Mechanism Design

6: Dynamic Mechanisms

Egor Starkov

Københavns Unversitet
Fall 2024

This slide deck:

Dynamic Problems

- Models considered so far were static: one report, one outcome.
  - Although we hinted towards dynamic incentives when discussing interim vs ex post IC/IR constraints.
- There are many dynamic problems in the real world:
  - Dynamic pricing when buyers' tastes evolve (e.g. experience goods) or buyers come and go over time;
  - Procurement from firms with changing costs;
  - Design of tax and social security systems;
  - Dynamic labor contracts
- How to develop dynamic mechanisms? Will see today.
- This lecture mostly follows Bergemann and Välimäki [2019].

What defines a dynamic problem? (1)

- Why can a dynamic problem not be seen as a sequence of independent static problems?
- Because there can be linkages across periods: (which ruin the independence)
  1. Information – future info evolves from (so depends on) past info and possibly past allocations.
  2. Preferences – usually evolve gradually. For our purposes, can see this as persistence in nformation.
  3. Allocations – set of feasible allocations today may depend on past outcomes (example: sale of fixed number of items over many periods).
- The same linkages mean that if we try to see the problem as a huge static problem (with same player in different periods seen as different players), the correlations in players' info and the set of feasible allocations will look weird and complicated.

Dynamic Model

- Periods $t \in \{0, 1, ..., T\}$; terminal time $T \leq \infty$; all players (incl. designer) have common discount factor δ.
- Players $i \in \{1, 2, ..., N\}$ have evolving types $\theta_{i,t} \in \Theta_i$, indep. across i.
  - Common prior $\theta_{i,0} \sim F_{i,0}$; types are Markov processes:

$$\theta_{i,t+1} \sim F_{i,t}(\theta_{i,t+1} | \theta_{i,t}, k_t).$$

- Every period: allocation $k_t \in K_t$ and payments $p_t \in \mathbb{R}^N$.
  Set of feasible allocations evolves as $K_{t+1} = g(K_t, k_t)$.
- Players' utilities: $u_i((k_t, p_t), \theta_t) = v_i(k_t, \theta_{i,t}) - p_{i,t}$.

Evolving Types

Possible interpretations of evolving types:

- Exogenous evolution ($\theta_{t+1} \perp k$);
  - Example: procuring goods over time from a firm with stochastically evolving costs $\theta_{i,t+1} = \gamma\theta_{i,t} + \varepsilon_{i,t+1}$.
- Endogenous evolution (depending on k_t);
  - Example: worker assigned to training by k_t will improve their future productivity $\theta_{i,t+1}$.
- Random arrival;
  - Players can arrive at the mechanism at random times.
  - Can model that by setting $\theta_{i,t} = \varnothing$ whenever i is not in the market/mechanism.

Dynamic Model: Assumptions

To fix ideas, assume the following for this class:

- The designer can commit to the whole future mechanism at $t = 0$.
- Contracts are binding – we ignore per-period IR constraints (except maybe IR at $t = 0$).
  - Justification: in quasilinear model, can ask players to put collateral at $t = 0$, to be repaid at a later date – this would eliminate incentives to quit mechanism after $t = 0$.
- All past reports and allocations are publicly observed.
- Player i at time t observes their type $\theta_{i,t}$ but not future types.

Direct Mechanisms

- As usual, we have the revelation principle, though there are caveats [Sugaya and Wolitzky, 2021].
- So can focus on mechanisms which ask players to report their types every period.
- Reporting strategies given by $\rho_i = \{r_{i,t}\}_{t=0}^{T}$, where $r_{i,t} : \Theta_i \times H_t \to \Theta_i$ and H_t is the set of public histories $h_t = \{k_s, (r_{1,s}, ..., r_{N,s})\}_{s<t}$.
- A dynamic direct mechanism is $(\kappa, \pi) = \{k_t, p_t\}_{t=0}^{T}$, where $k_t : \Theta \times H_t \to K_t$ and $p_t : \Theta \times H_t \to \mathbb{R}^N$.

