



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2010

Monopolistic screening under learning by doing

Gärtner, D

Abstract: This article investigates the design of incentives in a dynamic adverse selection framework where agents' production technologies display learning effects and agents' learning rates are private knowledge. In a simple two-period model with full commitment available to the principal, we show that whether learning effects are over- or underexploited crucially depends on whether more efficient agents also learn faster (so costs diverge through learning effects) or whether it is the less efficient agents who learn faster (so costs converge). We further show that an overexploitation of learning effects can occur also if the full-commitment assumption is relaxed.

DOI: <https://doi.org/10.1111/j.1756-2171.2010.00112.x>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-41744>

Journal Article

Accepted Version

Originally published at:

Gärtner, D (2010). Monopolistic screening under learning by doing. *RAND Journal of Economics*, 41(3):574-597.

DOI: <https://doi.org/10.1111/j.1756-2171.2010.00112.x>



University of Zurich

Socioeconomic Institute
Sozialökonomisches Institut

Working Paper No. 0718

Monopolistic Screening under Learning By Doing

Dennis Gärtner

December 2007

Socioeconomic Institute
University of Zurich

Working Paper No. 0718

Monopolistic Screening under Learning By Doing

December 2007

Author's address: Dennis Gärtner
E-mail: dennis.gaertner@soi.uzh.ch

Publisher Sozialökonomisches Institut
Bibliothek (Working Paper)
Rämistrasse 71
CH-8006 Zürich
Phone: +41-44-634 21 37
Fax: +41-44-634 49 82
URL: www.soi.uzh.ch
E-mail: soilib@soi.uzh.ch

Monopolistic Screening under Learning By Doing*

Dennis Gärtner[†]
University of Zürich

December 2007

Abstract

This paper investigates the design of incentives in a dynamic adverse selection framework when agents' production technologies display learning effects and agents' rate of learning is private knowledge. In a simple two-period model with full commitment available to the principal, we show that whether learning effects are over- or under-exploited crucially depends on whether learning effects increase or decrease the principal's uncertainty about agents' costs of production. Hence, what drives the over- or under-exploitation of learning effects is whether more efficient agents also learn faster (so costs diverge through learning effects) or whether it is the less efficient agents who learn faster (so costs converge). Furthermore, we show that if divergence in costs through learning effects is strong enough, learning effects will not be exploited at all, in a sense to be made precise.

JEL Classification: D82, L14, L43, L51, O31

Keywords: Asymmetric Information, Learning by Doing, Regulation

*This work heavily draws on Chapter 1 of my dissertation (Gärtner, 2007). It has benefited from comments by Bernard Salanié and seminar participants at the University of Zürich. Moreover, I am indebted to Armin Schmutzler and Jean Tirole for invaluable discussions based on earlier versions. Financial support for this project from the Swiss National Science Foundation under Grant No. PBZH1-108484 is gratefully acknowledged.

[†]University of Zürich, Socioeconomic Institute, Blümlisalpstr. 10, CH-8006 Zurich, Switzerland. email: dennis.gaertner@soi.uzh.ch.

1 Introduction

That private knowledge on technology is a key issue in regulatory practice has been a central theme in the literature for years (see Baron and Myerson, 1982 and the influential work by Laffont and Tirole, 1993). While we already have a fairly good understanding of optimal regulation in stationary settings, relatively little is still known about *dynamic* aspects of regulation and the interplay between regulation and innovation incentives when technology is endogenous. This is particularly unfortunate because innovation plays a prominent role in many traditional fields of government regulation such as telecommunications or electric utilities. In such industries, adequate regulatory practice must invariable take account of its dynamic impact on innovation.

This paper contributes to closing this gap by considering a well-documented kind of innovation: innovation through ‘learning by doing.’¹ Under learning by doing, the regulator’s task is to induce a level of production which takes careful account of future costs-savings through learning effects. In the tradition of the recent literature on regulation, this paper explores the challenge posed to the regulator when firms enjoy superior information also on this aspect of their technology.

We investigate this issue in a general principal-agent model where, in each period, a principal (the regulatory authority) procures a good from an agent (the regulated firm) in exchange for a monetary transfer. To capture learning effects, production costs are assumed to depend on the previous period’s level of production, where the extent of this learning effect is known only to the agent.

Our principal finding is that whether private information causes learning-effects to be under- or overexploited relative to the full-information benchmark crucially depends on how agents’ learning potential and their absolute level of efficiency are related: If learning permits inherently more efficient agents to expand their lead over less efficient agents, then learning effects will be underexploited. If, however, learning allows inherently less efficient agents to catch up, then learning effects will be overexploited.

The basic intuition for this result is simple: By the familiar rent-efficiency tradeoff, distortions in output are driven by the principal’s incentive to limit the rent payable to more efficient agents, which in turn corresponds to their cost advantage over less efficient agents. If learning

¹See Lewis and Yildirim (2002b) for an extensive list of industries where learning effects have been documented.

magnifies efficiency differences, the usual downward distortion results. However, if learning leads less efficient agents to catch up, the cost advantage enjoyed by more efficient agents decreases with first-period output, in which case upward distortions result. Moreover, we show that downward distortions in the former case may in fact be so strong as to induce a level of production which is inefficiently low even from a *static* point of view (i.e., ignoring dynamic learning effects).

To the best of our knowledge, this insight—particularly the possibility of learning effects being overexploited—is new to the regulation and procurement literature. In the most immediately related paper, Lewis and Yildirim (2002a) (see also Lewis and Yildirim, 2002b) also investigate the regulation of a privately informed monopolist who learns by doing. Two of their key findings are the following: First, learning effects will always be underexploited relative to a dynamically efficient benchmark. Second, however, learning effects will never be left unexploited altogether in the sense that output will nonetheless always exceed its *statically* efficient level.

The difference with our findings can be explained by thinking of learning-effects in terms of a cost and a benefit side—costs being higher output today, and benefits being lower production costs tomorrow. Lewis and Yildirim’s model focusses on private information only on the cost side: The regulator does not know the costs of raising output today, but he knows the impact this will have on tomorrow’s costs.² In contrast, our model introduces asymmetric information on the *benefit* side by supposing that the firm is privately informed on its learning rate, that is, on the extent to which higher output today translates into lower costs tomorrow. Our analysis shows that the aforementioned two key results of the previous literature are both sensitive to this aspect of private information.

On a more basic level, this paper’s contribution may also be understood as sharpening basic economic intuition concerning the connection between asymmetric information and the volume of trade: It has become a virtual commonplace to associate private information with inefficiently low trade. Our results bring back to mind that this intuition depends crucially on the presumption that increased trade exacerbates the value of private information and thereby informational rents. While this structure arises naturally in many models, we argue that learning by doing provides a case in point where it is just as natural for the reverse to be

²The same comment applies to learning effects as modeled in Gaudet et al. (1996).

true.

The rest of this paper is organized as follows. Section 2 sets up a basic two-period model of learning by doing and describes the full-information benchmark. Section 3 presents the optimal contract under asymmetric information and full commitment. We discuss both its efficiency properties, whether it even makes use of learning effects in the first place, and briefly discuss a number of extensions and generalizations. Section 4 in turn investigates contracts under spot commitment. Owing to the analytical complexity introduced by limited commitment, we essentially restrict ourselves to showing that the main insights of Section 3—particularly the possibility of inefficiently *high* trade—generalize to the case of spot commitment. Finally, Section 5 concludes with a brief discussion of the results and further possible applications of the general principal-agent structure considered in this paper.

2 A Simple Model of Learning By Doing

This section presents the basic model and, as a benchmark for our later analysis, characterizes the efficient allocation.

2.1 Setting up the Model

We consider a simple model in which a principal procures a good from an agent over two periods $t \in \{1, 2\}$. Let $q_t \in Q_t$ denote the amount of the good procured in period t and let $z_t \in \mathbb{R}$ denote the monetary transfer from the principal to the agent in that period. Unless stated differently, we let $Q_t \equiv \mathbb{R}_{\geq 0}$. In each period t , let the principal's utility be given by $v_t = S(q_t) - z_t$, where $S' > 0$ and $S'' < 0$. The principal discounts with $\delta \in (0, 1)$, overall utility from transactions over the two periods being $V = v_1 + \delta v_2$. The agent produces the good at a constant marginal cost $c_t > 0$ within each period, yielding a utility of $u_t = z_t - c_t q_t$ in each period t . The agent discounts with the same factor δ , leading to an overall utility of $U = u_1 + \delta u_2$. Both principal and agent are assumed risk neutral.³

The remaining assumptions detail the agent's production technology and its dependence on private information (we briefly discuss their relaxation in the context of contracts under full commitment in Section 3.5).

³Letting principal and agent share the same risk attitude and time preference focusses our analysis by avoiding motives to trade risk or intertemporal utility.

To focus our analysis, we assume the agent’s first-period marginal costs c_1 to be observable.⁴ However, we let the agent possess private information concerning the structure of second-period costs. Private information is represented by the scalar θ (the agent’s “type”), which is drawn from a commonly known distribution over Θ , and which the agent privately observes prior to first-period production (and prior to contracting). For simplicity, most of our analysis will assume two types $\Theta = \{\bar{\theta}, \underline{\theta}\}$ with $\bar{\theta} > \underline{\theta}$ and $\text{Prob}(\theta = \bar{\theta}) = \nu$.

Second-period marginal costs c_2 are a function of θ and first-period output q_1 . We model the presence of learning effects in production by letting $\partial c_2 / \partial q_1 < 0$: The higher first-period production, the lower the marginal costs of production in period 2 (for any given type θ).⁵

Next, we assume that c_2 is strictly decreasing in θ for all q_1 . Thus, an agent with a higher θ is more efficient in that he produces any output schedule $\mathbf{q} = (q_1, q_2)$ at a lower cost. Note that this assumption represents more than a mere normalization of the type space Θ : Since it is assumed to hold for all q_1 , it will provide a key sorting condition in our derivation of the optimal contract under asymmetric information.

