \theta^{t-1}, y^{t-1})$ is differentiable.

Assumption 4 For any $t \geq 2$, any $\varepsilon \in (0, 1)$ any $y^{t-1} \in Y^{t-1}$, the function $F_t^{-1}(\varepsilon | \cdot, y^{t-1})$ is differentiable.

Together, these conditions guarantee differentiability of the components of the z function.

Lemma 2 Let (z, G) be the canonical IS representation for the process corresponding to the kernels F . Assume that assumptions 3 and 4 hold. For any $t \geq 1$ and any $\tau \geq t$, the function $z_\tau((\cdot, \varepsilon_{-t}^\tau); y^{\tau-1}) : \mathcal{E}_t \rightarrow \Theta_\tau$ defined by (6) is differentiable with derivative

$$\frac{\partial z_\tau((\varepsilon_t, \varepsilon_{-t}^\tau); y^{\tau-1})}{\partial \varepsilon_t} = \hat{I}_t^t(\varepsilon^t, y^{t-1}) \hat{J}_t^\tau(\varepsilon^\tau, y^{\tau-1}), \quad (7)$$

where $\hat{J}_t^t(\varepsilon^t, y^{t-1}) \equiv 1$ and

$$\hat{J}_t^\tau(\varepsilon^\tau, y^{\tau-1}) \equiv \sum_{\substack{K \in \mathbb{N}, \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K \hat{I}_{l_{k-1}}^{l_k} \text{ for } \tau > t,$$

with

$$\hat{I}_t^t(\varepsilon^t, y^{t-1}) \equiv \frac{\partial F_t^{-1}(\varepsilon_t | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1})}{\partial \varepsilon_t}$$

and

$$\hat{I}_l^m(\varepsilon^m, y^{m-1}) \equiv \frac{\partial F_m^{-1}(\varepsilon_m | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{m-1})}{\partial \theta_l}, \quad m > l.$$

The proof follows again directly from the chain rule of differentiability. To apply the result in Proposition 2 one then simply needs to guarantee that, in addition to be differentiable, the (inverse of the) kernels be equi-Lipschitz continuous with appropriate bounds. Using the preceding lemma, this in turn can be guaranteed by assuming that the following hold.

Assumption 5 For any $t \geq 1$, there exists a $M_t \in \mathbb{R}_+$ such that, for any $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the function $F_t^{-1}(\cdot | \theta^{t-1}, y^{t-1})$ is M_t -Lipschitz continuous.

Assumption 6 For any $t, \tau, \tau > t$, there exists a $k_t^\tau \in \mathbb{R}_+$ such that, (a) $|\hat{J}_t^\tau(\varepsilon^\tau, y^{\tau-1})| \leq k_t^\tau$ for any $(\varepsilon^\tau, y^{\tau-1})$, and (b) for any $t \geq 1$, $\sum_{\tau=t+1}^T k_t^\tau < +\infty$.

It is easy to see that assumption 6 holds for example when, for any $t, \tau, \tau > t$, the function $F_\tau^{-1}(\varepsilon | (\theta_{-t}^{\tau-1}, \cdot), y^{\tau-1})$ is N_t^τ -equi-Lipschitz continuous in θ_t , and

$$\sum_{\tau=t+1}^T \left[\sum_{\substack{K \in \mathbb{N}, l \in \mathbb{N}^{K+1}: \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K N_{l_{k-1}}^{l_k} \right] < +\infty \text{ for any } t \geq 1. \quad (8)$$

This condition is easily satisfied when T is finite. When $T = \infty$, it is satisfied for example when θ follows an AR(k) process with coefficients $|\phi_j| < 1, j = 1, \dots, k$. We then have the following result.

Proposition 4 Suppose that assumptions 1, 3, 4, 5, and 6, hold. Fix the history $\hat{h}^{t-1} \equiv (\varepsilon^{t-1}, \hat{\varepsilon}^{t-1}, y^{t-1})$. The value function $V^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})$ is Lipschitz continuous in ε_t .

Furthermore, at any ε_t at which $V^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})$ is differentiable and $\hat{\Omega}$ is IC at $(\varepsilon_t, \hat{h}^{t-1})$,

$$\frac{\partial V^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})}{\partial \varepsilon_t} = \hat{I}_t^t(\varepsilon^t, y^{t-1}) \left\{ \mathbb{E}^{\hat{\mu}[\hat{\Omega}] | \varepsilon_t, \hat{h}^{t-1}} \left[\sum_{\tau=t}^T \hat{J}_t^\tau(\tilde{\varepsilon}^\tau, \tilde{y}^{\tau-1}) \frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_\tau} \right] \right\}.$$

Next, consider the result in Proposition 3. To apply this proposition one needs that the function $z(\cdot, \varepsilon_{-t}; y) : \mathcal{E}_t \rightarrow \Theta$ be equi-Lipschitz continuous and Frechet differentiable in ε_t . These properties can in turn be guaranteed by assuming that, in addition to conditions 3, 4, 5, the following conditions hold.

Assumption 7 For any t any $(\varepsilon_{-t}, \theta^{t-1}, y) \in \mathcal{E}_{-t} \times \Theta^{t-1} \times Y$, the function $W^t((\cdot, \varepsilon_{-t}), (\theta^{t-1}, \cdot), y) : (0, 1) \times \Theta_+^t \rightarrow \Theta$ defined by $W_s^t(\varepsilon, \theta, y) = \theta_s$ for all $s \leq t-1$ and

$$W_s^t(\varepsilon, \theta, y) = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1}) \quad \forall s \geq t$$

is Frechet differentiable, where $\Theta_+^t \equiv \prod_{s=t}^T \Theta_s$.

Assumption 8 For any $(\varepsilon, y) \in \mathcal{E} \times Y$,

$$\lim_{t \rightarrow T} \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t(\varepsilon^t, y^{t-1}) \right| < +\infty$$

Assumption 9 For any t , there exists a $K_t < \infty$ such that, for all $(\varepsilon, y) \in \mathcal{E} \times Y$,

$$\sup_{\tau \geq t} \delta^{\tau-t} \left| \hat{J}_t^\tau(\varepsilon^\tau, y^{\tau-1}) \right| \leq K_t.$$

Assumption 7 is an assumption of equi-differentiability of the kernels $F_s^{-1}(\varepsilon_s | \cdot, y^{s-1})$, $s \geq t$. Assumption 8 is a “backward-looking” analog of assumption 6; note that this condition is satisfied for example when, for any t, τ , $\tau > t$, the function $F_\tau^{-1}(\varepsilon | (\theta_{-t}^{\tau-1}, \cdot), y^{\tau-1})$ is N_t^τ -equi-Lipschitz continuous in θ_t , and

$$\lim_{t \rightarrow T} \sum_{\tau=1}^{t-1} \delta^\tau \left[\sum_{\substack{K \in \mathbb{N}, l \in \mathbb{N}^{K+1}: \\ t-\tau=l_0 < \dots < l_K=t}} \prod_{k=1}^K N_{l_{k-1}}^{l_k} \right] < +\infty. \quad (10)$$

Assumption 9, which is weaker than (in the sense of being implied by) assumption 6, is satisfied for example when there exists scalars $B \geq 0$ and $M \geq 0$ such that (i) $\delta(B + M) < 1$ and (ii) for all $t \geq 1$, $\tau > t$, $(\varepsilon, y) \in \mathcal{E} \times Y$,

$$\left| \hat{I}_t^\tau(\varepsilon^\tau, y^{\tau-1}) \right| \leq M^{\tau-1-t} B.$$

Combining the above results, leads to the following proposition.

Proposition 5 Suppose that assumptions 2, 3, 4, 5, 7, 8, and 9 hold. Then the same conclusions as in Proposition 4 hold.

Propositions 4 and 5 thus identify a set of conditions for the primitive environment (U, F) that guarantee that, in any IC mechanism, the agent’s expected payoff, when expressed using the canonical IS representation, satisfies the envelope formula of (1).

Note that, when applied to a *finite-horizon* setting, the conditions in the two propositions coincide; these conditions then reduce to assuming that the payoff U be differentiable and equi-Lipschitz continuous in θ and that the (inverse of the) kernels $F_t^{-1}(\varepsilon | \theta^{t-1}, y^{t-1})$ be differentiable and equi-Lipschitz continuous both in the “quantile” ε and in the past θ^{t-1} . Comparing these conditions to those in Proposition 1 in PST (2009) one can see that while the assumptions in that proposition rule out, for example, an atom at $\theta_t = \theta_t^\#$ that “shifts” with the past θ^{t-1} (e.g., fully persistent types), such a possibility is accommodated by the assumptions in Proposition 5 above. On the other hand, the assumptions in Proposition 5 rule out an atom at $\theta_t = \theta_t^\#$ whose measure grows with θ^{t-1} while such a possibility is allowed by the assumptions in Proposition 1 in PST (2009). When applied to finite-horizon environments, the result in (4) thus provides an alternative closed-form representation for the derivative of the value function that one can use, for example, when the assumptions in Proposition 1 in PST (2009) are violated. The most significant advantage of Proposition 5 however remains the fact that it also permits one to identify necessary conditions

for IC in *infinite horizon* settings.

