Working Paper

Unit-Contingent Power Purchase Agreement and Asymmetric Information about Plant Outage

Owen Q. Wu
Stephen M. Ross School of Business
University of Michigan

Volodymyr Babich
McDonough School of Business
Georgetown University

Ross School of Business Working Paper
Working Paper No. 1120
August 2011

Unit-Contingent Power Purchase Agreement and Asymmetric Information about Plant Outage

Owen Q. Wu
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan

Volodymyr Babich
McDonough School of Business, Georgetown University, Washington, DC

August 19, 2011

This paper analyzes a unit-contingent power purchase agreement between an electricity distributor and a power plant. Under such a contract the distributor pays the plant a fixed price if the plant is operational and nothing if plant outage occurs. Pricing a unit-contingent contract is complicated by the fact that the plant's true status is its private information. The difference between the electricity spot price and the unit-contingent contract price provides an incentive for the plant to misreport its status and earn profit at the distributor's expense. To prevent misreporting, the distributor may inspect the plant and levy penalties if misreporting is discovered. We show that some type of misreporting can actually benefit both the plant and the distributor, because it serves as a risk-allocation mechanism between the two parties, and we identify which type of misreporting under what circumstances is beneficial. We also explore the structural properties of the optimal contracts and design an implementation method using state-contingent options and analyze risk allocation over multiple periods.

1. Introduction

Electric power plants typically sell electricity to distributors pursuant to either "unit-contingent" or "financially firm" power supply contracts (Zaccaria et al. 2006). For example, Wabash Valley Power Association (WVPA), a generation and transmission cooperative that provides wholesale electricity to 28 distribution systems in the Midwest, has increased its unit-contingent power supply from about 24% of its total energy supply resources in 2005 to about 58% in 2009 (WVPA 2005, 2009). In the summer of 2009, WVPA's energy supply resources totaled 2,200 megawatts (MW), including 1,280 MW of unit-contingent contracts and 920 MW of firm contracts. In another example, California Independent System Operator (2003) reported that the net unit-contingent imports into California were as high as 3,000 MW, which accounted for about half of the total net imports.

In both firm and unit-contingent contracts, the distributor purchases a certain amount of power from a specific power generating unit at a predetermined price per megawatt hour (referred to as the *firm contract price* and *unit-contingent contract price*, respectively). Under the unit-contingent contract, the distributor pays the plant only when the power generating unit is operational (hence the transaction is unit-contingent), and in the event of an unplanned outage (e.g., generation facility failures or transmission facility failures that force the plant to be down), the plant is not financially responsible to the distributor and, consequently, the distributor has to purchase replacement power in the spot market. In contrast, under a firm contract, the plant must compensate the distributor for the cost of replacement power in the case of an unplanned outage. Thus, the financial risk associated with purchasing replacement power is borne by the power plant under the firm contract, and borne by the distributor under the unit-contingent contract. By shifting the risk of purchasing replacement power from the plant to the distributor, the unit-contingent contract stabilizes the power plant's income, which is often a necessary requirement for the plant to obtain financing.

To compensate the distributor for taking on the risk of purchasing replacement power, the price of the unit-contingent contract is typically lower than the firm contract price. It might not seem difficult to determine the unit-contingent contract price based on the unit-specific outage frequency and the spot price distribution. However, pricing the unit-contingent contract is complicated by the fact that the distributor cannot directly observe whether the plant indeed has an outage. In practice, the electricity flows are determined by the system operator and the distributor does not receive power directly from the plant. Each day, the plant reports to the distributor whether or not the power generating unit was operational on the previous operating day. The distributor pays the plant at the unit-contingent price if and only if the unit is *reported* to be on.

There is a potential incentive for the plant to misreport its status. When the electricity spot price is higher than the unit-contingent contract price, the plant that is up has an incentive to report to the distributor that the generating unit is down and sell electricity to the spot market. This hurts the distributor who has to buy replacement power at the higher spot price.

Conversely, when the spot price is lower than the unit-contingent contract price, the plant that is down has an incentive to buy electricity on the spot, supply it to the distributor, and report to the distributor that the generating unit is up. The distributor is also hurt because it could have bought the electricity at the lower spot price. Note that the plant reporting up when down is very different from a supplier's extra effort to fulfill the contract in the context of other manufacturing environments. It is typically costly to mitigate supply disruptions in a manufacturing environment due to higher procurement cost from alternative sources and logistic cost, whereas in the electricity industry, the grid provides easy access to the spot market for purchasing replacement power. Because of the electricity market's volatility, the spot price is frequently lower than the contract price. When that happens, plant outages allow the distributor to lower its procurement cost, but the plant's misreporting up deprives the distributor of that opportunity.

In the appendix, we present an example of possible misreporting behavior and its financial impact. We discuss related market regulations and explain why misreporting behavior in bilateral contracts is not monitored and why the system operator protects the plant's private information. Here we continue to discuss how the distributor views the misreporting problem.

It might appear that the distributor prefers the plant to always truthfully report its status. To enforce truthful reporting, the distributor could inspect the plant at a cost, e.g., administrative cost and legal fees for access to the plant's internal records or its account at the system operator. If misreporting is found, penalties can be levied. In practice, however, the distributor typically does not conduct inspections to recover potential damages. A manager at an electricity distributor firm revealed to us that they were aware of some plant's suspicious behavior but never conducted an inspection, because they found the unit-contingent contracts were lucrative after all: The risk-free profit during the "up" times outweighed the loss during the "down" times in the past. Furthermore, the distributor may have an impression that a little bit of misreporting by the plant brings extra profits to the plant, and thus the plant is willing to accept a lower unit-contingent contract price. However, the manager was not sure whether tolerating misreporting is really better or worse than prohibiting misreporting and offering a higher contract price.

We aim to address the following questions: Can the seemingly undesirable misreporting bring benefit to both parties? If so, under what circumstances should the distributor tolerate what type of misreporting? If tolerating some type of misreporting is not practical, what is a practical way to manage misreporting?

We show that, indeed, under some conditions the plant's misreporting benefits both the plant and the distributor because a certain type of misreporting reallocates the risk between the two parties. Specifically, misreporting is beneficial when the plant is down and the spot price falls into a region below the contract price. By purchasing electricity on the spot market to deliver to the distributor and misreporting up, the down plant gains a positive profit, which should belong to the distributor under truthful reporting. Anticipating this loss, the distributor lowers the contract price upfront. If the plant is risk-averse, the reduction in the contract price can exceed the expected increase in the plant's profit from misreporting, while keeping the plant's expected utility constant. This benefits the risk-neutral distributor. A risk-averse distributor would also benefit, because this type of misreporting increases the plant's apparent reliability, reducing the distributor's exposure to the spot price. If both parties are risk-neutral, risk allocations do not improve the expected profit, and a truth-telling contract is optimal.

The other type of misreporting – reporting down when the plant is up – will expose the distributor to the spot price risk while the plant earns a random profit. This increases the uncertainties of both firms' profits and is generally undesirable.

In general, we show that when the spot price falls into a certain region, misreporting behavior benefits both parties, but when spot price is not in that region, misreporting should be prohibited. In practice, the distributor can prohibit undesirable misreporting behavior by imposing a high penalty, but a clause encouraging misreporting is not likely to be part of a contract. Thus, we propose an implementation method using state-contingent options. Such options provide the plant with the same benefit as it would obtain from misreporting.

The contract gaming in this paper does not involve market manipulation (we assume the plant does not withhold power from the market). However, we show that contract gaming affects firms' cash flows and risk profiles, which in turn affect the firms' competitiveness. Our paper informs the regulators about the consequences of contract gaming in the unit-contingent contracts and proposes a method to improve risk allocation.

The rest of the paper is organized as follows. §2 reviews the relevant literature. §3 models the firms' interaction as a game with asymmetric information. §4 analyzes the role of misreporting in risk allocation. §5 extends the analysis to a two-period game. §6 discusses implementation issues and aggregate risk allocation. §7 presents numerical results. §8 concludes the paper with a summary of managerial insights and a discussion on model extensions.

2. Related Literature

We apply the theory of Bayesian games originally developed by Harsanyi (1967, 1968a,b) to model the interactions between the power plant and the distributor. Bayesian games have been applied to the electricity markets to model the suppliers' bidding processes in which each power plant's marginal cost is private information. Such a game has been analyzed in various market conditions by Ferrero et al. (1998), Shahidehpour et al. (2002), Li and Shahidehpour (2005), and Correia (2005), among others. Hortaçsu and Puller (2008) analyze bidding processes in which contract positions are private information. Unlike previous works, in this paper information asymmetry comes from the fact that the plant's status cannot be directly observed by the distributor, and the unit-contingent power purchase agreement introduces incentive conflicts into the system.

Several economics papers on contract theory are related to our work. For example, Laffont and Martimort (2002, Section 3.6) discuss an adverse selection problem with audits and costly state verification. The costly audit allows the principal to detect an untruthful agent's report and impose penalties. The Revelation Principle still applies, and under the truth-revealing mechanisms, punishments are never used, but the existence of punishments reduces agent's incentive to lie and, hence, reduces informational rents. Mookherjee and Png (1989) and Reinganum and Wilde (1985) apply the adverse selection problems with costly state verification to insurance and taxation. In contrast to these papers, our analysis is focused on a particular contract form commonly seen in practice. Because of the restriction on the contract set, instead of invoking the mechanism design approach (as in Myerson 1981, 1979, Guesnerie and Laffont 1984), we find the equilibrium of the Bayesian game directly. Within the unit-contingent contract space, we show the truth-revealing mechanism is not necessarily optimal.

Our work shares some similarities to the economics literature on contracting with costly state verification (e.g., Townsend 1979) and literature on incomplete contracts (e.g., Demski and Sappington 1991, Boot et al. 1993, Bernheim and Whinston 1998) in that by allowing flexibility to act to one party, a better equilibrium outcome can be achieved. However, there are a number of essential differences between our work and this literature. For instance, Townsend (1979) and related economics literature on bonding and insurance are concerned with the problem of signaling private information, when doing so involves a cost for the party that sends a signal – that is, the party with information can initiate a costly verification process, during which a part of the value of the firm is destroyed. In our problem, the firm without information can initiate a costly inspection to verify the other firm's report. Papers on incomplete contracts listed above use models where actions of one of the players are observable, but not contractible (this is the meaning of "non-verifiable" in their context). In our problem, actions of the plant, other than its report, are not observable (without an inspection). This is important, because analyses of incomplete contracts (e.g., Bernheim and Whinston 1998) rely on the second player knowing the actions of the first player when deciding whether or not to punish her for deviating from the desired equilibrium. Literature on incomplete contracts studies contracts that are not specified for all states of the world (by choice or due to non-verifiability). This affords tit-for-tat strategies as in Bernheim and Whinston (1998). Specifically, the main insight in Bernheim and Whinston (1998) is that if not all actions of the first player are contractible (but all actions are observable), some of the actions of the second player should be made non-contractible too (thus introducing strategic ambiguity), so that the second player has the flexibility of punishing the first player for "shirking". In contrast, our contract is not strategically ambiguous. It is fully specified based on all observable information. Thus, although there is some similarity between our and these economic literatures, the frameworks and the results are different.

Our game is similar to the inspection games used to study problems of arms control and treaty violations (e.g., Dresher 1962, Diamond 1982, Avenhaus and Canty 2005, among others). The settings of these games range from unobservable inspection and simultaneous moves (Diamond 1982) to observable inspection and sequential moves (Avenhaus and Canty 2005). Our problem differs from the inspection games in two aspects. First, the payoffs to the players in our problem depend on both a publicly observable stochastic process and a private information process observed only by the inspectee. The inspectee's incentive to violate depends on both processes, while the inspector's incentive to inspect depends on the public signal and the inspectee's report. Second, in our setting, before the game begins, the inspector offers a contract to the inspectee that affects players' incentives during the game.

Examples of inspection games in procurement are Reyniers and Tapiero (1995a,b), who study static quality-inspection games where the seller chooses product quality (unobservable by the buyer) and the buyer decides whether to inspect the product. They find non-cooperative and cooperative mixed strategy equilibria of the quality-inspection games. Reyniers and Tapiero (1995a) also design a contract that leads to a cooperative equilibrium. We design contracts that can benefit both players in a non-cooperative setting.

Heese and Kemahlioglu-Ziya (2009) study the misreporting of revenues by a retailer in a revenue-sharing contract between a supplier and the retailer, who enjoys private information about demand. Although the motivation of their research and the contracts considered are different from our paper, some results are similar: Heese and Kemahlioglu-Ziya (2009) also find that it might be beneficial to limit cheating by the retailer, but not eliminate it altogether.

3. Unit-Contingent Contract and Subsequent Inspection Game

A distributor buys electricity from a plant on a unit-contingent basis. The plant has two states: UP and DOWN (denoted in capital letters for visual convenience). In the UP state the plant produces power at its full capacity, which we normalize to be one unit, and the production cost is denoted as c. In the DOWN state the plant cannot produce anything. The two-status model is a reasonable approximation (e.g., in the example shown in appendix A, the plant's output falls between 0% and 90% of its capacity for only 5% of the time). The distributor sells the energy to satisfy a firm contract with a fixed price f. Unit-contingent contracts are often used for base-load plants, and the firm service contracts are part of the base-load.

The timeline of the interactions between the distributor and the plant is illustrated in Figure 1.

During the contracting stage, the distributor determines the contract parameters, which consist of a pair of functions: $(v(p), \phi(p))$, where p is the spot price to be realized in the contract execution stage, $v(p) > 0$ is the unit-contingent contract price (i.e., the payment from the distributor to the plant if the realized spot price is p and the plant delivers electricity to the distributor; otherwise, no payment is made), and $\phi(p) \geq 0$ is the penalty paid by the plant to the distributor if the plant is found to have misreported its status. In practice, unit-contingent contract prices are typically independent of the spot price, i.e., $v(p) = v$. The main insights hold regardless of whether the contract price is constant or dependent on the spot price. For a discussion on the importance of connecting the long-term contract prices with spot prices see Boyabatlı et al. (2011).

Figure 1: Timeline of the Unit-Contingent Contract

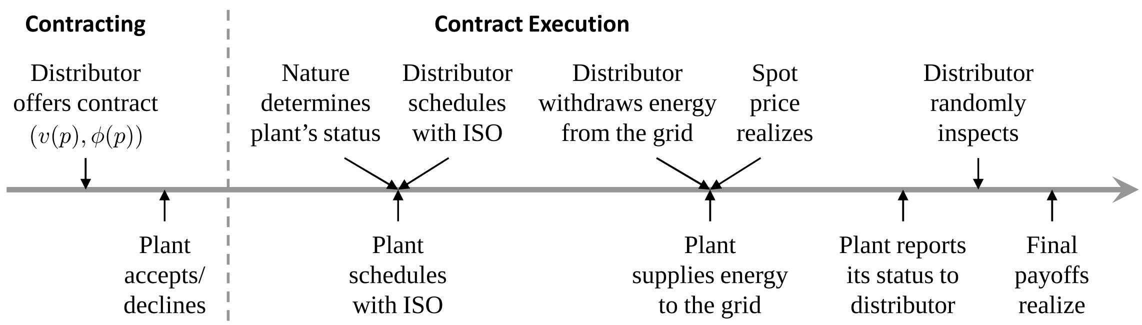

In this section, the contract execution stage is modeled as a single-period inspection game. Contract execution in a multi-period setting will be analyzed in §5 and §6. The differences between our game and the inspection games in the literature are reviewed in the previous section.

Nature's move: During the execution stage, it is common knowledge that nature lets the plant to be UP with probability $\gamma \in (0, 1)$ and DOWN with probability $1 - \gamma$. The plant's actual status is its private information. Nature also determines the spot price, p, based on the supply and demand balance in the market. It is common knowledge that the spot price follows a continuous distribution with probability density function $g(p)$, and the realized spot price is public information. We assume that the power generating unit is small so that its status has negligible influence on the distribution of the spot price.