We call $|\partial c_2 / \partial q_1|$ the agent’s *learning rate*, and say that an agent *learns faster* if he has a higher learning rate. Note that an agent may learn faster even though he is less efficient (i.e., has a lower θ). Indeed, key aspects of our analysis will crucially depend on whether learning rates increase or decrease in θ . To facilitate this, we assume that learning rates either increase or decrease in θ for all q_1 . Figure 1 illustrates the relevant constellations.

The following example not only provides an illustration of the setup, but—due to its analytical tractability—will prove useful for numerical examples given further below:

Example 1. Let second-period costs be given by $c_2(q_1, \theta) = c(\theta) - \gamma(\theta)q_1$ with c strictly increasing in θ , and with $\gamma > 0$, and let the principal’s objective function be given by $S(q_t) = aq_t - bq_t^2$, where $a, b > 0$. Then

⁴As we discuss in Section 3.5 in more detail, introducing private information on first-period costs c_1 introduces a second, competing motive for the principal to distort quantities in order to reduce *first-period* informational rents. This second motive in isolation already being well understood in the standard framework *without* learning, letting c_1 be observable serves to make this papers’ contribution more transparent.

⁵Note that we assume marginal costs to be constant *within* each period but change discontinuously from one period to the next. This assumption serves to isolate learning effects from simple scale economies. Indeed, what distinguishes the two is that learning by doing depends on both previous production volumes *and* on time.

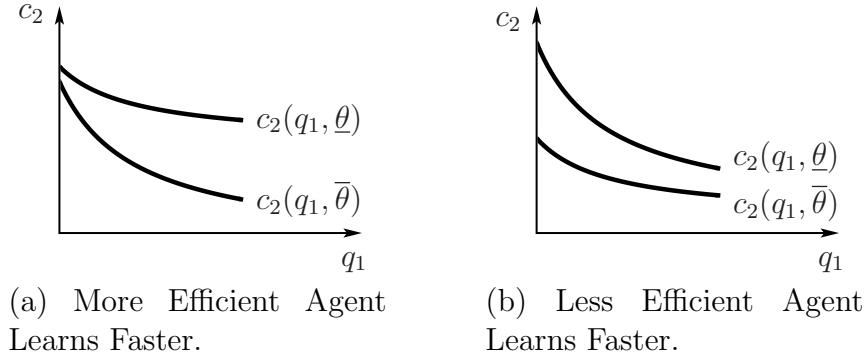


Figure 1: Types of Learning Effects.

more efficient agents learn faster if γ is increasing in θ , whereas less efficient agents learn faster if γ is decreasing in θ .⁶

In what follows, we will consider various settings for negotiating the exchanged quantities and transfers. Common to these settings, however, is the usual assumption of full bargaining power resting with the principal: The principal offers a contract (or a choice of contracts) to the agent, which the agent can decide to accept or reject. If the agent rejects, he obtains a (type-independent) reservation payoff of zero and negotiations end.

2.2 The Efficient Full-Information Benchmark

As a point of comparison for our later analysis, we first consider the efficient full-information benchmark. Assume for a moment that the agent's type θ is known to the principal. For any output schedule $\mathbf{q} = (q_1, q_2)$, the joint surplus of trade is then given by

$$W(\mathbf{q}; \theta) \equiv S(q_1) - c_1 q_1 + \delta[S(q_2) - c_2(q_1, \theta)q_2]. \quad (1)$$

Given his bargaining power and any known type θ , the informed principal will offer a contract $(\mathbf{q}^*(\theta), \mathbf{z}^*(\theta))$ which is efficient (i.e., first best), specifying production levels $\mathbf{q}^*(\theta) = (q_1^*(\theta), q_2^*(\theta))$ which maximize joint surplus $W(\mathbf{q}; \theta)$, and payments $\mathbf{z}^*(\theta) = (z_1^*(\theta), z_2^*(\theta))$ which leave the

⁶To make this example entirely compatible with our assumptions, the range of permissible q_1 and q_2 , (i.e., Q_1 and Q_2) must be bounded from above so as to ensure $S'(q_i) > 0$ and $c_2(q_1, \theta) > c_2(q_1, \bar{\theta}) > 0$.

agent his reservation utility.⁷

For later comparisons, we define *conditional* first-best output levels as follows. For any $(q_2; \theta)$, let $\hat{q}_1^*(q_2; \theta) \equiv \arg \max_{q_1} W(q_1, q_2; \theta)$ and, similarly, for any (q_1, θ) let $\hat{q}_2^*(q_1; \theta) \equiv \arg \max_{q_2} W(q_1, q_2; \theta)$.

Lemma 2.1. *The first-best output schedule $\mathbf{q}^*(\theta)$ and the contingent first-best output levels $\hat{q}_1^*(q_2; \theta)$ and $\hat{q}_2^*(q_1; \theta)$ have the following properties:*

- (a) \hat{q}_2^* is increasing in both q_1 and θ ;
- (b) \hat{q}_1^* is increasing in q_2 , and increasing (decreasing) in θ if more (less) efficient agents learn faster;
- (c) \mathbf{q}^* is increasing in θ if more efficient agents learn faster.

The proof of this and all later results is presented in the Appendix. As a general matter, we establish comparative static results such as Lemma 2.1 using supermodular analysis (cf. Milgrom and Roberts, 1990; Topkis, 1998). This approach exploits basic complementarity relations among arguments of the objective function and avoids imposing any unnecessary concavity assumptions on objectives.⁸ The latter is particularly valuable in our setting due to the concavity which learning effects naturally introduce into the cost function.⁹

These technical issues aside, the intuition for the above results is conceivably simple: Higher first-period output lowers the cost of additional second-period output, thereby raising incentives to expand the latter. Conversely, higher second-period output raises incentives to lower that output's costs through learning effects by expanding first-period output. Thus, each period's conditionally efficient output rises in the other period's output level. Moreover, a higher θ makes additional second-period

⁷Note that the first-best transfer schedule \mathbf{z}^* will never be unique. Indeed, if $(\mathbf{q}^*, \mathbf{z}^*)$ is a first-best contract, then any contract $(\mathbf{q}^*, \bar{\mathbf{z}})$ with the same discounted value of transfers (i.e., with $\bar{z}_1 + \delta \bar{z}_2 = z_1^* + \delta z_2^*$) will also be first-best.

⁸Nonetheless, readers unfamiliar with this technique will quickly verify the results under additional concavity assumptions by means of the first-order approach.

⁹Supermodular analysis extends also to situations where optimizers are not unique. In this case, comparative static results are interpretable in terms of ordering relations among *sets*. Although all our results permit such an interpretation, to avoid tedious notation, we will be somewhat loose in distinguishing between the *set* of optimizers and its individual elements.

Moreover, while comparative statics derived by this technique are very general in terms of covering also the possibility of corner solutions, this comes at the cost of all results applying only in a *weak* sense (i.e., “increasing” in Lemma 2.1 is to be read as “nondecreasing”). Since essentially all complementarity relations underlying our results are in fact *strict*, strict versions of our comparative static predictions are easily established for interior maximizers under mild additional conditions (see Edlin and Shannon, 1998, for technical details).

output less costly, which is why \hat{q}_2^* is increasing in θ . Incentives to raise first-period output in turn rise in the agent's learning rate, which is why \hat{q}_1^* 's response to a change in θ depends on how the learning rate changes in θ . Finally, if more efficient agents also learn faster, these effects complement each other, making \mathbf{q}^* rise in θ . Note that no such robust comparative result is available if the *less* efficient agent learns faster: a rise in θ then provides direct incentives to raise q_2 and lower q_1 , which are counteracted however by the complementarity between q_1 and q_2 , leaving the overall result ambiguous.

3 Contracts under Full Commitment

In contractual problems with investment characteristics, the outcome is generally sensitive to the level of intertemporal commitment available to the principal (see for instance Fudenberg et al., 1990). In this section, we investigate our problem of learning by doing under the most extreme form of commitment: We assume that at the start of period one (but after the agent has learned his type), the principal can offer a contract settling all future exchange which cannot be reneged on.

3.1 Characterizing the Optimal Contract

The full-commitment setting has the convenient property that, by the revelation principle and the stationarity of private information, we may equivalently restrict our attention to truth-revealing mechanisms of the type $\{\mathbf{q}(\tilde{\theta}), \mathbf{z}(\tilde{\theta})\}_{\tilde{\theta} \in \Theta}$ which specify contracts (i.e., exchanged quantities and transfers) for each type, and where these contracts are designed so as to make it optimal for the agent to truthfully reveal his type. For any such contract, we let

$$U(\theta) \equiv z_1(\theta) - c_1 q_1(\theta) + \delta \{z_2(\theta) - c_2 [q_1(\theta), \theta] q_2(\theta)\} \quad (2)$$

denote the θ -type's *equilibrium rent*.

To relax notation in this section, for any function of $\theta \in \{\bar{\theta}, \underline{\theta}\}$, we let an upper (lower) bar indicate that the function is evaluated at $\bar{\theta}$ ($\underline{\theta}$) and drop the argument θ . Thus, for instance, $\bar{U} \equiv U(\bar{\theta})$ and $\underline{U} \equiv U(\underline{\theta})$. Finally, we let

$$\Phi(\mathbf{q}) \equiv \delta q_2 [c_2(q_1) - \bar{c}_2(q_1)] \quad (3)$$

denote the *cost advantage* enjoyed by the $\bar{\theta}$ -agent over the $\underline{\theta}$ -agent for any output schedule $\mathbf{q} = (q_1, q_2)$. Intuitively, this cost advantage measures the *value* of private information enjoyed by the more efficient $\bar{\theta}$ -type.

With this notation in place, the optimal contract under full commitment can be characterized as follows:

Proposition 3.1. *The menu of contracts offered by the uninformed principal under full commitment is such that production schedules $\bar{\mathbf{q}}^{\text{SB}}$ and $\underline{\mathbf{q}}^{\text{SB}}$ solve*

$$\bar{\mathbf{q}}^{\text{SB}} \in \arg \max_{\mathbf{q}} \bar{W}(\mathbf{q}) \quad \text{and} \quad \underline{\mathbf{q}}^{\text{SB}} \in \arg \max_{\mathbf{q}} \left\{ \underline{W}(\mathbf{q}) - \frac{\nu}{1-\nu} \Phi(\mathbf{q}) \right\}. \quad (4)$$

Transfers $\bar{\mathbf{z}}^{\text{SB}}$ and $\underline{\mathbf{z}}^{\text{SB}}$ are chosen such that types' equilibrium rents are $\bar{U} = \Phi(\underline{\mathbf{q}}^{\text{SB}})$ and $\underline{U} = 0$.

