

The social planner promises to deliver lifetime utility  $\underline{U}$  to the agent and his objective is to minimize the expected present value of costs. The planner can freely borrow and save at the exogenous net interest rate  $r > 0$ , which is also his discount factor. Let  $\mathcal{K} : \mathbb{X} \rightarrow \mathbb{R}$  denote the planner's cost of allocating  $X \in \mathbb{X}$  to an agent, with  $\mathcal{K}$  continuous, increasing, and convex. For example, in an insurance model  $\mathcal{K}(X) = \mathcal{K}(c, y) = c - y$  is the planner's claims payout minus his income.

At the beginning of each date, agents observe their shocks (and if they die, they immediately leave the economy). Since the planner cannot observe  $\theta_t$ , agents then report their shock to the social planner, who determines  $X_t$ , and agents consume.

## 2.4 Planner's Problem

At  $t = 0$ , the planner selects an allocation  $x = \{x_t\}_{t=0}^\infty$ , which is a sequence of  $\mathcal{F}_t$ -progressively measurable<sup>9</sup> functions  $x_t : \Theta^t \rightarrow \mathbb{X}$ ; let  $\mathcal{X}$  denote the set of allocations. Then an agent's utility from an allocation  $x$  is

$$U(x) = \mathbb{E}_0 \left[ \int_0^\infty e^{-(\rho+\kappa)t} \tilde{u}(x_t(\theta^t); \theta_t) dt \right].$$

At each date agents report their shocks to the planner. A reporting strategy  $\sigma = \{\sigma_t\}_{t=0}^\infty \in \Sigma$  is a sequence of  $\mathcal{F}_t$ -progressively measurable functions  $\sigma_t : \Theta^t \rightarrow \Theta$ ; let  $\sigma^t \in \Sigma^t$  denote the history of reports up to date  $t$ .

Let  $\{x_t(\sigma^t(\theta^t))\}_{t=0}^\infty$  denote an agent's allocation given reporting strategy  $\sigma \in \Sigma$ . By the Revelation Principle, I can restrict attention to direct revelation mechanisms in which agents truthfully reveal their shocks to the planner. The truth-telling strategy specifies  $\sigma_t(\theta^t) = \theta_t$  for all  $t$ . Finally, let

$$\Sigma^* \equiv \{\sigma \in \Sigma \mid \sigma_t(\theta^t) = \hat{\theta}_t \leq \theta_t \forall t\}$$

---

<sup>9</sup>A function is progressively measurable if, for every  $t$ , the map  $[0, t] \times \Omega \rightarrow \mathbb{X}$  defined by  $(s, \omega) \mapsto x_s(\omega)$  is Borel  $([0, t]) \otimes \mathcal{F}_t$ -measurable.

denote the set of reporting strategies in which the agent cannot overreport his shock.

There are two important restrictions that I can impose on strategies. First, without loss of generality, I can focus on reporting strategies with  $\Sigma^t \subseteq \Theta^t$  since otherwise, the planner would immediately be able to detect a lie. Second, and more importantly, let  $\mathbb{Q}$  denote the probability measure induced by an arbitrary reporting strategy (defined formally via the Radon-Nikodym derivative); then  $\mathbb{Q}$  must be absolutely continuous with respect to the true measure  $\mathbb{P}$ .<sup>10</sup> This is important because it rules out jumps in agents' reports: since Brownian motions are almost surely continuous, an agent's skill process is almost surely continuous. Therefore, if an agent reports truthfully until date  $t$  but the report suddenly jumps, then the planner can be sure the agent is lying and punish him. Obviously, this restriction would not hold if the shock was driven by a jump-diffusion process.

Given a history  $\theta^t \in \Theta^t$  and a reporting strategy  $\sigma \in \Sigma$ , let  $v_t^\sigma(\theta^t)$  denote an agent's promised/continuation utility:

$$v_t^\sigma(\theta^t) \equiv \mathbb{E}_t \left[ \int_t^\tau e^{-\rho(s-t)} \tilde{u}(\sigma_s(\theta^s)) ds \right].$$

Let  $v_t$  denote the value of  $v_t^\sigma$  under truth-telling. An allocation is *incentive-compatible* if truth-telling yields weakly higher continuation utility than any other reporting strategy:

$$v_t(\theta^t) \geq v_t^\sigma(\theta^t) \quad \forall \sigma \in \Sigma, \theta^t \in \Theta^t. \quad (2.2)$$

Note that this is not a single constraint, but a doubly-continuum of constraints that rules out deviations following any history of shocks and any history of strategies.<sup>11</sup>

As mentioned above, an agent will not participate in the game unless his utility from

---

<sup>10</sup>The measure  $\mathbb{Q}$  is absolutely continuous with respect to  $\mathbb{P}$ , written  $\mathbb{Q} \ll \mathbb{P}$ , if for every measurable set  $A \subseteq \Omega$ ,  $\mathbb{P}(A) = 0$  implies  $\mathbb{Q}(A) = 0$ .

