

Volume 30, Issue 1

Bayesian-Nash vs dominant-strategy implementation with countervailing incentives: the two-type case

Annalisa Vinella
University of Bari

Abstract

We extend the principal/one-agent model with countervailing incentives to a framework in which the principal deals with two agents behaving non-cooperatively and protected by limited liability. Focusing on the two-type case, we show that, beside the situation in which first best is effected even without relying on type correlation, dominant-strategy implementation yields no penalty to the principal, with respect to Bayesian-Nash implementation, when the principal faces, on the opposite, very tight constraints.

I am indebted to David Martimort and Jérôme Pouyet for fruitful advice. I also wish to thank Jacques Crémer for further hints. The usual disclaimer applies.

Citation: Annalisa Vinella, (2010) "Bayesian-Nash vs dominant-strategy implementation with countervailing incentives: the two-type case", *Economics Bulletin*, Vol. 30 no.1 pp. 636-649.

Submitted: Jan 06 2010. **Published:** March 02, 2010.

1 Introduction

The literature about mechanism design in multi-agent environments relies widely on the notion of Bayesian-Nash (BN hereafter) implementation. This notion rests on important common knowledge assumptions, namely those regarding one agent's probability assessment about the other agents' private information. In practice, this may raise difficulties. Indeed, as evidenced by Wilson (1987) and reaffirmed thereafter by several authors hinging on the so-called Wilson Doctrine, BN incentive compatible mechanisms may yield unsatisfactory outcomes if the common knowledge hypotheses on which they stand are actually false.

One possibility for a principal to avoid dependence on critical common knowledge assumptions is to resort to the stronger notion of dominant-strategy (DS hereafter) implementation. DS mechanisms are simpler and thus practically more manageable¹. However, being based on a stronger incentive compatibility concept, they may induce less desirable achievements, as compared to BN mechanisms.

Following the Wilson Doctrine, the literature has identified circumstances under which a DS mechanism exactly replicates the outcome that would arise with a BN mechanism with regard to a rich variety of economic situations (within this domain of research, recall Mookherjee and Reichelstein 1992, Segal 2003, Chung and Ely 2007, for instance). Nonetheless, all existing studies focus on agency relationships in which agents display systematic incentives to over/under-report private information to the principal. Hence, the conclusions they achieve do not provide a clear clue about frameworks in which, by contrast, agents may face countervailing incentives to misrepresent information.

That an agent may face countervailing incentives means that, in the words of Lewis and Sappington (1989), he "may be tempted both to overstate and to understate his private information, depending upon its specific realization" (p.294). As Lewis and Sappington (1989) and Maggi and Rodriguez-Clare (1995) suggest, this may occur, for instance, in regulator-firm relationships when the firm/agent's production technology is such that the fixed cost depends negatively on the marginal cost, which is unknown to the regulator/principal. So far contractual performance in the presence of countervailing incentives has been explored with regard to single-agent environments². Of course, in the latter, no issue arises about the implementation notion to which the principal should refer.

In this note we extend the analysis of contract design in the presence of countervailing incentives to settings where the principal deals with two agents, each of whom may display countervailing incentives. As we describe in section 1, our focus is on the specific case in which the information structure is binary, the privately known types are positively correlated and the agents are protected by limited liability. Our contribution, to be presented in section 2 and 3, is twofold. We first evidence that, in the settings described so far, there are circumstances under which, contrary to standard literature results, information correlation is valueless for the principal's achievements. This occurs when the presence of countervailing incentives relaxes the principal's problem to the point that

¹Alternatively, a principal could ask agents to report not only their own information, but also their probability assessment about other agents' private information. However, this would make the mechanisms more complex. On the other hand, concerns about the principal's ability to commit would arise if, instead, the principal were to dictate herself beliefs to agents.

²Beside Lewis and Sappington (1989) and Maggi and Rodriguez-Clare (1995), to this domain of literature belong Lewis and Sappington (1989b), Brainard and Martimort (1996) and Jullien (2000), for instance.

the first-best outcome would be effected even with a single agent (or with independent types). Of course, under those circumstances, the implementation notion the principal adopts does not matter, provided it does not affect the equilibrium outcome. We then show that replacing the BN contract with a DS contract yields no penalty to the principal also in the opposite case in which her programme is so constrained that most important distortions are to be induced. By contrast, in this case, information correlation is useful to the principal precisely because of the presence of tight constraints.

2 The two-type two-agent model

We consider a risk-neutral principal who contracts with two risk-neutral agents for the provision of a good. The task of agent $i = 1, 2$ is to provide q_i units of the good. Following Lewis and Sappington (1989), production costs are given by

$$C_i(q_i; \theta_i) = \theta_i q_i + c(\theta_i). \quad (1)$$

This means that agent i produces at marginal cost θ_i and bears a fixed cost $c(\theta_i)$ that depends negatively on the realization of θ_i . That is, high (low) fixed cost is associated with low (high) variable unit cost.