Finally note that, while the formula in (4) describes the dynamics of the value function in the mechanism $\hat{\Omega}$ in which the agent reports the shocks ε instead of his payoff-relevant types θ , the same formula also permits one to express the derivative of the value function in the original mechanism Ω in which the agent reports θ instead of ε . To see this, it suffices to proceed as follows. Take any mechanism Ω for the primitive representation (U, F) and let $\hat{\Omega}$ be the corresponding mechanism in the independent-shock representation that is obtained from Ω using (3). Because, for any y , the agent's payoff in $\hat{\Omega}$ depends on ε only through $z(\varepsilon; y)$, we have that, for any t , y^{t-1} and ε^t the following identity holds:

$$\hat{V}^{\hat{\Omega}}(\varepsilon^t, \varepsilon^{t-1}, y^{t-1}) = V^{\Omega}(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1}). \quad (11)$$

Therefore, at any point of differentiability of $\hat{V}^{\hat{\Omega}}$ in ε_t ,

$$\frac{\partial \hat{V}^{\hat{\Omega}}(\varepsilon^t, \varepsilon^{t-1}, y^{t-1})}{\partial \varepsilon_t} = \frac{\partial V^{\Omega}(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} \frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t}. \quad (12)$$

While conditions (11) and (12) hold for *all* independent-shock representations, when (G, z) is the canonical IS representation of F ,

$$\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} = \hat{I}_t^t(\varepsilon^t, y^{t-1}).$$

Now suppose that the following two assumptions also hold.

Assumption 10 For all t , and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the c.d.f. $F_t(\cdot | \theta^{t-1}, y^{t-1})$ is strictly increasing on Θ_t .

Assumption 11 For all t and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the function $F_t(\cdot | \theta^{t-1}, y^{t-1})$ is absolutely continuous with density $f_t(\theta_t | \theta^{t-1}, y^{t-1}) > 0$ for a.e. $\theta_t \in \Theta_t$.

Combining (12) with (4), it is then easy to see that, when in addition to the assumptions in Proposition 5, assumptions 10 and 11 also hold, then $\hat{I}_t^t(\varepsilon^t, y^{t-1}) \neq 0$ and

$$\hat{I}_t^\tau(\varepsilon^\tau, y^{\tau-1}) = I_t^\tau(\theta^\tau, y^{\tau-1}) \Big|_{\theta^\tau = z^\tau(\varepsilon^\tau; y^{\tau-1})} \quad \text{and} \quad \hat{J}_t^\tau(\varepsilon^\tau, y^{\tau-1}) = J_t^\tau(\theta^\tau, y^{\tau-1}) \Big|_{\theta^\tau = z^\tau(\varepsilon^\tau; y^{\tau-1})}$$

where $J_t^t(\theta^t, y^{t-1}) \equiv 1$ and

$$J_t^\tau(\theta^\tau, y^{\tau-1}) \equiv \sum_{\substack{K \in \mathbb{N}, l \in \mathbb{N}^{K+1}; \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K I_{l_{k-1}}^{l_k}(\theta^{l_k}, y^{l_k-1}) \quad \text{for } \tau > t,$$

with

$$I_l^m(\theta^m, y^{m-1}) \equiv -\frac{\partial F_m(\theta_m|\theta^{m-1}, y^{m-1})/\partial\theta_l}{f_m(\theta_m|\theta^{m-1}, y^{m-1})} \text{ for } l < m.$$

The following is then an immediate implication of the aforementioned results.

Proposition 6 *Suppose that, in addition to the assumptions in Propositions 4 (or to those in Proposition 5), assumptions 10 and 11 hold. Then the conclusions of Proposition 2 in PST (2009) hold. That is, if Ω is IC at the truthful history $h^{t-1} \equiv (\theta^{t-1}, \theta^{t-1}, y^{t-1})$, then*

$$V^\Omega(\theta_t, h^{t-1}) \text{ is Lipschitz continuous in } \theta_t, \text{ and for a.e. } \theta_t, \quad (13)$$

$$\frac{\partial V^\Omega(\theta_t, h^{t-1})}{\partial\theta_t} = \mathbb{E}^{\mu[\Omega]|(\theta_t, h^{t-1})} \left[\sum_{\tau=t}^T J_t^T(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial\theta_\tau} \right].$$

Note that, while the conclusions in the two propositions are the same, the conditions in Proposition 6 that validate (13) are somewhat different from those in Proposition 2 in PST (2009). To better appreciate this, it is instructive to consider the case of a finite horizon. The key differences between the assumptions in the two propositions are then the following. While (for generic non-Markov settings) the backward-induction approach in PST (2009) requires the probability measures $F_t(\cdot|\theta^{t-1}, y^{t-1})$ to be continuous in θ^{t-1} in the total variation metric, such an assumption is not required under the IS approach in this paper. Furthermore, while the backward-induction approach requires the (absolute value of) the derivative of the Kernels $|\partial F_t(\theta_t|\theta^{t-1}, y^{t-1})/\partial\theta_s|$ to be bounded uniformly in (θ^{t-1}, y^{t-1}) by an integrable function $B_t : \Theta_t \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ such an assumption is not required under the IS approach. On the other hand, while the backward-induction approach only requires the payoff U and each kernel F_t to be partially differentiable in each θ_s , the IS approach requires these functions to be totally differentiable (in θ and θ^{t-1} respectively). The two propositions thus complement each other by identifying different sets of assumptions that validate the dynamic payoff formula given in (13) as a necessary condition for incentive compatibility.

3 Bandit Auctions

To illustrate how the results in the previous section can be put to work, we now consider a multi-agent quasilinear setting where buyers refine their valuations through consumption. For an illustration of how the multi-agent setting can be mapped into a single agent setting, we refer the reader to Section 4 in PST. The purpose of this section is to illustrate how the independent-shock approach outlined above can help characterize the properties of optimal mechanisms in infinite-horizon settings.

Setup. There is an auctioneer and $N \geq 1$ bidders. In each period, $t = 1, 2, \dots, T$, with $T = +\infty$,

the auctioneer has one indivisible, non-storable, object to sell. If he allocates the period- t object to agent i , he incurs a cost $c_{i,t} \in \mathbb{R}$. If the object goes unassigned, the auctioneer incurs no cost, and the object perishes. We then let $X_{i,t} \equiv \{0, 1\}$, $X_t = \left\{x_t \in \prod_{i=0}^N X_{i,t} : \sum_{i=0}^N x_{i,t} = 1\right\}$, $X = \prod_{t=1}^T X_t$ and $X_i^t = \prod_{t=1}^t X_{i,t}$ denote the relevant sets of feasible allocations (Here $i = 0$ refers to the seller).

Each bidder i 's Bernoulli function takes the form

$$U_i(\theta, x, p) = \sum_{t=1}^{\infty} \delta^t [\theta_{i,t} x_{i,t} - p_{i,t}],$$

whereas the auctioneer's Bernoulli function takes the form

$$U_0(\theta, x, p) = \sum_{t=1}^{\infty} \delta^t \left[\sum_{i=1}^N [p_{i,t} - x_{i,t} c_{i,t}] \right],$$

where $\theta_{i,t} \in \Theta_{i,t}$ and $p_{i,t} \in \mathbb{R}$ denote, respectively, the period- t valuation and the period- t payment for agent i , $\theta_t = (\theta_{1,t}, \dots, \theta_{N,t}) \in \prod_{i=1}^N \Theta_{i,t}$ and $p_t = (p_{1,t}, \dots, p_{N,t}) \in \mathbb{R}^N$ a profile of valuations and payments for period t , $\theta = (\theta_t)_{t=1}^T$ the complete state of the world, and $p = (p_t)_{t=1}^T \in \mathbb{R}^{NT}$ the payments received by the auctioneer over time.

The stochastic process governing the evolution of each $\theta_{i,t}$ is as follows. The first period valuation $\theta_{i,1}$ is drawn from an absolutely continuous c.d.f. $F_{i,1}$ strictly increasing over the interval $\Theta_{i,1} = (\underline{\theta}_{i,1}, \bar{\theta}_{i,1}) \subset \mathbb{R}$ with density $f_{i,1}(\theta_{i,1}) > 0$ if and only if $\theta_{i,1} \in \Theta_{i,1}$. For any $t > 1$, the period- t valuation $\theta_{i,t}$ is given by

$$\theta_{i,t} = \theta_{i,t-1} + \varepsilon_{i,t},$$

where the shock $\varepsilon_{i,t}$ is drawn from the c.d.f. $G_{i,t}(\cdot | x_i^{t-1})$. In other words,

$$F_{i,t}(\theta_{i,t} | \theta_i^{t-1}, x_i^{t-1}) = G_{i,t}(\theta_{i,t} - \theta_{i,t-1} | x_i^{t-1})$$

Given x^{t-1} , the shocks $\varepsilon_t = (\varepsilon_{1,t}, \dots, \varepsilon_{N,t})$ are independent across agents, i.e.

$$G(\cdot | x^{t-1}) = \times_{i=1}^N G_{i,t}(\cdot | x_i^{t-1})$$

Furthermore, for any $i = 1, \dots, N$, given x_i^T , the shocks $\varepsilon_i^T = (\varepsilon_{i,1}, \dots, \varepsilon_{i,T})$ are jointly independent, i.e.

$$G_i(\cdot | x_i^T) = \times_{t=1}^T G_{i,t}(\cdot | x_i^{t-1})$$

In addition to the aforementioned properties, the family of distributions $G_{i,t}(\cdot | x_i^{t-1})$ satisfies the following conditions: (a) for any x_i^{t-1} such that $x_{i,t-1} = 0$, $G_{i,t}(\cdot | x_i^{t-1})$ is a Dirac delta at 0; (b) for any t, τ , any x_i^{t-1} and $\hat{x}_i^{\tau-1}$ such that $x_{i,t-1} = \hat{x}_{i,\tau-1} = 1$ and $\sum_{s=1}^{\tau-1} \hat{x}_{i,s} = \sum_{s=1}^{t-1} x_{i,s}$,

$G_{i,t}(\cdot|x_i^{t-1}) = G_{i,\tau}(\cdot|\hat{x}_i^{t-1})$; (c) for any x_i^{t-1} such that $x_{i,t-1} = 1$, $G_{i,t}(\cdot|x_i^{t-1})$ is absolutely continuous with density strictly positive almost everywhere on \mathbb{R} .