Plant's move: Based on its realized status, the plant schedules with the ISO (system operator) and supplies energy to the grid if UP. The plant is prohibited to misreport to the ISO (see appendix B for regulations). After observing the realized spot price, the plant reports a status to the distributor. In practice, the price and status uncertainties are resolved when the plant reports its status. (See appendix B for more details on contract execution in practice and why the system operator protects the plant's private information from being accessible to the distributor.)

Distributor's move: The distributor schedules with the ISO and withdraws energy from the grid to meet its firm contract obligation. (The distributor cannot tell where the energy is actually produced.) After receiving the plant's status report, the distributor may inspect the truthfulness of the the plant's report at an inspection cost $k > 0$, which involves administrative cost and legal fees for access to the plant's internal records or its account at the system operator.

Table 1: Profits from Contract Execution

In each formula cell, the top expression is the plant's profit, the bottom expression is the distributor's profit.

		Nature determines the plant's status			
		UP		DOWN	
		Plant reports		**Plant reports**	
		UP	DOWN	UP	DOWN
Distributor	Inspect	$v(p) - c$ $f - v(p) - k$	$v(p) - c - \phi(p)$ $f - v(p) - k + \phi(p)$	$-\phi(p)$ $f - p - k + \phi(p)$	0 $f - p - k$
	Do Not Inspect	$v(p) - c$ $f - v(p)$	$p - c$ $f - p$	$v(p) - p$ $f - v(p)$	0 $f - p$

Given the moves by nature, the plant, and the distributor, the profits (or losses) of both firms are summarized in Table 1. For example, when the plant is UP but reports DOWN, if the distributor does not inspect, the plant earns a profit of $p - c$ (selling its output on the spot market), while the distributor's profit is $f - p$ (buying on the spot market to cover its firm contract). If the distributor inspects (at cost k), misreporting will be uncovered, and the cash flows are then corrected based on the true status (the plant effectively pays compensatory damages); furthermore, the plant must pay the punitive damages $\phi(p)$ specified in the contract. In practice, both penalties are specified in some contracts (see, e.g., Bachrach et al. 2003).

The spot price is set by the marginal generating unit on the grid and, therefore, is typically higher than the production cost c of a base-load plant. The spot price may sporadically drop below the production cost c, but the plant continues production due to high switching cost. Thus, when the plant is UP, it always produces energy at cost c.

The payoffs of the game are expressed as the utilities of the profits given in Table 1. In the multi-period settings in §5 and §6, the utility will be on the aggregate profit over multiple periods. The utility functions of the plant and the distributor are $U_P(\cdot)$ and $U_D(\cdot)$, respectively. We assume (i) both utility functions are concave and strictly increasing, (ii) $U_P(0) = 0$, and (iii) $U_D(x) \to \infty$ as $x \to \infty$. The last two assumptions are not crucial for the results, but help with the exposition.

3.1 Bayesian Equilibrium of the Inspection Game

In this section, we find the Bayesian equilibrium of the execution stage game for any given contract $(v(p), \phi(p)) \geq 0$ and any given realization of the spot price p. The mixed strategies of the plant and the distributor are characterized by the following probabilities:

x_{U} = Probability that the plant tells the truth (i.e., reports UP) when it is UP;

x_{D} = Probability that the plant tells the truth (i.e., reports DOWN) when it is DOWN;

y_{U} = Probability that the distributor inspects if the plant reports UP;

y_{D} = Probability that the distributor inspects if the plant reports DOWN.

How the game is played depends on whether the spot price p is higher or lower than the unit-contingent contract price $v(p)$, and in the case of $p > v(p)$, it further depends on whether the penalty is above or below the following threshold:

$$\widehat{\phi}(p) \stackrel{\text{def}}{=} v(p) + k - f + U_D^{-1}\left(\frac{1}{\gamma}U_D(f-p) - \frac{1-\gamma}{\gamma}U_D(f-p-k)\right), \quad \text{for } p > v(p). \qquad (1)$$

If the distributor is risk-neutral, (1) becomes $\widehat{\phi}(p) = v(p) - p + k/\gamma$. Then, $\phi(p) < \widehat{\phi}(p)$ is equivalent to $\gamma\big(\phi(p) + p - v(p)\big) < k$. When the distributor inspects, the expected compensation to the distributor does not exceed $\gamma\big(\phi(p) + p - v(p)\big)$, because $\phi(p) + p - v(p)$ is the total payment that the plant compensates the distributor if misreporting DOWN is found, and $\mathsf{P}\big\{\text{plant is UP} \mid \text{plant reports DOWN}\big\} = \frac{(1-x_{\mathsf{U}})\gamma}{(1-x_{\mathsf{U}})\gamma + (1-\gamma)} \leq \gamma$. Therefore, if $\gamma\big(\phi(p) + p - v(p)\big) < k$, the expected compensation from inspection cannot cover the inspection cost and the distributor will not conduct an inspection.

For a risk-averse distributor, $\phi(p) < \widehat{\phi}(p)$ is equivalent to $\gamma U_D(f - v(p) - k + \phi(p)) + (1-\gamma)U_D(f - p - k) < U_D(f - p)$, where the left side is the distributor's maximum expected utility if it inspects a DOWN report, and the right side is the distributor's expected utility if it does not inspect. Hence, the distributor will not conduct an inspection when the penalty $\phi(p)$ is below the threshold $\widehat{\phi}(p)$.

The Bayesian equilibrium under $p > v(p)$ is summarized in the following proposition, with the proof relegated to the online supplement.

Proposition 1a *When the spot price is above the contract price, $p > v(p)$, the Bayesian equilibrium of the contract execution stage game is as follows:*
(i) If the penalty $\phi(p)$ does not exceed the threshold $\widehat{\phi}(p)$ defined in (1), we have a pooling equilibrium where the plant always reports DOWN *($x_{\mathsf{U}}^* = 0, x_{\mathsf{D}}^* = 1$) and the distributor does not inspect ($y_{\mathsf{U}}^* = y_{\mathsf{D}}^* = 0$), and the equilibrium utilities of the plant and the distributor are:*

$$\mathsf{E}[U_P \mid p] = \gamma U_P(p-c), \qquad \mathsf{E}[U_D \mid p] = U_D(f-p). \qquad (2)$$

(ii) If $\phi(p) > \widehat{\phi}(p)$, we have a mixed strategy equilibrium:

$$x^*_{\mathsf{U}} = 1 - \frac{(1-\gamma)\widehat{\alpha}}{\gamma(1-\widehat{\alpha})}, \qquad \text{where} \quad \widehat{\alpha} = \frac{U_D(f-p) - U_D(f-p-k)}{U_D(f-v(p)-k+\phi(p)) - U_D(f-p-k)},$$

$$y^*_{\mathsf{D}} = \frac{U_P(p-c) - U_P(v(p)-c)}{U_P(p-c) - U_P(v(p)-c-\phi(p))},$$

$x^*_{\mathsf{D}} = 1$, and $y^*_{\mathsf{U}} = 0$. The firms' equilibrium utilities are:

$$\mathsf{E}[U_P \mid p] = \gamma U_P(v(p)-c), \qquad \mathsf{E}[U_D \mid p] = \gamma x^*_{\mathsf{U}} U_D(f-v(p)) + (1-\gamma x^*_{\mathsf{U}}) U_D(f-p). \tag{3}$$

Note that the firms' expected utilities in (3) do not directly involve the penalty. This is because the plant balances gains from misreporting and losses due to penalty, while the distributor balances the cost and benefit of inspection. In the resulting equilibrium, the plant's expected utility is the same whether it misreports or reports truthfully; the distributor's expected utility is the same whether it inspects or not. Note that the distributor uses the penalty to influence the equilibrium misreporting probability x^*_{U}, which in turn affects the distributor's expected utility.

Figure 2 illustrates a few useful properties of the above equilibrium. When the plant is UP, the probability of truthful reporting increases in the penalty $\phi(p)$ and approaches to one as $\phi(p) \to \infty$. If the plant reports DOWN and $\phi(p) > \widehat{\phi}(p)$, the distributor may conduct an inspection and the

Figure 2: Effect of Penalty $\phi(p)$ on Equilibrium Strategies and Utilities
Parameters: $\gamma = 0.8,\ k = 8,\ p = 75,\ v(p) = 70,\ U_D(x) = x,\ \widehat{\phi}(p) = v(p) - p + k/\gamma = 5$

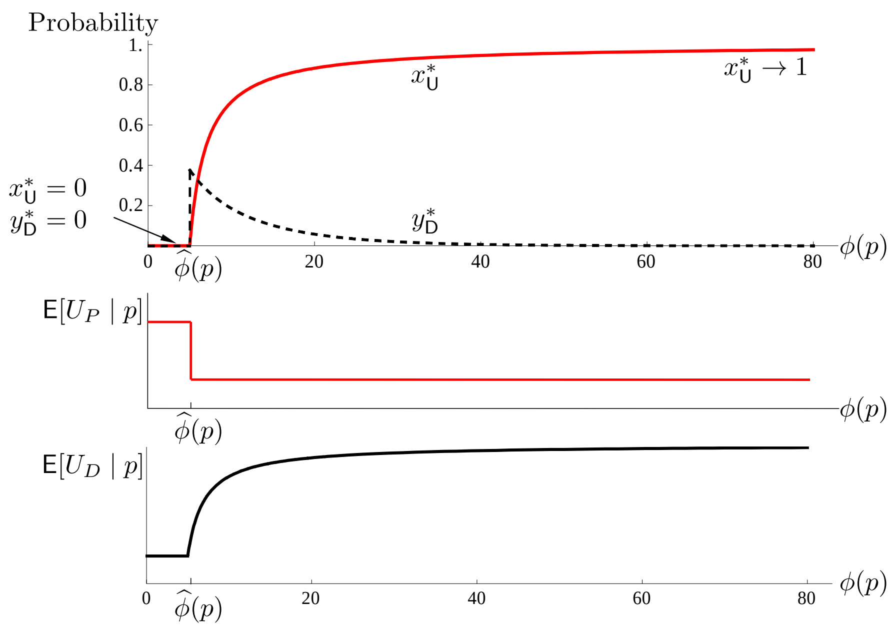

Note: For the case of $p < v(p)$, the effect of $\phi(p)$ is qualitatively the same, except that x^*_{U} and y^*_{D} are replaced by x^*_{D} and y^*_{U}, respectively.

probability of inspection, $\widehat{y}_{\mathsf{D}}$, decreases in $\phi(p)$ and approaches zero when $\phi(p) \to \infty$. Intuitively, a very large penalty with a very small inspection probability can effectively deter misreporting. When the penalty $\phi(p)$ increases, the plant's expected utility decreases (in a step-function fashion, see Figure 2) and the distributor's expected utility increases.

When the spot price p is below the contract price $v(p)$, the analysis is in parallel and we summarize the equilibrium below. The penalty threshold is

$$\widehat{\phi}(p) \overset{\text{def}}{=} p + k - f + U_D^{-1}\left(\frac{1}{1-\gamma}U_D(f - v(p)) - \frac{\gamma}{1-\gamma}U_D(f - v(p) - k)\right), \quad \text{for } p < v(p). \tag{4}$$

The Bayesian equilibrium under $p < v(p)$ is summarized in the following proposition.

Proposition 1b *When the spot price is below the contract price, $p < v(p)$, the Bayesian equilibrium of the contract execution stage game is as follows:*
(i) If the penalty $\phi(p)$ does not exceed the threshold $\widehat{\phi}(p)$ defined in (4), we have a pooling equilibrium where the plant always reports UP *($x^*_{\mathsf{U}} = 1, x^*_{\mathsf{D}} = 0$) and the distributor does not inspect ($y^*_{\mathsf{U}} = y^*_{\mathsf{D}} = 0$), and the equilibrium utilities of the plant and the distributor are:*

$$\mathsf{E}[U_P \mid p] = \gamma U_P(v(p) - c) + (1-\gamma)U_P(v(p) - p), \qquad \mathsf{E}[U_D \mid p] = U_D(f - v(p)). \tag{5}$$

(ii) If $\phi(p) > \widehat{\phi}(p)$, we have a mixed strategy equilibrium:

$$x^*_{\mathsf{D}} = 1 - \frac{\gamma\widehat{\beta}}{(1-\gamma)(1-\widehat{\beta})}, \qquad \textit{where} \quad \widehat{\beta} = \frac{U_D(f - v(p)) - U_D(f - v(p) - k)}{U_D(f - p - k + \phi(p)) - U_D(f - v(p) - k)},$$
$$y^*_{\mathsf{U}} = \frac{U_P(v(p) - p)}{U_P(v(p) - p) - U_P(-\phi(p))},$$

$x^*_{\mathsf{U}} = 1$, and $y^*_{\mathsf{D}} = 0$. The firms' equilibrium utilities are:

$$\mathsf{E}[U_P \mid p] = \gamma U_P(v(p) - c), \quad \mathsf{E}[U_D \mid p] = (1 - (1-\gamma)x^*_{\mathsf{D}})U_D(f - v(p)) + (1-\gamma)x^*_{\mathsf{D}}U_D(f - p). \tag{6}$$

4. Misreporting as a Risk Allocation Mechanism

Anticipating the equilibrium in the execution stage game, the distributor now chooses the contract parameters $(v(p), \phi(p))$ to influence the plant's misreporting behavior. Assume the plant's reservation utility is $\underline{U_P} > 0$ (recall $U_P(0) = 0$). The distributor's utility-maximization problem is:

$$\max_{\{v(p),\, \phi(p)\}} \mathsf{E}[U_D], \qquad s.t. \;\; \mathsf{E}[U_P] \;\geq\; \underline{U_P}. \tag{7}$$

Because we are analyzing a specific type of contract, the unit-contingent contract, we look for the optimal contract based on the explicit solution to the Bayesian game derived in the previous section, rather than taking the general mechanism design approach. In fact, general mechanism design leads to an impractical contract: Let $v(p, \mathsf{UP})$ and $v(p, \mathsf{DOWN})$ be the distributor's payment to the plant which reports UP and DOWN, respectively. The incentive compatibility requires $v(p, \mathsf{UP}) - v(p, \mathsf{DOWN}) = p$ for all p, because if this equality were $<$ (or $>$), it would provide incentive for misreporting DOWN (or UP). Any contract satisfying this equality exposes the plant to the spot price risk (outage leads to a payment reduction of p) and, therefore, does not possess the key feature of the unit-contingent contract.

4.1 Unit-Contingent Truth-Telling Contracts

We first analyze the unit-contingent contracts that enforce truth-telling. From the analysis in §3.1, we see that truth-telling is the equilibrium outcome when the penalty for misreporting is infinite in theory, but as shown in Figure 2, the truth-telling probability approaches to one quickly and the distributor's expected utility converges quickly when the penalty increases. Thus, a reasonably large penalty is sufficient to deter misreporting. For analytical tractability, we approximate the distributor's expected utility under a large penalty by its expected utility under infinite penalty. When penalty $\phi(p) = \infty$, $\forall p$, the equilibrium truth-telling probabilities $x^*_{\mathsf{U}} = x^*_{\mathsf{D}} = 1$. Therefore, (3) and (6) imply that the firms' expected utilities are:

$$\mathsf{E}[U_P] = \gamma \mathsf{E} U_P(v(p) - c), \qquad \mathsf{E}[U_D] = \gamma \mathsf{E} U_D(f - v(p)) + (1-\gamma)\mathsf{E} U_D(f - p). \tag{8}$$

Solving the problem in (7) with $\mathsf{E}[U_P]$ and $\mathsf{E}[U_D]$ defined in (8) gives the following lemma.