This result is easily understood by recognizing that, despite the presence of learning effects, sorting (i.e., the relevance of incentive constraints) is entirely driven by the assumption that c_2 is decreasing in θ . Since this unambiguously makes the $\bar{\theta}$ -agent more efficient, the relevant incentive problem is keeping *him* from falsely reporting $\underline{\theta}$ by granting him a rent equal to his cost advantage $\Phi(\underline{\mathbf{q}})$ for the corresponding production schedule—and leaving the $\underline{\theta}$ -agent a rent of zero. Deducting these rents from joint surplus, the principal is left with *reduced form profits* (i.e., incorporating the optimal choice of transfers $\bar{\mathbf{z}}$ and $\underline{\mathbf{z}}$) of

$$\Pi(\bar{\mathbf{q}}, \underline{\mathbf{q}}) \equiv \nu [\bar{W}(\bar{\mathbf{q}}) - \Phi(\underline{\mathbf{q}})] + (1 - \nu) \underline{W}(\underline{\mathbf{q}}), \quad (5)$$

maximization of which corresponds to condition (4).

The objective function (5) embodies the usual rent-efficiency tradeoff faced by an uninformed principal: His menu of contracts trades off expected joint surplus $\nu \bar{W}(\bar{\mathbf{q}}) + (1 - \nu) \underline{W}(\underline{\mathbf{q}})$ against the expected rent payments $\nu \Phi(\underline{\mathbf{q}})$ required to induce truthful reporting by the $\bar{\theta}$ -type. This tradeoff leads to inefficiencies whose precise nature we analyze next.

3.2 Partial Distortionary Incentives

A trivial implication of Proposition 3.1 is the usual ‘no distortion at the top’-result: In spite of information being private, the efficient $\bar{\theta}$ -type still produces first-best quantities in both periods, so $\bar{\mathbf{q}}^{\text{SB}} = \bar{\mathbf{q}}^*$. This leaves us with an investigation of the nature of distortions ‘at the bottom’, that is, of inefficiencies inherent in the contract offered to the inefficient $\underline{\theta}$ -type.

The rent-efficiency tradeoff responsible for this distortion involves the simultaneous use of *two* instruments, q_1^{SB} and q_2^{SB} . To clarify their individual roles and make the principal's motives more transparent, we first analyze what we shall call ‘partial’ distortionary incentives. In analogy

to the conditional first-best production schedules $\hat{q}_1^*(q_2)$ and $\hat{q}_2^*(q_1)$, we let $\hat{q}_1^{\text{SB}}(q_2)$ and $\hat{q}_2^{\text{SB}}(q_1)$ denote the levels of q_1 and q_2 , respectively, which maximize $\underline{W}(\mathbf{q}) - \frac{\nu}{1-\nu}\Phi(\mathbf{q})$ conditional on the other period's output level. With these definitions in place, the following partial distortionary motives can be identified:

Proposition 3.2. *Under full commitment, the uninformed principal faces the following (partial) distortionary incentives in designing the $\underline{\theta}$ -type's production schedule:*

- (a) *Conditional on any first-period output q_1 , the uninformed principal will distort second-period output downward, so $\hat{q}_2^{\text{SB}}(q_1) \leq \hat{q}_2^*(q_1)$ for all $q_1 \in Q_1$.*
- (b) *Conditional on any second-period output q_2 , the uninformed principal will distort first-period output*
 - (i) *downward if the more efficient agent learns faster, so $\hat{q}_1^{\text{SB}}(q_2) \leq \hat{q}_1^*(q_2)$ for all $q_2 \in Q_2$, and*
 - (ii) *upward if the less efficient agent learns faster, so $\hat{q}_1^{\text{SB}}(q_2) \geq \hat{q}_1^*(q_2)$ for all $q_2 \in Q_2$.*

Part (a) concerning second-period distortionary incentives is not surprising: Given a first-period output level, second-period marginal costs c_2 are a datum, and hence the principal's optimization problem is identical to the standard one-period model of procurement with privately known and constant marginal costs, for which downward distortion (i.e., inefficiently low trade) is a well-known result.

More interestingly, Part (b) identifies the distortionary incentives involved in choosing first-period output—and thereby the extent to which learning effects are under- or overexploited for any given second-period output. That the direction of these distortions crucially depends on which agent learns faster is quickly understood by referring back to our illustration of the two cases in Figure 1: If the more efficient agent learns faster, *decreasing* q_1 leads types' second-period costs c_2 to converge, thereby reducing the $\bar{\theta}$ -type's cost advantage (for fixed q_2) and hence the rent payable to him. Conversely, if the *less* efficient agent learns faster, the same effect is achieved by an *increase* in q_1 .

3.3 Overall Distortions in Trade

The partial distortions analyzed above provide a direct measure of the over- or underexploitation of learning effects by asking whether the uninformed principal's contracts can be pareto-improved upon by expanding




Figure 2: The Direction of Overall Distortions.

or reducing first-period output. Since the principal ultimately distorts both periods' output schedules simultaneously, however, attaining the full welfare optimum will also require simultaneous adjustments in both quantities.

As the next result shows, if the more efficient type learns faster, attaining the full welfare optimum requires *expanding* both periods' output: Partial distortions in this case are representative of overall distortions, so overall trade is inefficiently low.

Proposition 3.3. *Under full commitment, if the more efficient type learns faster, then private information causes an overall downward distortion in both first- and second-period output for the inefficient type, so $\mathbf{q}^{SB} \leq \mathbf{q}^*$.*

Figure 2(a) illustrates the results under the additional assumption that both \underline{W} and $\underline{W} - \nu\Phi$ are strictly concave. By the implied uniqueness of the maximizers and by Lemma 2.1, the conditional first-best outputs $\hat{q}_1^*(q_2)$ and $\hat{q}_2^*(q_1)$ are increasing functions. Moreover, concavity of \underline{W} implies that the \hat{q}_1^* -curve crosses the \hat{q}_2^* -curve from below at $\mathbf{q}^* = (q_1^*, q_2^*)$ in (q_1, q_2) -space. Now by Proposition 3.2(a), the \hat{q}_2^{SB} -curve will lie south of the \hat{q}_2^* -curve, and by Proposition 3.2(bi), the \hat{q}_1^{SB} -curve will lie west of the \hat{q}_1^* -curve. Hence, the equilibrium under private information—determined by the intersection of the \hat{q}_2^{SB} - and the \hat{q}_1^{SB} -curve—must lie in the shaded area in Figure 2(a). As illustrated, this area must lie in the southwest quadrant of \mathbf{q}^* .

Why an analogous argument fails when the *less* efficient agent learns faster is illustrated in Figure 2(b). Again, the \hat{q}_2^{SB} -curve must lie south of the \hat{q}_2^* -curve. However, Proposition 3.2(bii) in this case tells us that there will be an upward distortion in first-period output given any second-period output, so that the new equilibrium must lie east of the \hat{q}_1^* -curve. Hence, only equilibria with $\underline{q}_1^{\text{SB}} < \underline{q}_1^*$ and $\underline{q}_2^{\text{SB}} > \underline{q}_2^*$ can be excluded.¹⁰

Particularly, if the less efficient agent learns faster, it is possible for overall distortions to be upward in both periods, so that $\underline{q}_1^{\text{SB}} > \underline{q}_1^*$ and $\underline{q}_2^{\text{SB}} > \underline{q}_2^*$. This is illustrated by the following extension to Example 1:

Example 2. Assume the value of output to the principal is given by $S(q_t) = 100q_t - 80q_t^2$, cost structures are $c_1 = 75$, $\bar{c}_2(q_1) = 30 - 60q_1$, and $c_2(q_1) = 50 - 95q_1$, types are equally likely, and the common discount factor is $\delta = 0.7$. First-best production then entails $\bar{\mathbf{q}}^* = (0.30, 0.55)$ and $\underline{\mathbf{q}}^* = (0.38, 0.54)$, whereas contracts under asymmetric information and full commitment will entail $\underline{\mathbf{q}}^{\text{SB}} = (0.49, 0.58)$ for the $\underline{\theta}$ -type.

Alternative parameterizations of Example 1 will produce the other two possible directions in overall distortions.

3.4 Are Learning Effects Exploited At All?

Having gauged the outcome under incomplete information against efficient benchmarks, this section investigates whether downward distortionary incentives can be so severe as to eliminate the exploitation of learning effects altogether. For comparison, we consider the outcome which results if either first-period output has no impact on second-period marginal costs, so $\frac{\partial}{\partial q_1}c_2 \equiv 0$, or both principal and agent behave myopically, so $\delta = 0$. The resulting choice of q_1 in either case will maximize first-period surplus $S(q_1) - c_1q_1$ alone. Motivated by this, we introduce the following terminology:

Definition 3.4. Let $q_1^\circ \equiv \arg \max_q \{S(q) - c_1q\}$. A first-period output level q_1 *exploits* learning effects if $q_1 \geq q_1^\circ$; it *neglects* learning effects if $q_1 \leq q_1^\circ$.

¹⁰A simple generalization of this last argument (generalized beyond the graphical analysis' additional assumptions) runs as follows. Let $\underline{\mathbf{q}}^*$ denote the first-best output schedule (if the first-best output schedule is not unique, let $\underline{\mathbf{q}}^*$ denote any first-best schedule such that there exists no other first-best output schedule involving lower first- and higher second-period output). Then for any $\mathbf{q} = (q_1, q_2)$ with $q_1 < \underline{q}_1^*$ and $q_2 > \underline{q}_2^*$, we have $\underline{W}(\mathbf{q}) < \underline{W}(\underline{\mathbf{q}}^*)$. Moreover, $\Phi(\mathbf{q}) \leq \Phi(\underline{\mathbf{q}}^*)$ for any such \mathbf{q} because Φ is decreasing in q_1 and increasing in q_2 . But then $(1 - \nu)\underline{W}(\mathbf{q}) - \nu\Phi(\mathbf{q}) < (1 - \nu)\underline{W}(\underline{\mathbf{q}}^*) - \nu\Phi(\underline{\mathbf{q}}^*)$, so that no such \mathbf{q} can maximize the uninformed principal's objective in (4).

