Dynamic Mechanism Design: A Myersonian Approach – Supplementary Material

Alessandro Pavan
Northwestern University

Ilya Segal
Stanford University

Juuso Toikka
MIT

August 9, 2013

This document contains additional results and an omitted proof for the manuscript Dynamic Mechanism Design: A Myersonian Approach. Section S.1 contains the proof of the one-stage-deviation principle used in the proof of Theorem 2 in the main text. Section S.2 contains a detailed analysis of Example 5 from the main text. Section S.3 establishes conditions under which the allocation rule maximizing expected virtual surplus distorts allocations downwards compared to the first-best rule. Section S.4 discusses distortions in discrete type models using the language of impulse responses. Section S.5 considers optimal mechanisms in some classes of non-Markov environments.

All numbered items (i.e., sections, definitions, results, and equations) in this document contain the prefix S. Any numbered reference without a prefix refers to an item in the main text. Please refer to the main text for notation and definitions.

S.1 Proof of the One-Stage-Deviation Principle

The proof of Theorem 3 in the Appendix to the main text makes use of the following lemma:

Lemma S.1 *Suppose the environment is regular and Markov. Fix an allocation rule $\chi \in \mathcal{X}$ with belief system $\Gamma \in \mathbf{\Gamma}(\chi)$, and define the transfer rule ψ by (7). If a one-stage deviation from strong truthtelling is not profitable at any information set, then arbitrary deviations from strong truthtelling are not profitable at any information set.*

Proof. We first establish the result for the case where, for all $i = 1, \dots, n$, there exist a constant L_i^ψ and sequences $(K_{it})_{t=0}^\infty$ and $(M_i^\psi)_{t=0}^\infty$, with $L_i^\psi, ||M_i^\psi||, \|K_i\| < \infty$, such that for all $t \geq 0$ and $\theta^t \in \Theta^t$, (i) $\psi_{it}\left(\theta^t\right) \leq K_{it}$, and (ii) $\left|\psi_{it}\left(\theta^t\right)\right| \leq L_i^\psi \left|\theta_{it}\right| + M_{it}^\psi$. We then conclude the proof by showing that the transfer rule defined by (7) satisfies these bounds in any regular Markov environment.[1]

[1] Heurisitically, bound (i) ensures that for any possible deviation strategy, the net present value of the payments received by the agent from period T onwards is either small or negative when T is large. Bound (ii) together with the

The usual backward-induction argument establishes that, if no one-step deviation from strong truthtelling is profitable, then all *finite-stage* deviations from strong truthtelling are unprofitable. To establish that infinite-stage deviations are also unprofitable, suppose in negation that there exists an agent i and a period-s history $h_{is} = \left(\theta_i^s, \hat{\theta}_i^{s-1}, x_i^{s-1}\right)$ (not necessarily truthful for agent i) such that agent i raises his expected payoff at h_{is} by some $\varepsilon > 0$ by deviating from strong truthtelling to some strategy $\hat{\sigma}_i$, assuming that all other agents report truthfully. (All the expectations below will be conditional on history h_{is}.)

We then show that there exists some finite $T > s$ such that reversion from $\hat{\sigma}_i$ to strong truthtelling starting in period T cannot reduce the agent's expected payoff by more than $\varepsilon/2$. For this purpose, note first that under the bounds in part (ii) of the definition of Markov environments, the agent's time-s expectation of his type in each period $t \geq s$ under *any* strategy is bounded as follows:

$$\mathbb{E}[|\tilde{\theta}_{it}|] \leq \phi_i^{t-s} |\theta_{is}| + \textstyle\sum_{\tau=s}^{t} \phi^{t-\tau} E_{i\tau} \equiv S_{it}.$$

Note that letting $S_{it} \equiv 0$ for $t < s$ we have

$$\begin{aligned} \|S_i\| &= \delta^s \textstyle\sum_{t=s}^{\infty} (\delta\phi_i)^{t-s} |\theta_{is}| + \sum_{t=s}^{\infty} \delta^t \sum_{\tau=s}^{t} \phi_i^{t-\tau} |E_{i\tau}| \\ &= \frac{\delta^s}{1-\delta\phi_i} |\theta_{is}| + \textstyle\sum_{\tau=s}^{\infty} \delta^\tau |E_{i\tau}| \sum_{t=\tau}^{\infty} (\delta\phi_i)^{t-\tau} \\ &\leq \frac{\delta^s}{1-\delta\phi_i} |\theta_{is}| + \frac{1}{1-\delta\phi_i} \|E_i\| < \infty. \end{aligned}$$

Hence, by Condition U-SPR, the expected present value (EPV) of non-monetary utility flows starting in period T under any strategy is bounded in absolute value by $\sum_{t=T}^{\infty} \delta^t [L_i S_{it} + M_{it}]$, and so reversion from strategy $\hat{\sigma}_i$ to strong truthtelling in period T can reduce this EPV by at most twice this amount. Likewise, condition (ii) on payments implies that the EPV of payment flows from strong truthtelling starting in period T is bounded in absolute value by $\sum_{t=T}^{\infty} \delta^t \left[L_i^{\psi} S_{it} + M_{it}^{\psi}\right]$. As for the EPV of payment flows from period T onwards under strategy $\hat{\sigma}_i$, by condition (i) it is bounded from above by $\sum_{t=T}^{\infty} \delta^t K_{it}$. Adding up, we see that the total expected loss due to the reversion in period T does not exceed

$$\sum_{t=T}^{\infty} \delta^t \left[\left(2L_i + L_i^{\psi}\right) S_{it} + 2M_{it} + M_{it}^{\psi} + K_{it}\right].$$

Since $L_i, L_i^{\psi}, \|S_i\|, \|M_i\|, ||M_i^{\psi}||, \|K_i\| < \infty$, the loss converges to zero as $T \to \infty$, and so it falls below $\varepsilon/2$ for some T large enough. But then the finite-stage deviation to $\hat{\sigma}_i$ between periods s and T is profitable, a contradiction.

bounds from part (ii) of the definition of Markov environments ensure that, if the agent reverts to strong truthtelling in period T, the expected payments received after the reversion are small in absolute value. Thus, the reversion will not cost the agent much in terms of expected payments. Under Condition U-SPR, the same is true for the agent's expected non-monetary utility. Hence, when considering potential profitable deviations from strong truthtelling, it suffices to check only those deviations that revert to strong truthtelling at some finite period T.