Lemma 1 *The contract* $(v(p) = v_0,\ \phi(p) = \infty)$ *with* v_0 *satisfying* $\gamma U_P(v_0 - c) = \underline{U_P}$ *is optimal among all unit-contingent truth-telling contracts.*

Intuitively, when truth-telling is enforced, bilateral transactions occur only when the plant is up, and a constant contract price stabilizes both parties' income when the plant is up.

Note that the truth-telling contract in Lemma 1 is optimal among all contracts when both parties are risk-neutral. This is because when truth-telling is enforced, costly inspections will not occur. Therefore, the sum of both firms' expected profit is maximized, and keeping the plant's profit at its reservation level maximizes the distributor's expected profit.

4.2 Sub-Optimality of the Unit-Contingent Truth-Telling Contracts

If truth-telling is enforced, v_0 is the lowest contract price the plant is willing to accept, but is there a better contract with $v < v_0$ and a lower penalty? We consider a contract $(v(p), \phi(p))$ of the following form:

$$v(p) \equiv v < v_0, \qquad \phi(p) = \begin{cases} 0, & p \in [v-\delta, v], \\ \infty, & \text{otherwise}, \end{cases} \tag{9}$$

where v is slightly below v_0 and δ is small (we assume $\delta \in (0, k)$, where k is the inspection cost). Under the contract in (9), when the spot price $p \in [v-\delta, v]$, the equilibrium in Proposition 1b(i) is played.[1] In this equilibrium, facing zero penalty for misreporting, the plant would report UP when it is DOWN, obtaining a profit of $v - p$. Because of this benefit, the plant is willing to accept a contract price $v < v_0$. When $p \notin [v-\delta, v]$, the equilibrium is that in Proposition 1a(ii) or that in Proposition 1b(ii). In this equilibrium, due to large penalty, the plant is truthful: $x^*_{\mathsf{U}} = x^*_{\mathsf{D}} = 1$.

The following proposition shows that a contract of the form (9) can allocate the risk between the two parties more efficiently than the truth-telling enforcing contract.

Proposition 2a *If either the plant or the distributor or both are risk-averse, then there exists a unit-contingent contract of the form (9) that strictly dominates the truth-telling enforcing contract* $(v_0, \phi(\cdot) = \infty)$*, i.e., strictly improves* $\mathsf{E}[U_D]$ *in* (7) *while holding* $\mathsf{E}[U_P]$ *constant.*

Intuitively, a lower contract price $v < v_0$ reduces the plant's profit when the plant is UP, resulting in a utility loss of $\gamma U_P(v_0 - c) - \gamma U_P(v - c)$. To compensate for this loss, the distributor eliminates the misreporting penalty when the spot price $p \in [v-\delta, v]$. When the spot price falls in this range, the DOWN plant can still gain a small profit, $v - p$, by purchasing electricity on the spot market to deliver to the distributor and misreporting UP. This extra benefit can offset the loss, i.e., $(1-\gamma)\int_{v-\delta}^{v} U_P(v-p)g(p)dp = \gamma U_P(v_0 - c) - \gamma U_P(v - c)$, thereby keeping the plant's expected utility unchanged. If the plant is risk-averse, the expected value of this benefit, $(1-\gamma)\int_{v-\delta}^{v}(v-p)g(p)dp$, is smaller than the expected loss $\gamma(v_0 - v)$. Hence, the distributor's expected profit increases. Furthermore, the distributor now enjoys more certain profit: It obtains a fixed profit of $f - v$ (higher than $f - v_0$) with a higher probability than under the truth-telling enforcing contract.

We remark that when a plant reports UP when DOWN, it buys electricity from the spot market and sells to the distributor. Such a transaction can be easily executed because it is exactly what the plant has to do during the outages under a firm contract. Furthermore, the opportunities of misreporting UP when DOWN are not rare, because outages are not rare and the unit-contingent

[1] To see this, note that $\widehat{\phi}(p)$ defined in (4) satisfies $\dfrac{U_D(f-v) - U_D(f-v-k)}{U_D(f-p-k+\widehat{\phi}(p)) - U_D(f-v-k)} = 1 - \gamma < 1$, which implies that $\widehat{\phi}(p) > p + k - v$. Thus, for $p \in [v-\delta, v]$ and $\delta < k$, we have $\widehat{\phi}(p) > k - \delta > 0$. Hence, $\widehat{\phi}(p) > \phi(p) = 0$ for $p \in [v-\delta, v]$, which satisfies the condition in Proposition 1b(i).

contract price is typically close to the median spot price.

An important implication of Proposition 2a is that the seemingly undesirable action of the plant's reporting UP when DOWN could actually serve as a tacit risk-allocation mechanism between the plant and the distributor. A natural question is whether plant's reporting DOWN when UP could also serve the same purpose. Consider a contract $(v(p), \phi(p))$ that steers the plant to misreport DOWN when UP for spot price $p \in [v, v+\delta]$:

$$v(p) \equiv v < v_0, \qquad \phi(p) = \begin{cases} 0, & p \in [v, v+\delta], \\ \infty, & \text{otherwise}. \end{cases} \tag{10}$$

Proposition 2b *If either the plant or the distributor or both are risk-averse, then the unit-contingent contract of the form (10) is strictly dominated by the truth-telling enforcing contract.*

Intuitively, a truthful report of UP provides both firms with a fixed profit, whereas misreporting DOWN introduces additional variability to both parties' profits (the plant sells its output to the spot market and the distributor buys replacement power at the spot price). Thus, the contract in (10) performs worse than the truth-telling contracts. However, optimal contracts may contain both types of misreporting, because allowing misreporting DOWN would reduce the contract price v, which might affect the plant's extra profit obtained from misreporting UP.

4.3 Optimal Unit-Contingent Contracts

We first make the following observation based on Proposition 1 and Figure 2.

Corollary 1 *A unit-contingent contract $(v(p), \phi(p))$ with penalty $\phi(p) < \infty$ is strictly dominated by the contract $(v(p), \phi^\dagger(p))$, where $\phi^\dagger(p) = 0$ if $\phi(p) \le \widehat{\phi}(p)$, and $\phi^\dagger(p) > \phi(p)$ if $\phi(p) > \widehat{\phi}(p)$.*

From Figure 2, it can be seen that $(v(p), \phi^\dagger(p))$ defined in Corollary 1 will not change the plant's expected utility, but increase the distributor's expected utility. Intuitively, any penalty level below the threshold $\widehat{\phi}(p)$ induces misreporting without inspection, so does the zero penalty. When the penalty is above the threshold, a higher penalty leads to higher truth-telling probability, benefiting the distributor. Corollary 1 implies that, without loss of optimality, we can restrict our attention to only two penalty values: zero and very high penalty.

With the above observation, choosing a contract $(v(p), \phi(p))$ is equivalent to choosing a contract price function $v(p)$ and a price set L, within which the penalty is zero and beyond which the penalty is very high. Because there is no incentive for misreporting when $p = v(p)$, we can exclude the set

$\{p : p = v(p)\}$ from L. Let $L = L_l \cup L_r$, where

$$L_l \stackrel{\text{def}}{=} \{\ p:\ p < v(p),\ \phi(p) = 0\ \}, \qquad L_r \stackrel{\text{def}}{=} \{\ p:\ p > v(p),\ \phi(p) = 0\ \}.$$

From Corollary 1, zero penalty is never optimal when the threshold $\widehat{\phi}(p) < 0$. Thus, the two subsets above must be contained in the region where the threshold is non-negative:

$$L_l \subseteq S_l \stackrel{\text{def}}{=} \{\ p:\ p < v(p),\ \widehat{\phi}(p) \geq 0\ \}, \qquad L_r \subseteq S_r \stackrel{\text{def}}{=} \{\ p:\ p > v(p),\ \widehat{\phi}(p) \geq 0\ \}. \tag{11}$$

We refer to L_l and L_r as *zero-penalty regions* and S_l and S_r as *zero-penalty feasible regions.*

Using Proposition 1, we can derive the firms' expected utilities. When $p \in L_l$, the equilibrium utilities are given in (5):

$$\mathsf{E}[U_P \mid p] \ = \ \gamma U_P(v(p) - c) + (1-\gamma) U_P(v(p) - p), \qquad \mathsf{E}[U_D \mid p] \ = \ U_D(f - v(p)).$$

When $p \in L_r$, the equilibrium utilities are given in (2):

$$\mathsf{E}[U_P \mid p] \ = \ \gamma U_P(p - c), \qquad \mathsf{E}[U_D \mid p] \ = \ U_D(f - p).$$

When $p \notin (L_l \cup L_r)$, it is optimal for the distributor to set penalty $\phi(p) = \infty$. Both (3) and (6) imply that:

$$\mathsf{E}[U_P \mid p] \ = \ \gamma U_P(v(p) - c), \qquad \mathsf{E}[U_D \mid p] \ = \ \gamma U_D(f - v(p)) + (1-\gamma) U_D(f - p).$$

Integrating utilities over the above three price regions and rearranging terms, we can express the firms' expected utilities as follows:

$$\begin{aligned} \mathsf{E}[U_P] = {} & \gamma \mathsf{E} U_P(v(p) - c) + (1-\gamma) \int_{L_l} U_P(v(p) - p) g(p) dp \\ & + \gamma \int_{L_r} \big[U_P(p - c) - U_P(v(p) - c)\big] g(p) dp, \end{aligned} \tag{12}$$

$$\begin{aligned} \mathsf{E}[U_D] = {} & \Big[\gamma \mathsf{E} U_D(f - v(p)) + (1-\gamma) \mathsf{E} U_D(f - p)\Big] - (1-\gamma) \int_{L_l} \big[U_D(f - p) - U_D(f - v(p))\big] g(p) dp \\ & - \gamma \int_{L_r} \big[U_D(f - v(p)) - U_D(f - p)\big] g(p) dp. \end{aligned} \tag{13}$$

In (12), the first term $\gamma \mathsf{E} U_P(v(p) - c)$ is the plant's expected utility if it always reports the true status, and the other two terms are the additional utility the plant can get by misreporting its status in the zero-penalty regions L_l and L_r, respectively. In (13), the first term in the brackets is the distributor's expected utility when the plant is truthful, and the remaining two integrals account for the utility loss due to the plant's misreporting behavior.

Using (12) and (13), the distributor's problem in (7) now becomes:

$$\max_{\{v(p),\, L_i \subseteq S_i,\, i=l,r\}} \mathsf{E}[U_D], \qquad s.t. \;\; \mathsf{E}[U_P] \;\geq\; \underline{U_P}. \tag{14}$$

Solving the optimization problem in (14) is complicated in general, but we are able to identify some structural properties of the optimal contracts, as detailed in Proposition 3 and illustrated in Figure 3.

Proposition 3 *Suppose either the plant or the distributor or both are risk-averse. Then, the optimal contract* $(v^*(p), L_l^*, L_r^*)$ *has the following structural properties:*

a. Properties of the contract price $v^*(p)$*:*

(i) $v^*(p)$ *is constant outside of the zero-penalty regions:* $v^*(p) = v_1$*, for some* v_1 *and* $p \notin L_l^* \cup L_r^*$*.*

(ii) If the plant is risk-averse, $v^*(p)$ *increases in* p *for* $p \in L_l^*$ *at a rate no faster than one, and* $v^*(p) \geq v_1$ *for* $p \in L_l^* \cap [c, \infty)$*. If the plant is risk-neutral,* $v^*(p) = v_1$ *for* $p \in L_l^*$*.*

b. Properties of L_l^* *and* L_r^**:*

(i) If $L_r^* \neq \emptyset$*, then* L_r^* *is to the right of* L_l^**, i.e.,* $p_l < p_r$*,* $\forall p_l \in L_l^*$*,* $p_r \in L_r^*$*.*

(ii) If the distributor is risk-neutral, then L_l^* *is on the rightmost part of* S_l*:* $L_l^* = [p_l^*, \infty) \cap S_l$ *for some* p_l^**. If the plant is risk-neutral, then* L_l^* *is on the leftmost part of* S_l*:* $L_l^* = [0, p_l^*] \cap S_l$ *for some* p_l^**.*

(iii) L_r^* *is on the leftmost part of* S_r*:* $L_r^* = [0, p_r^*] \cap S_r$ *for some* p_r^**.*

(iv) If we restrict the contract price to be invariant with the spot price, i.e., $v(p) = v$*,* $\forall p$*, then all structural properties of* L_l^* *and* L_r^* *in parts (i)-(iii) still hold. In addition,* S_l *is an interval and, if one of the firms is risk-neutral,* L_l^* *is an interval.*

Figure 3: Structural Properties of the Optimal Unit-Contingent Contract
The case of risk-averse plant and risk-neutral distributor

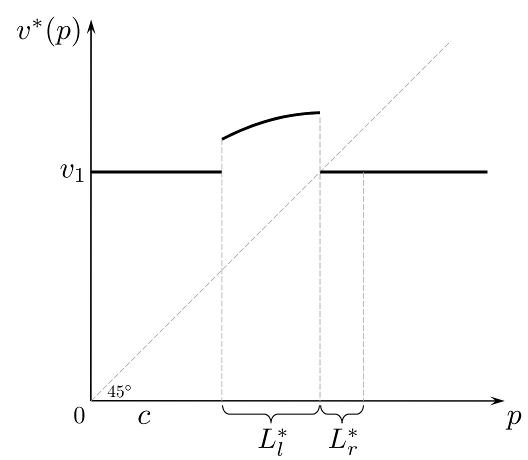

Property a(i) shows that in the price region where misreporting is prohibited ($p \notin L_l \cup L_r$), it is optimal to use a fixed contract price, because it stabilizes both firms' income when the plant is UP. When the price falls into the zero-penalty region L_l, the plant would report UP when DOWN to earn a profit of $v(p) - p$, and the plant gets $v(p) - c$ when it is actually UP. For a risk-averse plant, to reduce the plant's income variability, intuitively the distributor can select $v(p)$ so that the above incomes, $v(p) - p$ and $v(p) - c$, change mildly with p. This can be achieved if $v(p)$ increases in p at a rate no faster than one, which is exactly what part a(ii) prescribes. Note that for $p \in L_r^*$, any $v(p) < p$ is fine; no particular structure is necessary.

Property b(i) can be intuitively explained: If L_r is not to the right of L_l, then we would find $p_l \in L_l$ and $p_r \in L_r$ such that $v(p_r) < p_r < p_l < v(p_l)$, which implies $v(p_l) - v(p_r) > p_l - p_r$, i.e., $v(p)$ changes faster than the price changes. Such $v(p)$ is unlikely to be the best for stabilizing either firm's income.

For most settings, S_l and S_r are intervals and, consequently, properties b(ii) and b(iii) imply that L_l^* and L_r^* are also intervals. A typical optimal contract is illustrated in Figure 3.

We intuitively explain properties b(ii) and b(iii) under a fixed contract price v. (Part (iv) states that those properties hold when $v(p) \equiv v$.) Property b(ii) shows that if a risk-neutral distributor chooses to induce a risk-averse plant to report UP when DOWN to gain a profit of $v - p$, then p should be as large as possible. Intuitively, the risk-averse plant prefers high-probability small gains (corresponding to high p in the region $p < v$) over low-probability large gains (corresponding to low p) with the same expected value. Thus, by shifting the zero-penalty region L_l as close to v as possible, the distributor can lower the plant's expected gains from misreporting, thereby improving the distributor's expected profit while keeping the plant's expected utility constant. The intuition for the property of L_r^* in b(iii) is similar.

In reality, in terms of firm size, distributors are often much larger than plants. Thus, for a unit-contingent contract, we can consider the distributor is risk-neutral and the plant is risk-averse. Small distributors and large plants do exist, and the second statement in part b(ii) states that a risk-averse distributor who is contracting with a risk-neutral plant will try to move the zero-penalty region L_l to the left of the feasible region S_l. Intuitively, by inducing the plant to report UP when DOWN, the distributor essentially sacrifices a random profit of $v - p$ in return for a lower contract price v and lower profit variability (note that the risk-neutral plant's expected profit is held constant, so is the distributor's expected profit). To minimize the profit variability, the risk-averse distributor prefers sacrificing low-probability large profit (corresponding to low spot prices) over sacrificing high-probability small profit.