At the contracting stage, agent $i \in \{1, 2\}$ is privately informed about θ_i (his type). It is commonly known that, $\forall i \in \{1, 2\}$, $\theta_i \in \Theta = \{\theta_h, \theta_l\}$, with $\theta_h > \theta_l$, and that $c(\theta_l) = c_h$ and $c(\theta_h) = c_l$, with $c_h > c_l$. We denote $\Delta c = c_h - c_l$ and $\Delta\theta = \theta_h - \theta_l$. Prior beliefs are given by $\nu_{fg} = \Pr(\theta_1 = \theta_f, \theta_2 = \theta_g)$, $\forall f, g \in \{h, l\}$, with $\nu_{lh} = \nu_{hl}$ for simplicity. The degree of type correlation is $\rho \equiv \nu_{ll}\nu_{hh} - \nu_{lh}^2 > 0$.

Agents behave non-cooperatively. The principal ties them both in the offer of a unique grand-contract. Under the latter, agent i receives a transfer s_i for the production of q_i units of the good and obtains the profit $\pi_i(q_i, s_i) = s_i - [\theta_i q_i + c(\theta_i)]$. Production yields to the principal the gross benefit $\sum_{i=1,2} V(q_i)$, with $V(0) = 0$, $V' > 0$, $V'' < 0$, $V'(0) = +\infty$, $V'(+\infty) = 0$. The principal's utility is given by the gross benefit net of transfers, namely $U(q_1, q_2, s_1, s_2) = \sum_{i=1,2} [V(q_i) - s_i]$.

2.1 The principal's programme

The Revelation Principle applies. The principal focuses on truthful direct revelation mechanisms that include a lottery with two quantity-profit pairs for each possible type, namely $(\{q_{ff}, \pi_{ff}\}, \{q_{fg}, \pi_{fg}\})$, with $q_{ff} \equiv q_i(\theta_f, \theta_f)$ and $q_{fg} \equiv q_i(\theta_f, \theta_g) = q_j(\theta_g, \theta_f)$, $\forall f \neq g \in \{h, l\}$, $\forall i \neq j \in \{1, 2\}$, and similarly for profits. The specific mechanism to be chosen and, hence, the outcome to follow might depend on whether the principal prefers to rest on a weaker or tighter concept of incentive compatibility. To allow for both possibilities, we consider the two notions of BN and DS implementation.

2.1.1 Bayesian-Nash implementation

Suppose first that the principal wishes to induce truthtelling as a BN equilibrium. Then, for each agent, reporting truthfully is to yield the highest expected payoff, given his own type. Define $\Pr(\theta_j = \theta_f | \theta_i = \theta_f) = \nu_{ff} / (\nu_{ff} + \nu_{fg})$, $\forall i \neq j \in \{1, 2\}$, $\forall f \neq g \in \{h, l\}$, the (commonly known) posterior beliefs that agent i uses to evaluate his expected payoff when he has type θ_f . Further let $\Pi_f \equiv \nu_{ff}\pi_{ff} + \nu_{fg}\pi_{fg}$, $\forall f \neq g \in \{h, l\}$, the

expected rent of an agent of type θ_f . The principal selects the mechanism that maximizes her expected utility

$$\widehat{U} = 2 \sum_{f \neq g} \nu_{ff} [V(q_{ff}) - (c_g + \theta_f q_{ff}) - \pi_{ff}] + 2 \sum_{f \neq g} \nu_{fg} [V(q_{fg}) - (c_g + \theta_f q_{fg}) - \pi_{fg}]$$

subject to the incentive constraints

$$\begin{aligned} IC_l^{BN} &: \Pi_l \geq \nu_{ll}\pi_{hl} + \nu_{lh}\pi_{hh} + \Delta\theta(\nu_{ll}q_{hl} + \nu_{lh}q_{hh}) - \Delta c(\nu_{ll} + \nu_{lh}) \\ IC_h^{BN} &: \Pi_h \geq \nu_{lh}\pi_{ll} + \nu_{hh}\pi_{lh} - \Delta\theta(\nu_{lh}q_{ll} + \nu_{hh}q_{lh}) + \Delta c(\nu_{lh} + \nu_{hh}), \end{aligned}$$

the participation constraints

$$PC_f : \Pi_f \geq 0, \quad \forall f \in \{h, l\},$$

and the limited liability constraints

$$LL_{fg} : \pi_{fg} \geq -m, \quad \forall f, g \in \{h, l\},$$

with $0 \leq m < +\infty$. For future reference, denote this programme P^{BN} .