Equivalently (recall that $S'' < 0$), learning effects are exploited if the marginal benefit of first-period output S' (weakly) falls short of marginal costs c_1 , and neglected if the reverse holds.

Obviously, first-best quantities always exploit learning effects. This need not be true for q_1^{SB} , the first-period quantity procured from the $\underline{\theta}$ -type. The following result gives sufficient conditions for either case:

Proposition 3.5. *The contract offered to the $\underline{\theta}$ -agent exploits learning effects if*

$$\left| \frac{\partial}{\partial q_1} c_2(q_1) \right| \geq \nu \cdot \left| \frac{\partial}{\partial q_1} \bar{c}_2(q_1) \right| \quad (6)$$

for all $q_1 \in Q_1$; it neglects learning effects if (6) is reversed for all $q_1 \in Q_1$.

Thus, learning effects are exploited if either the efficient agent does not learn too much faster than the inefficient agent, or if efficient types are scarce enough. Intuitively, both ensure that the principal's rent-efficiency tradeoff is sufficiently in favor of efficiency—the former by reducing the efficient type's cost advantage and thereby his rent, the latter by making it less likely that such a rent will have to be paid in the first place.

In relation to our previous results in Section 3.3, Proposition 3.5 shows that even though downward distortions in q_1 may ensue if the *inefficient* agent learns faster, they will never be so strong as to eliminate the exploitation of learning effects altogether. In relation to the previous literature, Proposition 3.5 points out that this may however be the case if the *efficient* agent learns faster, depending on the distribution of types and how strongly learning rates differ. Particularly, the possibility of learning effects being neglected is absent in Lewis and Yildirim's (2002a) model.

3.5 Extensions and Limitations

Our analysis thus far has relied on several simplifying assumptions. Before the next section proceeds to relax what might seem the most serious and restrictive one—the assumption of full commitment—, we briefly discuss other possible extensions and the challenge they pose to our findings.

'Varying Learning Advantages': Assuming that one of the agents unambiguously learns faster has simplified our identification of first-period distortions, but has been immaterial to the derivation of the optimal contract itself in Proposition 3.1. Without this assumption, it is still true

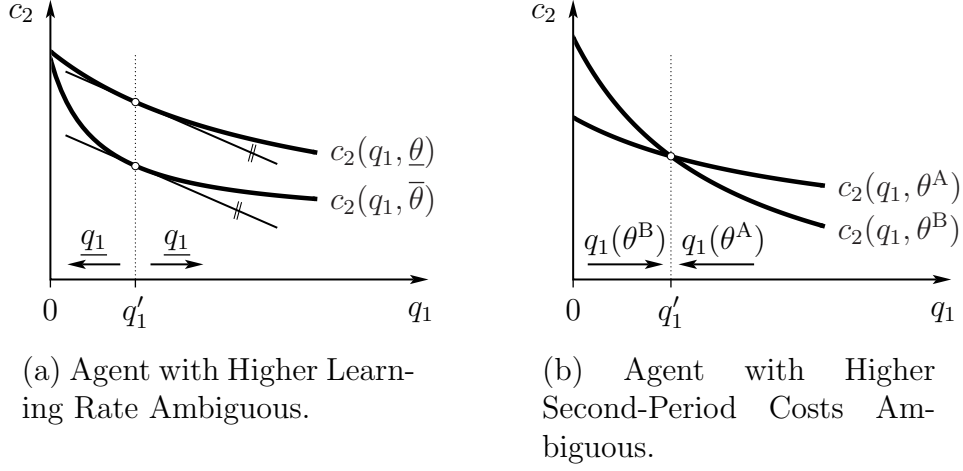


Figure 3: Examples of more General Second-Period Cost Functions.

that only the $\underline{\theta}$ -type's output is distorted, q_2 is distorted downward given q_1 , but the direction of the distortion in q_1 given q_2 will be ambiguous.

More specifically, consider the cost functions shown in Figure 3(a), where the more efficient $\bar{\theta}$ -agent learns faster for $q_1 < q'_1$ and slower for $q_1 > q'_1$. Our characterization of partial distortions in q_1 is then still valid to the extent that they will be downward if $q_1 < q'_1$, and upward if $q_1 > q'_1$. However, which regime is relevant is ambiguous and depends, *inter alia*, on the value of output relative to its costs.

Intersecting Cost Curves: In contrast, the assumption that the $\underline{\theta}$ -type has lower second-period costs c_2 for any q_1 has indeed been vital to our derivation of the optimal contract by determining which of the incentive constraints must bind. This assumption precludes, however, the possibility of one agent 'overtaking' the other due to learning effects.

Such a case is illustrated in Figure 3(b). Here, a straightforward extension to our previous analysis shows that (partial) distortions in $q_1(\theta^B)$ can only occur for $q_1(\theta^B) < q'_1$ and will be upward, whereas distortions in $q_1(\theta^A)$ can only occur for $q_1(\theta^A) > q'_1$ and will be downward. However, which of these distortions occurs in the optimum is again ambiguous, and both may in fact occur simultaneously.

The previous two examples generalize the main theme of our above analysis in the following straightforward way: Distortions, if they occur, aim to reduce the (locally) less efficient agent's cost-disadvantage. Whether this requires an increase or a decrease in first-period output depends on relative learning rates.

Private Information on First-Period Costs: We have assumed type-independent first-period costs c_1 to focus on the role of asymmetric information on the *returns* to learning by doing. If c_1 also depends on θ , we may generalize the $\bar{\theta}$ -type's comparative cost advantage to $\Phi(q_1, q_2) \equiv [c_1(\underline{\theta}) - c_1(\bar{\theta})]q_1 + \delta[c_2(q_1, \underline{\theta}) - c_2(q_1, \bar{\theta})]q_2$. Under the assumption that $c_1(\bar{\theta}) \leq c_1(\underline{\theta})$, the $\bar{\theta}$ -type is still unambiguously more efficient in *both* periods, and the derivation of the optimal contract goes through unchanged.¹¹ Particularly, the $\bar{\theta}$ -type's output schedule remains undistorted. There is now, however, an additional motive to distort q_1 downward in order to reduce the $\bar{\theta}$ -type's rent for first-period cost advantages, which—depending on which agent learns faster—will either reinforce or counteract the distortionary incentives identified above.

Finally, we note that if $c_2(q_1, \bar{\theta}) = c_2(q_1, \underline{\theta})$ for all q_1 , in addition, so that private information concerns the cost-side *exclusively*, then we are essentially in the setting considered by Lewis and Yildirim (2002a), and learning-effects will always be *underexploited* for the $\underline{\theta}$ -type.

Alternative Valuations of Output: We have assumed that the principal values output at $S(q_1) + \delta S(q_2)$. However, our key results concerning (partial) distortionary incentives in Proposition 3.2 are robust to more general valuation functions $\tilde{S}(\mathbf{q})$.¹² A particularly interesting extension involves letting \tilde{S} depend on q_2 alone: In this case, our model represents a pure investment problem with privately known returns, but where the *level* of investment is contractible. Reinterpreting our above results, underinvestment then ensues (in both a partial and an overall sense) if investment returns (in terms of cost savings) become *more* sensitive to private information with higher investment levels, whereas overinvestment (at least in a partial sense) occurs in the reverse case.

¹¹If $c_1(\bar{\theta}) > c_1(\underline{\theta})$, we again face a situation in which it is unclear which incentive constraint binds at the optimum.

¹²Except for our analysis in Section 3.4, all above results in fact easily generalize to cases in which the valuation function $\tilde{S}(\mathbf{q})$ displays complementarities in q_1 and q_2 in the sense that $\partial^2 \tilde{S} / \partial q_1 \partial q_2 \geq 0$. To understand this, note that the analysis has employed the particular form of valuation function only to the extent that it implies additive separability of valuations in q_1 and q_2 , which in turn implies that, due to cost-side effects alone, the surplus function W is complementary in q_1 and q_2 . Comparative static results reliant on this complementarity (i.e., Lemma 2.1 and Proposition 3.3) are thus robust to valuation functions which preserve this complementarity.

4 Contracts under Spot Commitment

The analysis thus far has assumed that, before first-period production takes place, the principal can commit to a contract spanning *both* periods of production. There are several reasons why it is interesting to relax this assumption. First, it is quite conceivable that the principal indeed cannot find a way to commit to not renegeing after the first period, be it due to restrictions imposed by the legal system or simply because the regulatory authority's commitment is limited to the current administration's life-span.¹³ Second, limited commitment is generally understood to be a deterrent to long-term investments (cf. Fudenberg et al., 1990). It should thus be interesting to see how our overinvestment result in particular stands up to limited commitment. Finally, limited commitment has been the focus in the immediately related literature (Lewis and Yildirim, 2002a, in particular), making the extension of our results to limited commitment desirable for reasons of comparison.

We therefore assume in this section that parties are limited to *spot* contracts: At the beginning of each period $t \in \{1, 2\}$, the principal can offer a contract specifying quantities q_t and transfers z_t only for this current period. The agent in turn can decline in each period, which yields him a reservation utility of zero and terminates the game.

In this dynamic setting with *stationary* private information, a central issue is the *rate* at which agents reveal this information over time. Particularly, by the well-known ratchet effect (see Freixas et al., 1985; Laffont and Tirole, 1987), constraints imposed by sequential rationality generally preclude equilibria in which agents fully reveal their type in the first-period, implying that there must be some degree of pooling in the first period.