The first condition says that valuations do not change in the absence of consumption. The second condition says that, conditional upon consumption in the preceding period, the distribution of the shock $\varepsilon_{i,t}$ is time-homogeneous in the sense that it depends on x_i^{t-1} only through $\sum_{s=1}^{t-1} x_{i,s}$. The last condition, which is not essential for the subsequent results, makes the analysis of IC simpler by establishing that the support of $\theta_{i,t}$ is independent of θ_i^{t-1} . Given these properties, for simplicity, hereafter we denote the distributions of the shocks by $G_{i,t}(\cdot|x_{i,t-1}, \sum_{s=1}^{t-1} x_{i,s})$.

Remark 1 *This kind of structure arises for example in a model with Normal learning. There θ_t is the posterior expectation of the true underlying valuation, and the impact of the k^{th} signal on the posterior is the same regardless of the period in which it arrives. More generally, this specification allows for learning by doing (or habit formation), for it does not require the mean of ε_t to be zero.*

Summarizing, the key assumptions are the following.

1. Valuations are independent across bidders.
2. Valuations change only upon consumption.
3. The environment is Markov: payoffs are additively time-separable and, conditional on allocations, the valuations follow a Markov process.
4. The valuation processes are time-homogenous: If bidder i wins the auction in period t , then the distribution of his period $t + 1$ valuation depends only on his valuation in period t and the total number of times he has won in the past.

IS representation, and Necessary IC conditions. The structure of the payoffs and of the process for the valuations suggest using a non-canonical IS representation where the distributions of the shocks depend on past decisions. For any $i = 1, \dots, N$, any $t \geq 1$, then let $z_{i,t}(\varepsilon_i^t, y_i^t) = \sum_{s=1}^t \varepsilon_{i,s}$, where for any i , any $t > 1$, the shock $\varepsilon_{i,t} \in \mathcal{E}_{i,t} = \mathbb{R}$ are the same shocks described above. Because the functions $z_{i,t}(\varepsilon_i^t, y_i^t)$ do not depend on $y_i^t = (x_i^{t-1}, p_i^{t-1})$, we simplify notation by denoting them by $z_{i,t}(\varepsilon_i^t)$, with $z_i^t(\varepsilon_i^t) = (z_{i,s}(\varepsilon_i^s))_{s=1}^t$. Each agent's reduced form payoff then takes the form

$$\begin{aligned} \hat{U}_i(\varepsilon, x, p) &= U_i(z_i^T(\varepsilon_i^T), x, p) \\ &= \sum_{t=1}^T \delta^t [z_{i,t}(\varepsilon_i^t) x_{i,t} - p_{i,t}] \\ &= \sum_{t=1}^T \delta^t \left[\left(\sum_{s=1}^t \varepsilon_{i,s} \right) x_{i,t} - p_{i,t} \right] \end{aligned}$$

Note that the flow payoffs $\hat{u}_{i,t}(\varepsilon_i^t, x_{i,t}) \equiv \delta^t [(\sum_{s=1}^t \varepsilon_{i,s}) x_{i,t} - p_{i,t}]$ is clearly differentiable in $\varepsilon_{i,1}$ with

$$\left| \frac{\partial \hat{u}_{i,t}(\varepsilon_i^t, x_{i,t})}{\partial \varepsilon_{i,1}} \right| = \delta^t x_{i,t} \leq \delta^t.$$

Thus, by the Lebesgue dominated convergence theorem, \hat{U}_i is differentiable in $\varepsilon_{i,1}$ with

$$\frac{\partial \hat{U}_i(\varepsilon, x, p)}{\partial \varepsilon_{i,1}} = \sum_{t=1}^T \frac{\partial \hat{u}_{i,t}(\varepsilon_i^t, x_{i,t})}{\partial \varepsilon_{i,1}} = \sum_{t=1}^T \delta^t x_{i,t}.$$

Furthermore, \hat{U}_i is equi-Lipschitz continuous in $\varepsilon_{i,1}$ since, for all (ε, x, p) ,

$$\left| \frac{\partial \hat{U}_i(\varepsilon, x, p)}{\partial \varepsilon_{i,1}} \right| \leq \sum_{t=1}^T |\delta^t x_{i,t}| \leq \frac{1}{1-\delta}.$$

From Proposition 1, we then have that, in any ex-ante IC mechanism $\hat{\Omega}$, each bidder i 's value function $\hat{V}_i^{\hat{\Omega}}(\varepsilon_{i,1})$ is Lipschitz continuous and for a.e. $\varepsilon_{i,t}$,

$$\frac{\partial \hat{V}_i^{\hat{\Omega}}(\varepsilon_{i,1})}{\partial \varepsilon_{i,1}} = \mathbb{E}^{\mathbb{E}^{\hat{\lambda}[\hat{\chi}]|\varepsilon_{i,1}}} \left[\sum_{t=1}^T \delta^t \tilde{x}_{i,t} \right].$$

Maximizing virtual surplus, Gittins indices. Given a deterministic allocation rule $\hat{\chi} := \langle \hat{\chi}_t : \mathbb{R}^{Nt} \rightarrow X_t \rangle_{t=1}^{\infty}$ the virtual surplus takes the form

$$\mathbb{E}^{\hat{\lambda}[\hat{\chi}]} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i=1}^N \left(\sum_{\tau=1}^t \tilde{\varepsilon}_{i,\tau} - c_{i,t} - \frac{1 - F_{i,1}(\tilde{\varepsilon}_{i,1})}{f_{i,1}(\tilde{\varepsilon}_{i,1})} \right) \hat{\chi}_{i,t}(\tilde{\varepsilon}^t) \right]$$

where $\hat{\lambda}[\hat{\chi}]$ is the unique probability measure over \mathcal{E} induced by the allocation rule $\hat{\chi}$, under truthtelling by all agents.

At this point, it is convenient to switch back to the original representation with θ as this makes the Markov structure more explicit. We may then write the virtual surplus as

$$\mathbb{E}^{\lambda[\chi]} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i=1}^N \left(\tilde{\theta}_{i,t} - c_{i,t} - \frac{1 - F_{i,1}(\tilde{\theta}_{i,1})}{f_{i,1}(\tilde{\theta}_{i,1})} \right) \chi_{i,t}(\tilde{\theta}^t) \right]$$

Once again, the notation $\lambda[\chi]$ stands for the unique probability measure over \mathcal{E} induced by the allocation rule $\hat{\chi}$, under truthtelling by all agents.

Now, given θ_1 , we have a standard $(N+1)$ -armed bandit problem (i.e., N bidders plus a safe arm with payoff 0 which corresponds to not selling) where the objective function is given by the above virtual surplus. If the auctioneer were to maximize social surplus instead of profits the objective function in the bandit problem would be

$$\mathbb{E}^{\lambda[\chi]} \left[\sum_{t=1}^{\infty} \delta^t \sum_{i=1}^N \left(\tilde{\theta}_{i,t} - c_{i,t} \right) \chi_{i,t}(\tilde{\theta}^t) \right].$$

The Gittins index for the safe arm is identically zero in both programs. The Gittins index for arm $i > 0$ in the profit maximization problem at time 1 given $\theta_{i,1}$ is by definition

$$\gamma_{i,1}(\theta_{i,1}) := \max_{\tau} \mathbb{E} \left[\frac{\sum_{t=1}^{\tau} \delta^t \left(\tilde{\theta}_{i,t} - c_{i,t} - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})} \right)}{\sum_{t=1}^{\tau} \delta^t} \mid \theta_{i,1} \right],$$

where τ is a stopping time (i.e., it is not just a scalar, but a policy rule that conditions on history), and the expectation is taken conditional on the first type being $\theta_{i,1}$ and the object being allocated to bidder i in each period until the stopping time τ . We have

$$\begin{aligned} \gamma_{i,1}(\theta_{i,1}) &= \max_{\tau} \mathbb{E} \left[\frac{\sum_{t=1}^{\tau} \delta^t \left(\tilde{\theta}_{i,t} - c_{i,t} - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})} \right)}{\sum_{t=1}^{\tau} \delta^t} \mid \theta_{i,1} \right] \\ &= \max_{\tau} \mathbb{E} \left[\frac{\sum_{t=1}^{\tau} \delta^t \left(\tilde{\theta}_{i,t} - c_{i,t} \right)}{\sum_{t=1}^{\tau} \delta^t} \mid \theta_{i,1} \right] - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})} \\ &= \gamma_{i,1}^E(\theta_{i,1}) - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})}, \end{aligned}$$

where $\gamma_{i,1}^E(\theta_{i,1})$ is the corresponding Gittins index in the efficient program. We can similarly calculate (and relate) the Gittins indices in an arbitrary period t following any bidder i history (i.e., any sequence of allocations to bidder i and a corresponding sequence of types θ_i^t). We have

$$\begin{aligned} \gamma_{i,t}(\theta_{i,1}, \theta_{i,t}, \sum_{u=1}^{t-1} x_{i,u}) &= \max_{\tau} \mathbb{E} \left[\frac{\sum_{s=t}^{\tau} \delta^{s-t} \left(\tilde{\theta}_{i,s} - c_{i,s} \right)}{\sum_{s=1}^{\tau} \delta^{s-t}} \mid \theta_{i,t}, \sum_{u=1}^{t-1} x_{i,u} \right] - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})} \\ &= \gamma_{i,t}^E(\theta_{i,t}, \sum_{u=1}^{t-1} x_{i,u}) - \frac{1 - F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})}. \end{aligned}$$

(Again the expectation is conditional on the optimal stopping policy τ .) Given the Markov structure the indices do not directly depend on $\theta_{i,-1}^{t-1}$ ($\gamma_{i,t}^E$ doesn't depend on $\theta_{i,1}$ either), and depend on x_i^{t-1} only through the sum of its terms. It is well-known that a Gittins index policy is an optimal policy

in a multi-armed bandit problem of the above form. This immediately implies the following result.