In practice, constant unit-contingent contract price is prevalent. Our analysis shows that the optimal contract contains non-constant contract price in some price region while a constant price in other regions. However, specifying the contract price as a function of the price is not easily implementable in practice. In the analysis of a two-period model in §5 and a multi-period model in §6, we focus on constant contract price.

5. Two-Period Interaction and Risk Allocation

In this section, we describe the main results from analyzing a two-period model. The analysis of the two-period game involves many technical details. We refer interested readers to the Technical Note by Wu and Babich (2011).

In the two-period model, the second period has the same sequence of events as in the single-period model. The two periods are linked in three ways: a) The firms' utilities are derived on the sum of cash flows in both periods; b) The penalty can be history-dependent, e.g., higher penalty on double misreports; c) Inspection in the second period detects misreports in both periods. These three aspects capture the reality: Firms typically care about profit and loss aggregated over an accounting period; penalty for repeated misreports tends to be more severe; inspection cost is a function of the effort, not the length of the data obtained.

Recall that in a single-period game, two types of misreports can happen:

Type 1: Plant reports $r_t = \mathsf{U}$ when its status is $s_t = \mathsf{D}$ and $p_t < v$;

Type 2: Plant reports $r_t = \mathsf{D}$ when its status is $s_t = \mathsf{U}$ and $p_t > v$.

In a two-period game, two more types of misreports that benefit the distributor are possible:

Type 3: Plant reports $r_t = \mathsf{D}$ when its status is $s_t = \mathsf{U}$ and $p_t < v$;

Type 4: Plant reports $r_t = \mathsf{U}$ when its status is $s_t = \mathsf{D}$ and $p_t > v$.

In a single-period game analyzed in §3, types 3 and 4 misreports are dominated strategies for the plant. Therefore, there is no need to consider them. However, in a two-period game, types 3 and 4 misreports may benefit the strategic plant in some circumstances. Because type 3 and 4 misreports benefit the distributor, we assume that there is no penalty for them.

In a perfect Bayesian equilibrium, we show that type 3 and 4 misreports can occur in the following situation. When the plant already misreported in the first period and the penalty on double misreports is high, the plant could make a type 3 or 4 misreport to the distributor in the second period, which indicates that it did not make double misreports of type 1 or 2. This reduces or eliminates the distributor's incentive to inspect and allows the plant to keep the misreporting benefit from the first period.

Our analysis also shows that whenever a penalty is imposed, a higher penalty level benefits both firms by raising both firms' expected utilities. As the penalty level increases, the inspection probability goes to zero. Because the inspection probability diminishes, the aforementioned benefit of type 3 or 4 misreport shrinks as well. In fact, we show that under very high penalty, type 3 or 4 misreport will no longer occur.

A more important strategic behavior present in the two-period model is that the plant may choose to give up the benefit from misreporting in the first period, so that it reserves the opportunity of misreporting in the second period without being penalized for double misreports. Such a strategy is desirable when the penalty on double misreports is high and the benefit from the first-period misreport is small. In our analysis, we rigorously derive conditions under which such strategic behavior will be played in the equilibrium (see Proposition TN.1).

Interestingly, through extensive numerical comparison of various contract structures, we show that an extra penalty on double misreports is undesirable for risk allocation, because it induces the plant to give up small benefits in the first period, and the single-period analysis tells us that small benefits allocated to the risk-averse plant are beneficial for risk allocation between the two parties.

In short, our theoretical and numerical analysis on the two-period game conclude that although richer strategic interactions between the contractual parties exist in the two-period game, they do not improve the risk allocation between the firms; employing the same penalty structure for both periods appears to be the best for risk allocation. The implication for a multi-period game is that a more complicated penalty structure may induce strategic behaviors that are undesirable for risk allocation.

6. Implementation Issues and Risk Allocation in Multiple Periods

In this section, we discuss how the misreporting problem can be managed in practice, and further study a multi-period model to analyze aggregate risk allocation. In this multi-period model, the same penalty structure (v, L_l, L_r) studied in §4.3 is used in every period. From the analysis of the two-period model in §5, we know that a more complicated penalty structure may induce strategic behaviors that are not desirable for risk allocation.

6.1 Implementation of the Optimal Contract

In practice, "no penalty for misreporting" is unlikely to be the type of language used in any contract. How should the distributor write the contract to manage the plant's misreporting behavior?

Recognizing that the purpose of the zero penalty is to give extra benefits to the plant in certain circumstances, the distributor can specify those extra benefits in the contract.

Specifically, any unit-contingent contract with constant contract price v and zero-penalty regions L_l and L_r can be implemented as a portfolio of the following three components:

1. A unit-contingent contract with price v and very high penalty for misreporting (effectively prohibiting plant's misreporting).

2. A put option to the plant with strike price v, which can be exercised *only when* the plant reports down and $p \in L_l$. The value of this option upon exercising is $v - p$.

3. A call option to the plant with strike price v, which can be exercised *only when* the plant reports up and $p \in L_r$. The value of this option upon exercising is $p - v$.

In the above portfolio, the put and call options are state-contingent and provide the plant with the same benefit as having the zero-penalty regions L_l and L_r. The option premia are already factored in the contract price, so no upfront payment is necessary.

In the analysis in §4.3, L_i $(i = l, r)$ must be contained in the feasible region S_i defined in (11), because when p falls outside of that region, the benefit of misreporting $|v-p|$ is large and, even if the penalty is zero, the distributor may still inspect to prevent the plant from getting the large benefit $|v - p|$. However, in the call and put options specified above, L_i is simply the price region where the distributor commits to pay $|v - p|$ to the plant that is in a certain status. Hence, L_i does not have to be contained in S_i. Nevertheless, a good power purchase agreement should provide the plant with an incentive to maintain its reliability. An L_l that contains a range of low prices would mean that the plant may get a large benefit even if it is DOWN, which reduces the incentive for maintaining reliability.

For the rest of this paper, we assume that L_l is at least above the plant's production cost c, so that a DOWN plant never obtains a profit larger than $v - c$. Let $\widetilde{S}_l = [\underline{p}, v]$, $\underline{p} \geq c$ and $\widetilde{S}_r = [v, \infty)$ denote the feasible regions for exercising the options. Because $\widetilde{S}_l$ and $\widetilde{S}_r$ are intervals, Proposition 3 b(ii) and b(iii) imply that L_l^* and L_r^* are also intervals. Proposition 4 identifies an additional property of the optimal contract under general utility functions for both firms.

Proposition 4 *Consider solving (14) with $v(p) \equiv v$ and S_i replaced by $\widetilde{S}_i$. If L_l^* is strictly contained in $\widetilde{S}_l$, then $L_r^* = \varnothing$.*

Proposition 4 essentially says if the portfolio contains a call option ($L_r^* \neq \varnothing$), then it must contain a put option whose exercise region cannot be further expanded (i.e., $L_l^* = \widetilde{S}_l$).

6.2 Multi-Period Contract Execution and Aggregate Risk Allocation

In practice, the execution of the unit-contingent contract lasts for many periods (thousands of periods typically). Both firms derive utilities from aggregate income over several months or years, and are concerned about the aggregate risk allocation.

If the firms' profits were independent across all periods, then improving the risk allocation in each period would be equivalent to improving the aggregate risk allocation. This is not the case in practice, where a period is an hour and profits in adjacent hours are correlated. In this section, we quantify the variability of the total profit over many periods and analyze aggregate risk allocation. We consider contracts that have constant contract price and structural properties in Propositions 3 and 4, and contract parameters do not change over multi-period execution. As discussed in §6.1, such contracts are straightforward to implement in practice. Finding the optimal contract with those structural properties is already non-trivial, as we will see below.

Assume the contract execution stage comprises N periods. The spot price process and the plant's status process are assumed to have stationary distributions. Let us redefine $g(p)$ as the stationary distribution of the spot price, and redefine γ as the long-run fraction of time that plant is UP. Let π_P and π_D be random variables representing the long-run stationary distribution of the single-period profit for the plant and the distributor, respectively.

Let Π_P and Π_D denote respectively the plant's and distributor's total profits over N periods. Because N is typically large in practice, we employ normal approximation for Π_P and Π_D. We apply the Central Limit Theorem for dependent random variables (Billingsley 1995, p. 363) because profits are serially correlated. Let the profit sequences be $\{\pi_{P1}, \pi_{P2}, \dots\}$ and $\{\pi_{D1}, \pi_{D2}, \dots\}$. For ease of exposition, assume π_{P1} and π_{D1} have stationary distributions ($\pi_{P1} \overset{\text{dist}}{=} \pi_P$, $\pi_{D1} \overset{\text{dist}}{=} \pi_D$) and both sequences are stationary. Then, we have the following relations:

$$N^{-1}\mathsf{E}[\Pi_P] = \mathsf{E}[\pi_P], \qquad N^{-1}\mathsf{Var}[\Pi_P] \;\to\; \mathsf{Var}[\pi_P] + 2\sum_{k=1}^{\infty} \mathsf{Cov}(\pi_{P1}, \pi_{P,1+k}), \qquad N \to \infty, \tag{15}$$

$$N^{-1}\mathsf{E}[\Pi_D] = \mathsf{E}[\pi_D], \qquad N^{-1}\mathsf{Var}[\Pi_D] \;\to\; \mathsf{Var}[\pi_D] + 2\sum_{k=1}^{\infty} \mathsf{Cov}(\pi_{D1}, \pi_{D,1+k}), \qquad N \to \infty. \tag{16}$$

The relations in (15)-(16) highlight that serial correlations of the profits contribute to the aggregate profit variability. Reducing the single-period profit variance, $\mathsf{Var}[\pi_P]$ and $\mathsf{Var}[\pi_D]$, only represents part of the total variance reduction. Do serial correlations also decline at the same time? Intuitively, if we reduce the profit variability in each period, serial correlations are likely to decrease as well – at least this is true if variability is reduced to zero.

Unfortunately but interestingly, for our problem serial correlations may sometimes increase faster than the decline in the single-period profit variance and, therefore, caution must be taken in analyzing the aggregate profit variability. Because serial correlations may increase or decrease aggregate profit variability (depending on how the spot price and plant's status evolve), we resort to numerical analysis in the next section to fully reveal the impact of serial correlations. Nevertheless, the following proposition summarizes some important findings.

Proposition 5 *Suppose $\gamma > 1/2$ and consider reducing the contract price v and expanding the put option's exercise region L_l, while holding $L_r = \varnothing$ and keeping $\mathsf{E}[\Pi_D]$ and $\mathsf{E}[\Pi_P]$ constant. Denote the lowest possible contract price as $v^\dagger$, at which L_l expands to the entire feasible region, i.e., $L_l = \widetilde{S}_l = [\underline{p}, v^\dagger]$. Let $\pi_P(v, L_l)$ and $\Pi_P(v, L_l)$ denote respectively the plant's single-period and aggregate profit under the contract $(v, L_l, \varnothing)$; similar notations apply to the distributor.*

(i) For any $v > v^\dagger$ and $L_l \subset \widetilde{S}_l$ that satisfy $\mathsf{E}[\Pi_P(v, L_l)] = \mathsf{E}[\Pi_P(v^\dagger, \widetilde{S}_l)]$, we have

$$\mathsf{Var}[\pi_P(v, L_l)] \ > \mathsf{Var}[\pi_P(v^\dagger, \widetilde{S}_l)] \quad \text{and} \quad \mathsf{Var}[\pi_D(v, L_l)] \ > \mathsf{Var}[\pi_D(v^\dagger, \widetilde{S}_l)]. \tag{17}$$

(ii) When v decreases, $\sum_{k=1}^{\infty} \mathsf{Cov}(\pi_{P1}, \pi_{P,1+k})$ may change in the opposite direction to the changes in $\mathsf{Var}[\pi_P]$ and, consequently, we may have

$$\mathsf{Var}[\Pi_P(v, L_l)] \ < \ \mathsf{Var}[\Pi_P(v^\dagger, \widetilde{S}_l)], \qquad \text{for some } v > v^\dagger.$$

Proposition 5(i) essentially says that whenever the put option's exercise region L_l is strictly smaller than the feasible region $\widetilde{S}_l = [\underline{p}, v]$, we can reduce both firms' single-period profit variability without changing the average profit by lowering v to $v^\dagger$ and setting $L_l = \widetilde{S}_l$.

If the inequalities in (17) could also hold for serial correlations, then from (15)-(16), any contract with $v > v^\dagger$ would not minimize the total profit variability. However, Proposition 5(ii) shows that the serial correlations may sometimes increase faster than the decline in the single-period profit variance.

7. Numerical Analysis

In this section, we simulate the firms' profit streams under various contracts of the form proposed in §6.1, quantify the profit variability, and identify superior risk allocations.

We run a 5,000-period simulation. We assume that the logarithm of the spot price follows a mean-reverting process:

$$d\log(p_t) = \eta(m - \log(p_t))dt + \sigma dW_t, \qquad m = 4.2, \quad \eta = 0.3, \quad \sigma = 0.387.$$

We let the initial price have the same distribution as the long-run stationary distribution. Thus, at any time t, $\log(p_t)$ has a mean of 4.2 and a standard deviation of $\sigma/\sqrt{2\eta} = 0.5$; hence, the average price is \$75.57/MWh. (The logarithm of the spot price at the PJM hub in July-August 2008 has a mean of 4.2 and a standard deviation of 0.5, estimated using the hourly price data.)

The plant's capacity is assumed to be 100 MW. Its status changes in a Markovian fashion: Being UP this period, it will be UP the next period with a probability of 0.95; being DOWN this period, it will return to UP the next period with a probability of 0.2. It can be calculated that the long-run fraction of time that the plant is up is $\gamma = 0.8$ (see http://www.nerc.com/page.php?cid=4|43|47 for statistics on generating unit availability). The production cost is $c =$ \$30/MWh (average operating cost of fossil steam engines is \$29.59/MWh, reported by Energy Information Administration 2007). The distributor's firm contract price is $f =$ \$80/MWh (PJM futures for 2008 summer electricity were traded around \$80/MWh in 2005; data available from Bloomberg).

Our numerical procedure can be used to find the optimal contract under specific utility functions. But a much more appealing way to demonstrate improvement in risk allocation is to compare the profit distributions under various contracts. Thus, in the following analysis, we do not impose any specific utility function for either firm, and we focus on comparing firms' profit variabilities while keeping the firms' average profits constant.

Aggregate Profit Distributions under Two Key Contracts

We illustrate the simulation results under the following two important contracts:

Contract A: Unit-contingent contract with contract price $v_0 =$ \$70/MWh and a high penalty to deter misreporting.

Contract B: Unit-contingent contract with contract price $v =$ \$67.93/MWh and a high penalty to deter misreporting, and a put option with strike price \$67.93/MWh, which can be exercised only when the plant is down and the spot price is between \$30 and \$67.93 per MWh, i.e., $L_l = \widetilde{S}_l = [c, v]$. (The parameters are such that the average profits are the same as in Contract A.)

We ran the simulation 10,000 times. For each simulation run and each contract above, we computed the plant's and distributor's total profit over 5,000 hours. The profit distributions under the two contracts are compared in Figure 4. We can see that normal distribution is a good approximation for the total profit.

Under Contract A, the plant's total profit averages at $0.8 \times (\$70/\text{MWh} - \$30/\text{MWh}) \times 100\text{MWh} \times 5000$ hours = \$16 million. Contract B yields the same average profits for both firms, but it also brings lower profit variability for both firms: The standard deviations of both firms' profits are reduced significantly (around 15% reduction) compared to Contract A.