This partial revelation of information can be formalized by means of an extension to the classical revelation principle due to Bester and Strausz (2001). By this extension, we may restrict ourselves to mechanisms where the principal makes the exchanged quantities dependent on type reports $\tilde{\theta}_t \in \Theta$ made in each period to date. In the second and final period, incentive constraints take the usual form, ensuring optimality of the agent reporting his true type. In contrast, the *partial* revelation of

¹³In the above two-period contract, renegeing will be mutually beneficial after first-period production due to the usual interim inefficiencies (see, for instance Laffont and Tirole, 1986). It should be borne in mind however that the principal has incentives to find a commitment device since the ability to commit will always make him better off.

information in the *first* period takes the form of each type $\theta \in \Theta$ reporting his true type with some probability $p_1(\theta) > 0$ which is strictly positive—but not necessarily equal to one.

4.1 Characterizing the Optimal Contract

Due to the above invocation of Bester and Strausz's revelation principle, an equilibrium of our two-period model under spot contracts therefore is described by quantities $q_1(\tilde{\theta}_1)$ and $q_2(\tilde{\theta}_1, \tilde{\theta}_2)$ and transfers $z_1(\tilde{\theta}_1)$ and $z_2(\tilde{\theta}_1, \tilde{\theta}_2)$ in each period for any history of reports $\tilde{\theta}_1, \tilde{\theta}_2 \in \Theta$, and probabilities $p_1(\theta) > 0$ with which each type $\theta \in \{\underline{\theta}, \bar{\theta}\}$ reports his true type in period one.

In a next step, in analogy to the full-commitment case, we may characterize rents optimally paid to each type in each period, thereby eliminating transfers from the principal's optimization program. Consider first the second period. At the start of the second period, previous gameplay has resulted in a report $\tilde{\theta}_1$ made by the agent. This report in turn determines (i) the agent's second-period marginal costs (via the first-period output invoked by the report), and (ii) an updated belief for the principal concerning the agent's true type, determined through Bayesian updating by comparing types' equilibrium reporting behavior $p_1(\cdot)$ with the actual report made. Given these type-dependent marginal costs and beliefs, the second-period subgame is identical to the standard one-period framework, implying that second-period rents will be $u_2(\tilde{\theta}_1, \underline{\theta}) = 0$ for the inefficient agent, and $u_2(\tilde{\theta}_1, \bar{\theta}) = \Phi[q_1(\tilde{\theta}_1), q_2(\tilde{\theta}_1, \underline{\theta})]$ for the efficient agent.

Concerning the first period, a specialty of our set-up is that first-period rents are type-independent since first-period costs are, so that we may simply denote the first-period rent by $u_1(\tilde{\theta}_1)$. Thus, incentive constraints become

$$\begin{aligned} u_1(\bar{\theta}) + \Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})] &\geq u_1(\underline{\theta}) + \Phi[q_1(\underline{\theta}), q_2(\underline{\theta}, \underline{\theta})], \\ u_1(\underline{\theta}) &\geq u_1(\bar{\theta}), \end{aligned}$$

whereas participation constraints read

$$\begin{aligned} u_1(\bar{\theta}) + \Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})] &\geq 0, \\ u_1(\underline{\theta}) &\geq 0. \end{aligned}$$

Reproducing the argument from the full-commitment case, the $\bar{\theta}$ -type's participation constraint can be neglected and the $\underline{\theta}$ -type's must bind.

Hence, letting $U(\theta) = u(\theta) + \delta u_2(\theta, \theta)$ denote the equilibrium rent of an agent of type θ , we have $U(\underline{\theta}) = 0$, and the remaining two incentive constraints collapse to

$$\Phi[q_1(\underline{\theta}), q_2(\underline{\theta}, \underline{\theta})] \leq U(\bar{\theta}) \leq \Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})]. \quad (7)$$

where the first inequality denotes the $\bar{\theta}$ -type's incentive constraint, and the second inequality denotes the $\underline{\theta}$ -type's incentive constraint.

But this is where parallels with the full-commitment case end and sequential rationality and the ratchet effect kick in. To see how, suppose that the $\underline{\theta}$ -type's incentive constraint (i.e., the second inequality in (7)) does *not* bind (as is the case under full commitment). This implies that $p_1(\underline{\theta}) = 1$, as the $\underline{\theta}$ -agent will strictly prefer truthful reporting in period 1. But this in turn implies that, having seen a first-period report of $\bar{\theta}$, the principal is *sure* of the agent's true type being $\bar{\theta}$. By sequential rationality, the principal will optimally respond to this by setting the $\bar{\theta}$ -agent's second-period rent $\Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})]$ to zero.¹⁴ But then, since $\Phi(\cdot, \cdot) \geq 0$, *both* inequalities in (7) must bind—a contradiction.

Hence, the *inefficient* type's incentive compatibility condition must bind, and the efficient type's equilibrium rent is thereby pinned down to $U(\bar{\theta}) = \Phi[q_1(\underline{\theta}), q_2(\underline{\theta}, \underline{\theta})]$. With these results in place, the principal's optimization program may be formulated as follows:

Proposition 4.1. *Under spot commitment, the principal optimally chooses $p_1(\cdot), p_2(\cdot) \in (0, 1]$, $q_1(\cdot)$ and $q_2(\cdot, \cdot)$ so as to maximize*

$$E_{(\bar{\theta}_1, \theta)} \{W[q_1(\bar{\theta}), q_2(\bar{\theta}_1, \theta); \theta]\} - \nu \cdot \Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})]$$

subject to

- (a) *the principal's second-period beliefs over types being rational given agents' first-period reporting behavior $p_1(\cdot)$,*
- (b) *second-period quantities $q_2(\cdot, \cdot)$ being sequentially optimal given first-period production and the updated beliefs.*
- (c) $\Phi[q_1(\underline{\theta}), q_2(\underline{\theta}, \underline{\theta})] \leq \Phi[q_1(\bar{\theta}), q_2(\bar{\theta}, \underline{\theta})]$, *with equality if $p_1(\bar{\theta}) < 1$,*¹⁵

Note that the clear characterization of the optimal spot contract in Proposition 4.1 stands in stark contrast to the usual ambiguity concerning which incentive constraint binds. Indeed, a key result in Laffont and

¹⁴He can always achieve this by setting $q_2(\bar{\theta}, \underline{\theta}) = 0$.

¹⁵Gärtner (2007) in fact shows that even if $p_1(\bar{\theta}) = 1$, it will never be optimal for the principal to leave this condition slack, implying that *both* incentive compatibility conditions must bind in the optimum.

Tirole’s (1987) two-type model of spot-contracting with *constant* (but privately known) marginal costs is that, depending on the specific parameterization, four distinct types of equilibria will result depending on which incentive constraint binds—an ambiguity which is reminiscent of the standard one-period analysis when the Spence-Mirrlees sorting condition is violated (cf. Guesnerie and Laffont, 1984). This ambiguity is circumvented in our setting by focussing on the *returns* to learning by doing: Assuming that first-period costs c_1 are type-independent implies that first-period rents u_1 are type-independent as well, so that first-period sorting (or ‘separation’) can be based only on *second* period rents u_2 . Thus, in contrast to Laffont and Tirole (1987), there is no conflict between first- and second-period sorting conditions.¹⁶

4.2 Numerical Results

As noted above, rather than derive results concerning the direction of distortions at the same level of generality as in Section 3 for the full-commitment case, we confine ourselves in this section to presenting results for two specific examples which illustrate that, also under spot commitment, both up- and downward distortions can occur.

4.2.1 Example 3: Inefficient Agent Learns Faster

We first consider an example for which the inefficient agent learns faster, which we know under full commitment causes the inefficient agent’s first-period output to be inefficiently high—at least given second-period output. Specifically, we use the same setting used in Example 2 in Section 3 to illustrate the possibility of overall upward distortions under full commitment:

Example 3. For the setting given in Example 2, equilibrium reporting strategies under spot commitment are given by $p_1(\bar{\theta}) = 1$ and $p_1(\underline{\theta}) = 0.74$ for equilibrium output menus given by $q_1(\bar{\theta}) = 0.34$, $q_1(\underline{\theta}) = 0.44$, $q_2(\bar{\theta}, \bar{\theta}) = 0.57$, $q_2(\bar{\theta}, \underline{\theta}) = 0.60$ ¹⁷, $q_2(\underline{\theta}, \underline{\theta}) = 0.58$, and $q_2(\underline{\theta}, \bar{\theta}) = 0.32$.

As noted, there are multiple ways to gauge distortions in the spot contract of Example 3 (comparisons are summarized in Table 1):

¹⁶As mentioned in Footnote 15, in this setting, it is in fact not only clear that the $\underline{\theta}$ -type’s incentive constraint must bind, but it can be shown that it will be optimal for the principal to have *both* incentive constraints bind. This implies essentially that there will be no first-period separation of types at all in our spot-commitment setup.

¹⁷Since $p_1(\bar{\theta}) = 1$, the equilibrium value of $q_2(\bar{\theta}, \underline{\theta})$ is not unique, as the principal attaches zero probability to observing such a sequence of reports in equilibrium.

	Period 1		Period 2			
	$q_1(\bar{\theta})$	$q_1(\underline{\theta})$	$q_2(\bar{\theta}, \bar{\theta})$	$q_2(\underline{\theta}, \bar{\theta})$	$q_2(\underline{\theta}, \underline{\theta})$	$q_2(\bar{\theta}, \underline{\theta})$
Probability	63%	37%	50%	0%	37%	13%
Spot Contract (SC)	0.341	0.444	0.565	0.604	0.576	0.319
(1) q_1^*, q_2^*	0.301	0.380	0.550	–	0.538	–
(2) $q_1^*, q_2^* p_1^{\text{SC}}$	0.314	0.380	0.555	0.580	0.538	0.500
(3) $q_1^* p_1^{\text{SC}}, q_2^{\text{SC}}$	0.301	0.396				
(4) $q_2^* p_1^{\text{SC}}, q_1^{\text{SC}}$			0.565	0.604	0.576	0.514

Note: Row ‘Probability’ shows probability of observing quantity in spot equilibrium; row ‘Spot Contract (SC)’ shows quantities in equilibrium spot contract; remaining rows are explained in the text.