We then finish the proof of Lemma S.1 by showing that the transfer rule ψ given by (7) satisfies bounds (i) and (ii) from above. Note that it suffices to establish the bounds separately for each of the three terms in (7) as summing up the bounds then gives the desired bound on the flow payment.

We begin with condition (ii). The first term in (7) satisfies this bound by (20) and $||\hat{\theta}_i|| < \infty$. For the second term, using (20) and the bounds from part (ii) of the definition of Markov environments,

$$\begin{aligned} \left|\delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[Q^{\chi,\Gamma}_{i,t+1}(\tilde{\theta}^t,\tilde{\theta}_{i,t+1})\right]\right| &\leq A_iB_i\left(\mathbb{E}^{F_{i,t+1}\left(\theta_{it},\chi_i^{t-1}(\theta^{t-1})\right)}\left[\left|\tilde{\theta}_{i,t+1}\right|\right]+\left|\hat{\theta}_{i,t+1}\right|\right) \\ &\leq A_iB_i\left(\phi_i\left|\theta_{it}\right|+E_{it}+\left|\hat{\theta}_{i,t+1}\right|\right), \end{aligned} \tag{S.1}$$

which satisfies the bound since $||E_i||, ||\hat{\theta}_i|| < \infty$. The third term satisfies the bound by U-SPR.

Next, we turn to condition (i). Using (6) and the fact that the payments (7) satisfy ICFOC by Theorem 2, we can write

$$Q^{\chi,\Gamma}_{i\tau}\left(\theta^{\tau-1},\theta_{i\tau}\right) = V^{\langle\chi,\Psi\rangle,\Gamma}_{i\tau}\left(\theta^{\tau-1},\theta_{i\tau}\right) - V^{\langle\chi,\Psi\rangle,\Gamma}_{i\tau}(\theta^{\tau-1},\hat{\theta}_{i\tau})$$

for $\tau = t, t+1$. Substituting in (7), omitting superscripts on V and Q to save space, and noting that

$$\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[V_{i,t+1}\left(\tilde{\theta}^t,\tilde{\theta}_{i,t+1}\right)\right] = V_{it}\left(\theta^{t-1},\theta_{it}\right),$$

we have

$$\begin{aligned} \psi_{it}\left(\theta^t\right) = \delta^{-t}&\left(\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[V_{i,t+1}(\theta_i^t,\tilde{\theta}^t_{-i},\hat{\theta}_{i,t+1})\right]-V_{it}(\theta^{t-1},\hat{\theta}_{it})\right) \\ &- \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[u_{it}\left(\theta_{it},\tilde{\theta}^t_{-i},\chi^t(\tilde{\theta}^t)\right)\right]. \end{aligned} \tag{S.2}$$

Now, the assumption that no agent has a profitable one-step deviation from strong truthtelling implies in particular that, given true history $(\theta^{t-1},\hat{\theta}_{it})$, agent i does not gain by reporting θ_{it} followed by strong truthtelling, or

$$V_{it}(\theta^{t-1},\hat{\theta}_{it}) \geq \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[\begin{array}{c} V_{i,t+1}\left(\theta_i^t,\tilde{\theta}^t_{-i},\tilde{\theta}_{i,t+1}\right)+\delta^t u_{it}(\hat{\theta}_{it},\tilde{\theta}^t_{-i},\chi^t(\theta_i^t,\tilde{\theta}^t_{-i})) \\ -\delta^t u_{it}(\theta_{it},\tilde{\theta}^t_{-i},\chi^t(\theta_i^t,\tilde{\theta}^t_{-i})) \end{array}\right],$$

where we have used the fact that in a Markov environment, once $\tilde{\theta}_{i,t+1}$ is realized following the misreport, the only payoff effect of the past true type being $\hat{\theta}_{it}$ rather than θ_{it} is through period-t utility. Using this inequality, and also noting that the distribution of $\tilde{\theta}^t_{-i}$ under the probability measure $\lambda_i\left[\chi,\Gamma\right]|\theta^{t-1},\theta_{it}$ does not depend on θ_{it}, so that we can replace the measure in (S.2) with $\lambda_i\left[\chi,\Gamma\right]|\theta^{t-1},\hat{\theta}_{it}$, the payments (S.2) are bounded above as follows:

$$\begin{aligned} \psi_{it}\left(\theta^t\right) \leq\ & \delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[V_{i,t+1}(\theta_i^t,\tilde{\theta}^t_{-i},\hat{\theta}_{i,t+1})-V_{i,t+1}(\theta_i^t,\tilde{\theta}^t_{-i},\tilde{\theta}_{i,t+1})\right] \\ &- \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[u_{it}(\hat{\theta}_{it},\tilde{\theta}^t_{-i},\chi^t(\theta_i^t,\tilde{\theta}^t_{-i}))\right]. \end{aligned}$$

Finally, use $\text{ICFOC}_{i,t+1}$ and (6) to rewrite the bound as

$$\psi_{it}\left(\theta^t\right) \leq -\delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[Q_{i,t+1}(\theta_i^t,\tilde{\theta}_{-i}^t,\tilde{\theta}_{i,t+1})\right] - \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[u_{it}(\hat{\theta}_{it},\tilde{\theta}_{-i}^t,\chi^t(\theta_i^t,\tilde{\theta}_{-i}^t))\right],$$

and note that both terms on the right are bounded from above by a term of a finite-norm sequence: the first term by (S.1) and $||E_i||, ||\hat{\theta}_i|| < \infty$, and the second by Condition U-SPR and $||\hat{\theta}_i|| < \infty$. ■