Figure 4: Aggregate Profit Distributions

Histograms are generated based on 10,000 simulated paths. Average profits are the same under both contracts. The plant's total profit variability (standard deviation) under Contract B is 16% (or \$48,000) lower than that under Contract A. The distributor's total profit variability under Contract B is 14% (or \$73,000) lower than that under Contract A.

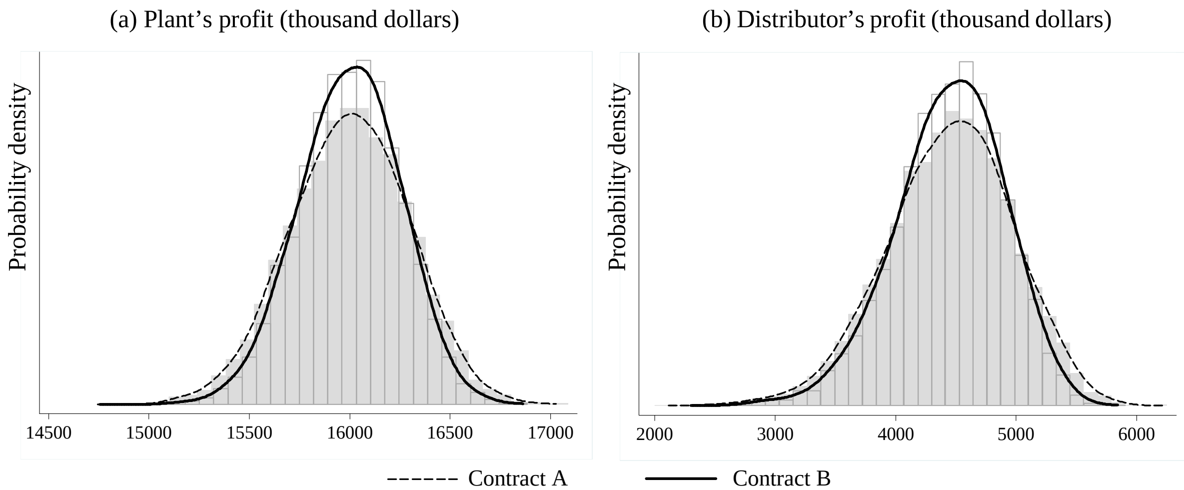

Impact of Serial Correlation on Total Profit Variability and Contract Choice

We now vary the contract price v in a wider range, adjust the state-contingent options to keep both firms' average profit constant, and measure profit variability reduction.

As discussed at the end of §6.1, the state-contingent put option should be used first, and the call option used only when the put option's exercise region L_l is equal to the entire feasible region $\widetilde{S}_l$. Thus, a kink point is $v^\dagger$ = \$67.93/MWh, at which $L_l = \widetilde{S}_l$ and $L_r = \varnothing$. This contract is exactly Contract B defined above.

For each contract price $v \in (v^\dagger, v_0)$, we let the put option's exercise region L_l be in the rightmost part of $\widetilde{S}_l$, i.e., $L_l = [p_1(v), v]$. We numerically find the price $p_1(v)$ that keeps both firms' average profit the same as under Contracts A and B. We also tried moving L_l to the leftmost of $\widetilde{S}_l$ or anywhere within $\widetilde{S}_l$, and found similar relations between serial correlation and profit variability.

Figure 5 illustrates the key numerical results. In Figure 5(a), the single-period profit variabilities (dashed curves) decline for both firms when the contract price v decreases from v_0 to $v^\dagger$, but increase if v drops below $v^\dagger$. Both firms' single-period profit variabilities are minimized at $v^\dagger$ = \$67.93/MWh, i.e., under Contract B.

In the multi-period setting, other contracts may allocate risk between the two firms better than Contract B, because serial correlation affects the total profit variability. For the distributor, the total profit variability (curves with triangular markers) has the same trend as the single-period profit variability. For the plant, however, serial correlation changes the trend of the total profit variability (curves with round markers), especially around the kink $v^\dagger$. With low serial correlation (or high mean-reverting speed, left panel in Figure 5(a)), total profit variability is still minimized at $v^\dagger$, but that is not the case with medium to high serial correlations. The kink is even reversed for high serial correlation (right panel Figure 5(a)), because when v decreases, the serial correlation may change in the opposite direction to the change in single-period variability (Proposition 5(ii)).

Figure 5: Impact of Serial Correlation on Total Profit Variability

The total profit variability $\mathsf{Var}(\Pi(v))$ is calculated based on 10,000 simulated paths. The single-period profit variability $\mathsf{Var}(\pi(v))$ is calculated based on the analytical expressions.

For either firm, $\dfrac{\mathsf{Var}(\Pi(v))}{\mathsf{Var}(\Pi(v_0))} > \dfrac{\mathsf{Var}(\pi(v))}{\mathsf{Var}(\pi(v_0))}$ if and only if $\dfrac{\sum_{k=1}^{\infty} \mathsf{Cov}(\pi_1(v), \pi_{1+k}(v))}{\sum_{k=1}^{\infty} \mathsf{Cov}(\pi_1(v_0), \pi_{1+k}(v_0))} > \dfrac{\mathsf{Var}(\pi(v))}{\mathsf{Var}(\pi(v_0))}$.

(a) Profit Variability Reductions as Functions of Contract Price $\boldsymbol{v}$

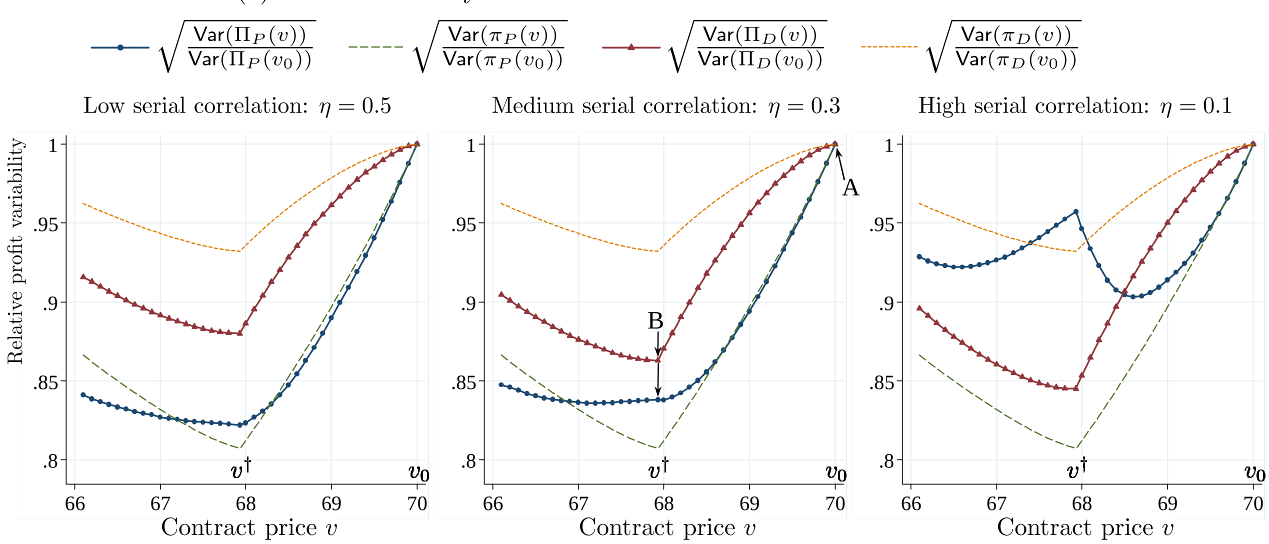

(b) Plant's Relative Total Profit Variability vs. Distributor's Relative Total Profit Variability

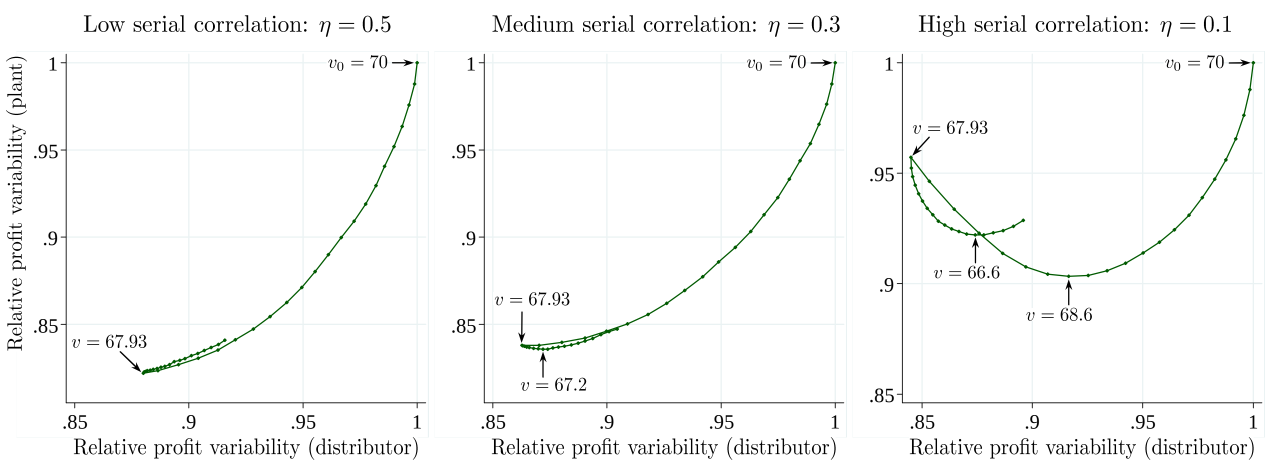

An intuitive explanation follows. When $v = v_0 = \$70$/MWh, the distributor bears all the spot price risk during the plant outage period. When v decreases from v_0 to $v^\dagger$, part of that spot price risk is shifted to the plant. Although this benefits both firms in a single period, the plant's profit becomes more strongly tied to the spot price and, therefore, the serial correlation of the plant's profit increases and the total profit variability may not decrease (evident in the right panel in Figure 5(a)). When v further decreases below $v^\dagger$, the call option needs to be included, and thus even more spot price risk is shifted to the plant, but the call option's payoff is negatively correlated with the put option's payoff, which reduces the serial correlation, and the plant's total profit variability may actually decline (seen in Figure 5(a) middle and right panels).

Figure 5(b) presents the information in Figure 5(a) in a different way to help with contract selection. We could identify a set of frontier contracts, for which there is no other contract that yields lower profit variabilities for both firms. Depending on the risk appetites of both parties, different contracts on the frontier can be optimal in terms of risk allocation. For example, in Figure 5(b) right panel, if the plant is risk-averse and the distributor is risk-neutral, then the contract with $v^* = 68.6$ best allocates the risk between the firms. If the distributor is also risk-averse, a different contract on the frontier would be desirable.

In the above analysis, we held the average profit at the same level as Contracts A and B, but once both firms' profit variabilities are reduced, a different average profit might be agreeable. We can then use the same procedure above to reduce variability under the new average profit level.

8. Conclusions and Extensions

Unit-contingent power supply contracts are widely used in the electricity industry. Because plants often possess private information about their operational status, the unit-contingent feature may provide incentives for a plant to misreport its true status. Industry's default belief is that the plant's misreporting reduces the distributor's profit (for any given contract) and, therefore, the distributor should either ban misreporting or lower the unit-contingent contract price to compensate for the losses due to misreporting. Our analysis shows that misreporting can serve as a risk-allocation mechanism between the plant and the distributor, and although some forms of misreporting are detrimental, others are beneficial for both the plant and the distributor. One of the managerial takeaways from this paper is that plant's reporting up when being down often improves risk allocation, whereas reporting down when being up typically impairs risk allocation.

To find how much misreporting (of each type) should be tolerated to maximize the distributor's benefit, we have identified structural properties of the optimal contract. These structural properties simplify the computational search for the optimal contract: The distributor can gradually reduce the contract price and expand the spot price region where misreporting is penalty-free. The region should be expanded below the contract price first (i.e., to tolerate misreporting up when down). In the typical case where the distributor is risk-neutral and the plant is risk-averse, this region is an interval with right endpoint being the contract price – a particularly convenient property in practice. Expanding the penalty-free region above the contract price should be considered when further expansion of the region below the contract price is not possible. Such a structural expansion of the penalty-free region essentially converts the problem into a one-dimensional search problem.

Richer strategic interactions exist in the two-period game, but they do not improve the risk allocation between the firms; employing the same penalty structure for both periods appears to be the best for risk allocation. The implication for a multi-period game is that a more complicated penalty structure may induce strategic behaviors that are undesirable for risk allocation.

When the penalty-free misreporting is difficult to implement as a contract clause, the distributor can adopt the following approach: prohibit misreporting (via high penalties) but provide the plant with state-contingent options. The cash flows from the plant's reporting up when down (down when up) are equivalent to the payoffs from a put (call) option. In the optimal portfolio, the distributor always uses the put option, but may or may not use the call option.

When implementing the above portfolio over multiple periods, the impact of serial correlations of the profits should be taken into account. A higher serial correlation of profit implies a higher total profit variability over the entire contract execution period. A practical takeaway from considering serial correlation is that the exercise region of the put (call) option will typically shrink (expand) because the put (call) option increases (decreases) the serial correlation of the plant's profit.

Our analysis can be extended in several directions. In practice, electricity distributors manage a portfolio of power supply contracts and service contracts. Large distributors may be treated as risk-neutral entities, and the contracting problem with each plant can be considered separately. For a risk-averse distributor, the portfolio selection problem should be analyzed using a model that contains multiple heterogeneous power plants and the distributor's demand profile (including customers' preferences). The analysis in this paper shows that both the distributor's and the plant's risk can be reduced when the misreporting issue is properly managed. With multiple unit-contingent contracts with multiple plants, if the contractual terms are similar, cash flows across multiple plants are positively correlated. Thus, our variance-reduction results will likely to be strengthened in the context of multiple unit-contingent contracts.

Power plants with multiple generating units may engage different units in different types of contracts. Concerns arise when a unit under a firm contract is down while another unit under a unit-contingent contract is running. In such a situation, the plant may wish to shift the output from the latter unit to satisfy the firm contract obligation first. Such a behavior can be analyzed using the approach in our paper, because the incentive for shifting the output is the same: The plant has an incentive to misreport down and shift output only if the spot price (the price of replacement power) is higher than the unit-contingent contract price.

This paper focuses on managing the incentive problem inherent in the unit-contingent contracts and the related risk allocation between a distributor and a plant. In addition to the physical contracts, firms also can use financial contracts to shape their risk profiles. A range of electricity derivatives can be used to shape market participants' risk profiles. Deng and Oren (2006) provide an excellent survey on electricity derivatives and risk management. The payoffs of the financial instruments depend on market prices that are realized before the plant reports its status, and thus these payoffs become the initial wealth of the firms at the beginning of the Bayesian game analyzed in the paper. Because the payoffs of the financial instruments do not depend on the plant status, the proposed state-contingent options are still beneficial, since they introduce payoffs dependent on the plant status. Thus, the risk-allocation mechanism analyzed in this paper is important for risk management.

Appendix

A. Preliminary Investigation of Possible Misreporting

Below we present a preliminary investigation into misreporting behavior based on limited data. Figure 6 shows the power delivery pattern of a coal-fired power generating unit. The data is hourly. Figure 6(a) shows the power delivery pattern before the plant was engaged in a unit-contingent contract: When the spot price was above \$150/MWh, the unit was producing power at its full capacity for about 98% of the time. (We could not access the information about the contract type prior to the unit-contingent contract, but most likely it sold to the spot market prior to that.) Figure 6(b) reveals that when the same unit was engaged in unit-contingent transactions, less power was delivered to the distributor at high spot prices. In fact, when the spot price was above \$150/MWh, outages were reported in almost a third of the time. The highlighted data points in Figure 6(b) represent 32 hours of reported outages. The average spot price of these 32 data points was \$180/MWh. Suppose the unit-contingent contract was 100 MW at \$75/MWh and the distributor had a firm service contract of 100 MW at \$80/MWh. Without those 32 hours of reported outage, the distributor could have earned ($80 − $75)/MWh × 100 MW × 32 hrs = $16,000, but instead it suffered a loss of ($180 − $80)/MWh × 100 MW × 32 hrs = $320,000. Thus, those 32 hours of reported outage reduced the distributor's profit by $336,000 over a four-month period.