Table 1: Evaluating Output Distortions in Example 3.

(1) *Comparison with Unconstrained First Best* (q_1^*, q_2^*): First-best quantities as described in Section 2 are shown in the third row of Table 1. Distortions are easy to evaluate for the $\bar{\theta}$ -type since he produces deterministic quantities in the spot equilibrium (recall $p(\bar{\theta}) = 1$): His output is higher under the spot contract than under the dynamic first-best. Comparisons are complicated for the $\underline{\theta}$ -agent by his equilibrium quantities under the spot contract being stochastic. However, his *average* first-period quantity $p_1(\underline{\theta})q_1(\underline{\theta}) + (1 - p_1(\underline{\theta}))q_1(\bar{\theta}) = 0.42$ exceeds $q_1^*(\underline{\theta})$, whereas his average second-period output of 0.51 falls short of $q_2^*(\underline{\theta})$. Average *overall* second-period output (i.e., averaged over both agents and reports) under the spot contract (0.54) is higher, however, than average first-best second-period output (0.52).

(2) *Comparison with First Best Conditional on First-Period Reporting* ($q_1^*, q_2^* | p_1^{\text{SC}}$): Next, efficient levels of $q_1(\cdot)$ and $q_2(\cdot)$ given agents reporting strategies under the spot contract are given in the fourth row of Table 1. Compared to this benchmark, there is a clear upward distortion in both agents’ period output and the $\bar{\theta}$ -type’s second-period output. The direction of the distortion in the $\underline{\theta}$ -type’s second-period output depends on his first-period report, but the average distortion under the spot contract is again upward.

(3) *Comparison with Efficient First-Period Output Conditional on First-Period Reporting and Second-Period Output* ($q_1^* | p_1^{\text{SC}}, q_2^{\text{SC}}$): Row five of Table 1 shows that given (i) agents’ reporting strategies and (ii) second-period output, both type’s first-period output is inefficiently high under

spot contracting.

(4) *Comparison with Efficient Second-Period Output Conditional on First-Period Reporting and First-Period Output* ($q_2^*|p_1^{\text{SC}}, q_1^{\text{SC}}$): The last line of Table 1 shows that there is no distortion in the $\bar{\theta}$ -agent's second-period output given his first-period output. Moreover, since $p_1(\bar{\theta}) = 1$ and therefore $\nu_2(\bar{\theta}) = 0$, the $\bar{\theta}$ -agent's second-period is also efficient under truthful reporting by this measure, whereas it is inefficiently low for a $\underline{\theta}$ -report.

4.2.2 Example 4: Efficient Agent Learns Faster

To round off this section of numerical examples, Example 4 below presents a setting in which the more efficient agent also learns faster. As one may expect, results in this example are less surprising in that the spot equilibrium entails unambiguous downward distortion in quantities traded:

Example 4. Assume the value of output to the principal is given by $S(q_t) = 100q_t - 80q_t^2$, cost structures are $c_1 = 75$, $c_2(q_1, \bar{\theta}) = 60 - 70q_1$, and $c_2(q_1, \underline{\theta}) = 65 - 60q_1$, $\nu = 0.5$, and $\delta = 0.7$. The spot equilibrium then involves no revelation of information in the first period with $q_1 = 0.237$, and second-period outputs (contingent only on second-period reports) of $q_2(\bar{\theta}) = 0.354$ and $q_2(\underline{\theta}) = 0.262$.

Comparisons of output distortions are shown in Table 2.¹⁸ Relative to all benchmarks, outputs are distorted downward in the equilibrium spot contract. Moreover, in this setting, the ‘benevolent first-period principal’ will implement the report-independent first-period quantity of $q_1 = 0.251$, resulting in second-period quantities of $q_2(\bar{\theta}) = 0.360$ and $q_2(\underline{\theta}) = 0.267$. Hence, also relative to this benchmark, quantities are unambiguously distorted downward.

5 Conclusion

Our model has analyzed how the introduction of privately known learning capabilities into the standard dynamic model of adverse selection influences incentive design. Contrary to previous work by Lewis and

¹⁸To present the results in the same format as Table 1, Table 2 equivalently presents the equilibrium in Example 4 as a mechanism *with* first-period message but where contracts are independent of this message and reporting strategies are completely uninformative (i.e., $p_1(\bar{\theta}) = 1 - p_1(\underline{\theta})$).

	<i>Period 1</i>		<i>Period 2</i>			
	$q_1(\bar{\theta})$	$q_1(\underline{\theta})$	$q_2(\bar{\theta}, \bar{\theta})$	$q_2(\underline{\theta}, \bar{\theta})$	$q_2(\underline{\theta}, \underline{\theta})$	$q_2(\bar{\theta}, \underline{\theta})$
Probability	50%	50%	25%	25%	25%	25%
Spot Contract (SC)	0.237	0.237	0.354	0.354	0.262	0.262
(1) q_1^*, q_2^*	0.267	0.237	0.368	–	0.308	–
(2) $q_1^*, q_2^* p_1^{\text{SC}}$	0.253	0.253	0.361	0.361	0.313	0.313
(3) $q_1^* p_1^{\text{SC}}, q_2^{\text{SC}}$	0.245	0.245				
(4) $q_2^* p_1^{\text{SC}}, q_1^{\text{SC}}$			0.354	0.354	0.308	0.308

Note: See the explanation in Table 1 and in the text.

Table 2: Evaluating Output Distortions in Example 4.

Yildirim (2002a), we have considered a setting in which agents are privately informed about the *rate* at which they learn rather than just the cost side.

The focus of our investigation has been on whether this information being private leads to an under- or an overexploitation of learning effects relative to the efficient level. Under full commitment, we have shown that this crucially depends on whether learning effects let inherently more efficient agents expand their lead, or whether they enable inherently less efficient agents to catch up. In the first case, we obtain results similar to Lewis and Yildirim’s in that learning effects will be under-exploited—the only difference being that in our case, this distortion can be so strong as to eliminate the exploitation of learning effects altogether. In the second case, we obtain entirely new results: If learning effects let inherently less efficient agents catch up, the principal has an incentive to *overexploit* learning effects. Moreover, we have shown that this effect is not driven by the full-commitment assumption: An overexploitation of learning effects may result also under spot commitment, despite the general notion that limited commitment tends to deter rather than encourage long-term investments.

More generally, our analysis has shown that in order to predict an under- or overexploitation of learning effects in dynamic adverse selection settings, it is important to identify whether these learning effects serve to magnify or to diminish existing differences in efficiency between types. Concerning vertical procurement relationships, for instance, we may seek to categorize supplying industries along these lines according

to their technology. For instance, consider rather simple low-tech inputs produced in more traditional ‘bread-and-butter’ industries where there is little scope for large technological improvements. Even if there is originally some scope for improvements through learning by doing, we would eventually expect all agents to ‘catch on to the trick’ (some types sooner, some later), after which there is little scope for further improvement. Thus, we would expect learning effects to quickly subside and to equalize agents’ productivity. In such industries, our model would predict learning effects to be over- rather than under-exploited. In contrast, consider suppliers of more high-tech products such as the computer chip industry. Here, we would expect significant scope for long-run improvements in production technologies. Further, we would expect inherently more innovative and creative suppliers to ever increase their lead over less efficient suppliers through accumulated learning effects. For such suppliers, our model predicts learning effects to be under- rather than overexploited. Similar technological arguments may be applied to the regulation of monopolistic suppliers.

One may also imagine applications of our model outside of the realm of pure procurement and regulation settings. Take, for instance, labor contracts of the type considered in Miyazaki’s (1977) ‘internal labor market rat race’, where employees’ productivity on the job is privately known, and labor contracts specify how hard an agent is expected to work on the job. Assume, in addition, that how hard an agent works today influences his future productivity on the job. If we expect hard work to make a less efficient worker catch up with the more efficient worker’s productivity, we should expect the employer to ask agents to work inefficiently hard on the job—essentially, aggravating the ‘rat race’. If, on the other hand, we expect harder work today to magnify productivity differences between workers (as might be the case on more creative jobs), we should expect the employer to relax workers’ workload below the efficient level.

Other potential fields of application may include credit market problems in the style of Freixas and Laffont (1990), where the borrower is privately informed about the returns to his project. If the project’s future returns systematically depend on the size of the loan today, then our analysis would predict inefficiently high first-period loans if raising loans has a larger impact on productivity in inherently less productive projects, and vice versa if productivity in inherently more productive projects is more strongly affected. Finally, the insights may be applied to models of discrimination in quantity or quality by a monopolistic supplier of a

consumption good (see Maskin and Riley, 1984; Mussa and Rosen, 1978) if we assume that consumers get used to or even addicted to the good, so that consuming more (or a higher quality) of the good today increases consumers' willingness to pay tomorrow.¹⁹ The learning-speed criterion of our model then pertains to whether customers with a higher willingness to pay for the good also get used to the good faster, or whether it is the customers who value the good less who get used to it at a faster rate. Our analysis predicts an underexploitation of the 'addiction factor' in the former case, and an overexploitation in the latter.

¹⁹The idea of such 'rational addiction' has been introduced by Becker and Murphy (1988), albeit in the context of a competitive market. Boone and Shapiro (2006) have more recently investigated a setting quite similar to the application we have in mind here, the main difference being that the discriminating monopolist sells on 'anonymous' spot markets, that is, any information revealed to him in the first-period is useless in the second-period). However, in many settings, even if the seller has a sizeable number of clients, he can nonetheless design a 'personalized' sequence of contracts, as is done for instance by video rental chains by means of membership cards.