S.2 Details for Example 5

We provide here the details omitted in the discussion of Example 5 in the main text. We verify first that the environment is regular and Markov. Since X_t is bounded, the buyer's payoff function satisfies U-D, U-ELC, and U-SPR. The assumptions on the process imply that $0 \leq \theta_t \leq 1$ for all t, and hence the bounds from part (ii) of the definition of Markov environments are satisfied as well. For F-BIR, note that a state representation for the process is obtained by putting $z_t(\theta^{t-1}, x^{t-1}, \varepsilon_t) = \phi(\theta_{t-1}) + \varepsilon_t$. Hence, by the chain rule and formula (4), $\partial Z_{(s),t}(\theta^s, \varepsilon^t)/\partial\theta_s$ exists, and the impulse responses take the form $I_{(s),t}(\theta^t) = \prod_{\tau=s}^{t-1} \phi'(\theta_\tau)$. Furthermore, we have $\left|\partial Z_{(s),t}/\partial\theta_s\right| \leq b^{t-s}$ for all $t \geq s \geq 0$. Thus, recalling that $b \geq 1$, we may put $C_{(s),t}(\varepsilon) \equiv b^t$ for all s, t, and ε to obtain the bounding functions $C_{(s)}$, since then $\left\|C_{(s)}\right\| \leq \sum_{t=0}^{\infty}(\delta b)^t < \infty$. Thus F-BIR is satisfied.

The expected virtual surplus from allocation rule χ takes the form

$$\mathbb{E}^{\lambda}\left[\sum_{t=0}^{T}\delta^t\left(\chi_t(\tilde{\theta}^t)\left(a + \tilde{\theta}_t - \frac{1}{\eta_0(\tilde{\theta}_0)}\prod_{\tau=0}^{t-1}\phi'(\tilde{\theta}_\tau)\right) - \frac{(\chi_t(\tilde{\theta}^t))^2}{2}\right)\right],$$

and hence it can be maximized pointwise in time t and type history θ^t to obtain the allocation rule χ given in the main text. Since the kernels satisfy F-AUT, we can use ex-post monotonicity from Corollary 1 to verify that χ is implementable. Here, it requires that, for all $t = 0, \ldots, T$ and $\theta \in \Theta$,

$$\sum_{s=t}^{T}\delta^{s-t}\prod_{\tau=t}^{s-1}\phi'(\theta_\tau)\chi_s(\hat{\theta}_t, \theta_{-t}^s)$$

be nondecreasing in $\hat{\theta}_t$.

We restrict attention to the case where χ prescribes an interior solution at every history by assuming that $b^T \leq a \min_{\theta_0} \eta_0(\theta_0)$. (Note that if $b > 1$, this requires T to be finite. Otherwise, it suffices to assume that the hazard rate η_0 be bounded away from 0 and take a large enough.) Then it is enough to verify that (i) $\partial\chi_t(\theta)/\partial\theta_t \geq 0$ for all $0 \leq t \leq T$ and $\theta \in \Theta^\infty$, and (ii) $\partial\chi_s(\theta)/\partial\theta_t \geq -\frac{1-b\delta}{b\delta}\partial\chi_t(\theta)/\partial\theta_t$ for all $0 \leq t < s \leq T$, and $\theta \in \Theta^\infty$.[2] If η_0 is nondecreasing, then $\partial\chi_t(\theta)/\partial\theta_t \geq 1$ for all t and $\theta \in \Theta^\infty$ so Condition (i) holds. As for (ii), note that for $s > t$ we have

$$\frac{\partial\chi_s(\theta)}{\partial\theta_t} = \frac{\partial}{\partial\theta_t}\left[a + \theta_s - \frac{1}{\eta_0(\theta_0)}\prod_{\tau=0}^{s-1}\phi'(\theta_\tau)\right] \geq -\frac{b^{s-1}}{\eta_0(\theta_0)}\phi''(\theta_t).$$

[2] To see this, note that since $0 \leq \phi' \leq b < \frac{1}{\delta}$, conditions (i) and (ii) imply

$$\frac{\partial}{\partial\hat{\theta}_t}\left[\sum_{s=t}^{T}\delta^{s-t}\prod_{\tau=t}^{s-1}\phi'(\theta_\tau)\chi_s(\hat{\theta}_t,\theta_{-t}^s)\right] \geq \left(1 - \frac{1-b\delta}{b\delta}\sum_{s=t+1}^{T}\delta^{s-t}\prod_{\tau=t}^{s-1}\phi'(\theta_\tau)\right)\frac{\partial\chi_t(\hat{\theta}_t,\theta^{t-1})}{\partial\hat{\theta}_t} \geq (b\delta)^T\frac{\partial\chi_t(\hat{\theta}_t,\theta^{t-1})}{\partial\hat{\theta}_t} \geq 0.$$

Hence, recalling that $b^T \leq a \min_{\theta_0} \eta_0(\theta_0)$, we see that (ii) holds if

$$\phi'' \leq \frac{1-b\delta}{b\delta}\frac{1}{a}.$$

In particular, if $\phi'' \leq 0$, then this condition always holds, and χ is in fact strongly monotone. However, if $\phi'' < \frac{1-b\delta}{b\delta}\frac{1}{a}$ with $\phi''(y) \in (0, \frac{1-b\delta}{b\delta}\frac{1}{a}]$ for some $y \in (0,1)$, then χ fails strong monotonicity, but it is still ex-post monotone, and thus implementable.

S.3 Downward Distortions

One property that is often encountered in applications is that of downward distortions. Here we establish sufficient conditions for the allocation rule that maximizes expected virtual surplus to satisfy this property. In view of our discussion of non-Markov environments below, we state the result for the general environment.

Intuitively, the principal introduces downward distortions to reduce the agents' information rents when these rents are increasing in the allocations x. To ensure this, we make the usual assumption that higher types have higher marginal utilities, extending it to allow for multidimensional allocations and types:

Condition (U-SCP) ***Utility Single-Crossing Property:*** *X is a partially ordered set, and for all $i = 1, \ldots, n$ and $\theta_{-i} \in \Theta_{-i}$, $U_i(\theta, x)$ has increasing differences in (θ_i, x).*

We also need to assume that the different dimensions of the allocations are (weakly) complementary to each other in the objective function. This is done by invoking Condition U-COMP from the main text, which was used there as a condition for strong monotonicity.