Dynamic Implementation

- Looking for a truthful equilibrium in a direct mechanism.
- "Equilibrium" is a sketchy term in dynamic incomplete-info games.
  - There is at least a dozen different equilibrium concepts and refinements in use.
  - Main concern in general: off-equilibrium-path beliefs. What should a player believe after observing an event they considered impossible? Different answers can strongly affect the predicted outcome.
  - Not a big problem in mechdesign – players do not observe any actions until it's too late to act.
- Look for Perfect Bayesian Equilibria.
  - Each player chooses report to maximize expected util, expecting others to report truthfully.
  - Beliefs are updated using Bayes' rule whenever possible (i.e., on equilibrium path).
  - In general in PBE: We can assume anything we want about off-path beliefs to sustain eqm. In our problem: won't need to.

This slide deck:

Efficient Allocation

- Suppose we want to implement the efficient allocation κ^*.
- But what is κ^* in a dynamic problem?

$$\kappa^* \in \arg\max_{\{k_t\}_{t=0}^T} \mathbb{E}\left\{\sum_{t=0}^{T} \delta^t \sum_{i=0}^{N} v_i(k_t, \theta_{i,t})\right\}$$

- Must optimize over the whole path $\{k_t\}_{t=0}^T$ rather than period-by-period.
  - Today's allocation k_t may affect tomorrow's types θ_{t+1} and set of alternatives K_{t+1}.
- Also remember that $k_t : \Theta \times H_t \to K_t$ is a highly-dimensional object in itself.
- So simply finding κ^* is in general a difficult optimal control problem.
- **Remark**: ex post efficiency is unattainable in dynamics – k_t must be chosen before θ_{t+s} learned. Interim efficiency is the best we can hope for.

Efficient Implementation

- Ok, suppose we found κ^*, what next?
- In static setting we used VCG aka the pivot mechanism: each player had to pay the externality they imposed on everyone else:

$$p_i(\theta) = -\sum_{j \neq i} v_j\left(k^*(\theta), \theta_j\right) + \sum_{j \neq i} v_j\left(k^*_{-i}(\theta_{-i}), \theta_j\right)$$

- The idea translates almost verbatim to the dynamics.
  - Problem: the externality that i imposes on others via report $\theta_{i,t}$ may manifest in other periods – not necessarily at t.
- Enter dynamic pivot mechanism! [Bergemann and Välimäki, 2010]

Dynamic Pivot Mechanism

Flow social surplus	$w_t(k_t, \theta_t) \equiv \sum_{i=1}^{N} v_i(k_t, \theta_{i,t}).$
Welfare	$W_t(\theta_t, K_t) \equiv \max_{k_t \in K_t} \left\{ w_t(k_t, \theta_t) + \delta \mathbb{E} W_{t+1}(\theta_{t+1}, K_{t+1}) \right\}.$
i's marginal contribution	$M_{i,t}(\theta_t, K_t) \equiv W_t(\theta_t, K_t) - W_{-i,t}(\theta_t, K_t)$
can be written recursively as	$M_{i,t}(\theta_t, K_t) = m_{i,t}(\theta_t, K_t) + \delta \mathbb{E} M_{i,t+1}(\theta_{t+1}, K_{t+1}).$
Payments	$p^*_{i,t} \equiv v_i(k^*_t, \theta_{i,t}) - m_{i,t}(\theta_t, K_t).$

- The dynamic pivot mechanism is given by $\kappa = \kappa^*$ and $\rho = \{p^*_{i,t}\}_{t=0}^{T}$.
- Note that i must pay his flow marginal contribution rather than simply $w(k^*) - w(k^*_{-i})$.
- This is because i by influencing today's allocation k_t, i will also affect future types of other players and the set of available allocations – have to account for that.

This slide deck:

Dynamic Revenue Maximization

- Second canonical question: what is the optimal mechanism?
  - Main example: dynamic pricing (there's huge literature, more or less related to DMD).
  - With binding contracts: mobile service, loans, insurance

Question

There is one buyer with time-changing valuation $\theta_t \in \Theta \subset \mathbb{R}$ for the item.
What is the seller-optimal mechanism for $\{$repeated purchases, one-time purchase$\}$?

- Again, insights from static models carry over after reasonable modifications.
  - Now we want to distinguish between info that the buyer has before signing up for a mechanism
  - and which they acquire after signing the contract.