Appendix

A.1 Proofs of Results in Section 2

Proof of Lemma 2.1. Using elementary robust comparative statics and supermodular analysis (cf. Topkis, 1998), the results require identifying suitable complementarity relations among q_1 , q_2 and θ in the objective function W . Observe first that $\frac{\partial^2}{\partial q_1 \partial q_2} W = -\delta \frac{\partial}{\partial q_1} c_2 > 0$, so that W has increasing differences in \mathbf{q} , implying that \hat{q}_1^* and \hat{q}_2^* are both increasing in the quantity produced in the other period. Next, $\frac{\partial}{\partial q_2} W = -\delta c_2$, which is increasing in θ since we have assumed c_2 to decrease in θ . Hence, W has increasing differences in (q_2, θ) , implying that \hat{q}_1^* is increasing in θ . Finally, $\frac{\partial}{\partial q_1} W = -\delta q_2 \frac{\partial}{\partial q_1} c_2$. Thus, W has increasing differences in (q_1, θ) if more efficient agents learn faster ($|\partial c_2 / \partial q_1|$ increasing in θ), and decreasing differences if less efficient agents learn faster, yielding the comparative static result for \hat{q}_1^* in θ . Finally, if more efficient agents learn faster, W will thereby have increasing differences in all pairs of arguments, implying that \mathbf{q}^* is increasing in θ . \square

A.2 Proofs of Results in Section 3

Proof of Proposition 3.1. Using the Φ -function defined in (3), the incentive constraints may be written compactly as

$$\bar{U} \geq \underline{U} + \Phi(\underline{\mathbf{q}}) \quad (\text{A.1})$$

$$\underline{U} \geq \bar{U} - \Phi(\bar{\mathbf{q}}). \quad (\text{A.2})$$

Using the surplus functions \bar{W} and \underline{W} and agents' rents \bar{U} and \underline{U} , the principal's payoff from any contract may be written as

$$\nu[\bar{W}(\bar{\mathbf{q}}) - \bar{U}] + (1 - \nu)[\underline{W}(\underline{\mathbf{q}}) - \underline{U}]. \quad (\text{A.3})$$

The principal maximizes this payoff by choice of $\{\bar{\mathbf{q}}, \bar{U}\}$ and $\{\underline{\mathbf{q}}, \underline{U}\}$, subject to incentive constraints (A.1) and (A.2), and subject to the participation constraints, which we restate here for easy reference:

$$\bar{U} \geq 0 \quad (\text{A.4})$$

$$\underline{U} \geq 0. \quad (\text{A.5})$$

Observe first that only allocations satisfying the *implementability condition*

$$\Phi(\bar{\mathbf{q}}) \geq \Phi(\underline{\mathbf{q}}) \quad (\text{A.6})$$

can be realized. This condition follows from combining (A.1) and (A.2).

Next, we argue that for any menu of allocations $\{\bar{\mathbf{q}}, \underline{\mathbf{q}}\}$ satisfying (A.6), the principal will optimally set $\underline{U} = 0$ and $\bar{U} = \Phi(\underline{\mathbf{q}})$. To see this, note first that

constraint (A.4) may be neglected: Since $\Phi \geq 0$, it is implied by (A.1) and (A.5). Given this insight, the principal must optimally set $\underline{U} = 0$: Otherwise, he could decrease \underline{U} and \bar{U} by the same small amount without violating any of the remaining constraints, and thereby strictly increase his payoff (A.3). But then the remaining constraints, (A.1) and (A.2), simplify to $\Phi(\underline{\mathbf{q}}) \leq \bar{U} \leq \Phi(\bar{\mathbf{q}})$, so that the principal must optimally set $\bar{U} = \Phi(\underline{\mathbf{q}})$: If not, we could decrease \bar{U} by a small amount without violating any of the remaining constraints. Finally, this implies that the only remaining constraint, (A.2), simply becomes (A.6).

Hence, the principal's optimization problem may be restated as choosing $\bar{\mathbf{q}}$ and $\underline{\mathbf{q}}$ so as to maximize

$$\nu[\bar{W}(\bar{\mathbf{q}}) - \Phi(\underline{\mathbf{q}})] + (1 - \nu)W(\underline{\mathbf{q}}) \quad (\text{A.7})$$

subject to (A.6). Due to the additively separable structure and since $\nu \in (0, 1)$, this is equivalent to condition (4). To complete the proof, it therefore remains to be shown that any $\{\bar{\mathbf{q}}, \underline{\mathbf{q}}\}$ satisfying (4) also satisfy (A.6). The former implies in particular that

$$\bar{W}(\bar{\mathbf{q}}) \geq \bar{W}(\underline{\mathbf{q}}) \quad \text{and} \quad W(\underline{\mathbf{q}}) - \frac{\nu}{1 - \nu}\Phi(\underline{\mathbf{q}}) \geq W(\bar{\mathbf{q}}) - \frac{\nu}{1 - \nu}\Phi(\bar{\mathbf{q}}). \quad (\text{A.8})$$

Since $\bar{W}(\underline{\mathbf{q}}) - W(\underline{\mathbf{q}}) = \Phi(\underline{\mathbf{q}})$ for any $\underline{\mathbf{q}}$ by the definition of \bar{W} and W , the inequalities in (A.8) may be added to yield $\Phi(\bar{\mathbf{q}})/(1 - \nu) \geq \Phi(\underline{\mathbf{q}})/(1 - \nu)$, which implies (A.6). \square

Proof of Proposition 3.2. We again employ supermodular analysis to derive the results. To this end, define the real-valued function $g(\mathbf{q}; \tau)$ such that

$$g(\mathbf{q}; \tau) = \begin{cases} W(\underline{\mathbf{q}}), & \text{for } \tau = 0, \\ W(\underline{\mathbf{q}}) - \frac{\nu}{1 - \nu}\Phi(\underline{\mathbf{q}}), & \text{for } \tau = 1. \end{cases} \quad (\text{A.9})$$

(a) For any \mathbf{q} ,

$$\frac{\partial}{\partial q_2} g(\mathbf{q}; 1) - \frac{\partial}{\partial q_2} g(\mathbf{q}; 0) = -\frac{\nu}{1 - \nu} \frac{\partial}{\partial q_2} \Phi(\underline{\mathbf{q}}), \quad (\text{A.10})$$

which is strictly negative (see the definition of Φ). Thus, for any q_1 , g has strictly decreasing differences in q_2 and τ , implying that $\arg \max_{q_2} g(q_1, q_2; \tau)$ is decreasing in τ for any q_1 . By definition of g , this set corresponds to the conditional second-period first-best for $\tau = 0$ and to the conditional second-period second-best, which proves the claim.

(b) For any \mathbf{q} ,

$$\frac{\partial}{\partial q_1} g(\mathbf{q}; 1) - \frac{\partial}{\partial q_1} g(\mathbf{q}; 0) = -\frac{\nu}{1 - \nu} \frac{\partial}{\partial q_1} \Phi(\underline{\mathbf{q}}), \quad (\text{A.11})$$

the sign of which depends on whether the more efficient $\bar{\theta}$ -agent also learns faster ($\frac{\partial}{\partial q_1} \Phi(\underline{\mathbf{q}}) \geq 0$) or whether the $\underline{\theta}$ -agent learns faster ($\frac{\partial}{\partial q_1} \Phi(\underline{\mathbf{q}}) \leq 0$), with

each of these inequalities being strict for $q_2 > 0$. Thus, whenever the more efficient agent learns faster, g has decreasing differences in q_1 and τ for any q_2 (and strictly so for any $q_2 > 0$), which proves claim (bi). On the other hand, whenever the less efficient agent learns faster, g has *increasing* differences in q_1 and τ for any q_2 , thereby proving part (bii). \square

Proof of Proposition 3.3. We again employ the auxiliary function $g(\mathbf{q}; \tau)$ defined in (A.9) in the proof of Proposition 3.2, so that $\underline{\mathbf{q}}^* \in \arg \max_{\mathbf{q}} g(\mathbf{q}; 0)$ and $\underline{\mathbf{q}}^{\text{SB}} \in \arg \max_{\mathbf{q}} g(\mathbf{q}; 1)$. As established there, for the case at hand, g has strictly decreasing differences in q_2 and τ for any q_1 , and strictly decreasing differences in q_1 and τ for any $q_2 > 0$. Hence, g has strictly decreasing differences in (\mathbf{q}, τ) . Moreover, g is supermodular in \mathbf{q} for $\tau = 0$ since W is supermodular in \mathbf{q} (see the proof of Proposition 2.1). Hence, the set of maximizers of g is decreasing in τ , which proves the claim.²⁰ \square

Proof of Proposition 3.5. Observe first that $\hat{q}_1^{\text{SB}}(0) = q_1^\circ$ since the principal's optimal choice of q_1 conditional on $q_2 = 0$ simply maximizes first-period surplus.²¹ Moreover, $\hat{q}_1^{\text{SB}}(q_2^{\text{SB}}) = q_1^{\text{SB}}$ by definition of the full and conditional optima. Since $q_2^{\text{SB}} \geq 0$, learning effects will thus be exploited if $\hat{q}_1^{\text{SB}}(q_2)$ is increasing in q_2 , and unexploited if it is decreasing. This in turn depends on whether the principal's objective $\underline{W} - \nu\Phi$ has increasing or decreasing differences in (q_1, q_2) . Using the definitions of \underline{W} and Φ , we have

$$\frac{\partial^2}{\partial q_1 \partial q_2} \{ \underline{W}(\mathbf{q}) - \frac{\nu}{1-\nu} \Phi(\mathbf{q}) \} = \frac{\delta}{1-\nu} \left[\left| \frac{\partial}{\partial q_1} c_2(q_1) \right| - \nu \cdot \left| \frac{\partial}{\partial q_1} \bar{c}_2(q_1) \right| \right],$$

so that the former will be the case whenever condition (6) holds (and the latter whenever (6) is reversed), which completes the proof. \square

²⁰Note that supermodularity of g in \mathbf{q} for $\tau = 1$ is *not* required for the proof (and not generally satisfied). The interested reader is invited to verify that, given a *binary* parameter space, Theorem 2.8.1 in Topkis (1998) in fact requires supermodularity of the objective for only *one* of the two parameter values: Together with increasing differences, this is easily seen to imply Topkis' condition (2.8.1), which produces the result by Lemma 2.8.1.

²¹The same is true of the conditional first-best output, i.e. $\hat{q}_1^*(0) = q_1^\circ$, which is why the \hat{q}_1^{SB} - and \hat{q}_1^* -curves meet at $q_2 = 0$ in Figure 2.