We also invoke F-FOSD to ensure that an agent's type in each period $t > 1$ is positively linked to his period-0 type, which implies that making the agent's utility less sensitive to his future types reduces his ex-ante information rent. F-FOSD also implies that the impulse responses are nonnegative.

Finally, recall that F-AUT assumes that the process is autonomous, and hence the impulse responses can be written as $I_{i,(0),t}(\theta_i^t)$.

We then have the following result.

Proposition S.1 (Downward distortions) *Suppose that Conditions F-AUT, F-FOSD, U-SCP, and U-COMP hold. If the allocation rule χ maximizes expected virtual surplus and the allocation rule χ^* is efficient, then the allocation rule given by $\chi_t(\theta^t) \wedge \chi_t^*(\theta^t)$ for all $t \geq 0$, all $\theta^t \in \Theta^t$ maximizes expected virtual surplus and the allocation rule given by $\chi_t(\theta^t) \vee \chi_t^*(\theta^t)$ for all $t \geq 0$, all $\theta^t \in \Theta^t$ is efficient.*

Proof. Since X is a lattice, the set $\mathcal{X}$ of feasible allocation rules is also a lattice with the meet and join operations defined pointwise (i.e. , for each θ). Define $g : \mathcal{X} \times \{-1, 0\} \to \mathbb{R}$ as

$$g(\chi, q) \equiv \mathbb{E}^{\lambda}\left[\sum_{i=0}^{n} U_i(\tilde{\theta}, \chi(\tilde{\theta})) + q \sum_{i=1}^{n} \frac{1}{\eta_{i0}(\tilde{\theta}_{i0})} \sum_{t=0}^{\infty} I_{i,(0),t}(\tilde{\theta}_i^t) \frac{\partial U_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}}\right].$$

Then $g(\chi, 0)$ is the expected total surplus and $g(\chi, -1)$ is the expected virtual surplus. Condition F-AUT ensures that the stochastic process $\lambda[\chi]$ doesn't depend on χ and that each $I_{i,(0),t}\left(\theta_i^t, x_i^{t-1}\right)$ does not depend on x_i^{t-1}, which is reflected in the formula. Condition F-FOSD ensures that each $I_{i,(0),t}(\theta_i^t) \geq 0$. Together with Condition U-SCP, this ensures that g has increasing differences in (χ, q). Together with U-COMP, this ensures that g is supermodular in χ. The result then follows from Topkis's Theorem (see, e.g., Topkis, 1998). ■

In particular, Proposition S.1 implies that if either the allocation χ that maximizes expected virtual surplus, or the efficient allocation rule χ^* is uniquely defined with probability 1, then $\chi(\theta) \leq \chi^*(\theta)$ with probability 1. (More generally, the set of allocation rules solving the relaxed problem is below the set of efficient allocation rules in the strong set order.)

Downward distortions are typical in applications featuring a single agent, where Conditions F-AUT, F-FOSD, U-SCP, and U-COMP are often implicitly assumed. (As discussed in the main text, Courty and Li (2000) provide a single-agent example that violates F-FOSD and where distortions are upward.). On the contrary, Condition U-COMP is less likely to be satisfied in settings featuring multiple-agents, because of possible capacity constraints that prevent the choice set X from being a lattice. As a result, distortions need not be downward in these applications, even if the other conditions in Proposition S.1 are met. For example, in the bandit auctions considered in Section 5 of the main text, Conditions F-AUT and U-COMP are violated, and distortions are upward at some histories.

S.4 Distortions with Discrete Types

We comment here on how the logic of impulse responses can be used to understand distortions in the optimal allocation rule in models where the agent's types are discrete. For concreteness, we consider a buyer-seller setting similar to Example 5 from the main text. That is, the buyer's payoff is given by $U_1(\theta, x) = \sum_{t=0}^{\infty} \delta^t \theta_t x_t$, with $X_t = [0, \bar{x}]$ for some $\bar{x} >> 0$, and the seller's payoff is given by $U_0(\theta, x) = -\sum_{t=0}^{\infty} \delta^t \frac{x_t^2}{2}$. We assume now, however, that the type spaces Θ_t, $t \geq 0$, are discrete. Our first-order approach is not directly applicable to this setting, but it can be adapted to it by focusing on local downward incentive constraints instead of ICFOC, and using discrete versions of the impulse responses.

For simplicity, we begin with the setting of Battaglini (2005), who considers a Markov process over the binary type space $\{L, H\}$ (i.e., $\Theta_t = \{L, H\}$ for all t). For this setting, we can use the following state representation $\langle \mathcal{E}, G, z \rangle$: in each period $t \geq 1$, G_t is the uniform distribution on $\mathcal{E}_t = (0, 1)$, and $z_t\left(\theta^{t-1}, \varepsilon_t\right) = H$ if $\varepsilon_t > 1 - q_{\theta_{t-1}}$ and $z_t\left(\theta^{t-1}, \varepsilon_t\right) = L$ otherwise. This induces a Markov process on the types with transition probabilities $\Pr\{\tilde{\theta}_t = H | \theta_{t-1}\} = q_{\theta_{t-1}}$, and the assumption $q_H > q_L$ ensures positive serial correlation. In this setting, the discrete one-period-ahead impulse response can be defined as $I(\theta_{t-1}, \theta_t) = \frac{1}{H-L}\mathbb{E}[z(H, \tilde{\varepsilon}_t) - z(L, \tilde{\varepsilon}_t) | z(\theta_{t-1}, \tilde{\varepsilon}_t) = \theta_t]$, i.e., the expected effect of the previous type being H rather than L on the current type, given the observed previous and current types (θ_{t-1}, θ_t). It is then easy to see that $I(\theta_{t-1}, \theta_t) = 0$ whenever the type switches, i.e., $\theta_t \neq \theta_{t-1}$. For example, when type switches from $\theta_{t-1} = L$ to $\theta_t = H$, this means that $\varepsilon_t > 1 - q_L$, and therefore $\varepsilon_t > 1 - q_H$, hence the new type would also have been H had the previous type been H. Similarly, when the type switches from H to L, the new type would also have been L had the previous type been L. Because of the Markov nature of the process, the impulse response of the period-t type to the period-0 type takes the form $I_{(0),t}(\theta^t) = \prod_{\tau=0}^{t-1} I(\theta_\tau, \theta_{\tau+1})$, and therefore as soon as the type switches, the causal effect of the initial type is severed, ensuring efficiency from that point onward.