Proposition 7 *Let χ^* be the allocation rule such that for all i, t, θ^t , and x^{t-1} ,*

$$\chi_{i,t}^*(\theta^t, x^{t-1}) = 1 \quad \Rightarrow \quad i \in \arg \max_{j \in \{0, \dots, N\}} \gamma_{j,t}(\theta_{j,1}, \theta_{j,t}, \sum_{\tau=1}^{t-1} x_{j,\tau}).$$

Then χ^ maximizes dynamic virtual surplus.*

Obviously social surplus is also maximized by using an index policy.

Incentive compatibility. The efficient policy can be implemented for example with the Team Mechanism. Consider then the profit maximizing policy.

Proposition 8 *Assume that first period hazard rates are nondecreasing. Then the allocation rule of Proposition 7 is sustained under an optimal mechanism.*

The proof in the Appendix proceeds as follows. First we show that, because at any $t \geq 2$, χ^* is efficient given the period-1 reports θ_1 (with the seller's adjusted cost of serving bidder i in period t set equal to $c_{i,t} - \frac{1-F_{i,1}(\theta_{i,1})}{f_1(\theta_{i,1})}$), there always exist a system of payments that induce each bidder to report truthfully from period two onwards, at any period- t history, $t \geq 2$, irrespective of whether the history is truthful or not. Next, we use the fact χ^* is weakly monotone to show that there exists a payment scheme p^* that, in addition to the property described above, it also induces the bidders to report truthfully in period 1. Lastly, we verify that under the mechanism $\Omega^* = [\chi, p^*]$ each bidder finds it optimal to participate and each type $\theta_{i,1}$ obtains a payoff equal to the outside option. That the mechanism Ω^* is optimal then follows from the same arguments that establish Proposition 5 in PST.

Appendix

Proof of Proposition 1.

Fix the history $\hat{h}^{t-1} = h_t = (\varepsilon^{t-1}, \hat{\varepsilon}^{t-1}, y^{t-1})$ and for any $\varepsilon_t, m_t \in \mathcal{E}_t$, let $\mu[\hat{\Omega}, \hat{\sigma}]|m_t, \varepsilon_t, \hat{h}^{t-1}$ denote the (unique) measure over $\mathcal{E} \times \mathcal{E} \times Y$ that is obtained by assuming that, after history $h_t = (\varepsilon^t, \hat{\varepsilon}^{t-1}, y^{t-1})$ is reached, in period t the agent sends the message m_t and then starting from period $t+1$ onwards he follows an arbitrary strategy

$$\hat{\sigma} \equiv \langle \hat{\sigma}_t : \mathcal{E}^t \times \mathcal{E}^{t-1} \times Y^{t-1} \rightarrow \mathcal{E}_t \rangle_{t=1}^T.$$

The key observation here is that, because of the independence of the shocks, the restriction of the measure $\mu[\hat{\Omega}, \hat{\sigma}]|m_t, \varepsilon_t, \hat{h}^{t-1}$ on future shocks, current and future reports, and current and future

decisions, i.e. on $\mathcal{E}_{t+1} \times \dots \times \mathcal{E}_T \times \mathcal{E}_t \times \dots \times \mathcal{E}_T \times Y_t \times \dots \times Y_T$, does not depend on the true shock ε_t .⁹ Formally, let $P(m_t, \hat{h}^{t-1}, \hat{\sigma})$ denote such restriction and $\delta_{\varepsilon_t, \hat{h}^{t-1}}$ denote the Dirac measure at $(\varepsilon_t, \hat{h}^{t-1})$ over past and current shocks, past reports, and past decisions, i.e. over $\mathcal{E}_1 \times \dots \times \mathcal{E}_t \times \mathcal{E}_1 \times \dots \times \mathcal{E}_{t-1} \times Y_1 \times \dots \times Y_{t-1}$. Then the measure $\mu[\hat{\Omega}, \hat{\sigma}]|_{m_t, \varepsilon_t, \hat{h}^{t-1}}$ on $\mathcal{E} \times \mathcal{E} \times Y$ can be decomposed as

$$\mu[\hat{\Omega}, \hat{\sigma}]|_{m_t, \varepsilon_t, \hat{h}^{t-1}} = \delta_{\varepsilon_t, \hat{h}^{t-1}} \times P(m_t, \hat{h}^{t-1}, \hat{\sigma}).$$

By implication,

$$\mathbb{E}^{\mu[\hat{\Omega}, \hat{\sigma}]|_{m_t, \varepsilon_t, \hat{h}^{t-1}}} \left[\hat{U}(\tilde{\varepsilon}, \tilde{y}) \right] = \mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\hat{U}(\varepsilon^t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T) \right].$$

Now, because for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$ the function $\hat{U}(\cdot, \varepsilon_{-t}, y)$ is A_t -Lipschitz continuous, we have that, for any $\varepsilon_t, \varepsilon'_t \in \mathcal{E}_t$ any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$,

$$\left| \frac{\hat{U}((\varepsilon_t, \varepsilon_{-t}), y) - \hat{U}((\varepsilon'_t, \varepsilon_{-t}), y)}{\varepsilon_t - \varepsilon'_t} \right| \leq A_t.$$

On the other hand, because $P(m_t, \hat{h}^{t-1}, \hat{\sigma})$ is a probability measure, $\mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} [A_t] = A_t$. Hence by the Lebesgue dominated convergence theorem,

$$\begin{aligned} & \lim_{\varepsilon'_t \rightarrow \varepsilon_t} \frac{\mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\hat{U}(\varepsilon^{t-1}, \varepsilon_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T) \right] - \mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\hat{U}(\varepsilon^{t-1}, \varepsilon'_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T) \right]}{\varepsilon_t - \varepsilon'_t} \\ &= \lim_{\varepsilon'_t \rightarrow \varepsilon_t} \mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\frac{\hat{U}(\varepsilon^{t-1}, \varepsilon_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T) - \hat{U}(\varepsilon^{t-1}, \varepsilon'_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T)}{\varepsilon_t - \varepsilon'_t} \right] \\ &= \mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\lim_{\varepsilon'_t \rightarrow \varepsilon_t} \frac{\hat{U}(\varepsilon^{t-1}, \varepsilon_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T) - \hat{U}(\varepsilon^{t-1}, \varepsilon'_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T)}{\varepsilon_t - \varepsilon'_t} \right] \\ &= \mathbb{E}^{P(m_t, \hat{h}^{t-1}, \hat{\sigma})} \left[\frac{\partial \hat{U}(\varepsilon^{t-1}, \varepsilon_t, \tilde{\varepsilon}_{t+1}, \dots, \tilde{\varepsilon}_T, y^{t-1}, \tilde{y}_t, \dots, \tilde{y}_T)}{\partial \varepsilon_t} \right] \in [-A_t, A_t], \end{aligned}$$

which implies that, for any plan of action $\hat{\sigma}$, the expected payoff $\mathbb{E}^{\mu[\hat{\Omega}, \hat{\sigma}]|_{m_t, \cdot, \hat{h}^{t-1}}} \left[\hat{U}(\tilde{\varepsilon}, \tilde{y}) \right]$ is A_t -Lipschitz continuous and differentiable in ε_t . The result then follows from essentially the same arguments that establish Theorem 2 in Milgrom and Segal (2002):¹⁰ the value function $\hat{V}^{\hat{\Omega}}(\cdot, \hat{h}^{t-1})$

⁹To be precise, it also does not depend on the true shocks experienced prior to period t ; that is, it depends on the history $(\varepsilon_t, \hat{h}^{t-1}) = (\varepsilon^t, \hat{\varepsilon}^{t-1}, y^{t-1})$ only through the reported shocks $\hat{\varepsilon}^{t-1}$ and the past decisions y^{t-1} .

¹⁰Theorem 2 in Milgrom and Segal (2002) establishes only that the value function is absolutely continuous; this is because that theorem assumes that the payoff is differentiable with an integrable bound instead of differentiable and equi-Lipschitz continuous. It is however immediate to see that the same arguments that establish Theorem 2 in Milgrom and Segal also establish that the value function is equi-Lipschitz continuous under the stronger assumptions considered here.

is Lipschitz continuous in ε_t and at any history $(\varepsilon_t, \hat{h}^{t-1})$ at which $\hat{\Omega}$ is IC and $\hat{V}^{\hat{\Omega}}(\cdot, \hat{h}^{t-1})$ is differentiable

$$\frac{\partial \hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})}{\partial \varepsilon_t} = \mathbb{E}^{\hat{\mu}[\hat{\Omega}]|\varepsilon_t, \hat{h}^{t-1}} \left[\frac{\partial \hat{U}(\tilde{\varepsilon}, \tilde{y})}{\partial \varepsilon_t} \right]$$

where $\hat{\mu}[\hat{\Omega}]|\varepsilon_t, \hat{h}^{t-1}$ is the measure over $\mathcal{E} \times \mathcal{E} \times Y$ induced by assuming that, starting from period t the agent follows a truthful strategy at all current and future information sets. ■

Proof of Proposition 2. For any t and any $(\varepsilon^t, y^t) \in \mathcal{E}^t \times Y^t$, let

$$\hat{u}_t(\varepsilon^t, y^t) \equiv u_t(z^t(\varepsilon^t; y^{t-1}), y^t),$$

so that, under assumption 1,

$$\hat{U}(\varepsilon, y) \equiv U(z(\varepsilon; y), y) = \sum_{t=1}^T \hat{u}_t(\varepsilon^t, y^t).$$

The result follows from combining the following two lemmas below.