2.1.2 Dominant-strategy implementation

Suppose next that the principal wishes to induce truthtelling as a DS. Then, for each type of either agent, reporting truthfully is to yield the highest payoff for each possible report of the other agent. The principal's programme is analogous to P^{BN} , except that IC_l^{BN} and IC_h^{BN} are replaced by

$$\begin{aligned} IC_l^{DS} &: \pi_{ll} \geq \pi_{hl} + \Delta\theta q_{hl} - \Delta c \\ IC_{lh}^{DS} &: \pi_{lh} \geq \pi_{hh} + \Delta\theta q_{hh} - \Delta c \\ IC_{hl}^{DS} &: \pi_{hl} \geq \pi_{ll} - \Delta\theta q_{ll} + \Delta c \\ IC_{hh}^{DS} &: \pi_{hh} \geq \pi_{lh} - \Delta\theta q_{lh} + \Delta c. \end{aligned}$$

Denote this programme P^{DS} .

2.2 The first-best outcome

At the first-best outcome (FB hereafter), quantities are pinned down such that the marginal benefit equals the marginal cost:

$$V'(q_f^*) = \theta_f, \quad \forall f \in \{h, l\}, \quad (3)$$

with $q_f^* \equiv q_{ff}^* = q_{fg}^*$, $f \neq g$. Profits are set so as to leave no expected rent to agents:

$$\Pi_f^* = 0, \quad \forall f \in \{h, l\}. \quad (4)$$

3 Results

An established result in the economic literature is that the principal benefits from the presence of a second agent who shares with the first agent, even only stochastically, some information that is unknown to her. This is not necessarily the case in two-type settings in which agents display countervailing incentives.

Proposition 1. *As long as $\Delta c \in [\Delta\theta q_h^*, \Delta\theta q_l^*]$, FB is implemented even if the principal faces a single agent (or two agents with independent types) so that correlated information yields no benefit to the principal.*

To state Proposition 1, one does not even need to solve P^{BN} and P^{DS} . It suffices to look at the situation in which the principal faces one sole agent who exhibits countervailing incentives, which is a two-type version of the problem studied by Lewis and Sappington (1989). As specified in Proposition 1, in the single-agent (or uncorrelated information) setting with binary information structure, the optimal contract effects FB for all values of Δc that belong to the interval $[\Delta\theta q_h^*, \Delta\theta q_l^*]$ ³. Actually, for such values, any gain that could be obtained by misreporting the value of θ would be offset by an associated loss in terms of fixed cost. Incentives to cheat are thus costlessly removed. Of course, this outcome can always be replicated, for the same values of Δc , in a correlated information environment, even when agents are protected by limited liability. It follows that, under these circumstances, the ability to contract with a second agent with correlated type yields no benefit to the principal, contrary to standard literature findings.

Another established result in the literature is that, in correlated information settings, *a priori*, BN implementation makes the principal weakly better off as compared to DS implementation. This occurs because the former requires a weaker standard of incentive compatibility. The principal thus enjoys a greater discretion at specifying payments and can take a bigger advantage of type correlation. It follows that, with a BN mechanism, the principal can (at least) replicate any outcome that she can achieve in DS. In the specific framework here explored, whether BN implementation brings an extra gain to the principal depends on two elements, namely the value of m and that of Δc .

The result that BN implementation makes the principal weakly better off carries over as long as m is sufficiently large. In that case, limited liability constraints are not binding and, whatever the value of Δc , the principal can design a BN contract that effects FB. Actually, this possibility is not specific to environments characterized by the presence of countervailing incentives. In a correlated information setting with systematic incentives to over-report type, Gary-Bobo and Spiegel (2006) show that the principal can recommend the efficient production and extract surplus with an appropriate BN contract, provided that liability is sufficiently large⁴. Analogous possibility is not at hand with a DS contract. In the specific situation we model, FB is implemented in DS thanks to the presence of countervailing incentive, precisely as in single-agent frameworks.

A more interesting case arises when limited liability constraints are tight (m small). Then, the peculiar nature of agents' incentives imposes more structure on the optimal

³Details about the single-agent setting and the optimal single-agent contract are reported in Appendix A. In particular, the case of FB implementation appears in P_3^S in Appendix A.2.

⁴In turn, this outcome rests on the classical finding that FB is implemented in BN settings with correlated information and no limited liability concerns (Cr mer and McLean 1988 and Riordan and Sappington 1988).

contract and thus affects the comparison between implementation concepts⁵.