Flashback: Static Model

- In the static optimal mechanism, seller's expected revenue was

$$\mathbb{E}R = \mathbb{E}_\theta \left[v(k(\theta), \theta) - \frac{1 - F(\theta)}{\phi(\theta)} \frac{\partial v(k(\theta), \theta)}{\partial \theta} \right]$$

(we derived this for $v(k, \theta) = k\theta$: $\quad \mathbb{E}R = \mathbb{E}_\theta \left[k(\theta) VS^{static}(\theta) \right]$)

- Had to trade off max social surplus $v(k, \theta)$ (i.e., efficiency) against information rents.
    - Had to leave some money to the buyer to incentivize truthful reporting.

Static Model – Posterior Information

Example

- Consider static optimal mechanism setting (1 period, 1 item, 1 buyer),
- except: buyer only learns θ after signing up for the mechanism.
- What is the optimal contract?

Static Model – Posterior Information

Example

- Consider static optimal mechanism setting (1 period, 1 item, 1 buyer),
- except: buyer only learns θ after signing up for the mechanism.
- What is the optimal contract?
  - Designer's problem is

$$\max_{(k,p)} \{\mathbb{E}_\theta p(\theta)\}$$

$$\text{s.t. } (IC): \; v(k(\theta),\theta) - p(\theta) \geq v(k(\theta),\hat{\theta}) - p(\hat{\theta}) \quad \forall \theta, \hat{\theta},$$

$$(eaIR): \; \mathbb{E}_\theta \left[v(k,\theta) - p \right].$$

  - Only real difference from Myerson: ex ante IR instead of interim IR.
  - Solution: choose efficient k^* and charge $p(\theta) \equiv p = \mathbb{E}_\theta \left[v(k^*(\theta),\theta) \right]$
  - Perfect information extraction; no information rents to the buyer; full efficiency.
  - Remark: this solution would not work with $N > 1$ bidders competing for 1 item (why?)

Dynamics and Information (i)

Statement (Future Extraction)

Designer can extract all of buyer's future info at no cost.

- Same idea: "sell" the item (subscription) to the buyer at ex ante expected value.
- Then only buyer's initial info θ_0 matters for IC:
  - in future periods use buyer-optimal allocation rule $\Rightarrow$ buyer's IC is satisfied without any extra transfers.
- (FE) sounds reasonable, but it is not a formal theorem.
- The literature is currently at the stage "let's hope that (FE) holds".

Dynamics and Information (ii)

Statement (Future Extraction)

Designer can extract all of buyer's future info at no cost.

- The literature is currently at the stage "let's hope that (FE) holds".
- In particular, the protocol is:
  1. Solve the dynamic problem as if all future info is public.
  2. Get some allocation and transfers.
  3. Check whether the resulting mechanism satisfies dynamic IC (at $t > 0$).
  4. Pray that it does.
- Pavan, Segal, and Toikka [2014] provide some sufficient conditions for this to work, but these are considered by some as too restrictive.
- We today take the "pray that (FE) holds" approach and only worry about extracting the buyer's initial type θ_0 – we are back to the static problem.

Caveat

Statement (Future Extraction)

Designer can extract all of buyer's future info at no cost.

Caveat

"Ignore future information" is not the same as "ignore future types"!

- Type θ_0 is (in general) correlated with future θ_t,
- so θ_0 contains some information about θ_t,
- so we **cannot** work as if know θ_t for $t \geq 1$.

Caveat

Statement (Future Extraction)

Designer can extract all of buyer's future info at no cost.

Caveat

"Ignore future information" is not the same as "ignore future types"!

- Solution: separate types from information through orthogonalization.
  - Suppose $\theta_{t+1} \sim F_{t+1}(\theta_{t+1}|\theta_t, k_t)$.
  - Let $\varepsilon_{t+1} \equiv F_{t+1}(\theta_{t+1}|\theta_t, k_t)$. Then $\varepsilon_{t+1} \sim U[0,1]$ and independent of θ_t.
  - In a direct mechanism, ask player to report θ_0 in period 0 and ε_t in period t, then recover θ_{t+1} from these reports.