Lemma 3 *Fix t . Suppose that, for any $\tau \geq t$, there exists a $D_{t,\tau} \in \mathbb{R}_+$ such that (a) for all $(\varepsilon_{-t}^\tau, y^\tau) \in \mathcal{E}_{-t}^\tau \times Y^\tau$, the function $\hat{u}_\tau(\cdot, \varepsilon_{-t}^\tau, y^\tau)$ is $D_{t,\tau}$ -Lipschitz and differentiable, and (b) $\sum_{\tau=t}^T D_{t,\tau} < +\infty$. Then there exists an $A_t \in \mathbb{R}_+$ such that, for any $(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$, the function $\hat{U}((\cdot, \varepsilon_{-t}), y)$ is A_t -Lipschitz continuous and differentiable with*

$$\frac{\partial \hat{U}((\varepsilon_t, \varepsilon_{-t}), y)}{\partial \varepsilon_t} = \sum_{\tau=t}^T \frac{\partial \hat{u}_\tau(\varepsilon^\tau, y^\tau)}{\partial \varepsilon_t}.$$

Proof of the Lemma. Under the assumptions of the Lemma we have that

$$\begin{aligned} \lim_{\varepsilon'_t \rightarrow \varepsilon_t} \frac{\hat{U}((\varepsilon_t, \varepsilon_{-t}), y) - \hat{U}((\varepsilon'_t, \varepsilon_{-t}), y)}{\varepsilon_t - \varepsilon'_t} &= \lim_{\varepsilon'_t \rightarrow \varepsilon_t} \sum_{\tau=t}^T \frac{\hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau) - \hat{u}_\tau((\varepsilon'_t, \varepsilon_{-t}^\tau), y^\tau)}{\varepsilon_t - \varepsilon'_t} \\ &= \sum_{\tau=t}^T \lim_{\varepsilon'_t \rightarrow \varepsilon_t} \frac{\hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau) - \hat{u}_\tau((\varepsilon'_t, \varepsilon_{-t}^\tau), y^\tau)}{\varepsilon_t - \varepsilon'_t} \\ &= \sum_{\tau=t}^T \frac{\partial \hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau)}{\partial \varepsilon_t} \end{aligned}$$

where the second equality is by the Lebesgue dominated convergence theorem, since, for any

$(\varepsilon_{-t}, y) \in \mathcal{E}_{-t} \times Y$, any $\varepsilon_t, \varepsilon'_t \in \mathcal{E}_t$,

$$\sum_{\tau=t}^T \left| \frac{\hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau) - \hat{u}_\tau((\varepsilon'_t, \varepsilon_{-t}^\tau), y^\tau)}{\varepsilon_t - \varepsilon'_t} \right| \leq \sum_{\tau=t}^T D_{t,\tau} < +\infty.$$

■

Lemma 4 *Suppose the assumptions in Proposition 2 hold. Then for all $\tau \geq t$ there exists $D_{t,\tau} \in \mathbb{R}_+$ such that (a) for all $(\varepsilon_{-t}^\tau, y^\tau) \in \mathcal{E}_{-t}^\tau \times Y^\tau$, $\hat{u}_\tau((\cdot, \varepsilon_{-t}^\tau), y^\tau) : \mathcal{E}_t \rightarrow \mathbb{R}$ is $D_{t,\tau}$ -Lipschitz continuous and differentiable with*

$$\frac{\partial \hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau)}{\partial \varepsilon_t} = \sum_{l=t}^{\tau} \frac{\partial u_\tau(z^\tau(\varepsilon^\tau, y^{\tau-1}), y^\tau)}{\partial \theta_l} \frac{\partial z_l(\varepsilon^l, y^{l-1})}{\partial \varepsilon_t},$$

and (b) $\sum_{\tau=t}^T D_{t,\tau} < +\infty$.

Proof of the Lemma. Fix $(\varepsilon_{-t}^\tau, y^\tau) \in \mathcal{E}_{-t}^\tau \times Y^\tau$ and let $z^\tau((\cdot, \varepsilon_{-t}^\tau); y^{\tau-1}) : \mathcal{E}_t \rightarrow \mathbb{R}^\tau$ denote the vector-valued function defined by

$$z^\tau((\varepsilon_t, \varepsilon_{-t}^\tau); y^{\tau-1}) = (z_s(\varepsilon^s; y^{s-1}))_{s=1}^\tau \quad \forall \varepsilon_t \in \mathcal{E}_t$$

Because each component function z_s is differentiable in ε_t so is $z^\tau((\cdot, \varepsilon_{-t}^\tau); y^{\tau-1})$. The function $\hat{u}_\tau((\cdot, \varepsilon_{-t}^\tau), y^\tau) : \mathcal{E}_t \rightarrow \mathbb{R}$ defined by

$$\hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau) \equiv u_\tau(z^\tau((\varepsilon_t, \varepsilon_{-t}^\tau); y^{\tau-1}), y^\tau) \quad \forall \varepsilon_t \in \mathcal{E}_t$$

is thus the composition of two differentiable functions and hence, by the chain rule, it is itself differentiable and its derivative satisfies the formula in the statement of the lemma. Furthermore,

$$\begin{aligned} \left| \frac{\partial \hat{u}_\tau((\varepsilon_t, \varepsilon_{-t}^\tau), y^\tau)}{\partial \varepsilon_t} \right| &\leq \sum_{l=t}^{\tau} \left| \frac{\partial u_\tau(z^\tau(\varepsilon^\tau, y^{\tau-1}), y^\tau)}{\partial \theta_l} \right| \left| \frac{\partial z_l(\varepsilon^l, y^{l-1})}{\partial \varepsilon_t} \right| \\ &\leq B_\tau \sum_{l=t}^{\tau} C_{t,l}. \end{aligned}$$

Thus $\hat{u}_\tau((\cdot, \varepsilon_{-t}^\tau), y^\tau)$ is Lipschitz continuous with constant $D_{t,\tau} = B_\tau \sum_{l=t}^{\tau} C_{t,l}$. Finally we have

$$\begin{aligned} \sum_{\tau=t}^T D_{t,\tau} &= \sum_{\tau=t}^T \left(B_\tau \sum_{l=t}^{\tau} C_{t,l} \right) \leq \sum_{\tau=t}^T \left(B_\tau \sum_{l=t}^T C_{t,l} \right) \\ &= \left(\sum_{l=t}^T C_{t,l} \right) \sum_{\tau=t}^T B_\tau < +\infty. \end{aligned}$$

■ ■

Proof of Proposition 4. The proof consists in showing that assumptions ??, ??, 3, 4, 5, and 6 imply that the mapping z satisfies the properties of Proposition 2. Once this is established, the result follows from Proposition 2.

Part (1). Using Lemma 2, for any $t \geq 1$ any $(\varepsilon^t, y^{t-1}) \in \mathcal{E}^t \times Y^{t-1}$,

$$\left| \frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} \right| = \left| \hat{I}_t^t(\varepsilon^t, y^{t-1}) \right| \leq M_t \quad (15)$$

(where the inequality follows from assumption 5). Furthermore, for any $\tau \geq t \geq 1$ any $(\varepsilon^\tau, y^{\tau-1}) \in \mathcal{E}^\tau \times Y^{\tau-1}$,

$$\left| \frac{\partial z_\tau(\varepsilon^\tau; y^{\tau-1})}{\partial \varepsilon_t} \right| = \left| \hat{I}_t^t(\varepsilon^t, y^{t-1}) J_t^\tau(\varepsilon^\tau, y^{\tau-1}) \right| \leq M_t k_t^\tau \quad (16)$$

(where the inequality follows from assumptions 5 and 6).

The function $z_\tau((\cdot, \varepsilon_{-t}^\tau); y^{\tau-1}) : \mathcal{E}_t \rightarrow \Theta_\tau$ is thus $C_{t,\tau}$ -Lipschitz continuous with constant $C_{t,\tau}$ equal to the RHS of (15) if $\tau = t$ and to the RHS of (16) if $\tau > t$. That $\sum_{\tau=t}^T C_{t,\tau} < +\infty$ follows directly by assumption 6. ■

Proof of Proposition 5. The proof shows that under the assumptions in the proposition the mapping z satisfies the properties of Proposition 3. Once this is established, the result then follows from Proposition 3.

For simplicity, we prove the result here for $t = 1$. Similar arguments establish the result for any arbitrary $t > 1$. Part (1) establishes (Frechet) differentiability of $z(\cdot, \varepsilon_{-1}; y) : \mathcal{E}_1 \rightarrow \Theta$. Part (2) establishes equi-Lipschitz continuity.

Part (1): Differentiability. Frechet differentiability of $z(\cdot, \varepsilon_{-1}; y) : \mathcal{E}_1 \rightarrow \Theta$ is established using the implicit function theorem (IFT) for Banach spaces. We simplify notation by treating z as a function of ε_1 only. That is, fix (ε_{-1}, y) and drop it and let $\varepsilon := \varepsilon_1$. Furthermore, to simplify the derivation, assume that $\mathcal{E}_1 = \mathbb{R}$. It is immediate to see that the formula in (4) can be obtained from the results below by multiplying everything by $\hat{I}_1^1(\varepsilon_1) = dF_1^{-1}(\varepsilon_1)/d\varepsilon_1$.