Proposition 2. *With m sufficiently small, BN implementation yields no additional benefit to the principal as compared to DS implementation:*

(i) *whenever $\Delta c < \Delta\theta Q_{hl}$, with Q_{hl} as pinned down in (5a) below, in which case the optimal grand-contract entails:*

(i/a) *for type θ_h , surplus extraction and downward production distortions*

$$V'(Q_{hl}) = \theta_h + \frac{\nu_{ll}}{\nu_{lh}} \Delta\theta \quad (5a)$$

$$V'(Q_{hh}) = \theta_h + \frac{\nu_{lh}}{\nu_{hh}} \Delta\theta; \quad (5b)$$

(i/b) *for type θ_l , efficient production and information rent*

$$\Pi_l = \Delta\theta (\nu_{ll} Q_{hl} + \nu_{lh} Q_{hh}) - \Delta c (\nu_{ll} + \nu_{lh}) - \frac{\rho m}{\nu_{hh}}; \quad (6)$$

(ii) *whenever $\Delta c \in [\Delta\theta q_h^*, \Delta\theta q_l^*]$, in which case the optimal grand-contract effects FB;*

(iii) *whenever $\Delta c > \Delta\theta Q_{lh}$, with Q_{lh} as pinned down in (7b) below, in which case the optimal grand-contract entails:*

(iii/a) *for type θ_l , surplus extraction and upward production distortions*

$$V'(Q_{ll}) = \theta_l - \frac{\nu_{lh}}{\nu_{ll}} \Delta\theta \quad (7a)$$

$$V'(Q_{lh}) = \theta_l - \frac{\nu_{hh}}{\nu_{lh}} \Delta\theta; \quad (7b)$$

(iii/b) *for type θ_h , efficient production and information rent*

$$\Pi_h = \Delta c (\nu_{lh} + \nu_{hh}) - \Delta\theta (\nu_{hh} Q_{lh} + \nu_{lh} Q_{ll}) - \frac{\rho m}{\nu_{ll}}. \quad (8)$$

The circumstance that, in multi-agent environments with correlated types and limited liability, FB is at hand for $\Delta c \in [\Delta\theta q_h^*, \Delta\theta q_l^*]$ whether the principal resorts to BN or DS implementation, follows directly from Proposition 1. Provided full efficiency entails for such values even with a single agent (or with two agents whose types are independent), there is nothing more the principal can do in a correlated information setting. Therefore, the possibility to better exploit type correlation with a BN mechanism is valueless to the principal.

Proposition 2 further evidences that, as long as m is small enough, the optimal grand-contract is robust to variations in the equilibrium concept not only when Δc takes intermediate values, but also when it takes very small and very large values. Noticeably, the latter are both situations in which it is particularly difficult to induce truth-telling because agents' types display a clear incentive to misreport in one specific direction. Indeed, the gain that could be obtained by misrepresenting θ in that direction would more

⁵See Appendix B.1 and B.2 for a full presentation of the optimal BN and DS mechanism with m small.

than compensate the associated loss in terms of fixed cost. Specifically, type θ_l is tempted to over-report in situation (i), whereas type θ_h is tempted to under-report in situation (iii). Under these circumstances, the principal faces a (nearly) standard adverse selection problem, in which quantities are optimally distorted to decrease information rents⁶.

Figure 1 and 2 provide a graphical representation of the path that quantities follow as Δc takes different values respectively in the optimal BN and DS grand-contract when m is small. To better illustrate the content of Proposition 1, in Figure 1 the BN quantity path (the black line) is contrasted with that in the optimal single-agent contract (the red line). A comparison of the two graphs evidences that the BN quantities coincide with the DS ones for very small, intermediate and very large values of Δc , as previously pointed out.

A neat message ensues from Proposition 2. In multi-agent environments with a binary information structure and severe liability limits, tightening the equilibrium concept from BN to DS yields no loss of generality to the principal as long as agents display either sufficiently important countervailing incentives (case (ii)) or sufficiently strong systematic incentives to misreport type (case (i) and (iii)).

This conclusion is no longer true for case (i) and (iii) if m is still sufficiently large that, although FB is beyond reach for those values of Δc , participation constraints are binding. Then, in the optimal BN contract, quantities are distorted just enough to ensure participation without compromising information release. That is, the principal takes for case (i) and (iii) the same approach that she adopts for all the other values of Δc for which FB is not effected⁷. Noticeably, this approach is only at hand in environments characterized by the presence of countervailing incentives. Under DS, even if agents can be punished enough to remove all rents, the principal is compelled to pick one among the allocations that are feasible in the BN framework. Specifically, the optimal allocation is such that distortions are contained in the quantities that are more likely to be assigned (namely, with positive correlation, q_{hh} and q_{ll} respectively). Yet, distortions remain more important in the other quantities so as to warrant incentive compatibility. The BN allocation is thus at least as efficient as the DS allocation. The result that BN implementation weakly dominates is restored.