References

- Baron, David P., and Roger B. Myerson (1982) 'Regulating a monopolist with unknown costs.' *Econometrica* 50(4), 911–930
- Becker, Gary S., and Kevin M. Murphy (1988) 'A theory of rational addiction.' *Journal of Political Economy* 96(4), 675–700
- Bester, Helmut, and R. Strausz (2001) 'Contracting with imperfect commitment and the revelation principle: The single agent case.' *Econometrica* 69(4), 1077–1098
- Boone, Jan, and Joel Shapiro (2006) 'Selling to consumers with endogenous types.' Working paper, Universitat Pompeu Fabra, Barcelona
- Edlin, Aaron S., and Chris Shannon (1998) 'Strict monotonicity in comparative statics.' *Journal of Economic Theory* 81(1), 201–219
- Freixas, Xavier, and Jean-Jacques Laffont (1990) 'Optimal banking contracts.' In *Essays in Honor of Edmond Malinvaud, Vol. 2, Macroeconomics*, ed. Paul Champsaur (Cambridge, Mass.: MIT Press)
- Freixas, Xavier, Roger Guesnerie, and Jean Tirole (1985) 'Planning under incomplete information and the ratchet effect.' *Review of Economic Studies* 52(2), 173–191
- Fudenberg, Drew, Bengt Holmstrom, and Paul Milgrom (1990) 'Short-term contracts and long-term agency relationships.' *Journal of Economic Theory* 51(1), 1–31
- Gaudet, Gérard, Pierre Lasserre, and Ngo Van Long (1996) 'Dynamic incentive contracts with uncorrelated private information and history dependent outcomes.' *Japanese Economic Review* 47(4), 321–334
- Gärtner, Dennis (2007) 'Essays in industrial organization and mechanism design.' Ph.D. Thesis, University of Zürich, available at <http://www soi.uzh.ch/research/dissertations/gaertner.pdf>
- Guesnerie, Roger, and Jean-Jacques 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(3), 329–369
- Laffont, Jean-Jacques, and Jean Tirole (1986) 'Using cost observation to regulate firms.' *Journal of Political Economy* 94(4), 614–641

- (1987) ‘Comparative statics of the optimal dynamic incentive compatible contract.’ *European Economic Review* 31(4), 901–926
- (1993) *A Theory of Incentives in Procurement and Regulation* (Cambridge, Mass.: MIT Press)
- Lewis, Tracy R., and Huseyin Yildirim (2002a) ‘Learning by doing and dynamic regulation.’ *RAND Journal of Economics* 33(1), 22–36
- (2002b) ‘Managing dynamic competition.’ *American Economic Review* 92(4), 779–797
- Maskin, Eric, and John Riley (1984) ‘Monopoly with incomplete information.’ *RAND Journal of Economics* 15(2), 171–196
- Milgrom, Paul, and John Roberts (1990) ‘The economics of modern manufacturing: Technology, strategy and organization.’ *American Economic Review* 80(3), 511–528
- Miyazaki, Hajime (1977) ‘The rat race and internal labor markets.’ *Bell Journal of Economics* 8(2), 394–418
- Mussa, Michael, and Sherwin Rosen (1978) ‘Monopoly and product quality.’ *Journal of Economic Theory* 18(2), 301–317
- Topkis, Donald M. (1998) *Supermodularity and Complementarity* (Princeton, New Jersey: Princeton University Press)

Working Papers of the Socioeconomic Institute at the University of Zurich

The Working Papers of the Socioeconomic Institute can be downloaded from http://www soi.uzh.ch/research/wp_en.html

- 0718 Monoplistic Screening under Learning by Doing, Dennis Gärtner, December 2007, 29 p.
- 0717 An analysis of the Swiss vote on the use of genetically modified crops, Felix Schläpfer, November 2007, 23 p.
- 0716 The relation between competition and innovation – Why is it such a mess? Armin Schmutzler, November 2007, 26 p.
- 0715 Contingent Valuation: A New Perspective, Felix Schläpfer, November 2007, 32 p.
- 0714 Competition and Innovation: An Experimental Investigation, Dario Sacco, October 2007, 36p.
- 0713 Hedonic Adaptation to Living Standards and the Hidden Cost of Parental Income, Stefan Boes, Kevin Staub, Rainer Winkelmann, October 2007, 18p.
- 0712 Competitive Politics, Simplified Heuristics, and Preferences for Public Goods, Felix Schläpfer, Marcel Schmitt, Anna Roschewitz, September 2007, 40p.
- 0711 Self-Reinforcing Market Dominance, Daniel Halbheer, Ernst Fehr, Lorenz Goette, Armin Schmutzler, August 2007, 34p.
- 0710 The Role of Landscape Amenities in Regional Development: A Survey of Migration, Regional Economic and Hedonic Pricing Studies, Fabian Waltert, Felix Schläpfer, August 2007, 34p.
- 0709 Nonparametric Analysis of Treatment Effects in Ordered Response Models, Stefan Boes, July 2007, 42p.
- 0708 Rationality on the Rise: Why Relative Risk Aversion Increases with Stake Size, Helga Fehr-Duda, Adrian Bruhin, Thomas F. Epper, Renate Schubert, July 2007, 30p.
- 0707 I'm not fat, just too short for my weight – Family Child Care and Obesity in Germany, Philippe Mahler, May 2007, 27p.
- 0706 Does Globalization Create Superstars?, Hans Gersbach, Armin Schmutzler, April 2007, 23p.
- 0705 Risk and Rationality: Uncovering Heterogeneity in Probability Distortion, Adrian Bruhin, Helga Fehr-Duda, and Thomas F. Epper, July 2007, 29p.
- 0704 Count Data Models with Unobserved Heterogeneity: An Empirical Likelihood Approach, Stefan Boes, March 2007, 26p.
- 0703 Risk and Rationality: The Effect of Incidental Mood on Probability Weighting, Helga Fehr, Thomas Epper, Adrian Bruhin, Renate Schubert, February 2007, 27p.
- 0702 Happiness Functions with Preference Interdependence and Heterogeneity: The Case of Altruism within the Family, Adrian Bruhin, Rainer Winkelmann, February 2007, 20p.
- 0701 On the Geographic and Cultural Determinants of Bankruptcy, Stefan Buehler, Christian Kaiser, Franz Jaeger, June 2007, 35p.
- 0610 A Product-Market Theory of Industry-Specific Training, Hans Gersbach, Armin Schmutzler, November 2006, 28p.
- 0609 Entry in liberalized railway markets: The German experience, Rafael Lalive, Armin Schmutzler, April 2007, 20p.
- 0608 The Effects of Competition in Investment Games, Dario Sacco, Armin Schmutzler, April 2007, 22p.

- 0607 Merger Negotiations and Ex-Post Regret,
Dennis Gärtner, Armin Schmutzler, September 2006, 28p.
- 0606 Foreign Direct Investment and R&D offshoring,
Hans Gersbach, Armin Schmutzler, June 2006, 34p.
- 0605 The Effect of Income on Positive and Negative Subjective Well-Being,
Stefan Boes, Rainer Winkelmann, May 2006, 23p.
- 0604 Correlated Risks: A Conflict of Interest Between Insurers and Consumers and Its
Resolution,
Patrick Eugster, Peter Zweifel, April 2006, 23p.
- 0603 The Apple Falls Increasingly Far: Parent-Child Correlation in Schooling and the
Growth of Post-Secondary Education in Switzerland,
Sandra Hanslin, Rainer Winkelmann, March 2006, 24p.
- 0602 Efficient Electricity Portfolios for Switzerland and the United States,
Boris Krey, Peter Zweifel, February 2006, 25p.
- 0601 Ain't no puzzle anymore: Comparative statics and experimental economics,
Armin Schmutzler, December 2006, 45p.
- 0514 Money Illusion Under Test,
Stefan Boes, Markus Lipp, Rainer Winkelmann, November 2005, 7p.
- 0513 Cost Sharing in Health Insurance: An Instrument for Risk Selection?
Karolin Becker, Peter Zweifel, November 2005, 45p.
- 0512 Single Motherhood and (Un)Equal Educational Opportunities: Evidence for Germany,
Philippe Mahler, Rainer Winkelmann, September 2005, 23p.
- 0511 Exploring the Effects of Competition for Railway Markets,
Rafael Lalive, Armin Schmutzler, April 2007, 33p.
- 0510 The Impact of Aging on Future Healthcare Expenditure;
Lukas Steinmann, Harry Telser, Peter Zweifel, December 2006, 23p.
- 0509 The Purpose and Limits of Social Health Insurance;
Peter Zweifel, September 2005, 28p.
- 0508 Switching Costs, Firm Size, and Market Structure;
Simon Loertscher, Yves Schneider, August 2005, 29p.
- 0507 Ordered Response Models;
Stefan Boes, Rainer Winkelmann, March 2005, 21p.
- 0506 Merge or Fail? The Determinants of Mergers and Bankruptcies in Switzerland, 1995-
2000; Stefan Buehler, Christian Kaiser, Franz Jaeger, March 2005, 18p.
- 0505 Consumer Resistance Against Regulation: The Case of Health Care
Peter Zweifel, Harry Telser, Stephan Vaterlaus, February 2005, 23p.
- 0504 A Structural Model of Demand for Apprentices
Samuel Mühlemann, Jürg Schweri, Rainer Winkelmann and Stefan C. Wolter,
February 2005, 25p.
- 0503 What can happiness research tell us about altruism? Evidence from the German
Socio-Economic Panel
Johannes Schwarze, Rainer Winkelmann, September 2005, 26p.
- 0502 Spatial Effects in Willingness-to-Pay: The Case of Two Nuclear Risks
Yves Schneider, Peter Zweifel, September 2007, 31p.
- 0501 On the Role of Access Charges Under Network Competition
Stefan Buehler, Armin Schmutzler, January 2005, 30p.
- 0416 Social Sanctions in Interethnic Relations: The Benefit of Punishing your Friends
Christian Stoff, February 2006, 51p.
- 0415 Single Motherhood and (Un)equal Educational Opportunities: Evidence from
Germany,
Philippe Mahler, Rainer Winkelmann, October 2005, 23p.