For any $t > 1$, also let

$$f_t(\theta^{t-1}) := F_t^{-1}(\varepsilon_t | \theta^{t-1}, y^{t-1}),$$

where ε_t and y^{t-1} are fixed and hence dropped. Notice that, given $(f_t)_{t \geq 2}$, $z : \mathbb{R} \rightarrow \Theta$ is defined

implicitly as the solution to the system:

$$\begin{aligned} \varepsilon - \theta_1 &= 0 \\ f_2(\theta_1) - \theta_2 &= 0 \\ &\vdots \\ f_t(\theta^{t-1}) - \theta_t &= 0 \\ &\vdots \end{aligned}$$

Equivalently,

$$\begin{aligned} \varepsilon - z_1(\varepsilon) &= 0 \\ f_2(z_1(\varepsilon)) - z_2(\varepsilon) &= 0 \\ &\vdots \\ f_t(z^{t-1}(\varepsilon)) - z_t(\varepsilon) &= 0 \\ &\vdots \end{aligned}$$

Motivated by this observation we endow $\mathbb{R} \times \Theta$ with the norm $\|h\|_{\mathbb{R} \times \Theta} = \sup_t \delta^t |h_t|$, where $t = 0, 1, 2, \dots$ with $h_0 \in \mathbb{R}$ and $h_{-0} \in \Theta$. (That is, $\|\cdot\|_{\mathbb{R} \times \Theta}$ is just like the norm $\|\cdot\|_\delta$ except that t starts from 0 rather than 1.) This turns $\mathbb{R} \times \Theta$ into a Banach space (recall that $\|\theta\|_\delta < \infty$ for all $\theta \in \Theta$). Also note that while the results here assume that $\mathcal{E}_1 = \mathbb{R}$ and that $\Theta = \mathbb{R}_\delta^\infty$ where \mathbb{R}_δ^∞ is the set of all real sequences $x \in \mathbb{R}^\infty$ such that $\|x\|_\delta < \infty$, all the subsequent results extend to the case that $\mathcal{E}_1 \subsetneq \mathbb{R}$ and that $\Theta \subsetneq \mathbb{R}_\delta^\infty$.

We then define $T : \mathbb{R} \times \Theta \rightarrow \Theta$ by

$$T(\varepsilon, \theta) = \begin{pmatrix} \varepsilon - \theta_1 \\ f_2(\theta_1) - \theta_2 \\ \vdots \\ f_t(\theta^{t-1}) - \theta_t \\ \vdots \end{pmatrix}.$$

So z is now implicitly defined as the unique solution to

$$T(\varepsilon, z(\varepsilon)) = 0.$$

By the IFT for Banach spaces if

- (1) $T(\varepsilon', z(\varepsilon')) = 0$ all $\varepsilon' \in \mathbb{R}$;
- (2) T is Frechet differentiable;
- (3) $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is a Banach space isomorphism from Θ onto Θ ,

then z is Frechet differentiable in a neighborhood of ε (Above the bounded linear operator $DT(\varepsilon, z(\varepsilon))$ is the derivative of T at $(\varepsilon, z(\varepsilon))$.)

Condition 1 is true by construction. As for 2 and 3, define $f_1 : \mathbb{R} \rightarrow \Theta_1$ by setting $f_1(\varepsilon) = \varepsilon$. Let $f = (f_1, f_2, \dots) : \mathbb{R} \times \Theta \rightarrow \Theta$. We then have

$$T(\varepsilon, \theta) = f(\varepsilon, \theta) - \theta.$$

Throughout we adopt the convention of indexing an element x of $\mathbb{R} \times \Theta$ starting from zero so that $x = (x_0, x_1, \dots)$ with $x_0 \in \mathbb{R}$ and $x_{-0} \in \Theta$.

Simply writing out the definitions we have the following preliminary observation.

Lemma 5 *Let f be Frechet differentiable at (ε, θ) . Assume that each $f_t : \Theta^{t-1} \rightarrow \mathbb{R}$, $t \geq 2$, is differentiable. Then the derivative of f at (ε, θ) is the bounded linear operator $Df(\varepsilon, \theta) : \mathbb{R} \times \Theta \rightarrow \Theta$ defined by*

$$Df(\varepsilon, \theta)\alpha = \begin{pmatrix} \alpha_0 \\ \nabla f_2(\theta^1) \cdot \alpha_{-0}^1 \\ \vdots \\ \nabla f_t(\theta^{t-1}) \cdot \alpha_{-0}^{t-1} \\ \vdots \end{pmatrix}$$

where, for each $t \geq 1$, $\nabla f_t(\theta^{t-1})$ denotes the gradient of f_t at θ^{t-1} .

■

Proof. Proof of the Lemma. The proof is standard and hence omitted.

Note that the assumption that f is Frechet differentiable is guaranteed by Assumption 7, while that each $f_t : \Theta^{t-1} \rightarrow \mathbb{R}$, $t \geq 2$, is differentiable is guaranteed by Assumption 4. We then have the following result. ■

Lemma 6 *Assume that f is Frechet differentiable at (ε, θ) . Then T is Frechet differentiable at (ε, θ) , and its derivative at (ε, θ) is the bounded linear operator $DT(\varepsilon, \theta) : \mathbb{R} \times \Theta \rightarrow \Theta$ defined by*

$$DT(\varepsilon, \theta)h = Df(\varepsilon, \theta)h - h_{-0}.$$

Proof. Proof of the Lemma. Note first that T maps from a Banach space into a Banach space.

Since $T(\varepsilon, \theta) = f(\varepsilon, \theta) - \theta$, the obvious candidate for its Frechet derivative is

$$DT(\varepsilon, \theta)h = Df(\varepsilon, \theta)h - h_{-0},$$

which is clearly linear. In order to show that $DT(\varepsilon, \theta)$ is a bounded operator, note that for all $h \in \mathbb{R} \times \Theta$,

$$\begin{aligned} \|DT(\varepsilon, \theta)h\|_\delta &\leq \|Df(\varepsilon, \theta)h\|_\delta + \|h_{-0}\|_\delta \\ &\leq B \|h\|_{\mathbb{R} \times \Theta} + \|h\|_{\mathbb{R} \times \Theta} \\ &= (1 + B) \|h\|_{\mathbb{R} \times \Theta} \\ &< +\infty \end{aligned}$$

since $B < +\infty$ and $h \in \Theta$, which implies that $\|h\|_{\mathbb{R} \times \Theta} < +\infty$. The existence of the constant B follows from the fact that the Frechet derivative of f is by definition a bounded operator. It remains to show that $DT(\varepsilon, \theta)$ is the derivative of T at (ε, θ) . We have

$$\begin{aligned} 0 &\leq \lim_{h \rightarrow 0} \frac{\|T((\varepsilon, \theta) + h) - T(\varepsilon, \theta) - DT(\varepsilon, \theta)h\|_\delta}{\|h\|_{\mathbb{R} \times \Theta}} \\ &= \lim_{h \rightarrow 0} \frac{\|f((\varepsilon, \theta) + h) - \theta - h_{-0} - f(\varepsilon, \theta) + \theta - Df(\varepsilon, \theta)h + h_{-0}\|_\delta}{\|h\|_{\mathbb{R} \times \Theta}} \\ &= \lim_{h \rightarrow 0} \frac{\|f((\varepsilon, \theta) + h) - f(\varepsilon, \theta) - Df(\varepsilon, \theta)h\|_\delta}{\|h\|_{\mathbb{R} \times \Theta}} = 0, \end{aligned}$$

where the last equality follows by the assumed Frechet differentiability of f . ■

We now turn to Condition (3). Let \hat{J}_s^t denote the (inverse) impulse response functions, as defined in the main text and note that, with the notation used in this proof, for any $\tau < t$,

$$\frac{\partial f_t(z^{t-1}(\varepsilon))}{\partial \theta_\tau} := \frac{\partial F_t^{-1}(\varepsilon_t | z^{t-1}(\varepsilon), y^{t-1})}{\partial \theta_\tau} = \hat{I}_\tau^t((\varepsilon, \varepsilon^{t-1}), y^{t-1}).$$

Because (ε_{-1}, y) is fixed, we then let $\hat{I}_\tau^t(\varepsilon) := \hat{I}_\tau^t((\varepsilon, \varepsilon^{t-1}), y^{t-1})$ and

$$\hat{J}_\tau^t(\varepsilon) = \sum_{\substack{K \in \mathbb{N}, \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K \hat{I}_{l_{k-1}}^{l_k}(\varepsilon).$$

with $\hat{J}_t^t(\varepsilon) \equiv 1$. We then have the following result. ■

Lemma 7 *Let f be Frechet differentiable. Fix $\varepsilon \in \mathbb{R}$. If*

$$\lim_{t \rightarrow T} \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t(\varepsilon) \right| < \infty,$$

then $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is a Banach space isomorphism from Θ onto Θ .