<sup>11</sup>Discrete time models such as Fernandes and Phelan (2000) and Kapička (2013) often use a nearly equivalent concept called “temporary incentive compatibility” that explicitly rules out only one-period deviations.

an allocation exceeds  $\underline{U}$ , which yields the participation constraint

$$U(x) \geq \underline{U}. \quad (2.3)$$

Let  $\mathcal{X}^{\text{IC}}$  denote the set of incentive-compatible allocations that induce the agent to play the game, i.e., those allocations that satisfy the incentive-compatibility constraint (2.2) and participation constraint (2.3). This set is convex since  $\tilde{u}_t$  is affine.

The present value of the cost to the planner of delivering an allocation  $x$  is

$$\Psi(x) = \mathbb{E}_0 \left[ \int_0^\infty e^{-rt} \mathcal{K}(x_t(\theta^t)) dt \right].$$

Since the planner's goal is to find the cheapest allocation  $x \in \mathcal{X}^{\text{IC}}$ , his problem, written sequentially, is

$$K(\underline{U}) \equiv \min_{x \in \mathcal{X}^{\text{IC}}} \Psi(x).$$

An allocation is *efficient* if it attains this minimum and by Berge's Theorem of the Maximum, a solution necessarily exists and is continuous.

## 2.5 Relaxed Problem and First-Order Approach

The planner's problem as constructed above is hopelessly complicated because of the set of incentive constraints. Instead, I follow the first-order approach and replace (2.2) with a single first-order condition that is both necessary and sufficient for optimality. All proofs are in the Appendix.

I follow the approach outlined in Fernandes and Phelan (2000), Kapička (2013), and Pavan, Segal, and Toikka (2014) in discrete time, and Williams (2011) and San-nikov (2014) in continuous time, by reformulating the planner's problem recursively. To do so, I need to derive the law of motion for each of the planner's state variables. The first state variable is the agent's private information,  $\theta_t$ , whose law of motion is

(2.1). The second state variable is an agent's continuation utility,  $v_t$ , as is standard in mechanism design problems.<sup>12</sup> The planner uses continuation utility in order to track promise-keeping through time. My first result, which is standard in the continuous time optimal contracting literature, establishes the law of motion for  $v_t$  under an arbitrary contract.

**Proposition 2.1.** *Fix a contract and a reporting strategy  $\sigma$  with finite expected payoff to the agent. Then the process  $\{v_t\}$  corresponds to the agent's continuation utility if and only if there exists a process  $\hat{\Delta}_t \in H^2$  such that<sup>13</sup>*

$$dv_t = (\rho v_t - \tilde{u}_t) dt + \sigma_\theta \theta_t \hat{\Delta}_t dZ_t - v_t (dR_t - \kappa dt) \quad (2.4)$$

and the transversality condition  $\mathbb{E}_t [e^{-\rho(T \wedge \tau)} v_{T \wedge \tau}] \rightarrow 0$  as  $T \rightarrow \infty$  holds.<sup>14</sup>

The drift of (2.4) tracks promise-keeping:  $\tilde{u}_t$  is just-delivered utility and  $\rho v_t$  is everything owed going forward. The process  $\hat{\Delta}_t$  is the sensitivity of  $v_t$  to variations in the shock process. This also has a natural interpretation in terms of options. Thinking of the contract between the principal and the agent as a package of call options, then  $v_t$  represents the value of the options and  $\hat{\Delta}_t$  is the “delta” of the options, i.e., the sensitivity of the value to changes in the underlying shock process.

Since this is a dynamic mechanism design problem, however, I need a third state variable to track threat-keeping.<sup>15</sup> This is because misreporting changes the planner's beliefs about the distribution of the shock process, which affects what the agent can report to the planner at all future dates. In other words, violating the incentive

---

<sup>12</sup>See Thomas and Worrall (1988) and Sannikov (2008), for example.