Observe also that the solution to the relaxed program is efficient when the initial type is H, since only type H's incentive constraint is considered in the relaxed program. These arguments yield Battaglini's "Generalized No Distortion at the Top Principle" (GNDTP): any switch to H yields efficiency from that period onward. Battaglini's other conclusion is the "Vanishing Distortions at the Bottom Principle" (VDBP): the distortions for histories $(L, \ldots, L)$ shrink with time. This conclusion can be understood by noting that the impulse response at $(\theta_{t-1}, \theta_t) = (L, L)$ is less than 1: indeed, for some of the shocks that leave type L unchanged, had the previous type been H instead of L, the current type would have been H.

This logic also demonstrates that GNDTP extends to any discrete type process satisfying FOSD: indeed, when type θ_{t-1} switches to the highest possible type $\theta_t = \bar{\theta}_t$, this implies that any previous type $\theta'_{t-1} > \theta_{t-1}$ would have also switched to $\bar{\theta}_t$. Since the relaxed program considers only downward adjacent incentive constraints, this means a period-$(t-1)$ type θ'_{t-1} who pretended to be θ_{t-1} is no longer distinguishable from θ_{t-1}, in which case distortions should be eliminated forever. However, the distortions need not be immediately eliminated by switching from θ_{t-1} to the lowest possible type $\underline{\theta}_t$. Indeed, since it is the *downward* local IC constraints that bind in the relaxed program, and since types *above* θ_{t-1} may not have switched to $\underline{\theta}_t$ by experiencing the same shocks thus remaining distinguishable, this creates a reason to distort after reporting $\underline{\theta}_t$. We conclude that VDBP also generalizes to multiple discrete types.

Nonetheless, as the number of discrete types increases, the results become qualitatively closer to the continuous-type case than to Battaglini's two-type case, since the binding downward adjacent IC constraints start approximating the ICFOC constraints. For example, while GNDTP and VDPB still hold, their relevance decreases as the number of types grows, for these properties apply only to histories that occur with a small probability. Furthermore, as the distance between the two lowest types vanishes, distortions "at the bottom" become arbitrarily small immediately from $t = 1$. For intermediate types, distortions can in general be nonmonotone in type as in Example 6 in the main text.

S.5 Optimal Mechanisms in Non-Markov Environments

We now discuss how our characterization of implementability for Markov environments can be used to derive sufficient conditions for implementability in some classes of non-Markov environments, and then extend our results on optimal mechanisms to these environments.

The difficulty in establishing sufficient conditions for implementability in non-Markov environments stems from the fact that an agent who lied in the past may find it optimal to continue to lie in subsequent periods. Not knowing how the agent will optimally behave following a lie makes it hard to identify conditions that guarantee the suboptimality of the first lie. This problem does not arise in Markov environments as there restricting attention to strongly truthful equilibria is without loss of generality, but in a non-Markov environment this restriction entails a loss. To see this, assume for simplicity, that there are only two periods, that there is a single-agent, and that there are only two types in each period. Below we prove, by means of a counterexample, that the following result is *false*.

Result. Given any indirect mechanism Ω and any strategy σ that is sequentially optimal for the agent in Ω, there exists a direct mechanism Ω^D with message space equal to Θ_t, $t = 1, 2$, where strongly truthful strategy is optimal and implements the same outcome as σ in Ω.

Counterexample. The agent's private information is $\theta_t \in \Theta \equiv \{\theta_L, \theta_H\}$, with $\theta_L, \theta_H \in (0, 1)$. The set of possible decisions is $X_0 = \varnothing$ and $X_1 = [0, 1]$. The agent's payoff is $U = \theta_0\theta_1 x - p$, where $x \in X$ and $p \in R$. Then consider the following choice rule: $\chi_1(\theta_0, \theta_1) = 1$ if $(\theta_0, \theta_1) = (\theta_H, \theta_H)$ and $\chi_1(\theta_0, \theta_1) = 0$ otherwise; $\psi(\theta_0, \theta_1) = p \in (\theta_H\theta_L, \theta_H^2)$ if $(\theta_0, \theta_1) = (\theta_H, \theta_H)$ and $\psi(\theta_0, \theta_1) = 0$ otherwise. This choice rule can be implemented for example by offering the agent the direct mechanism corresponding to the choice rule. The following strategy is then trivially sequentially optimal for the agent: $\sigma_0(\theta_0) = \theta_0$, for any θ_0, $\sigma_1(\theta_0, \theta_1, \hat{\theta}_0) = \theta_1$ if $\hat{\theta}_0 = \theta_0$ and $\sigma_1(\theta_0, \theta_1, \hat{\theta}_0) = \theta_L$ if $\hat{\theta}_0 \neq \theta_0$. It is easy to see that there exists no direct mechanism that induces the agent to report truthfully in each period after any possible history (including non truthful ones) and that implements the above choice rule. $\square$

The problem with strong truthtelling stems from the fact that a direct mechanism with message space $M_t = \Theta_t$ does not permit the mechanism to treat differently in period t an agent with type history θ^t who reported truthfully in all periods (including the current one) from an agent with type history $\mathring{\theta}^t = (\mathring{\theta}^{t-1}, \theta_t)$, with $\mathring{\theta}^{t-1} \neq \theta^{t-1}$, who reported truthfully in period t but who lied in the past by reporting θ^{t-1}. When the environment is not Markov, as in the above example, requiring the strongly truthful strategy to be optimal then precludes the possibility of implementing certain choice rules.[3]