This example does not provide conclusive evidence of misreporting, but it shows the significant financial impact of misreporting if it occurs.

Figure 6: Electricity Delivery from a Coal-Fired Power Generating Unit

Disclaimer: We do not claim any conclusive evidence that the plant has misreported; the data presented here is for research purposes only, and should not be used for any other purpose. The data excludes pre-scheduled maintenance periods.

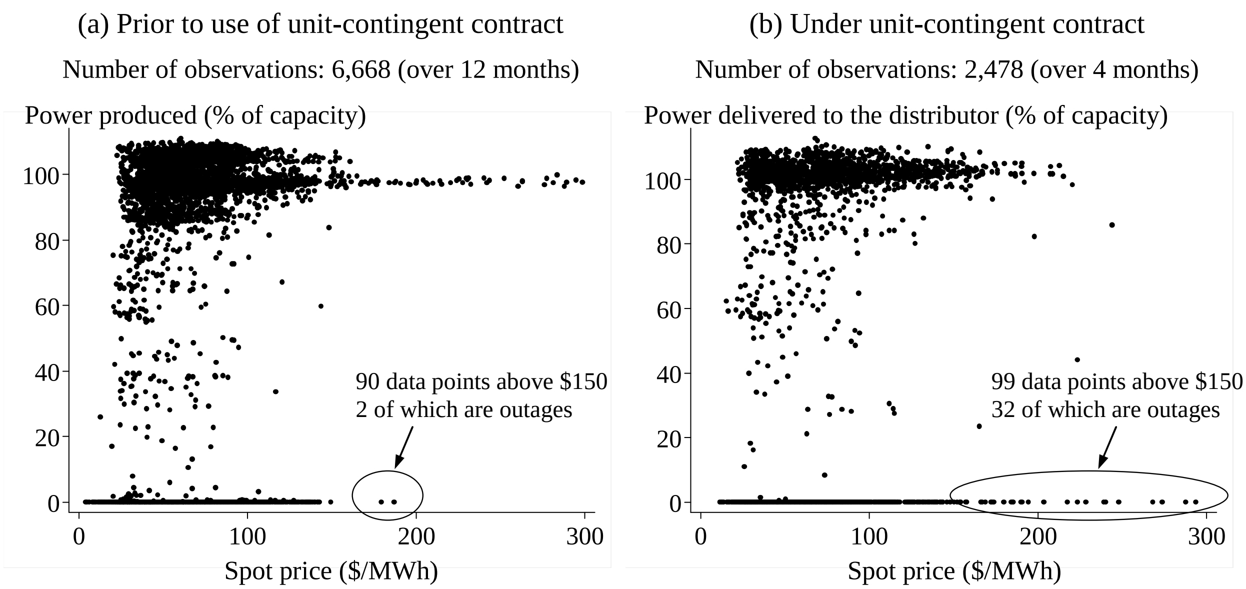

B. Related Market Background

This appendix provides information on market regulations and discusses why misreporting behavior in bilateral contracts is not monitored and cannot be directly observed by the distributor.

There are ten electricity wholesale markets in North America. Each market is organized by an independent system operator (ISO), a nonprofit organization providing market mechanisms to coordinate and control the operation of the electrical power system.

Electricity market rules are quite complex. At a high level, the following is the sequence of events in a unit-contingent contract:

1) The plant and the distributor schedule with the ISO.

2) The plant supplies electricity to the grid, and the distributor withdraws electricity from the grid. If the plant's status changes in real-time, it updates the schedule with the ISO.

3) The day after the operating day (or within 1 to 6 days after the operating day, depending on the contracts), the plant reports to the distributor its status on the operating day. This report is used to calculate the financial settlement between the plant and the distributor.

4) Actual payment is due on a set day (e.g., the 20th day) of the next month.

Note that the plant's report to the distributor, based on which the plant and the distributor settle their bilateral contracts (over which ISO has no supervision), comes the day after the operating day, when electricity prices have already been realized.

It is important to note that misreporting to an ISO is prohibited by Federal Energy Regulatory Commission (FERC), as stated in §5.1.1 and §8.2 of the Market Monitoring and Mitigation Manual of Midwest ISO (2009). The Manual also details the calculation of penalty charges for the misconduct. Due to the explicit market regulations and severe consequences, it is reasonable to assume that the plant always truthfully reports its status to the ISO. On the other hand, the plant reports to the distributor, because the distributor does not have the right to access the plant's true status recorded by the ISO for confidentiality reasons, as stated in §A.4 of the Market Settlements Calculation Guide of Midwest ISO (2010):

> *"Determinants are calculation components shown on Settlement Statements that enable verification of Charge Types. There are two types of determinants provided on Settlement Statements: 1) public, and 2) private. ... Private determinants represent confidential participant data related only to individual participants. Cleared transactions and cleared virtual schedule data are examples of private determinants. Participants have access to all public data, and only individual private data. The Midwest ISO provides as many determinants as possible while maintaining each participant's confidentiality. ... Although the participant knows their own load volume, they do not have access to the volume of all other participants individually."*

Furthermore, the ISO only acts as a transmission coordinator needed for the execution of bilateral contracts. Actually, monitoring bilateral contracts is beyond the ISO's responsibility, as described in §3.5 of the Market Monitoring and Mitigation Manual of Midwest ISO (2009):

> *"MMM (Market Monitoring and Mitigation) is generally concerned with any MP (market participant) behavior that affects the competitiveness of the Energy and Operating Reserve Markets and increases LMPs (Locational Marginal Prices), MCPs (Market Clearing Prices) or ORSGPs (Offer Revenue Sufficiency Guarantee Payments). The MMM process, however, is not directly concerned with Internal or External Bilateral Transaction Schedules, with bilateral Capacity, or*

with private transmission rights that are not under Midwest ISO administration, unless they affect the Energy and Operating Reserve Markets and services provided by the Midwest ISO."

A third party does not exist to monitor or regulate bilateral contracts, e.g., unit-contingent contracts. In short, the distributor does not have free access to the plant's private status.

References

Avenhaus, R., M. J. Canty 2005. Playing for time: A sequential inspection game. *European Journal of Operational Research* **167** 475–492.

Bachrach, D., R. Wiser, M. Bolinger, W. Golove 2003. Comparing the risk profiles of renewable and natural gas electricity contracts. Lawrence Berkeley National Laboratory, LBNL-50965.

Bernheim, B. D., M. D. Whinston 1998. Incomplete contracts and strategic ambiguity. *American Economic Review* **88**(4) 902–932.

Billingsley, P. 1995. *Probability and Measure*. 3rd edn. John Wiley & Sons. New York.

Boot, A. W. A., S. I. Greenbaum, A. V. Thakor 1993. Reputation and discretion in financial contracting. *American Economic Review* **83**(5) 1165–1183.

Boyabatlı, O., P. R. Kleindorfer, S. R. Koontz 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science*. Forthcoming.

California Independent System Operator 2003. Five-Year Assessment (2004-2008). Available at http://www.caiso.com/docs/09003a6080/28/5b/09003a6080285b79.pdf.

Correia, P. 2005. Games with incomplete and asymmetric information in poolco markets. *Power Systems, IEEE Transactions on* **20**(1) 83–89.

Demski, J. S., D. E. M. Sappington 1991. Resolving double moral hazard problems with buyout agreements. *RAND Journal of Economics* **22**(2) 232–240.

Deng, S.-J., S. S. Oren 2006. Electricity derivatives and risk management. *Energy* **31**(6-7) 940–953.

Diamond, H. 1982. Minimax policies for unobservable inspections. *Mathematics of Operations Research* **7**(1) 139–153.

Dresher, M. 1962. A sampling inspection problem in arms control agreements: A game-theoretic analysis. Memorandum for Advanced Research Projects Agency. RM-2972-ARPA.

Energy Information Administration 2007. Electric Power Annual 2006. Available at http://www.eia.doe.gov/cneaf/electricity/epa/epa_sum.html.

Ferrero, R., J. Rivera, S. Shahidehpour 1998. Application of games with incomplete information for pricing electricity in deregulated power pools. *Power Systems, IEEE Transactions on* **13**(1) 184–189.

Fudenberg, D., J. Tirole 1996. *Game Theory*. The MIT Press. Cambridge, Massachusetts.

Guesnerie, R., J.-J. Laffont 1984. A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm. *Journal of Public Economics* **25** 329–369.

Harsanyi, J. C. 1967. Games with incomplete information played by "Bayesian" players, I-III: Part I. The basic model. *Management Science* **14**(3) 159–182.

Harsanyi, J. C. 1968a. Games with incomplete information played by "Bayesian" players, I-III: Part II. Bayesian equilibrium points. *Management Science* **14**(5) 320–334.

Harsanyi, J. C. 1968b. Games with incomplete information played by "Bayesian" players, I-III: Part III. The basic probability distribution of the game. *Management Science* **14**(7) 486–502.

Heese, H. S., E. Kemahlioglu-Ziya 2009. Don't ask, don't tell: Revenue sharing with dishonest retailers. Working paper, Indiana University, University of North Carolina.

Hortaçsu, A., S. L. Puller 2008. Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market. *RAND Journal of Economics* **39**(1) 86–114.

Laffont, J.-J., D. Martimort 2002. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press. Princeton, New Jersey.

Li, T., M. Shahidehpour 2005. Strategic bidding of transmission-constrained GENCOs with incomplete information. *Power Systems, IEEE Transactions on* **20**(1) 437–447.

Midwest ISO 2009. Business Practices Manual: Market Monitoring and Mitigation. Available at http://www.midwestiso.org/publish.

Midwest ISO 2010. Business Practices Manual: Market Settlements Calculation Guide. Available at http://www.midwestiso.org/publish.

Mookherjee, D., I. Png 1989. Optimal auditing, insurance, and redistribution. *The Quarterly Journal of Economics* **104**(2) 399–415.

Myerson, R. B. 1979. Incentive compatibility and the bargaining problem. *Econometrica* **47** 61–73.

Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* **6**(1) 58–73.

Reinganum, J. F., L. L. Wilde 1985. Income tax compliance in a principal-agent framework. *Journal of Public Economics* **26**(1) 1–18.

Reyniers, D. J., C. S. Tapiero 1995a. Contract design and the control of quality in a conflictual environment. *European Journal of Operational Research* **82** 373–382.

Reyniers, D. J., C. S. Tapiero 1995b. The delivery and control of quality in supplier-producer contracts. *Management Science* **41**(10) 1581–1589.

Shahidehpour, M., H. Yamin, Z. Li 2002. *Market Operations in Electric Power Systems*. John Wiley & Sons. New York.

Townsend, R. M. 1979. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* **21**(2) 265–293.

Wabash Valley Power Association 2005. 2005 Summer Capacity Survey. Available at http://www.in.gov/iurc/files/05summer_wvpa.pdf.

Wabash Valley Power Association 2009. 2009 Summer Capacity Survey. Available at http://www.in.gov/iurc/files/WVPA_Summer_Reliability_Presentation_2009(1).pdf.

Wu, O. Q., V. Babich 2011. Technical Note for unit-contingent power purchase agreement and asymmetric information about plant outage. University of Michigan and Georgetown University.

Zaccaria, E., D. Hoog, D. Fromer, M. Mayers, D. Kane, K. Husar, G. Hawk, P. O'Neill 2006. Method and computerized system for reducing risk in an energy industry. *United States Patent*. No. 7027992. Available at http://www.patentstorm.us/patents/7027992/description.html.

Online Supplement

Proof of Proposition 1a [Bayesian equilibrium for $p > v(p)$]. If the plant is DOWN, reporting DOWN (truth-telling) is the dominating strategy for the plant. Therefore, $x^*_{\mathsf{D}} = 1$.

If the plant is UP, it obtains $U_P(v(p) - c)$ by reporting UP (truth-telling), and an expected utility of $y_{\mathsf{D}} U_P(v(p) - c - \phi(p)) + (1 - y_{\mathsf{D}}) U_P(p - c)$ by reporting DOWN, where y_{D} is the probability that the distributor inspects the plant that reports DOWN. Reporting DOWN is preferred by the plant if

$$y_{\mathsf{D}} < \frac{U_P(p-c) - U_P(v(p)-c)}{U_P(p-c) - U_P(v(p)-c-\phi(p))} \stackrel{\text{def}}{=} \widehat{y}_{\mathsf{D}} \in (0,1],$$

where $\widehat{y}_{\mathsf{D}} \in (0,1]$ because $p > v(p)$ and $\phi(p) \geq 0$. Thus, the plant's optimal response expressed as the truth-telling probabilities x_{U} and x_{D}, given the inspection probability y_{D}, is

$$\text{When } \mathsf{UP}, \begin{cases} x^*_{\mathsf{U}} = 0 \text{ (report } \mathsf{DOWN}\text{)}, & \text{if } y_{\mathsf{D}} < \widehat{y}_{\mathsf{D}}; \\ x^*_{\mathsf{U}} \in [0,1], & \text{if } y_{\mathsf{D}} = \widehat{y}_{\mathsf{D}}; \\ x^*_{\mathsf{U}} = 1 \text{ (report } \mathsf{UP}\text{)}, & \text{if } y_{\mathsf{D}} > \widehat{y}_{\mathsf{D}}; \end{cases} \tag{A.1}$$

$$\text{When } \mathsf{DOWN},\ x^*_{\mathsf{D}} = 1 \text{ (report } \mathsf{DOWN}\text{)}.$$

Next, we analyze the distributor's best response, given the plant's truth-telling probability. If the plant reports UP, the distributor knows that the plant is indeed UP (because pretending to be UP while DOWN is a dominated strategy for the plant when $p > v(p)$), so the distributor will not inspect: $y^*_{\mathsf{U}} = 0$. If the plant reports DOWN, the distributor forms a belief about the plant's actual status according to the Bayes rule:

$$\mathsf{P}\big\{\text{plant is } \mathsf{UP} \mid \text{plant reports } \mathsf{DOWN}\big\} = \frac{(1-x_{\mathsf{U}})\gamma}{(1-x_{\mathsf{U}})\gamma + (1-\gamma)} \stackrel{\text{def}}{=} \alpha(x_{\mathsf{U}}) \in [0,\gamma].$$

Based on the above belief about misreporting, if the distributor chooses to inspect, its expected utility is

$$\alpha(x_{\mathsf{U}}) U_D(f - v(p) - k + \phi(p)) + \big(1 - \alpha(x_{\mathsf{U}})\big) U_D(f - p - k).$$

Without inspection, the distributor can obtain $U_D(f-p)$. Thus, the distributor will not inspect if its belief about misreporting is below a certain level:

$$\alpha(x_{\mathsf{U}}) < \frac{U_D(f-p) - U_D(f-p-k)}{U_D(f - v(p) - k + \phi(p)) - U_D(f-p-k)} \stackrel{\text{def}}{=} \widehat{\alpha} \in (0,\infty). \tag{A.2}$$

Because $\alpha(x_{\mathsf{U}}) \leq \gamma$, if $\widehat{\alpha} > \gamma$, then, regardless of the plant's strategy x_{U}, inequality (A.2) holds and the distributor does not inspect. The condition $\widehat{\alpha} > \gamma$ is equivalent to $\phi(p) < \widehat{\phi}(p)$, where $\widehat{\phi}(p)$ is given below and also defined in (1) in the paper:

$$\widehat{\phi}(p) \stackrel{\text{def}}{=} v(p) + k - f + U_D^{-1}\left(\frac{1}{\gamma} U_D(f-p) - \frac{1-\gamma}{\gamma} U_D(f-p-k)\right), \qquad \text{for } p > v(p).$$

In other words, regardless of the plant's strategy, a penalty payment that is below $\widehat{\phi}(p)$ will not provide sufficient incentive for the distributor to inspect at cost $k > 0$. If $\phi(p) \geq \widehat{\phi}(p)$, then the distributor's inspection decision will depend on the plant's truth-telling probability. The distributor will not inspect if (A.2) holds, which is equivalent to $x_{\mathsf{U}} > \alpha^{-1}(\widehat{\alpha})$, or

$$x_{\mathsf{U}} > 1 - \frac{(1-\gamma)\widehat{\alpha}}{\gamma(1-\widehat{\alpha})} \stackrel{\text{def}}{=} \widehat{x}_{\mathsf{U}},$$

where $\widehat{x}_{\mathsf{U}} \in [0,1)$ because $\widehat{\alpha} \in (0,\gamma]$ when $\phi(p) \geq \widehat{\phi}(p)$.