Virtual Surplus

- Optimal allocation κ maximizes virtual surplus $=$ real surplus – information rents.
  - This pins down optimal mechanism $(\kappa, \pi + C)$ up to the constant C.
  - C is determined from IR at $t = 0$ – skip the step of finding it.
- In static model, virtual surplus is (*note inconsistency in how VS is defined here vs in past lectures!*)

$$VS(k,\theta) = v(k(\theta),\theta) - \frac{1 - F(\theta)}{\phi(\theta)} \frac{\partial v(k(\theta),\theta)}{\partial \theta}$$

- Now in dynamics, real surplus is

$$S(\kappa,\theta) \equiv \sum_{t \geq 0} \delta^t v(k_t(\theta_t),\theta_t).$$

Calculating $VS(\kappa,\theta) = S(\kappa,\theta) - \frac{1-F_0(\theta_0)}{\phi_0(\theta_0)} \frac{\partial S(\kappa,\theta)}{\partial \theta_0}$ requires understanding how S depends on θ_0 (the only source of inforents for the buyer).

Virtual Surplus

$$\frac{\partial S(\kappa,\theta)}{\partial \theta_0} = \sum_{t\geq 0} \delta^t \frac{\partial v(k_t(\theta_t),\theta_t)}{\partial \theta_t} \frac{\partial \theta_t}{\partial \theta_0}$$

- Let $I_t(\theta_t|\theta^{t-1}, k_{t-1}) \equiv \frac{\partial \theta_t}{\partial \theta_0}$ be impulse response function, where $\theta^t \equiv (\theta_0, \theta_1, ..., \theta_t)$.
- I_t shows the effect of θ_0 on θ_t given fixed realization of uncertainty $\{\varepsilon_s\}_{s\leq t}$.
- Can compute that

$$I_t(\theta_t|\theta^{t-1}, k_{t-1}) = -\prod_{s=1}^{t} \frac{\frac{\partial F_s(\theta_s|\theta^{s-1}, k_{s-1})}{\partial \theta_{s-1}}}{\phi_s(\theta_s|\theta^{s-1}, k_{s-1})}.$$

Virtual Surplus

Then

$$\frac{\partial S(\kappa,\theta)}{\partial\theta_0} = \sum_{t\geq 0}\delta^t \frac{\partial v(k_t(\theta_t),\theta_t)}{\partial\theta_t}\frac{\partial\theta_t}{\partial\theta_0}$$
$$= \sum_{t\geq 0}\delta^t I_t(\theta_t|\theta^{t-1},k_{t-1})\frac{\partial v(k_t(\theta_t),\theta_t)}{\partial\theta_t}$$

so the whole virtual surplus as a function of the whole $\theta = (\theta_1,\theta_2,...)$ is

$$VS(\kappa,\theta) = S(\kappa,\theta) - \frac{1-F_0(\theta_0)}{\phi_0(\theta_0)}\frac{\partial S(\kappa,\theta)}{\partial\theta_0}$$
$$= S(\kappa,\theta) - \frac{1-F_0(\theta_0)}{\phi_0(\theta_0)}\sum_{t\geq 0}\delta^t I_t\frac{\partial v_t}{\partial\theta_t}$$

(again, definition slightly different than in static opt.mech.; this one is more general)

Optimal Mechanism

- To find optimal allocation, take expectation of $VS(\kappa, \theta)$ over $\{\varepsilon_t\}$ to get $VS(\kappa, \theta_0)$ and maximize over κ. (Still a difficult problem, for the same reasons as for efficient κ^*.)

$$\max_{\kappa} \mathbb{E}_{\varepsilon}\left[S(\kappa, \theta) - \frac{1 - F_0(\theta_0)}{\phi_0(\theta_0)} \sum_{t \geq 0} \delta^t I_t \frac{\partial v_t}{\partial \theta_t} \mid \theta_0\right]$$

- Then find expected (as of $t = 0$) payments from the envelope representation of the buyer's expected utility:

$$\frac{dU_{b,0}(\theta_0)}{d\theta_0} = \mathbb{E}\left[\sum_{t=0}^{T} \delta^t I_t(\theta_t|\theta^{t-1}, k_{t-1}) \frac{\partial v(k_t, \theta_t)}{\partial \theta_t} \mid \theta_0\right].$$

- Note that this will pin down the "expected-at-time-0" payments $\mathbb{E}_{\varepsilon}[\sum_t \delta^t p_t(\theta^t)|\theta_0]$. These payments can be redistributed across periods and histories since both seller and buyer are risk-neutral.
  - Will usually have to do this redistribution to ensure IC at $t > 0$. No good recipe here.