Proof. Proof of the Lemma. Fix ε and then drop it from all \hat{I} and \hat{J} functions. By Lemma 6, T is Frechet differentiable at $(\varepsilon, z(\varepsilon))$ with derivative $DT(\varepsilon, z(\varepsilon))\alpha$. We show first that $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is onto. Fix $u \in \Theta$. By Lemmas 5 and 6 we have

$$DT(\varepsilon, z(\varepsilon))(0, \theta) = Df(\varepsilon, z(\varepsilon))(0, \theta) - \theta = \begin{pmatrix} 0 - \theta_1 \\ \hat{I}_1^2 \theta_1 - \theta_2 \\ \vdots \\ \sum_{s=1}^{t-1} \hat{I}_s^t \theta_s - \theta_t \\ \vdots \end{pmatrix}.$$

Hence to find the preimage of u , we define θ recursively by

$$\begin{aligned} \theta_1 &= -u_1, \\ \theta_t &= -u_t + \sum_{s=1}^{t-1} \hat{I}_s^t \theta_s \text{ for } t > 1. \end{aligned} \tag{17}$$

It remains to show that if $u \in \Theta$, that is, if $\|u\|_\delta < \infty$, then θ so constructed has $\|\theta\|_\delta < \infty$, and hence $\theta \in \Theta$. We first use the above recursion to show by induction on t that

$$\theta_t = - \sum_{l=1}^t \hat{J}_l^t u_l.$$

For $t = 1$ this is clearly true. Suppose it holds for all $\tau < t$. We have

$$\theta_t = -u_t + \sum_{k=1}^{t-1} \hat{I}_k^t \theta_k = -\hat{J}_t^t u_t - \sum_{k=1}^{t-1} \hat{I}_k^t \sum_{l=1}^k \hat{J}_l^k u_l = - \sum_{k=1}^t \hat{J}_k^t u_k,$$

which establishes the claim. Hence for all t ,

$$\delta^t |\theta_t| \leq \sum_{k=1}^t \delta^{t-k} \left| \hat{J}_k^t \right| \delta^k |u_k| \leq \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t \right| \|u\|_\delta,$$

implying that

$$\|\theta\|_\delta \leq \|u\|_\delta \sup_t \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t \right|.$$

We claim that if $\lim_{t \rightarrow \infty} \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t \right| < \infty$, then $\sup_t \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t \right| < \infty$ and the right-hand side is finite. This can be seen as follows. Since f is Frechet differentiable, $Df(\varepsilon, z(\varepsilon))$ is a bounded operator. Thus there exists $B < \infty$ such that for all $\alpha \in \mathbb{R} \times \Theta$, all $t \geq 2$,

$$B \|\alpha\|_{\mathbb{R} \times \Theta} \geq \|Df(\varepsilon, z(\varepsilon))\alpha\|_\delta \geq \delta^t \left| \sum_{s=1}^{t-1} \hat{I}_s^t \alpha_s \right|.$$

Now choose α such that $\exists! \tau \geq 2 : \alpha_\tau = 1$ and $\alpha_s = 0$ for $s \neq \tau$. Then the above inequality implies that for all $t > \tau$,

$$\left| \hat{I}_\tau^t \right| \leq B \delta^{\tau-t} < \infty.$$

By implication $\left| \hat{J}_\tau^t \right| < \infty$ for all t and all $\tau \leq t$. Thus for all t the sum $\sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t \right|$ is finite. Hence the sup can be infinite only if it is approached as $t \rightarrow \infty$. But the limit is finite by Assumption 9.

We finish the proof by showing that $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is an isomorphism. Recall that every one-to-one bounded linear operator from a Banach space onto a Banach space is an isomorphism (see, e.g., Corollary 1.6.6. in Megginson's "An Introduction to Banach Space Theory"). Above it was shown that $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is onto, and it is obviously bounded. Hence it suffices to show that $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is one-to-one. Let θ, θ' be such that $DT(\varepsilon, z(\varepsilon))(0, \theta) = DT(\varepsilon, z(\varepsilon))(0, \theta')$. By inspection of the first line in the formula for $DT(\varepsilon, z(\varepsilon))(0, \theta)$ above, we have $\theta_1 = \theta'_1$. But then the second line gives $\theta_2 = \theta'_2$, and so on. Thus $\theta = \theta'$ and hence $\theta \mapsto DT(\varepsilon, z(\varepsilon))(0, \theta)$ is one-to-one as desired. ■

The property in Lemma 7 that $\lim_{t \rightarrow T} \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t(\varepsilon) \right| < \infty$ is guaranteed by Assumption 8. Lemmas 5-7 together with the IFT thus establish Frechet differentiability of $z(\cdot, \varepsilon_{-1}; y) : \mathcal{E}_1 \rightarrow \Theta$.

Part (2): Lipschitz continuity. Next we address Lipschitz continuity of z . Let $\hat{J}_1 := (\hat{J}_1^1, \hat{J}_1^2, \dots, \hat{J}_1^t, \dots)$. (Recall that $\hat{J}_1^1 \equiv 1$.) ■

Lemma 8 *Let $\varepsilon \in \mathbb{R}$. Under the assumptions in the Proposition, the Frechet derivative of z at ε , $z'(\varepsilon) : \mathbb{R} \rightarrow \Theta$, is given by*

$$z'(\varepsilon)\alpha = \hat{J}_1(\varepsilon)\alpha.$$

Proof. Proof of the Lemma. Fix ε . Since z is Frechet differentiable at ε , there exists a bounded linear operator $z'(\varepsilon) : \mathbb{R} \rightarrow \Theta$ such that

$$\lim_{\alpha \rightarrow 0} \frac{\|z(\varepsilon + \alpha) - z(\varepsilon) - z'(\varepsilon)\alpha\|_\delta}{|\alpha|} = 0.$$

We will show by induction on t that $z'(\varepsilon)\alpha = \hat{J}_1(\varepsilon)\alpha$. Note first that the operator $\alpha \mapsto \hat{J}_1(\varepsilon)\alpha$ is obviously linear. It is bounded, since for all $\alpha \in \mathbb{R}$,

$$\left\| \hat{J}_1(\varepsilon)\alpha \right\|_\delta = \sup_t \delta^t \left| \hat{J}_1^t(\varepsilon)\alpha \right| \leq |\alpha| \sup_t \sum_{\tau=0}^{t-1} \delta^\tau \left| \hat{J}_{t-\tau}^t(\varepsilon) \right| < +\infty$$

where the last inequality follows from the proof in the previous lemma. Now, by induction, suppose that, for all $\tau < t$, $z_\tau : \mathbb{R} \rightarrow \Theta_\tau$ is differentiable at ε with $z'_\tau(\varepsilon) = \hat{J}_1^\tau(\varepsilon)$. Note that for $t = 2$ this is trivially true since $f_1(\varepsilon)$ is linear. Next note that $z^{t-1} := (z_1, \dots, z_{t-1}) : \mathbb{R} \rightarrow \Theta^{t-1}$ is differentiable at ε with gradient $\nabla z^{t-1}(\varepsilon) = (\hat{J}_1^1(\varepsilon), \dots, \hat{J}_1^{t-1}(\varepsilon))$. Since

$$0 = \lim_{\alpha \rightarrow 0} \frac{\|z(\varepsilon + \alpha) - z(\varepsilon) - z'(\varepsilon)\alpha\|_\delta}{|\alpha|} \geq \delta^t \lim_{\alpha \rightarrow 0} \frac{|z_t(\varepsilon + \alpha) - z_t(\varepsilon) - z'_t(\varepsilon)\alpha|}{|\alpha|} \geq 0,$$

z_t is differentiable at ε with derivative $z'_t(\varepsilon)$. Furthermore, f_t is differentiable by Assumption 4 so that by the chain rule

$$\begin{aligned} z'_t(\varepsilon) &= \lim_{\alpha \rightarrow 0} \frac{|z_t(\varepsilon + \alpha) - z_t(\varepsilon)|}{|\alpha|} \\ &= \lim_{\alpha \rightarrow 0} \frac{|f_t(z^{t-1}(\varepsilon + \alpha)) - f_t(z^{t-1}(\varepsilon))|}{|\alpha|} \\ &= \sum_{\tau=1}^{t-1} \frac{\partial f_t(z^{t-1}(\varepsilon))}{\partial \theta_\tau} \hat{J}_1^\tau(\varepsilon) \\ &= \sum_{\tau=1}^{t-1} \hat{I}_\tau^t(\varepsilon) \hat{J}_1^\tau(\varepsilon) \\ &= \hat{J}_1^t(\varepsilon), \end{aligned}$$

where the last two equalities follow by the definitions of \hat{I} and \hat{J} . This establishes the inductive step and concludes the proof. ■

Assumption 9 then guarantees that there exists $K_1 < \infty$ such that for all ε ,

$$\left\| \hat{J}_1(\varepsilon) \right\|_\delta \leq K_1.$$

That z is Lipschitz continuous then follows from Lemma 8 together with Proposition 2 on p. 176 of Luenberger (1969). ■

Proof of Proposition 6. The initial steps of the proof are in the main text. Here we simply prove that, under the assumptions in the proposition, the formula in (4) reduces to the one in (13).