In summary, the distributor's best response expressed as inspection probabilities y_{U} and y_{D}, given the plant's strategy x_{U}, is

$$\text{When plant reports UP, } y^*_{\mathsf{U}} = 0 \text{ (do not inspect);}$$

$$\text{When plant reports DOWN, } \begin{cases} \text{If } \phi(p) < \widehat{\phi}(p), \text{ then } y^*_{\mathsf{D}} = 0 & \text{(do not inspect);} \\ \text{If } \phi(p) \geq \widehat{\phi}(p), \text{ then} & \\ \quad y^*_{\mathsf{D}} = 1 \text{ (inspect)}, & \text{if } x_{\mathsf{U}} < \widehat{x}_{\mathsf{U}}; \\ \quad y^*_{\mathsf{D}} \in [0,1], & \text{if } x_{\mathsf{U}} = \widehat{x}_{\mathsf{U}}; \\ \quad y^*_{\mathsf{D}} = 0 \text{ (do not inspect)}, & \text{if } x_{\mathsf{U}} > \widehat{x}_{\mathsf{U}}. \end{cases} \tag{A.3}$$

Combining the best response functions in (A.1) and (A.3), we derive the Bayesian equilibrium as follows. The plant truthfully reports its status if it is DOWN, and the distributor never inspects when the plant reports UP:

$$x^*_{\mathsf{D}} = 1, \qquad y^*_{\mathsf{U}} = 0.$$

If $\phi(p) < \widehat{\phi}(p)$, the plant reports DOWN when it is UP, and the distributor does not inspect:

$$x^*_{\mathsf{U}} = 0, \qquad y^*_{\mathsf{D}} = 0. \tag{A.4}$$

If $\phi(p) = \widehat{\phi}(p)$, there are many equilibria: $x^*_{\mathsf{U}} = 0$, $y^*_{\mathsf{D}} \in [0, \widehat{y}_{\mathsf{D}}]$. The distributor is indifferent among all these equilibria, but the plant would prefer the distributor not to inspect. Thus, we assume the equilibrium played in this scenario is the same as in (A.4).

If $\phi(p) > \widehat{\phi}(p)$ and $\phi(p) > 0$, we have a mixed strategy equilibrium:

$$x^*_{\mathsf{U}} = \widehat{x}_{\mathsf{U}}, \qquad y^*_{\mathsf{D}} = \widehat{y}_{\mathsf{D}}.$$

If $\phi(p) > \widehat{\phi}(p)$ and $\phi(p) = 0$, there are many equilibria: $x^*_{\mathsf{U}} \in [0, \widehat{x}_{\mathsf{U}}]$, $y^*_{\mathsf{D}} = 1$. The plant is indifferent among all these equilibria, but the distributor would prefer the plant to truthfully report with probability $\widehat{x}_{\mathsf{U}}$. Thus, we assume that the equilibrium played in this scenario is:

$$x^*_{\mathsf{U}} = \widehat{x}_{\mathsf{U}}, \qquad y^*_{\mathsf{D}} = 1 = \widehat{y}_{\mathsf{D}}.$$

The above Bayesian equilibria can be summarized into two cases $\phi(p) \leq \widehat{\phi}(p)$ and $\phi(p) > \widehat{\phi}(p)$, which are exactly what Proposition 1a describes.

Next, we prove the firms' expected utilities are those in (2) and (3) in the paper for any given spot price realization $p > v(p)$. If $\phi(p) \leq \widehat{\phi}(p)$, the plant's profit is zero when DOWN, and it reports DOWN when UP to obtain a profit of $p - c$. Thus, the plant's expected utility is $\gamma U_P(p-c)$ (recall that the plant is UP with probability γ and $U_P(0) = 0$). Since the plant always reports DOWN in this case, the distributor always purchases from the spot and obtains a utility of $U_D(f-p)$.

If $\phi(p) > \widehat{\phi}(p)$, the plant's profit is zero when DOWN. When the plant is UP, it plays a mixed strategy, so its expected utility is equal to the utility it obtains by playing any pure strategy, e.g., reporting UP gives the plant a utility of $U_P(v(p)-c)$. Thus, the plant's expected utility is $\gamma U_P(v(p)-c)$.

When the plant reports UP (which happens with probability γx^*_{U}), the distributor does not inspect and gets a profit of $f-v(p)$. When the plant reports DOWN, the distributor plays a mixed strategy, so its expected utility is equal to the utility it obtains by playing any pure strategy, e.g., not conducting inspection, which gives the distributor a utility of $U_D(f-p)$. Thus, the distributor's expected utility is $\gamma x^*_{\mathsf{U}} U_D(f-v(p)) + (1-\gamma x^*_{\mathsf{U}})U_D(f-p)$, as shown in (3) in the paper. ∎

Proof of Proposition 1b. The proof parallels that for Proposition 1a.

Proof of Lemma 1. We set $\phi(p) = \infty$ and look for the optimal contract price function $v(p)$. Consider $\max_{v(p)} \mathsf{E}[U_D]$, *s.t.* $\mathsf{E}[U_P] \geq \underline{U_P}$, where $\mathsf{E}[U_P]$ and $\mathsf{E}[U_D]$ are defined in (8) in the paper. The Lagrangian is

$$\gamma \mathsf{E} U_D(f-v(p)) + (1-\gamma)\mathsf{E} U_D(f-p) + \lambda\big(\gamma \mathsf{E} U_P(v(p)-c) - \underline{U_P}\big),$$

where $\lambda \geq 0$ is the Lagrangian multiplier. The Euler-Lagrange equation for this problem is:

$$U_D'(f-v(p)) = \lambda U_P'(v(p)-c), \quad \forall p.$$

The above Euler-Lagrange equation requires the contract price to be constant: $v(p) \equiv v$. To determine the optimal contract price, note that when v increases, $\mathsf{E}[U_P]$ increases while $\mathsf{E}[U_D]$ decreases. Therefore, the distributor should set the contract price at v_0, determined by

$$\gamma U_P(v_0 - c) = \underline{U_P}.$$

The distributor's expected utility under the optimal unit-contingent truth-telling contract is

$$\mathsf{E}[U_D] = \gamma U_D(f-v_0) + (1-\gamma)\mathsf{E} U_D(f-p). \tag{A.5}$$

Proof of Proposition 2a. We consider a contract of the form (9) in the paper: $v(p) \equiv v < v_0$, $\phi(p) = 0$ if $p \in [v-\delta, v]$, otherwise $\phi(p) = \infty$. Using the equilibrium utilities given in (3), (5), and (6) in the paper, the expected utilities of the plant and the distributor under the above contract can be expressed as:

$$\mathsf{E}[U_P] \;=\; \gamma U_P(v-c) + (1-\gamma)\int_{v-\delta}^{v} U_P(v-p)g(p)dp, \tag{A.6}$$

$$\begin{aligned} \mathsf{E}[U_D] &= \gamma U_D(f-v) + (1-\gamma)\int_{p\notin[v-\delta,v]} U_D(f-p)g(p)dp + (1-\gamma)\int_{v-\delta}^{v} U_D(f-v)g(p)dp \\ &= \gamma U_D(f-v) + (1-\gamma)\mathsf{E} U_D(f-p) - (1-\gamma)\int_{v-\delta}^{v} \big[U_D(f-p) - U_D(f-v)\big]g(p)dp. \end{aligned} \tag{A.7}$$

We choose the parameters v and δ such that $\delta \in (0,k)$ and $\delta < v-c$, and that the plant's expected utility in (A.6) is maintained at the reservation level $\underline{U_P} = \gamma U_P(v_0 - c)$. That is, v and δ also satisfy:

$$\gamma\big[U_P(v_0-c) - U_P(v-c)\big] - (1-\gamma)\int_{v-\delta}^{v} U_P(v-p)g(p)dp \;=\; 0. \tag{A.8}$$

To see the existence of such v and δ, note that (A.8) defines an implicit function $\delta(v)$ with $\delta(v_0) = 0$. The function $\delta(v)$ is continuous because the left side of (A.8) is continuously differentiable in (δ, v) and its partial derivative with respect to δ is non-zero. Therefore, there exists v such that $\delta(v) \in (0, k)$ and $\delta(v) < v - c$.

By the concavity of the utility functions, for $v < v_0$ and $p \in [v - \delta, v]$, we have

$$\begin{aligned} U_D(f-v) - U_D(f-v_0) &\geq U_D'(f-v)(v_0-v) \;=\; U_P'(v-c)(v_0-v)\frac{U_D'(f-v)}{U_P'(v-c)} \\ &\geq \big[U_P(v_0-c) - U_P(v-c)\big]\frac{U_D'(f-v)}{U_P'(v-c)}, \end{aligned} \tag{A.9}$$

$$\begin{aligned} U_D(f-p) - U_D(f-v) &\leq U_D'(f-v)(v-p) \;=\; U_P'(v-c)(v-p)\frac{U_D'(f-v)}{U_P'(v-c)} \\ &\leq U_P(v-p)\frac{U_D'(f-v)}{U_P'(v-c)}, \end{aligned} \tag{A.10}$$

where the last inequality follows from $U_P(0) = 0$ and $p \geq v - \delta > c$.

We now compare the distributor's expected utility under the truth-telling contract in (A.5) with its expected utility under the contract (9) expressed in (A.7). Taking the difference and employing the relations in (A.9) and (A.10), we have

$$\begin{aligned} \text{(A.7)} - \text{(A.5)} &= \gamma\big[U_D(f-v) - U_D(f-v_0)\big] - (1-\gamma)\int_{v-\delta}^{v}\big[U_D(f-p) - U_D(f-v)\big]g(p)dp \\ &\geq \left[\gamma\big[U_P(v_0-c) - U_P(v-c)\big] - (1-\gamma)\int_{v-\delta}^{v} U_P(v-p)g(p)dp\right]\frac{U_D'(f-v)}{U_P'(v-c)} \;=\; 0, \end{aligned} \tag{A.11}$$

where the last equality is due to (A.8). Notice that the inequality (A.11) will be strict as long as one of the inequalities in (A.9) and (A.10) is strict. This means that as long as one of the distributor and the plant is risk-averse, there exists a contract in which the distributor can obtain a strictly higher utility than it would get under a truth-telling contract, while the plant's utility is kept constant. ∎

Proof of Proposition 2b. The proof is included in the Technical Note.

Proof of Proposition 3 [Structural properties of the optimal contract].

a. Properties of $v^*(p)$***.*** Consider the problem (14) in the paper. The Lagrangian is $\mathsf{E}[U_D] + \lambda(\mathsf{E}[U_P] - \underline{U_P})$, where $\mathsf{E}[U_P]$ and $\mathsf{E}[U_D]$ are defined in (12) and (13) in the paper, and $\lambda \geq 0$ is the Lagrangian multiplier. The Euler-Lagrange equations for this problem are:

$$U_D'(f - v^*(p)) = \lambda\Big[\gamma U_P'(v^*(p) - c) + (1-\gamma)U_P'(v^*(p) - p)\Big], \quad p \in L_l^*, \tag{A.12}$$

$$U_D'(f - v^*(p)) = \lambda\, U_P'(v^*(p) - c), \quad p \notin L_l^* \cup L_r^*. \tag{A.13}$$

There is no equation for $v(p)$ when $p \in L_r^*$, because the plant always reports DOWN when $p \in L_r^*$ and $v(p)$ does not affect the expected utilities.

(i) In the price region where misreporting is prohibited (i.e., $p \notin L_l^* \cup L_r^*$), eq. (A.13) implies that the contract price is a constant, denoted as $v^*(p) \equiv v_1$, $p \notin L_l^* \cup L_r^*$, where v_1 satisfies $U_D'(f - v_1) = \lambda\, U_P'(v_1 - c)$.

(ii) In the price region $p \in L_l^*$ where misreporting is not penalized, (A.12) implicitly determines $v^*(p)$.

Assuming that $U_P(\cdot)$ and $U_D(\cdot)$ are twice differentiable, we have

$$\frac{dv^*(p)}{dp} = \frac{\lambda(1-\gamma)U_P''(v^*(p)-p)}{\lambda\big[\gamma U_P''(v^*(p)-c)+(1-\gamma)U_P''(v^*(p)-p)\big]+U_D''(f-v^*(p))} \in [0,1].$$

If $c \in L_l^*$, (A.12) becomes $U_D'(f-v^*(c)) = \lambda U_P'(v^*(c)-c)$, which implies $v^*(c) = v_1$. Because $v^{*\prime}(p) \geq 0$, we have $v^*(p) \geq v_1$ for $p \geq c$. If the plant is risk-neutral, then $v^{*\prime}(p) = 0$ and $v^*(p) = v_1$ for $p \in L_l^*$.

b. Properties of L_l^* and L_r^*.

(i) L_r^* is to the right of L_l^*. The proof is included in the Technical Note.

(ii) Position of L_l^*. When $U_D(x) = x$, $\mathsf{E}[U_D]$ in (13) in the paper is the distributor's expected profit:

$$f - \gamma \mathsf{E}v(p) - (1-\gamma)\mathsf{E}p - (1-\gamma)\int_{L_l}(v(p)-p)g(p)dp - \gamma\int_{L_r}(p-v(p))g(p)dp. \tag{A.14}$$

Consider a contract $(v(p), L_l, L_r)$ that has the properties in parts (i) and (ii), and $p \geq c, \forall p \in L_l$. Suppose there exist $L_a \subseteq S_l \backslash L_l$ and $L_b \subseteq L_l$ such that $p_a > p_b$ for any $p_a \in L_a, p_b \in L_b$, and

$$\int_{L_a} U_P(v(p)-p)g(p)dp \;=\; \int_{L_b} U_P(v(p)-p)g(p)dp.$$

The above equality implies that if we adjust L_l to include L_a but exclude L_b, the plant's expected utility in (12) in the paper remains unchanged. We now show that this adjustment improves (A.14). Since L_a and L_b are disjoint, there exists p_o such that $p_b \leq p_o \leq p_a$, $\forall p_a \in L_a, p_b \in L_b$.