Dynamic Revenue Maximization: Conclusions

$$\max_{\kappa} \mathbb{E}_{\varepsilon} \left[S(\kappa, \theta) - \frac{1 - F_0(\theta_0)}{\phi_0(\theta_0)} \sum_{t \geq 0} \delta^t I_t \frac{\partial v_t}{\partial \theta_t} \mid \theta_0 \right]$$

- Insight: if $|I_t|$ decreasing with t, i.e., θ_0 contains little information about θ_t for large t then optimal k_t converges to the efficient k_t^*.
- **Distortions vanish over time.**
- See Bergemann and Välimäki (2019, ch.5) for applications.

This slide deck:

What now?

- Will look at dynamic mechanisms within some special settings.
- Beyond the models we looked at, not within.
- Will go very quickly: no solving models, just setup and results.
- Will see a common theme emerging.

This slide deck:

Dynamic Insurance [Thomas and Worrall, 1990]

- One risk-neutral lender (designer), one risk-averse borrower (agent), common discount factor β.
- Time $t = 0, 1, \ldots$.
- Agent receives random exogenous income $\theta_t \sim i.i.d. F[\underline{\theta}, \bar{\theta}]$.
  - Concave utility $u(c)$, so would like to insure.
  - Special assumption: $u(\underline{c}) = -\infty$, where $\underline{c} > 0$ is subsistence level.
- Principal designs insurance contract.
  - Goal: minimize cost of providing (ex ante expected) util V_0 to agent.
  - Agent reports θ_t in every period, mechanism pays him $b_t(\theta_t, \theta_{t-1}, \ldots)$
  - Perfect commitment on both sides – no IR.
  - But must incentivize truthful reporting of income θ_t – IC.

Agent's incentives

- At all t, agent maximizes lifetime utility

$$V_t \equiv \sum_{s=t}^{\infty} \beta^s u(\theta_s + b_s).$$

- Let $g^t = (\hat{\theta}_0, ..., \hat{\theta}_t)$ be history of past reports.
- Then agent's IC at g^{t-1} is:

$$u(\theta_t + b_t(g^{t-1}, \theta_t)) + \beta V_{t+1}(g^{t-1}, \theta_t) \geq$$
$$\geq u(\theta_t + b_t(g^{t-1}, \hat{\theta}_t)) + \beta V_{t+1}(g^{t-1}, \hat{\theta}_t).$$

for all $\theta_t, \hat{\theta}_t$.

Relation to standard model

- Note that there are no allocations, only money across periods.
- One way to relate to our standard quasilinear model:

usual model	*this model*
allocation k	today's transfer b_t
transfer t	continuation util V_{t+1}

- The main intertemporal linkage comes from the need to deliver on promised V_{t+1}.

Efficient contract

- Moving on to the results.
- In the optimal contract, at every g^{t-1}:
  - b_t is decreasing in $\hat{\theta}_t$ (insurance);
  - V_{t+1} is increasing in $\hat{\theta}_t$ (IC).
  - In particular, $b_t(\underline{\theta}) > 0 > b_t(\overline{\theta})$; $V_{t+1}(\underline{\theta}) < V_t < V_{t+1}(\overline{\theta})$.
- First-best (cheapest way to deliver util V_t) would be to provide full insurance, but have to trade efficiency against info rents, so incomplete insurance in the optimum. (Standard opt.mech logic)

Immiseration

Theorem (Immiseration)

$$\lim_{t \to \infty} V_t \overset{a.s.}{=} -\infty$$

- In the limit, consumption and utility converge to $\bar{c}$ and $-\infty$ resp.
- Neat mathematical result, but I haven't found any good intuitive explanations of where it comes from and after some thorough thinking cannot offer any correct intuition of my own.
- Popular paper, has quite some citations and influential follow-up papers.

This slide deck:

Dynamic Allocation without Money [Guo and Hörner, 2018]

- One principal, one agent.
- Time $t = 0, 1, \dots$
- In each period: agent's type $v \in \{L, H\}$, principal chooses $a \in \{0, 1\}$. Utilities (P,A):

(u_P, u_A)	$v = H$	$v = L$
$a = 0$	$(0, 0)$	$(0, 0)$
$a = 1$	$(H - c, H)$	$(L - c, L)$

  with $H > c > L > 0$.
- Idea: principal can provide funding for agent's project, it is costly for the principal, but agent always wants more funding.
- Persistence: $\mathbb{P}(v_{t+1} = v_t) = \rho \geq 1/2$.
- Principal's goal: max own discounted util subject to IC.