[3] To bypass this difficulty, Townsend (1988) suggested mechanisms where at each period each agent is asked to report his complete type history rather than just the latest type, as in Myerson (1986). (In other words, Townsend's approach amounts to Markovizing the model by enlargening the state space.) With each report, an agent is then

While general results for non-Markov environments are difficult (if not impossible) to obtain for the above reason, there are nevertheless a few special classes of non-Markov environments, of interest for applications, where our results can be used to establish the implementability of the allocation rule maximizing expected virtual surplus:

First, there are environments whose primitives are not Markov, but where each agent i's payoff depends on the history of received signals θ_i^t only through a *one-dimensional* summary statistics $\varphi_{it}(\theta_i^t)$ whose evolution is Markov and which enters into i's payoff in a Markov way. As far as agent i's incentives are concerned, one may then simply redefine his type to be $\varphi_{it}(\theta_i^t)$ and apply Theorem 3 and its corollaries to the Markovised environment. An example is provided by the Gaussian updating foundations for our bandit auction application discussed in footnote 37. If the signals s_i are taken as primitives (i.e., as types), then the model is not Markov. However, each agent i's period-t payoff depends only on his posterior expectation θ_{it} of his unknown true valuation, which can be shown to follow a Markov process by routine projection results.

Second, our results readily extend to a class of time-separable non-Markov environments that can fit many applications, including sequential auctions, procurement, and regulation. The two key assumptions are that payoffs and decisions separate over time, and that impulse responses depend only on current and initial types and are nonincreasing functions thereof.

Formally, the *separable environment* is defined as follows. Let $X = \prod_{t=1}^{\infty} X_t$ with each $X_t \subset \mathbb{R}^{n+1}$. Each agent i's payoff takes the separable form $U_i(\theta, x) = \sum_{t=0}^{\infty} \delta^t u_{it}(\theta_{it}, x_{it})$, where the flow payoff functions u_{it} are differentiable, non-decreasing functions of the current type θ_{it} that satisfy the following conditions: (i) there exists a constant K_i such that for all $t \geq 0$, $\partial u_{it}/\partial\theta_{it} < K_i$, (ii) u_{it} is strictly increasing in x_{it} and strictly submodular in (θ_{it}, x_{it}), and (iii) the bounds for Condition U-SPR are satisfied. The principal's payoff takes the separable form $U_0(\theta, x) = \sum\limits_{t=0}^{\infty} \delta^t u_{0t}(x_t)$.

The agents' types are assumed to follow processes that satisfy Conditions F-AUT and F-FOSD. In addition, we assume that the first-period hazard rate $\eta_{i0}(\theta_{i0})$ of each agent i is nondecreasing function of θ_{i0} and that, for all $i = 1, \ldots, n$ and $t \geq 0$, the impulse response function $I_{it}(\theta_i^t)$ depends only on $(\theta_{i0}, \theta_{it})$ and is nonincreasing in $(\theta_{i0}, \theta_{it})$.

The above assumptions on the utility functions imply Conditions U-D and U-ELC. Hence for the separable environment to be regular, it suffices that the processes satisfy conditions F-BE and F-BIR, which we assume henceforth. An example of a process satisfying all of the above assumptions is given by the linear AR(k) process consider in Example 2 of the main text, provided that the impulse responses and shocks satisfy $\left\|I_{i,(s)}\right\| < B_i, \|\mathbb{E}[|\tilde{\varepsilon}_i|]\| < \infty$ for some B_i for all $i = 1, \ldots, n$ and $s \geq 0$ (more generally, any ARIMA process satisfies the above assumptions given appropriate restrictions on coefficients).

given a chance to "correct" his past lies, in which case restricting attention to strong truthtelling is always without loss of generality by the argument used to establish the revelation principle in static models. However, this approach entails multidimensional reports in each period, which creates its own set of difficulties.

We then have the following result:

Proposition S.2 (Optimal mechanism for separable environments) *Consider the separable environment described above. Suppose that the expected virtual surplus attains its supremum on the set of feasible allocation rules.*[4] *Then there exists an optimal mechanism* $\langle \chi, \psi \rangle$ *that has the following properties:*

(i) The allocation rule χ *is strongly monotone, and for all* $t \geq 0$ *and* $\theta^t \in \Theta^t$, $\chi_t(\theta^t)$ *depends only on* (θ_0, θ_t)*, and for all* $t \geq 0$ *and* λ*-almost every* θ^t,

$$\chi_t(\theta^t) \in \arg\max_{x_t \in X_t} \left\{ u_{0t}(x_t) + \sum_{i=1}^{n} \left(u_{it}(\theta_{it}, x_{it}) - \frac{1}{\eta_{i0}\left(\theta_{i0}\right)} I_{i,(0),t}(\theta_{i0}, \theta_{it}) \frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}} \right) \right\}. \quad \text{(S.3)}$$

(ii) Flow payments are given by

$$\psi_{it}(\theta_0, \theta_t) = -u_{it}(\theta_{it}, \chi_{it}(\theta_0, \theta_t)) + \int_{\underline{\theta}_{it}}^{\theta_{it}} \frac{\partial u_{it}(r, \chi_{it}(\theta_0, (r, \theta_{-i,t})))}{\partial \theta_{it}} dr \quad \textit{for all } t \geq 1, \quad \text{(S.4)}$$

and

$$\psi_{i0}(\theta_0) = -\mathbb{E}^{\lambda|\theta_0} \left[U_i(\tilde{\theta}_i, \chi(\tilde{\theta})) + \sum_{t=1}^{\infty} \delta^t \psi_{it}(\theta_0, \tilde{\theta}_t) \right] + \int_{\underline{\theta}_{i0}}^{\theta_{i0}} \mathbb{E}^{\lambda_i|r} \left[\sum_{t=0}^{\infty} \delta^t I_{i,(0),t}(\tilde{\theta}_{i0}, \tilde{\theta}_{it}) \frac{\partial u_{it}(\tilde{\theta}_{it}, \chi_{it}(\tilde{\theta}_t))}{\partial \theta_{it}} \right] dr. \quad \text{(S.5)}$$

(iii) Strongly truthful strategies form a subgame perfect equilibrium of the complete information version of the model (where agents observe each others' types).[5]