4 Concluding remarks

Overall, our analysis predicts that, in correlated two-type settings where agents may display countervailing incentives and are protected by limited liability, the optimal grand-contract is robust to variations in the implementation notion (BN *vs* DS) both when the principal's programme is so relaxed that FB would be at hand even with a single agent (or with two agents whose types are independent) and when it is so constrained that the principal is to induce most important distortions. Actually, under both these circumstances, the principal can run a mechanism that induces truth-telling in DS and still secure the outcome that she would achieve with the BN optimal contract.

The lesson our study delivers appears to be useful with regard to situations in which the principal is concerned with the robustness of the BN contract. This concern may arise because BN incentive-compatible mechanisms may perform in unsatisfactory ways

⁶In situation (i) (resp. (iii)) the rent that type θ_l (resp. θ_h) obtains is decreased by distorting *downward* (resp. *upward*) the quantities assigned to type θ_h (resp. θ_l).

⁷Check P_2^{BN} and P_4^{BN} in Appendix B.1.

if the common knowledge assumptions on the agents' posterior beliefs, on which they critically rests, turn out to be incorrect⁸. Although restricted to the simplest possible information structure, our investigation identifies circumstances under which the principal can costlessly escape reliance on common knowledge assumptions in frameworks in which agents display countervailing incentives to misreport type.

The observations made so far evidence a similarity between correlated information contexts characterized by the presence of countervailing incentives and independent information contexts. Still referring to standard situations with systematic incentives to over/under-report, Mookherjee and Reichelstein (1992) show that implementation in DS entails no loss of generality, with respect to BN implementation, provided agents' types are independent.

The existence of this similarity with independent information settings should not convey the idea that, in environments in which countervailing incentives arise and agents are protected by limited liability, correlated information is equally (ir)relevant for the principal whenever BN and DS achievements coincide (cases (i) to (iii) in Proposition 2). Because countervailing incentives tend to relax the principal's programme, whereas limited liability operates in the opposite direction, the role of correlation depends ultimately on how intense those incentives are. Correlation is of no value as long as the presence of countervailing incentives relaxes the principal's programme to the point that FB entails even if only small deficits can be imposed *ex post* (Proposition 1 and case (ii) in Proposition 2). That is, sufficiently important countervailing incentives can work as a substitute for correlated information. It is thus rather intuitive that correlated information comes back to be useful when such incentives are not in place. Actually, it does improve contractual performance as soon as systematic incentives to misreport prevail, in which case the constraints the principal faces are especially severe (case (i) and (iii) in Proposition 2)⁹.

⁸Compare, for instance, Chung and Ely (2007) who provide foundations for a principal to focus on DS mechanisms whenever she is unwilling to make strong common knowledge hypotheses about agents' conjectures.

⁹To check this point it suffices to compare the parts of the optimal BN and DS grand-contract presented in the Proposition with P_1^S and P_5^S in Appendix A.2.




Figure 1: Quantity path in the BN grand-contract with m small (black), as contrasted with that in the single-agent optimal contract (red), with $\alpha \equiv \Delta\theta \left(\frac{\nu_{ll}Q_{hl} + \nu_{lh}Q_{hh}}{\nu_{ll} + \nu_{lh}} \right) - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})}$, $\beta \equiv \Delta\theta q_h^* - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})}$, $\gamma \equiv \Delta\theta q_l^* + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})}$ and $\delta \equiv \Delta\theta \left(\frac{\nu_{hh}Q_{lh} + \nu_{lh}Q_{ll}}{\nu_{lh} + \nu_{hh}} \right) + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})}$




Figure 2: Quantity path in the DS grand-contract with m small and $\eta \equiv \Delta\theta Q_{hl}$, $\iota \equiv \Delta\theta Q_{hh} - \frac{\rho m}{\nu_{lh}\nu_{hh}}$, $\kappa \equiv \Delta\theta q_h^* - \frac{\rho m}{\nu_{lh}\nu_{hh}}$, $\mu \equiv \Delta\theta q_h^*$, $\varpi \equiv \Delta\theta q_l^*$, $\vartheta \equiv \Delta\theta q_l^* + \frac{\rho m}{\nu_{ll}\nu_{lh}}$, $\sigma \equiv \Delta\theta Q_{ll} + \frac{\rho m}{\nu_{ll}\nu_{lh}}$, $\zeta \equiv \Delta\theta Q_{lh}$.