First note that, under the assumptions in the proposition, from the implicit function theorem

applied to the identity

$$F_s^{-1}(F_s(\theta_s|\theta^{s-1}, y^{s-1})|\theta^{s-1}, y^{s-1}) = \theta_s$$

for any $\theta_s \in \Theta_s$, $s > t$, any (θ^{s-1}, y^{s-1}) , $F_s(\theta_s|\theta^{s-1}, y^{s-1})$ is differentiable in θ^{s-1} . Next note that the implicit function theorem applied to the identity

$$F_s(F_s^{-1}(\varepsilon_s|\theta^{s-1}, y^{s-1})|\theta^{s-1}, y^{s-1}) = \varepsilon_s$$

implies that, for any $t < s$, ε_s , (θ^{s-1}, y^{s-1}) ,

$$\frac{\partial F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^s)}{\partial \theta_t} = - \frac{\frac{\partial F_s(\theta_s | \theta^{s-1}, y^{s-1})}{\partial \theta_t} \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})}}{f_s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})}}.$$

It follows that

$$\begin{aligned} \hat{I}_t^s(\varepsilon^s, y^{s-1}) &\equiv - \frac{\partial F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^s)}{\partial \theta_t} \Big|_{\theta^{s-1} = z^{s-1}(\varepsilon^{s-1}, y^{s-2})} \\ &= - \frac{\partial F_s(\theta_s | \theta^{s-1}, y^{s-1}) / \partial \theta_t}{f_s(\theta_s | \theta^{s-1}, y^{s-1})} \Big|_{\theta^{s-1} = z^{s-1}(\varepsilon^{s-1}, y^{s-2})} \\ &\equiv I_t^s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta^s = z^s(\varepsilon^s, y^{s-1})}. \end{aligned}$$

and hence that

$$\hat{J}_t^s(\varepsilon^s, y^{s-1}) = J_t^s(z^s(\varepsilon^s; y^{s-1}), y^{s-1}).$$

By the definition of independent-shock representation, we then have that

$$\begin{aligned} \frac{\partial V^\Omega(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} &= \\ &\mathbb{E}^{\hat{\mu}_{[\hat{\Omega}]|\varepsilon^t, \varepsilon^{t-1}, y^{t-1}}} \left[\sum_{\tau=t}^T J_t^\tau(z^\tau(\tilde{\varepsilon}^\tau; y^{\tau-1}), y^{\tau-1}) \frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_\tau} \right] \\ &= \mathbb{E}^{\mu[\Omega]|z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1}} \left[\sum_{\tau=t}^T J_t^\tau(\tilde{\theta}^\tau, y^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_\tau} \right], \end{aligned}$$

which is the same formula as in (13). ■

Proof of Proposition 8. First note that, because the environment is quasilinear and Markov (in the sense of Definition 1), whether or not a bidder finds it optimal to report truthfully at any given private history $(\theta_i^t, m_i^{t-1}, x_i^{t-1}, p_i^{t-1})$ depends only on his current type $\theta_{i,t}$, the history of messages he sent m_i^t and the number of times he consumed in the past, $\sum_{s=1}^{t-1} x_{i,s}$.

Second note that, starting from period two onwards, and irrespective of whether the period-one announcements were truthful or not, the allocation rule χ of Proposition 7 coincides with the one that maximizes the sum of the bidders' continuation payoffs and of the seller's adjusted continuation payoff, where the latter is obtained by replacing each cost $c_{i,t}$ with the cost

$$\hat{c}_{i,t} := c_{i,t} - \frac{1 - F_{i,1}(m_{i,1})}{f_1(m_{i,1})}.$$

Furthermore, because each player's continuation payoff (including the seller's adjusted continuation payoff) depends only on his own types and decisions, it is immediate to see that truth-telling at any period- t history, $t \geq 2$, can be obtained by using for example the Team Mechanism payments of Athey and Segal:

$$p_{i,t}^*(\theta^t) = - \sum_{\substack{j=0, \dots, N \\ j \neq i}} u_{j,t}(\chi^{*t}(\theta^t), \theta^t) \quad (18)$$

for all $i = 1, \dots, N$, all $t \geq 2$, where $u_{j,t}(\chi^{*t}(\theta^t), \theta^t) = \theta_{j,t} \chi_{j,t}^*(\theta^t)$ if $j \neq 0$ and $u_{0,t}(\chi^{*t}(\theta^t), \theta^t) = - \sum_{i=1}^N \hat{c}_{i,t} \chi_{i,t}^*(\theta^t)$ if $j = 0$.

To establish incentive compatibility at $t = 1$, we use the analog of the weak monotonicity result of Proposition 8 in PST (note that, while the result there is for the case where T is finite, it is immediate to see that the same arguments apply to the case $T = \infty$). We start with the following result.

Lemma 9 *Let χ^* be the allocation rule of Proposition 7. Then for all i and $\theta_{i,1}$*

$$\mathbb{E}^{\lambda[\chi^*]|\theta_{i,1}, m_{i,1}} \left[\sum_{t=1}^{\infty} \delta^t \chi_{i,t}^* \left(\left(m_{i,1}, \tilde{\theta}_{i,-1}^t \right), \tilde{\theta}_{-i}^t \right) \right]$$

is nondecreasing $m_{i,1}$.

Proof of the Lemma. To prove the lemma, it is convenient to note that, without loss, we can think of the valuation process of bidder i as being generated as follows. First, $\theta_{i,1}$ is drawn according to $F_{i,1}$. Next, a sequence $\eta_i = (\eta_{i,k})_{k=1}^{\infty}$ is drawn according to the product measure $\times_{k=1}^{\infty} G_i(\cdot|1, k)$. Then we take the innovation term following the k^{th} time i wins the object to be $\eta_{i,k}$. Now, while the distribution of θ (and even that of ε) depends on $m_{i,1}$ through χ , the distribution of $\eta = (\eta_1, \dots, \eta_N)$ is independent of $m_{i,1}$.

Fix i , $\theta_{i,1}$, $m_{i,1}$ and $m'_{i,1} > m_{i,1}$. We establish the result by showing, by induction on k , that for any realizations of $\theta_{-i,1}$ and of η and for all k , if bidder i 's k^{th} win comes in period t given $m_{i,1}$, then it comes at some $s \leq t$ given $m'_{i,1}$. To this end, fix an arbitrary realizations $\theta_{-i,1}$ and η . Let $k = 1$ and suppose to the contrary that given $m_{i,1}$ bidder i 's first win comes in period t , and given

$m'_{i,1}$ it comes in period $s > t$. (If bidder i never wins given $m_{i,1}$, then $t = +\infty$ and the claim is trivially true.) Note that, since (θ_1, η) is fixed, the allocations are the same in both cases in periods $1, \dots, t-1$. Hence also the other bidders' realized Gittins indices in period t are the same in both cases. Since i wins at t , his Gittins index $\gamma_{i,1}(m_{i,1})$ must be the highest. But given the hazard rate assumption and the form of the process for $\theta_{i,t}$, the index $\gamma_{i,1}(m_{i,1})$ is nondecreasing in $m_{i,1}$. Hence i must also win at t with $m'_{i,1}$, a contradiction.

Next, suppose that the claim is true for some $k \geq 1$. Assume to the contrary that the $k+1^{\text{th}}$ win given $m_{i,1}$ comes in period t , and given $m_{i,1}$ it comes in period $s > t$. Note that during periods $1, \dots, t-1$ bidder i wins k times in both cases. Also, since (θ_1, η) is fixed, the remaining $t-1-k$ wins go to the same bidders in both cases. In particular, since the Gittins index only depends on the first message, the most recent valuation, and the number of times the bidder has won the object, it must be that the other bidders' period- t Gittins indices are the same in both cases. Now, since $(\theta_{i,1}, \eta_i)$ is fixed, bidder i 's realized period t valuation $\theta_{i,t}$ is the same in both cases since in each case he won k times prior to period t . Thus we know that his period- t Gittins indices in the two cases satisfy

$$\begin{aligned} \gamma_{i,t}(m'_{i,1}, \theta_{i,t}, k) &= \max_{\tau} \mathbb{E} \left[\frac{\sum_{s=t}^{\tau} \delta^{s-t} (\tilde{\theta}_{i,s} - c_{i,s})}{\sum_{s=t}^{\tau} \delta^{s-t}} \mid \theta_{i,t}, k \right] - \frac{1 - F_{i,1}(m'_{i,1})}{f_1(m'_{i,1})} \\ &\geq \max_{\tau} \mathbb{E} \left[\frac{\sum_{s=t}^{\tau} \delta^{s-t} (\tilde{\theta}_{i,s} - c_{i,s})}{\sum_{s=t}^{\tau} \delta^{s-t}} \mid \theta_{i,t}, k \right] - \frac{1 - F_{i,1}(m_{i,1})}{f_1(m_{i,1})} \\ &= \gamma_{i,t}(m'_{i,1}, \theta_{i,t}, k), \end{aligned}$$

where the inequality is by the hazard rate assumption. Thus, if i wins at t given $m_{i,1}$, he must also win at t given $m'_{i,1}$. But this contradicts $s > t$. ■

Now, consider the payment scheme such that $p_{i,t}^*$ is as in (18) for all $t > 1$, while for $t = 1$

$$p_{i,1}^*(\theta_{i,1}) \equiv \mathbb{E}^{\lambda[\chi^*]|\theta_{i,1}} \left[\sum_{t=1}^{\infty} \tilde{\theta}_{i,t} \chi_{i,t}^*(\tilde{\theta}^t) - \sum_{t=2}^{\infty} p_{i,t}^*(\tilde{\theta}^t) \right] - \int_{\underline{\theta}_{i,1}}^{\theta_{i,1}} D_i^{[\chi^*]}(z, z) dz,$$

where

$$D_i^{[\chi^*]}(z, z) \equiv \mathbb{E}^{\lambda[\chi^*]|\theta_{i,1}, m_{i,1}} \left[\sum_{t=1}^{\infty} \delta^t \chi_{i,t} \left(\left(m_{i,1}, \tilde{\theta}_{i,-1}^t \right), \tilde{\theta}_{-i}^t \right) \right]$$

From the same arguments as in Proposition 8 in PST, one can then verify that, under the mechanism $\Omega^* = [\chi^*, p^*]$, the agent finds it optimal to report truthfully at all histories and to participate in period one. One can also check that each bidder's expected payoff when his period-one type is the lowest, i.e., at $\theta_{i,1} = \underline{\theta}_{i,1}$ is zero. That the mechanism Ω is optimal then follows from the same

arguments that establish Proposition 5 in PST. ■