Before presenting the proof, we note that Proposition 1, which provides conditions on the primitives that guarantee strong monotonicity of an allocation rule maximizing expected virtual surplus, extends to non-Markov environments by slightly modifying the statement to account for the more general payoff functions and processes:

Proposition S.3 (Primitive conditions for strong monotonicity) *Suppose Conditions F-AUT and F-FOSD hold, and for all* $i = 0, \ldots, n$*, and* $t \geq 0$*,* X_{it} *is a subset of an Euclidean space. Suppose that either of the following conditions is satisfied:*

(i) Condition U-COMP holds, and for all $i = 1, \ldots, n$*, agent* i*'s virtual utility*

$$U_i(\theta, x) - \frac{1}{\eta_{i0}(\theta_{i0})} \sum_{t=0}^{\infty} I_{i,(0),t}(\theta_i^t) \frac{\partial U_i\left(\theta, x\right)}{\partial \theta_{it}}$$

has increasing differences in (θ, x)*, and the same is true of the principal's utility* $U_0\left(x, \theta\right)$.

[4] This can be ensured, for example, by assuming that, for all $t \geq 0$, X_t is compact and that, for all $i = 1, \ldots, n$, $t \geq 0$ and $\theta \in \Theta$, $u_{it}(\theta, \cdot)$ is continuous on X_t.

[5] Because valuations are private, the reader may wonder whether truthful reporting is actually dominant in the proposed mechanism. The answer is negative; under arbitrary strategies, the opponents' future actions may depend on their observed allocations in a way that may induce the agent to depart from truthful reporting.

(ii) Condition U-DSEP holds, and for all $i = 1, \ldots, n$, and $t \geq 0$, $X_{it} \subset \mathbb{R}$ and there exists a nondecreasing function $\varphi_{it} : \Theta_i^t \to \mathbb{R}^m$, with $m \leq t$, such that agent i's virtual flow utility

$$u_{it}\left(\theta^t, x_t\right) - \frac{1}{\eta_{i0}(\theta_{i0})} \sum_{\tau=0}^{t} I_{i,(0),\tau}(\theta_i^\tau) \frac{\partial u_{it}\left(\theta^t, x_t\right)}{\partial \theta_{i\tau}}$$

depends only on $\varphi_{it}\left(\theta_i^t\right)$ and x_{it} and has strictly increasing differences in $(\varphi_{it}\left(\theta_i^t\right), x_{it})$, while the principal's flow utility depends only on x_t.

Then, if the problem of maximizing expected virtual surplus has a solution, it has a solution χ such that, for all $i = 1, \ldots, n$ and $\theta_{-i} \in \Theta_{-i}$, $\chi_i\left(\theta_i, \theta_{-i}\right)$ is nondecreasing in θ_i.

The proof is identical to that of Proposition 1 and hence omitted.

Proof of Proposition S.2. Because the environment is time-separable, allocation rule χ maximizes expected virtual surplus if, and only if, for all $t \geq 0$ and λ-almost every $\theta^t \in \Theta^t$, $\chi_t(\theta^t)$ satisfies (S.3). Furthermore, note that the separable environment satisfies the conditions in part (ii) of Proposition S.3 with the functions φ given by $\varphi_{it}(\theta_i^t) = (\theta_{i0}, \theta_{it})$ for all $i = 1, \ldots, n$ and $t \geq 0$. Hence, there exists then a maximizer that is strongly monotone, and by inspection of (S.3), $\chi_t(\theta^t)$ clearly depends only on (θ_0, θ_t). This establishes property (i).

To see that the payments given in (ii) implement the maximizer χ and that property (iii) holds, note first that, under the proposed mechanism, *incentives separate over time*, starting from $t = 1$ onwards. That is, in every period $t \geq 1$, and for any history, each agent i maximizes his payoff by simply choosing the current message $\hat{\theta}_{it}$ so as to maximize his flow payoff $u_{it} + \psi_{it}$. This follows because agent i's period-t message has no effect on the allocations in periods $\tau > t$, and because of the separable payoffs and F-AUT the allocation in period t has no effect on future payoffs or types. That under the payments given by (ii), the agent finds it optimal to report truthfully, irrespective of his beliefs about the other agents' types and messages, and irrespective of whether or not he has been truthful in the past, then follows from standard results from static mechanism design by observing that (i) the allocation $\chi_{it}(\theta_0, \theta_t)$ is monotone in θ_{it} for all $(\theta_{-i,t}, \theta_0)$, and (ii) values are private, i.e., each u_{it} depends on θ only through θ_{it}. In other words, it is as if each agent i were facing a single-agent static decision problem indexed by $(\theta_{-i,t}, \theta_0)$. In particular, note that reporting truthfully remains optimal even if each agent were able to observe both his own future types, the other agents' past, current, and future types, and the messages sent by the other agents, thus making his beliefs completely irrelevant.

As for period zero, we first establish the following lemma, which is similar to Theorem 3 in the main text. However, because of the environment being non-Markov, the restriction to strongly truthful strategies is in general with loss, and hence the result only provides us with a sufficient condition for implementability.[6]

[6] Note that the lemma establishes a result which holds more generally than in the context of the separable environments. In particular, it can be useful in finite-horizon settings where incentive compatibility can be established by solving each agent's optimization problem by backward induction.

Lemma S.2 *Suppose the environment is regular. Fix $t \geq 0$, and consider a choice rule $\langle \chi, \psi \rangle$. Suppose there exists a belief system $\Gamma \in \mathbf{\Gamma}(\chi)$ such that, for all $i = 1, \ldots, n$, (i) $\langle \chi, \psi \rangle$ with belief system Γ satisfies ICFOC$_{i,t}$, (ii) for all $\theta^{t-1} \in \Theta^{t-1}$ and $\theta_{it}, \hat{\theta}_{it} \in \Theta_{it}$,*

$$\int_{\hat{\theta}_{it}}^{\theta_{it}} \left[D_{it}^{\chi,\Gamma}(\theta^{t-1}, r) - D_{it}^{\chi \circ \hat{\theta}_{it}, \Gamma}(\theta^{t-1}, r) \right] dr \geq 0,$$

where the functions D are as defined in equation (6), and (iii) strong truthtelling is an optimal continuation strategy for agent i. Then, for all agents, a one-step deviation from strong truthtelling is unprofitable at every period-t truthful history.