References

- Brainard, S.L. and D. Martimort (1996), "Strategic Trade Policy Design with Asymmetric Information and Public Contracts", *The Review of Economic Studies* 63(1), 81-105
- Chung, K.S. and J.C. Ely (2007), "Foundations of Dominant Strategy Mechanisms", *Review of Economic Studies* 74(2), 447-476
- Crémer, J. and R. McLean (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions", *Econometrica* 56(6), 1247-1258
- Gary-Bobo, R. and Y. Spiegel (2006), "Optimal State-Contingent Regulation under Limited Liability", *RAND Journal of Economics* 37, 431-448
- Jullien, B. (2000), "Participation Constraints in Adverse Selection Models", *Journal of Economic Theory* 93(1), 1-47
- Lewis, T.R. and D.E.M. Sappington (1989), "Countervailing Incentives in Agency Problems", *Journal of Economic Theory* 49, 294-313
- Lewis, T.R. and D.E.M. Sappington (1989b), "Inflexible Rules in Incentive Problems", *American Economic Review* 79(1), 69-84
- Maggi, G. and A. Rodriguez-Clare (1995), "On Countervailing Incentives", *Journal of Economic Theory* 66, 238-263
- Mookherjee, D. and S. Reichelstein (1992), "Dominant strategy implementation of Bayesian incentive compatible allocation rules", *Journal of Economic Theory* 56(2), 378-399
- Riordan, M. and D.E.M. Sappington (1988), "Optimal Contracts with Public Ex Post Information", *Journal of Economic Theory* 45, 189-199
- Segal, I. (2003), "Optimal Pricing Mechanisms with Unknown Demand", *American Economic Review* 93(3), 509-529
- Wilson, R. (1987), "Game-Theoretic Approaches to Trading Processes," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, 33-77, Cambridge University Press, Cambridge

A The single-agent environment

A.1 The principal's programme

In the single-agent setting, in which type is θ_l and θ_h with probability v and $(1 - v)$ respectively, the principal picks the menu of contracts $\{q_f, \pi_f\}$, $\forall f \in \{h, l\}$, that maximizes expected utility

$$U^S = v [V(q_l) - (c_h + \theta_l q_l)] + (1 - v) [V(q_h) - (c_l + \theta_h q_h)] - [v\pi_l + (1 - v)\pi_h]$$

subject to the agent's incentive constraints

$$\begin{aligned} IC_l^S &: \pi_l \geq \pi_h + \Delta\theta q_h - \Delta c \\ IC_h^S &: \pi_h \geq \pi_l - \Delta\theta q_l + \Delta c \end{aligned}$$

and participation constraints

$$PC_f^S : \pi_f \geq 0, \quad \forall f \in \{h, l\}.$$

Denote this programme P^S .

A.2 The optimal contract

For each agent's type, the first-best outcome (FB hereafter) is characterized by:

$$V'(q_f^*) = \theta_f \tag{9a}$$

$$\pi_f^* = 0. \tag{9b}$$

The contract Γ^S that solves P^S is characterized as in $P_1^S - P_5^S$ below.

P_1^S) For $\Delta c \in [0, \Delta\theta\tilde{q}_h)$, with $\tilde{q}_h = q_h$ as pinned down by (10a) below, Γ^S entails (9a) for $f = l$, (9b) for $f = h$ and

$$V'(q_h) = \theta_h + \frac{v}{1-v}\Delta\theta \tag{10a}$$

$$\pi_l = \Delta\theta\tilde{q}_h - \Delta c. \tag{10b}$$

P_2^S) For $\Delta c \in (\Delta\theta\tilde{q}_h, \Delta\theta q_h^*)$, Γ^S entails (9a) for $f = l$, (9b) for $f = l, h$ and

$$q_h = \frac{\Delta c}{\Delta\theta} \tag{11}$$

P_3^S) For $\Delta c \in [\Delta\theta q_h^*, \Delta\theta q_l^*)$, Γ^S implements FB.

P_4^S) For $\Delta c \in (\Delta\theta q_l^*, \Delta\theta\tilde{q}_l)$, with $\tilde{q}_l = q_l$ as pinned down by (13a) below, Γ^S entails (9a) for $f = h$, (9b) for $f = l, h$ and

$$q_l = \frac{\Delta c}{\Delta\theta}. \tag{12}$$

P_5^S) For $\Delta c > \Delta\theta\tilde{q}_l$, Γ^S entails (9a) for $f = h$, (9b) for $f = l$ and

$$V'(q_l) = \theta_l - \frac{1-v}{v}\Delta\theta \quad (13a)$$

$$\pi_h = \Delta c - \Delta\theta\tilde{q}_l. \quad (13b)$$

B The multi-agent environment

B.1 The Bayesian-Nash grand-contract (with m small)

The grand-contract that solves P^{BN} , to be denoted Γ^{BN} , is characterized as in $P_1^{BN} - P_5^{BN}$ below.