<sup>13</sup>A process  $\{X_t\}$  is in the space  $H^2[0, T]$  if it is  $\mathcal{F}_t$ -progressively measurable and is square-integrable,  $\mathbb{E} \left[ \int_0^T X_t^2 dt \right] < \infty$ .

<sup>14</sup>In a finite-horizon economy, the transversality condition is replaced with a terminal boundary condition  $v_T = U(T)$  and (2.4) is a backwards stochastic differential equation with a known terminal condition but unknown initial condition.

<sup>15</sup>If the shock process was driven by a jump-diffusion process instead, then I would need another state variable to keep track of jumps and my results probably would not hold anymore.

constraint today affects future incentive constraints. To deter such deviations, the planner has to provide the agent with some information rent. Formally, following DeMarzo and Sannikov (2015), an agent’s information rent is the sensitivity of his continuation value to his report, evaluated at truth-telling:

$$\Delta_t \equiv \left. \frac{\partial v_t}{\partial \theta_t} \right|_{\hat{\theta}_t = \theta_t} = \frac{\partial v_t}{\partial \theta_t}. \quad (2.5)$$

This state variable is present in Fernandes and Phelan (2000), Williams (2011), Kapička (2013), and Pavan, Segal, and Toikka (2014) in settings with private information, and Sannikov (2014) in a setting in which private actions have long-run effects.<sup>16</sup> Pavan, Segal, and Toikka (2014) argue that this variable is like an impulse response function that “describe(s) how a change in the agent’s current type propagates through his type process.” Also, this derivative is a Malliavin derivative instead of an ordinary one. The Malliavin derivative is a way to formalize impulse response functions in continuous time by defining what it means to “differentiate” a Brownian motion, since the standard one-shot deviation principle does not have bite in these settings.<sup>17</sup> In fact, as the proof of Proposition 2.1 shows,  $\hat{\Delta}_t$  is the “black box” process in the Martingale Representation Theorem. Meanwhile, the Clark-Ocone-Haussman Theorem shows that this process is precisely the Malliavin derivative, thus explaining the link.

**Proposition 2.2.** *A necessary condition for truth-telling to be optimal is that,*

$$\hat{\Delta}_t = \Delta_t \quad (2.6)$$

---

<sup>16</sup>It is not a state variable, however, in Sannikov (2008) and other settings where hidden actions have one-time effects. There, it is well-known that the first-order approach is valid under standard monotonicity assumptions.

<sup>17</sup>See Di Nunno, Øksendal, and Proske (2009) for an introduction to Malliavin calculus and Borovička, Hansen, and Scheinkman (2014) for more on how to use Malliavin derivatives to compute impulse response functions; I use these methods on an optimal taxation problem in Chapter 1 of this dissertation.

for all  $t$ , where  $\widehat{\Delta}_t$  is the process in Proposition 2.1.

This is the first-order necessary condition for optimality and is quite intuitive: the sensitivity of continuation utility must equal the sensitivity evaluated at truth-telling. Let  $\mathcal{X}^{\text{FOA}}$  denote the set of allocations that satisfy (2.6), which necessarily satisfies  $\mathcal{X}^{\text{IC}} \subseteq \mathcal{X}^{\text{FOA}}$ . This is precisely the condition identified in discrete time by Pavan, Segal, and Toikka (2014) (see their Theorem 1) and in continuous time by Williams (2011) (see his equation (10)). However, Williams (2011) and I arrive at this result via different methods. He first transforms the state space using the Girsanov Theorem and then uses a version of the stochastic maximum principle to derive this expression directly from an agent's Hamiltonian. On the other hand, I follow Sannikov (2014) and use the Giransov Theorem to directly maximize an agent's utility function under small deviations from truth-telling, using Malliavin calculus to compute the derivatives.<sup>18</sup> It is worth noting that in static mechanism design settings such as Sannikov (2008), the derivation of the first-order condition establishes both necessity and sufficiency. That proof considers strategies that deviate up until date  $t$  and play truthfully after. However, in that setting, strategies played before  $t$  do not affect strategies played after  $t$ , which is clearly not the case here. Finally, if I restrict strategies so that  $\sigma \in \Sigma^*$ , then (2.6) becomes  $\widehat{\Delta}_t \leq \Delta_t$ .