If $p_o \in L_l$, from part (i) we have $v(p_o) \geq v_1$ and $p_o - p_b \geq v(p_o) - v(p_b)$ because $v(p)$ increases in p at a rate no faster than one. If $p_o \notin L_l$, then $v(p_o) = v_1 \leq v(p_b)$. In either case, we have:

$$v(p_a) - p_a \;=\; v_1 - p_a \;\leq\; v(p_o) - p_o \;\leq\; v(p_b) - p_b.$$

Since $U_P(\cdot)$ is strictly concave and $U_P(0) = 0$, we have

$$\frac{U_P(v(p_b)-p_b)}{v(p_b)-p_b} \;\leq\; \frac{U_P(v(p_o)-p_o)}{v(p_o)-p_o} \;\stackrel{\text{def}}{=}\; C_3 \;\leq\; \frac{U_P(v(p_a)-p_a)}{v(p_a)-p_a},$$

where the equalities hold only when p_a or p_b coincides with p_o. Hence,

$$\begin{aligned}
\int_{L_a}(v(p)-p)g(p)dp \;&=\int_{L_a}\frac{v(p)-p}{U_P(v(p)-p)}U_P(v(p)-p)g(p)dp \\
&< C_3^{-1}\int_{L_a}U_P(v(p)-p)g(p)dp \qquad = C_3^{-1}\int_{L_b}U_P(v(p)-p)g(p)dp \\
&< \int_{L_b}\frac{v(p)-p}{U_P(v(p)-p)}U_P(v(p)-p)g(p)dp \;=\int_{L_b}(v(p)-p)g(p)dp.
\end{aligned}$$

Therefore, we can improve (A.14) by replacing L_b with L_a, which in effect shifts L_l toward the right.

If $p < c$ for some $p \in L_l$, we can still prove the above structure for L_l using similar arguments. The difference is that the adjustment not only involves replacing L_b with L_a, but also updating $v(p)$ for $p \in L_a \cup L_b$. The proof is omitted due to the length of this document.

Using similar method as above, we can show that when $U_P(x) = x$, the distributor can improve its own expected utility by shifting L_r toward the left.

(iii) Position of L_r^*. The following proof does not rely on the risk-neutrality of any firm. Consider a contract $(v(p), L_l, L_r)$ that has the properties in parts (i) and (ii). Suppose there exist $L_a \subseteq S_r \backslash L_r$ and

$L_b \subseteq L_r$ such that $p_a < p_b$ for any $p_a \in L_a, p_b \in L_b$, and

$$\int_{L_a} \big[U_P(p-c) - U_P(v(p)-c)\big] g(p)dp \quad = \quad \int_{L_b} \big[U_P(p-c) - U_P(v(p)-c)\big] g(p)dp.$$

Thus, if we adjust L_r to include L_a but exclude L_b, the plant's expected utility in (12) in the paper remains unchanged. We now show that this adjustment improves (13) in the paper.

Without loss of generality, we assume $v(p)$ is non-decreasing in p for $p > v(p)$. This can be achieved by setting $v(p) = v_1$ for $p > v_1$ and setting $v(p) = p - \epsilon$ when $p \le v_1$, where $\epsilon > 0$ is small. This

Since L_a and L_b are disjoint, there exists p_o such that $p_a \le p_o \le p_b$, $\forall p_a \in L_a, p_b \in L_b$. By the concavity of the utility functions, we have

$$\begin{aligned}
\frac{U_P(p_b-c) - U_P(v(p_b)-c)}{p_b - v(p_b)} &\le \frac{U_P(p_o-c) - U_P(v(p_o)-c)}{p_o - v(p_o)} \overset{\text{def}}{=} C_P \le \frac{U_P(p_a-c) - U_P(v(p_a)-c)}{p_a - v(p_a)},\\
\frac{U_D(f-v(p_a)) - U_D(f-p_a)}{p_a - v(p_a)} &\le \frac{U_D(f-v(p_o)) - U_D(f-p_o)}{p_o - v(p_o)} \overset{\text{def}}{=} C_D \le \frac{U_D(f-v(p_b)) - U_D(f-p_b)}{p_b - v(p_b)},
\end{aligned}$$

where the equalities hold only when p_a or p_b coincides with p_o. Hence,

$$\begin{aligned}
& \int_{L_a} \big[U_D(f-v(p)) - U_D(f-p)\big] g(p)dp \\
\le & \int_{L_a} (p - v(p)) C_D \; g(p)dp && = \int_{L_a} (p-v(p)) C_P \; g(p)dp \frac{C_D}{C_P} \\
\le & \int_{L_a} \big[U_P(p-c) - U_P(v(p)-c)\big] g(p)dp \frac{C_D}{C_P} && = \int_{L_b} \big[U_P(p-c) - U_P(v(p)-c)\big] g(p)dp \frac{C_D}{C_P} \\
\le & \int_{L_b} (p-v(p)) C_P \; g(p)dp \frac{C_D}{C_P} && = \int_{L_b} (p-v(p)) C_D \; g(p)dp \\
\le & \int_{L_b} \big[U_D(f-v(p)) - U_D(f-p)\big] g(p)dp.
\end{aligned}$$

Notice that if at least one of the two parties is risk-averse, then at least one of the above inequalities will hold strictly. Therefore, we can improve (13) by replacing L_b with L_a, which in effect shifts L_r toward the left. Thus, the optimal L_r should be contained in the leftmost part of S_r.

(iv) Constant contract price case. When $v(p) \equiv v$, by definition, S_l must be to the left of S_r, and part (i) holds. Parts (ii) and (iii) can be proven using the same approach detailed above. With a constant contract price, the definition of S_l in (11) in the paper and the definition of $\widehat{\phi}(p)$ in (4) in the paper imply that S_l is an interval. ∎

Proof of Proposition 4. Throughout this proof, $v(p) \equiv v$. Suppose both $\widetilde{S}_l \backslash L_l$ and L_r have positive measures. We consider shrinking L_r and expanding L_l to improve the objective. Let $L_a \subseteq \widetilde{S}_l \backslash L_l$ and $L_b \subseteq L_r$, such that L_a and L_b have positive measures and

$$(1-\gamma)\int_{L_a} U_P(v-p) g(p)dp \quad = \quad \gamma \int_{L_b} \big[U_P(p-c) - U_P(v-c)\big] g(p)dp.$$

That is, $\mathsf{E}[U_P] \ge \underline{U_P}$ remains satisfied if we add L_a to L_l and subtract L_b from L_r. We now show that this adjustment improves (13). For any $p_a \in L_a \backslash \{v\}$, $p_b \in L_b \backslash \{v\}$, we have $p_a < v < p_b$. Because $U_D(\cdot)$ is concave,

$$\frac{U_D(f-p_a) - U_D(f-v)}{v - p_a} \;\le\; U_D'(f-v) \;\le\; \frac{U_D(f-v) - U_D(f-p_b)}{p_b - v}.$$

Because $U_P(\cdot)$ is concave, $U_P(0)=0$, and $p_a > c$ (since $p_a \in \widetilde{S}_l$), we have

$$\frac{U_P(v-p_a)}{v-p_a} \;\geq\; U_P'(v-c) \;\geq\; \frac{U_P(p_b-c)-U_P(v-c)}{p_b-v}.$$

Hence,

$$\begin{aligned}
& (1-\gamma)\int_{L_a} \big[U_D(f-p)-U_D(f-v)\big]g(p)dp \\
= \;& (1-\gamma)\int_{L_a} \frac{U_D(f-p)-U_D(f-v)}{v-p}\cdot\frac{v-p}{U_P(v-p)}\,U_P(v-p)\,g(p)dp \\
\leq \;& (1-\gamma)\frac{U_D'(f-v)}{U_P'(v-c)}\int_{L_a} U_P(v-p)g(p)dp \;=\; \gamma\frac{U_D'(f-v)}{U_P'(v-c)}\int_{L_b}\big[U_P(p-c)-U_P(v-c)\big]g(p)dp \\
\leq \;& \gamma\int_{L_b}\frac{U_D(f-v)-U_D(f-p)}{p-v}\cdot\frac{p-v}{U_P(p-c)-U_P(v-c)}\big[U_P(p-c)-U_P(v-c)\big]g(p)dp \\
= \;& \gamma\int_{L_b}\big[U_D(f-v)-U_D(f-p)\big]g(p)dp.
\end{aligned}$$

Therefore, we can improve the objective by expanding L_l and shrinking L_r until that is not possible any more. Hence, if L_r has a positive measure, then $L_l = \widetilde{S}_l$. ∎

Proof of Proposition 5. For any contract $(v, L_l, \varnothing)$, $L_l \subset \widetilde{S}_l$, consider the following two contacts: $(v^e, \varnothing, \varnothing)$ and $(v^\dagger, L_l^\dagger, \varnothing)$ with $L_l^\dagger = \widetilde{S}_l = [\underline{p}, v^\dagger]$, that yield the same average profit:

$$\gamma(v-c)+(1-\gamma)\int_{L_l}(v-p)g(p)dp \;=\; \gamma(v^e-c), \tag{A.15}$$

$$\gamma(v^\dagger-c)+(1-\gamma)\int_{L_l^\dagger}(v^\dagger-p)g(p)dp \;=\; \gamma(v^e-c). \tag{A.16}$$

Note that v^e is uniquely determined from (A.15), and $v^\dagger$ is uniquely determined from (A.16). The latter is because $\gamma(v-c)+(1-\gamma)\int_{\underline{p}}^{v}(v-p)g(p)dp$ is strictly increasing in v. Furthermore, $v^\dagger < v < v^e$.

Summing up (A.15) and (A.16), we have

$$\gamma(v+v^\dagger-2c)+(1-\gamma)\left[\int_{L_l\cap L_l^\dagger}(v+v^\dagger-2p)g(p)dp+\int_{L_l\setminus L_l^\dagger}(v-p)g(p)dp+\int_{L_l^\dagger\setminus L_l}(v^\dagger-p)g(p)dp\right] = 2\gamma(v^e-c). \tag{A.17}$$

Taking the difference between (A.15) and (A.16) gives

$$\gamma(v-v^\dagger)+(1-\gamma)\left[(v-v^\dagger)\int_{L_l\cap L_l^\dagger}g(p)dp+\int_{L_l\setminus L_l^\dagger}(v-p)g(p)dp-\int_{L_l^\dagger\setminus L_l}(v^\dagger-p)g(p)dp\right] \;=\; 0. \tag{A.18}$$

We now show that the contract $(v^\dagger, L_l^\dagger, \varnothing)$ yields a lower profit variability for the plant than does the contract $(v, L_l, \varnothing)$. Consider the difference of the profit variance:

$$\begin{aligned}
\mathsf{E}[\pi_P^2]-\mathsf{E}[\pi_P^{\dagger 2}] &= \left[\gamma(v-c)^2+(1-\gamma)\int_{L_l}(v-p)^2g(p)dp\right]-\left[\gamma(v^\dagger-c)^2+(1-\gamma)\int_{L_l^\dagger}(v^\dagger-p)^2g(p)dp\right] \\
&= \gamma(v+v^\dagger-2c)(v-v^\dagger)+(1-\gamma)\int_{L_l\cap L_l^\dagger}(v+v^\dagger-2p)(v-v^\dagger)g(p)dp \\
&\quad +(1-\gamma)\int_{L_l\setminus L_l^\dagger}(v-p)^2g(p)dp-(1-\gamma)\int_{L_l^\dagger\setminus L_l}(v^\dagger-p)^2g(p)dp.
\end{aligned}$$

Substituting the first two terms on the right-side above using (A.17) and then combining integrals, we have:

$$\begin{aligned}\mathsf{E}[\pi_P^2] - \mathsf{E}[\pi_P^{\dagger 2}] &= 2\gamma(v^e - c)(v - v^\dagger) - (1-\gamma)\int_{L_l\backslash L_l^\dagger}(p - v^\dagger)(v-p)g(p)dp - (1-\gamma)\int_{L_l^\dagger\backslash L_l}(v-p)(v^\dagger - p)g(p)dp \\ &> 2\gamma(v^e - c)(v - v^\dagger) - (1-\gamma)\int_{L_l\backslash L_l^\dagger}(p - v^\dagger)(v-p)g(p)dp - (1-\gamma)(v - \underline{p})\int_{L_l^\dagger\backslash L_l}(v^\dagger - p)g(p)dp.\end{aligned}$$

Substituting the last term on the right-side above using (A.18), we have:

$$\begin{aligned}\mathsf{E}[\pi_P^2] - \mathsf{E}[\pi_P^{\dagger 2}] &> 2\gamma(v^e - c)(v - v^\dagger) - \gamma(v - v^\dagger)(v - \underline{p}) - (1-\gamma)(v - \underline{p})(v - v^\dagger)\int_{L_l\cap L_l^\dagger} g(p)dp \\ &\quad - (1-\gamma)\int_{L_l\backslash L_l^\dagger}\big[(p - v^\dagger)(v-p) + (v - \underline{p})(v-p)\big]g(p)dp.\end{aligned} \tag{A.19}$$

Since $p > v^\dagger > \underline{p}$ for $p \in L_l\backslash L_l^\dagger$, we can enlarge the last integrand in (A.19) as follows:

$$(p - v^\dagger)(v-p) + (v - \underline{p})(v-p) = (p - v^\dagger)(\underline{p} - p) + (v - \underline{p})(v - v^\dagger) \;<\; (v - \underline{p})(v - v^\dagger).$$

Then, we can combine the two integrals in (A.19) using $(L_l \cap L_l^\dagger) \cup (L_l\backslash L_l^\dagger) = L_l$, and factor out $(v - v^\dagger)$:

$$\begin{aligned}\mathsf{E}[\pi_P^2] - \mathsf{E}[\pi_P^{\dagger 2}] &> (v - v^\dagger)\left[\gamma(2v^e - 2c - v + \underline{p}) - (1-\gamma)(v - \underline{p})\int_{L_l} g(p)dp\right] \\ &> (v - v^\dagger)\left[\frac{1}{2}(2v^e - 2c - v + \underline{p}) - \frac{1}{2}(v - \underline{p})\int_{L_l} g(p)dp\right] \\ &> (v - v^\dagger)\left[\frac{1}{2}(2v^e - 2c - v + \underline{p}) - \frac{1}{2}(v - \underline{p})\right] \\ &= (v - v^\dagger)(v^e - c - v + \underline{p}) \;>\; 0.\end{aligned} \tag{A.20}$$

where the inequality (A.20) is due to $\gamma > 1/2$ and $2v^e - 2c - v + \underline{p} = (v^e - v) + (v^e - c) + (\underline{p} - c) > 0$. Hence, the plant's profit has a lower variance under contract $(v^\dagger, L_l^\dagger, \varnothing)$. For the distributor, we have

$$\begin{aligned}\mathsf{E}[\pi_D^2] - \mathsf{E}[\pi_D^{\dagger 2}] &= \gamma(2f - v - v^\dagger)(v^\dagger - v) + (1-\gamma)\left[\int_{L_l\cap L_l^\dagger}(2f - v - v^\dagger)(v^\dagger - v)g(p)dp\right. \\ &\quad \left. + \int_{L_l^\dagger\backslash L_l}(2f - p - v^\dagger)(v^\dagger - p)g(p)dp - \int_{L_l\backslash L_l^\dagger}(2f - p - v)(v - p)g(p)dp\right] \\ &> \gamma(2f - v - v^\dagger)(v^\dagger - v) + (1-\gamma)\left[\int_{L_l\cap L_l^\dagger}(2f - v - v^\dagger)(v^\dagger - v)g(p)dp\right. \\ &\quad \left. + \int_{L_l^\dagger\backslash L_l}(2f - v - v^\dagger)(v^\dagger - p)g(p)dp - \int_{L_l\backslash L_l^\dagger}(2f - v^\dagger - v)(v - p)g(p)dp\right] \\ &= (2f - v - v^\dagger)\big(\mathsf{E}[\pi_D] - \mathsf{E}[\pi_D^\dagger]\big) \;=\; 0,\end{aligned}$$

where the inequality follows from $p < v^\dagger < v$ for any $p \in L_l^\dagger\backslash L_l$, and $p > v^\dagger$ for any $p \in L_l\backslash L_l^\dagger$. In the final step, we used $\mathsf{E}[\pi_D] = \gamma(f - v) + (1-\gamma)(f - \mathsf{E}p) - (1-\gamma)\int_{L_l}(v-p)g(p)dp$ and a similar expression for $\mathsf{E}[\pi_D^\dagger]$. $\mathsf{E}[\pi_D] - \mathsf{E}[\pi_D^\dagger] = 0$ is because the two contracts yield the same average profit. ∎