Proof of Lemma S.2. Since strong truthtelling is assumed to be an optimal continuation strategy, the result follows from the same arguments as in the proof of Theorem **??** in the main text. ■

In order to apply Lemma S.2 to establish optimality of truthtelling in period 0, we observe that the choice rule $\langle \chi, \psi \rangle$ given in the statement of the proposition satisfies ICFOC by construction, and that strong truthtelling is an optimal continuation strategy from period 1 onwards by the argument preceding the lemma. Hence it remains to verify the integral monotonicity condition in the lemma. But this follows by Corollary 1, since χ is strongly monotone and the environment satisfies the conditions needed for strong monotonicity to imply integral monotonicity. (Note that Corollary 1 does not invoke the Markov property.) Hence strongly truthful strategies from a PBE. Moreover, because the beliefs used to define the period-0 transfer ψ_{i0} in part (ii) are given by the exogenous measure $\lambda|\theta_0$, truthful reporting in period 0 is in fact optimal regardless of the agent's beliefs. Together with the fact that truthful reporting is also optimal in each subsequent period irrespective of the beliefs, this establishes that strongly truthful strategies form a subgame perfect equilibrium of the complete information version of the model.

Finally, because each $u_{it}(\theta_{it}, x_{it})$ is nondecreasing in θ_{it} and F-FOSD holds, it is immediate to verify that, under the proposed mechanism, participating is optimal for all period-0 types. ■

We conclude by discussing possible implementations of the profit-maximizing rule. First, consider the special case where flow payoffs are linear (i.e., $u_{it}(\theta_{it}, x_{it}) = \theta_{it} x_{it}$) as in auctions, and where types evolve according to an AR(k) process (or, more generally, any process for which the impulse responses I_{it} depend only on the initial types). The implementation is then particularly simple. Suppose that there is no allocation in the first period and consider implementation in a PBE rather than in a periodic ex-post equilibrium (both assumptions simplify the discussion but are not essential for the argument). In period zero, each agent i is then asked to choose from a menu of "handicaps" $(I_{i,(0),t}(\theta_{i0})/\eta_{i0}(\theta_{i0}))_{t=1}^{\infty}$, indexed by θ_{i0}, with each handicap costing $\psi_{i0}(\theta_{i0})$ as defined in (S.5) but with the measure $\lambda|\theta_0$ replaced by the measure $\lambda_i|\theta_{i0}$. Then in each period $t \geq 1$, a "handicapped" VCG mechanism is played with transfers as in (S.4). (Eső and Szentes (2007) derived this result in the special case of a two-period model with allocation only in the second period. See Kakade et al (2011) for an extension to a different class of non-Markov environments.)

This logic extends to nonlinear payoffs and more general processes, in the sense that in the first period the agents still choose from a menu of future plans (indexed by the first-period type). However, in general, in the subsequent periods the distortions will depend also on the current reports through the partial derivatives $\partial u_{it}(\theta_{it}, x_{it})/\partial\theta_{it}$ and through the impulse responses $I_{i,(0),t}(\theta_{i0}, \theta_{it})$. However, as long as payoff separate over time and the impulse responses depend only on initial and current types, the intermediate reports (i.e., reports in periods $1, \ldots, t-1$) remain irrelevant both for the period-t allocation and for the period-t payments.

Finally, note that the argument used to establish sufficiency for this family of non-Markov problems ("virtual" VCG payments in each period $t > 0$ along with monotonicity of the allocation rule and payments at $t = 0$ given by (S.5)) extends to a few environments where payoffs are not time-separable, but where virtual payoffs continue to be an affine transformation of the true payoffs with constants that depend only on the initial reports. Consider, for example, an economy where all processes are AR(1) and where payoffs are given by $U_i = \sum_{t=0}^{\infty}\delta^t\theta_{it}x_{it} - c_{it}(x^t)$ for all $i = 0, \ldots, n$. The dynamic virtual surplus then coincides with the true surplus of a fictitious economy where all agents' payoffs are as in the original economy and where the principal's payoff is given by $U_0 - \sum_{i=1}^{n}\frac{1}{\eta_{i0}(\theta_{i0})}\sum_{t=0}^{\infty}\delta^t I_{i,(0),t}x_{it}$. Then note that, irrespective of whether or not the agents reported truthfully in period zero, the allocation rule χ that solves the relaxed program maximizes the surplus of this fictitious economy from period $t = 1$ onwards. This property, together with the fact that values are private from $t = 1$ onwards in this fictitious economy, then implies that incentives for truthful reporting at any period $t \geq 1$ can be provided by using either the team payments of Athey and Segal (2012) or the pivot payments of Bergemann and Välimäki (2010). Furthermore, as long as the rule χ satisfies integral monotonicity in period 0, then incentives can also be provided in period 0 by adding to the aforementioned payments an initial payment given by (S.5). For an application of similar ideas to a family of problems where the dynamic virtual surplus takes the form of a multiplicative transformation of the true surplus, with the scale depending only on the initial reports, see the recent paper by Kakade et al (2011).

References

[1] Athey, S., and I. Segal (2012), "An Efficient Dynamic Mechanism," Mimeo, Stanford University.

[2] Bergemann, D., and J. Välimäki (2010), "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771-789.

[3] Courty, P. and H. Li (2000), "Sequential Screening," *Review of Economic Studies*, 67, 697-717.

[4] Eső, P. and B. Szentes (2007), "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74(3), 705-731.

[5] Kakade, S., Lobel, I. and H. Nazerzadeh (2011), "Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism," mimeo Wharton School, University of Pennsylvania.

[6] Milgrom, P. and C. Shannon (1994), "Monotone Comparative Statics," *Econometrica*, 62, 157-180.

[7] Topkis, D. (1998), *Supermodularity and Complementarity*, Princeton University Press.