P_1^{BN}) For $\Delta c < \Delta\theta \left(\frac{\nu_{ll}Q_{hl} + \nu_{lh}Q_{hh}}{\nu_{ll} + \nu_{lh}} \right) - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})}$, with Q_{hl} and Q_{hh} as pinned down by (14a) below, Γ^{BN} entails (3) for $f = l$, $g = h$, (4) for $f = h$ and

$$V'(Q_{hl}) = \theta_h + \frac{\nu_{ll}}{\nu_{lh}}\Delta\theta; \quad V'(Q_{hh}) = \theta_h + \frac{\nu_{lh}}{\nu_{hh}}\Delta\theta \quad (14a)$$

$$\Pi_l = \Delta\theta(\nu_{ll}Q_{hl} + \nu_{lh}Q_{hh}) - \Delta c(\nu_{ll} + \nu_{lh}) - \frac{\rho m}{\nu_{hh}}. \quad (14b)$$

Profits are defined as

$$\pi_{hl} = -m; \quad \pi_{hh} = \frac{\nu_{lh}}{\nu_{hh}}m. \quad (15)$$

P_2^{BN}) For $\Delta c \in \left[\Delta\theta \left(\frac{\nu_{ll}Q_{hl} + \nu_{lh}Q_{hh}}{\nu_{ll} + \nu_{lh}} \right) - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})}, \Delta\theta q_h^* - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})} \right)$, Γ^{BN} entails (3) for $f = l$, $g = h$, (4) $\forall f \in \{l, h\}$ and

$$\hat{q}_h \equiv \nu_{ll}q_{hl} + \nu_{lh}q_{hh} = \frac{1}{\Delta\theta} \left[\Delta c(\nu_{ll} + \nu_{lh}) + \frac{\rho m}{\nu_{hh}} \right]. \quad (16)$$

Profits are defined as in (15).

P_3^{BN}) For $\Delta c \in \left[\Delta\theta q_h^* - \frac{\rho m}{\nu_{hh}(\nu_{ll} + \nu_{lh})}, \Delta\theta q_l^* + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})} \right]$, Γ^{BN} implements FB.

P_4^{BN}) For $\Delta c \in \left(\Delta\theta q_l^* + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})}, \Delta\theta \left(\frac{\nu_{hh}Q_{lh} + \nu_{lh}Q_{ll}}{\nu_{lh} + \nu_{hh}} \right) + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})} \right]$, with Q_{ll} and Q_{lh} as pinned down by (19a) below, Γ^{BN} entails (3) for $f = h$, $g = l$, (4) $\forall f \in \{l, h\}$ and

$$\hat{q}_l \equiv \nu_{hh}q_{lh} + \nu_{lh}q_{ll} = \frac{1}{\Delta\theta} \left[\Delta c(\nu_{lh} + \nu_{hh}) - \frac{\rho m}{\nu_{ll}} \right]. \quad (17)$$

Profits are defined as

$$\pi_{ll} = \frac{\nu_{lh}}{\nu_{ll}}m; \quad \pi_{lh} = -m. \quad (18)$$

P_5^{BN}) For $\Delta c > \Delta\theta \left(\frac{\nu_{hh}Q_{lh} + \nu_{lh}Q_{ll}}{\nu_{lh} + \nu_{hh}} \right) + \frac{\rho m}{\nu_{ll}(\nu_{lh} + \nu_{hh})}$, Γ^{BN} entails (3) for $f = h$, $g = l$, (4) for $f = l$ and

$$V'(Q_{ll}) = \theta_l - \frac{\nu_{lh}}{\nu_{ll}}\Delta\theta; \quad V'(Q_{lh}) = \theta_l - \frac{\nu_{hh}}{\nu_{lh}}\Delta\theta \quad (19a)$$

$$\Pi_h = \Delta c(\nu_{lh} + \nu_{hh}) - \Delta\theta(\nu_{hh}Q_{lh} + \nu_{lh}Q_{ll}) - \frac{\rho m}{\nu_{ll}}. \quad (19b)$$

Profits are defined as in (18).

B.2 The dominant-strategy grand-contract (with m small)

The grand-contract that solves P^{DS} , to be denoted Γ^{DS} , is characterized as in $P_1^{DS} - P_9^{DS}$ below.