With the first-order necessary condition in hand, I can derive the law of motion of  $\Delta_t$  under the first-order approach, i.e., the law of motion of information rent under the optimal contract.

**Proposition 2.3.** *The finite process  $\{\Delta_t\}$  is characterized by (2.5) if and only if, for some  $\Gamma_t \in H^2$ ,*

$$d\Delta_t = ((\rho - \mu'(\theta_t)) \Delta_t - \widetilde{u}_{\theta,t} - \sigma_{\theta}^2 \theta_t \Gamma_t) dt + \sigma_{\theta} \theta_t \Gamma_t dZ_t - \Delta_t (dR_t - \kappa dt) \quad (2.7)$$

---

<sup>18</sup>He, Wei, Yu, and Gao (2014) apply this method in a moral hazard model with learning, gaining considerable tractability by working with CARA utility.

*It follows that overreporting is never an issue because consumption, and hence utility, does not depend on the report.*

Unfortunately, outside of very special cases such as those above, it is difficult (if not impossible) to prove that upward incentive constraints are slack; if one could prove this, then Assumption 2.2 obviously would not be necessary.

**Theorem 2.1.** *Under Assumptions 2.1 and 2.2 and the transversality condition of Proposition 2.1, the optimal contract under the first-order approach is globally optimal.*

To my knowledge, this is the first sufficiency result for the first-order approach entirely in terms of model primitives and strategies. Pavan, Segal, and Toikka (2014) establish a sufficient condition in Markov environments but their condition is extremely complicated (see their Theorem 3) while Kapička (2013) derives a simpler, but still endogenous integral monotonicity condition that needs to be verified ex post (see his Section 6.3).<sup>21</sup> Finally, Williams (2011) shows that the first-order approach is sufficient for optimality if  $\theta_t \Gamma_t$  is appropriately bounded by an expression that depends on the utility function and other endogenous objects. This essentially amounts to limiting the volatility of agents' reports relative to the volatility of the shock process.<sup>22</sup> He shows that his condition can sometimes be simplified but that this simplification (and even his original condition) might be overly stringent and that the first-order approach can be sufficient even though the conditions fail.

So why do I obtain a sharper result than Williams (2011) and Kapička (2013) even though our settings are essentially identical? Since I work in continuous time, I have a more powerful set of tools than does the latter. Working in continuous time also imposes a lot more structure on reports. As for the former, he constructs

---

<sup>21</sup>Farhi and Werning (2013) and Golosov, Troshkin, and Tsyvinski (2016b) numerically check that his condition holds in their settings.

<sup>22</sup>I thank Drew Fudenberg for this interpretation.

a linear approximation to the utility gain from deviating, and uses the incentive constraint and his condition to ultimately show that the gain is negative. In contrast, I follow Sannikov (2014) and construct a *quadratic* approximation to the gain from deviating.<sup>23</sup> Using the law of motion (2.7), I show that Assumptions 2.1 and 2.2 are enough to guarantee that this gain is negative. He, Wei, Yu, and Gao (2014) and Di Tella and Sannikov (2016) use a similar verification method in their settings, and it is important to note that the appropriate approximation is very setting-specific and requires some guesswork.

The results in this section established that the planner's relaxed problem is equivalent to the original problem. For completeness' sake, I will complete the characterization of the first-order approach by explicitly writing out the planner's recursive problem. Recall that  $K(x)$  is the cost to the planner of delivering an allocation; rewrite this as  $K(\theta, v, \Delta)$  to emphasize the dependence on the state variables. Then the solution to the planner's recursive problem is obtained by solving his Hamilton–Jacobi–Bellman equation:

$$(r + \kappa) K = \inf_{\{x, \Gamma\}} \left\{ \mathcal{K}(x) + K_\theta \mu(\theta) + K_v [(\rho + \kappa)v - \tilde{u}] + K_\Delta [(\rho + \kappa - \mu'(\theta))\Delta - \tilde{u}_\theta - \sigma_\theta^2 \theta \Gamma] \right. \\ \left. + \frac{\sigma_\theta^2}{2} [K_{\theta\theta} + K_{vv} \theta^2 \Delta^2 + K_{\Delta\Delta} \theta^2 \Gamma^2] + \sigma_\theta^2 \theta^2 [K_{\theta v} \Delta + K_{\theta\Delta} \Gamma + K_{v\Delta} \Delta \Gamma] \right\}. \quad (2.8)$$

In general, this equation is very difficult to solve because the directions of the second derivatives depend on the choice of  $\Gamma$ . While Williams (2011) gives an example with a closed-form solution, in general one can simplify the problem by using duality to formulate an equivalent problem over the space of Lagrange multipliers. I defer to Sannikov (2014), DeMarzo and Sannikov (2015), and Chapter 1 of this dissertation (where I use duality to solve a dynamic Mirrlees model) for more on this procedure.