Connection

- Like Thomas and Worrall, but there had only transfers, no allocations. Here only allocations, no transfers.
- Same idea behind IC: induce truthtelling today by varying future utility promises.

usual model	*this model*
allocation k	today's allocation a_t
transfer t	continuation util V_{t+1}

- Opt. mech: if agent does not require funding today, allow to claim more funding in the future. For $v = H$ agent, funding today is more valuable than in the future (since $\mathbb{E}v_{t+s} < H$), for $v = L$ future funding is more valuable than today $\Rightarrow$ IC.

Polatization

- Let $U_t \equiv (1-\delta)\mathbb{E}\left[\sum_{s \geq t} \delta^{s-t} a_t v_t\right]$ denote agent's util.

  Note $U_t \in [0, \bar{U}]$ for some $\bar{U}$.

Theorem (Polarisation)

Under the optimal mechanism, $U_t \to \{0, \bar{U}\}$ as $t \to \infty$.

- U_t is (not really, but similar for our purposes to) a martingale bounded on both sides – both boundaries are absorbing, and U_t hits one of them sooner or later.

This slide deck:

Power Dynamics in Organizations [Li et al., 2017]

- One principal, one agent.
- Time $t = 0, 1, \ldots$
- In each period: principal chooses $a \in \{0, 1, 2, 3\}$. Utilities (P,A):

		principal	agent
$a = 0$	(default)	0	0
$a = 1$	(agent-preferred)	b	B
$a = 2$	(principal-preferred)	B	b
$a = 3$	(nuke humanity)	$-\infty$	$-\infty$

  with $B > b > 0$.
- Principal-preferred project only available at any t with probability p. Only the agent knows whether $a = 2$ is available at a given t. Agent suggests a project to principal at every t.
- Principal's goal: maximize expected util subject to agent's IC.

Possible Modes

- Centralization
  - Principal always chooses the default $a = 0$.
- Cooperative Empowerment
  - Agent suggests and principal implements principal-preferred $a = 2$ when available, agent-preferred $a = 1$ otherwise.
  - The "best" outcome.
- Restricted Empowerment
  - Agent suggests and principal implements principal-preferred $a = 2$ when available, default $a = 0$ otherwise.
- Unrestricted Empowerment
  - Agent suggests and principal implements agent-preferred $a = 1$ always.
- Total Annihilation
  - Principal implements $a = 3$; only used as off-path threat.

Polatization

Theorem

In the optimal relational contract, the principal chooses cooperative empowerment for the first τ periods, where τ is random and finite with probability one.
For $t > \tau$, the relationship results in unrestricted empowerment, restricted empowerment, or centralization forever

- The relationship inevitably slips out of the cooperative mode into one of the uncooperative ones:
  - either the agent gets unlimited power,
  - or the principal loses trust in him.
  - Although convergence to restricted empowerment (semi-cooperative outcome) is possible...

Conclusion

- Lessons from the three papers:
  - relying on promises of future utility for incentive provision leads to huge asymptotic inefficiencies.
- Drastically different from the quasilinear setting we considered,
  - where inefficiencies vanished over time...

References I

D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010. URL `https://doi.org/10.3982/ECTA7260`. Publisher: Wiley Online Library.

D. Bergemann and J. Välimäki. Dynamic Mechanism Design: An Introduction. *Journal of Economic Literature*, 57(2):235–274, June 2019. ISSN 0022-0515. doi: 10.1257/jel.20180892. URL `https://www.aeaweb.org/articles?id=10.1257/jel.20180892`.

Y. Guo and J. Hörner. Dynamic allocation without money. Technical report, Working Paper, 2018.

J. Li, N. Matouschek, and M. Powell. Power Dynamics in Organizations. *American Economic Journal: Microeconomics*, 9(1):217–41, 2017. doi: 10.1257/mic.20150138. URL `http://pareto-optimal.com/s/LMP-Submission.pdf`.

References II

A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: a myersonian approach. *Econometrica*, 82(2):601–653, 2014. doi: 10.3982/ECTA10269.

T. Sugaya and A. Wolitzky. The revelation principle in multistage games. *The Review of Economic Studies*, 88(3):1503–1540, 2021. URL `https://doi.org/10.1093/restud/rdaa041`. Publisher: Oxford University Press.

J. Thomas and T. Worrall. Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, 51(2):367–390, 1990.