P_1^{DS}) For $\Delta c \in [0, \Delta\theta Q_{hl}]$, Γ^{DS} entails (3) for $f = l, g = h$, (14a), (4) for $f = h$ and (14b). Profits are defined as in (15) and as

$$\pi_{ll} = \Delta\theta Q_{hl} - \Delta c - m; \quad \pi_{lh} = \Delta\theta Q_{hh} - \Delta c + \frac{\nu_{lh}}{\nu_{hh}}m. \quad (20)$$

P_2^{DS}) For $\Delta c \in \left(\Delta\theta Q_{hl}, \Delta\theta Q_{hh} - \frac{\rho m}{\nu_{lh}\nu_{hh}}\right)$, Γ^{DS} entails (3) for $f = l, g = h$, (4) for $f = h$ and

$$q_{hl} = \frac{\Delta c}{\Delta\theta}; \quad q_{hh} = Q_{hh} \quad (21a)$$

$$\Pi_l = \nu_{lh}(\Delta\theta Q_{hh} - \Delta c) - \frac{\rho m}{\nu_{hh}}. \quad (21b)$$

Profits are defined as in (15) and as

$$\pi_{ll} = -m; \quad \pi_{lh} = \Delta\theta Q_{hh} - \Delta c + \frac{\nu_{lh}}{\nu_{hh}}m. \quad (22)$$

P_3^{DS}) For $\Delta c \in \left(\Delta\theta Q_{hh} - \frac{\rho m}{\nu_{lh}\nu_{hh}}, \Delta\theta q_h^* - \frac{\rho m}{\nu_{lh}\nu_{hh}}\right)$, Γ^{DS} entails (3) for $f = l, g = h$, (4) for $f = l, h$, and

$$q_{hl} = \frac{\Delta c}{\Delta\theta}; \quad q_{hh} = \frac{1}{\Delta\theta} \left(\Delta c + \frac{\rho m}{\nu_{lh}\nu_{hh}} \right). \quad (23)$$

Profits are defined as in (15) and as

$$\pi_{ll} = -m; \quad \pi_{lh} = \frac{\nu_{ll}}{\nu_{lh}}m. \quad (24)$$

P_4^{DS}) For $\Delta c \in \left(\Delta\theta q_h^* - \frac{\rho m}{\nu_{lh}\nu_{hh}}, \Delta\theta q_l^*\right)$, Γ^{DS} entails (3) for $f = l, g = h$, (4) for $f = l, h$, and

$$q_{hl} = \frac{\Delta c}{\Delta\theta}; \quad q_{hh} = q_h^*. \quad (25)$$

Profits are defined as in (15) and (24).

P_5^{DS}) For $\Delta c \in (\Delta\theta q_h^*, \Delta\theta q_l^*)$, Γ^{DS} implements FB.

P_6^{DS}) For $\Delta c \in \left(\Delta\theta q_l^*, \Delta\theta q_l^* + \frac{\rho m}{\nu_{ll}\nu_{lh}}\right)$, Γ^{DS} entails (3) for $f = h, g = l$, (4) for $f = l, h$ and

$$q_{ll} = q_l^*; \quad q_{lh} = \frac{\Delta c}{\Delta\theta} \quad (26)$$

Profits are defined as in (18) and as

$$\pi_{hl} = \frac{\nu_{hh}}{\nu_{lh}}m; \quad \pi_{hh} = -m. \quad (27)$$

P_7^{DS}) For $\Delta c \in \left(\Delta\theta q_l^* + \frac{\rho m}{\nu_u \nu_{lh}}, \Delta\theta Q_u + \frac{\rho m}{\nu_u \nu_{lh}} \right)$, Γ^{DS} entails (3) for $f = h, g = l$, (4) for $f = l, h$ and

$$q_{ul} = \frac{1}{\Delta\theta} \left(\Delta c - \frac{\rho m}{\nu_u \nu_{lh}} \right); \quad q_{lh} = \frac{\Delta c}{\Delta\theta}. \quad (28)$$

Profits are defined as in (18) and (27).

P_8^{DS}) For $\Delta c \in \left(\Delta\theta Q_u + \frac{\rho m}{\nu_u \nu_{lh}}, \Delta\theta Q_{lh} \right)$, Γ^{DS} entails (3) for $f = h, g = l$, (4) for $f = l$ and

$$q_{ul} = Q_u; \quad q_{lh} = \frac{\Delta c}{\Delta\theta} \quad (29a)$$

$$\Pi_h = \nu_{lh}(\Delta c - \Delta\theta Q_u) - \frac{\rho m}{\nu_u}. \quad (29b)$$

Profits are defined as in (18) and as

$$\pi_{hl} = \Delta c - \Delta\theta Q_u + \frac{\nu_{lh}}{\nu_u} m; \quad \pi_{hh} = -m. \quad (30)$$

P_9^{DS}) For $\Delta c > \Delta\theta Q_{lh}$, Γ^{DS} entails (19a), (3) for $f = h, g = l$, (4) for $f = h$ and (19b). Profits are defined as in (18) and as

$$\pi_{hl} = \Delta c - \Delta\theta Q_u + \frac{\nu_{lh}}{\nu_u} m; \quad \pi_{hh} = \Delta c - \Delta\theta Q_{lh} - m. \quad (31)$$