---

<sup>23</sup>He derives a very simple sufficient condition, but it still depends on a term similar to  $\Gamma_t$ . That said, his model has moral hazard instead of persistent private information.



## 2.6 Conclusion

This paper studies the first-order approach in a widely-used class of dynamic mechanism design models. My main results characterize the first-order necessary condition for optimality, showing that as long as an agent cannot overreport the value of his shock each period, this first-order condition is sufficient, too. This restriction is very reasonable for many types of dynamic mechanism design problems such as optimal taxation models (see Farhi and Werning (2013) and Golosov, Troshkin, and Tsyvinski (2016b), for example). My condition improves on those derived by Williams (2011) and Kapička (2013) in a nearly identical setting. I do so by utilizing a new set of techniques developed by Sannikov (2014).

It would be interesting to investigate whether the results in this paper are unique to the continuous time setting, or if they in fact do hold in discrete time. On one hand, the model is the same in both cases so the economics should not change but on the other hand, taking the time-step to zero and hence reducing the impact of any single report could have tangible impacts. In addition, in discrete time the process  $\{\theta_t\}$  can “jump” while here the reporting process must be sufficiently well-behaved. Second, while the class of models covered in this paper is widely-applicable, it does not cover settings such as career concerns models à la Holmström (1999), taste shock models like Atkeson and Lucas (1992), or buyer-seller games like in Battaglini and Lamba (2015) where agents have an incentive to overreport. Future work should look into extending the techniques/proofs in this setting to others and, in particular, determining whether the first-order approach is sufficient as long as agents are restricted to behaving as one would expect them to anyway.

- (2014): “House Price Gains and U.S. Household Spending from 2002 to 2006,” NBER Working Paper No. 20152.
- MIRRELES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38(2), 175–208.
- MOREIRA, A., AND A. SAVOV (2016): “The Macroeconomics of Shadow Banking,” NYU and Yale University Working Paper.
- NADAULD, T. D., AND S. M. SHERLUND (2013): “The Impact of Securitization on the Expansion of Subprime Credit,” *Journal of Financial Economics*, 107(2), 454–76.
- NUÑO, G., AND B. MOLL (2015): “Controlling a Distribution of Heterogeneous Agents,” Princeton University Working Paper.
- NUÑO, G., AND C. THOMAS (2014): “Bank Leverage Cycles,” European Central Bank Working Paper.
- ORDOÑEZ, G. (2015): “Sustainable Shadow Banking,” University of Pennsylvania Working Paper.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82(2), 601–53.
- PIKETTY, T., AND E. SAEZ (2003): “Income Inequality in the United States, 1913–1998,” *Quarterly Journal of Economics*, 118(1), 1–39.
- PISKORSKI, T., AND A. TCHISTYI (2010): “Optimal Mortgage Design,” *Review of Financial Studies*, 23(8), 3098–3140.
- PLANTIN, G. (2014): “Shadow Banking and Bank Capital Regulation,” *Review of Financial Studies*, 28(1), 146–75.
- POZSAR, Z., T. ADRIAN, A. ASHCRAFT, AND H. BOESKY (2010): “Shadow Banking,” Discussion Paper, Federal Reserve Bank of New York Staff Report No. 458.
- PRAT, J., AND B. JOVANOVIĆ (2014): “Dynamic Contracts When the Agent’s Quality is Unknown,” *Theoretical Economics*, 9(3), 865–914.
- PURNANANDAM, A. (2011): “Originate-to-Distribute Model and the Subprime Mortgage Crisis,” *Review of Financial Studies*, 24(6), 1881–1915.
- REED, W. J., AND M. JORGENSEN (2004): “The Double Pareto-Lognormal Distribution - A New Parametric Model for Size Distributions,” *Communications in Statistics - Theory and Methods*, 33(8), 1733–53.
- REVUZ, D., AND M. YOR (1999): *Continuous Martingales and Brownian Motion*. Springer.