DISSERTATION

Titel der Dissertation

"Stochastic Modelling:
Applications to Managerial Problems"

Verfasser

Mag. David Hirschmann

angestrebter akademischer Grad

Doctor of Philosophy (PhD)

Wien, 2014

Studienkennzahl lt. Studienblatt:	A 094 151 146
Studienrichtung lt. Studienblatt:	Betriebswirtschaft
Betreuer:	Univ.-Prof. Dr. Franz Wirl

Acknowledgments

It is fair to say that this thesis would not exist if it had not been for Prof. Wirl and Prof. Wozabal. Both gave me the chance to work with them at university while granting me a lot of freedom for my own research. They were always available for discussion and I have learned a lot from them, not only about research but in many different aspects.

No less important were my family and friends. My parents, my grandparents as well as my sister and my brothers have been encouraging and supporting me ever since I can remember and I am very grateful for having them. Also, I want to thank Myra and Dani for helping me - knowingly or not - through difficult times during the PhD, Magdalena for tipping the scales when I took the decision to enroll in the PhD program, and Susanne for her support and understanding during the last months.

There are a lot of other people who offered their help and contributed in one way or another. In particular, I would like to mention Prof. Vetschera, Prof. Fabel, Prof. Müller and Prof. Wälde, as well as Prof. Krattenthaler, Prof. Steinbauer, Prof. Hofbauer and Prof. Teschl from the Faculty of Mathematics.

Finally, I want to thank my PhD colleagues Christoph Graf, Lea Six and Philipp Fuchs. Thanks for the discussions and for accompanying me throughout the challenging and instructive time of writing this thesis.

Vienna, August 2014 David Hirschmann

Contents

Stochastic Modelling: Applications to Managerial Problems

Introduction

0.1 Genesis

If everything had worked out as planned on the first day in my PhD program, the title of the thesis would probably read "stochastic control theory in energy markets". But things do not always work out as planned. Instead, my research area drifted apart into different directions.

Now, the thesis covers two different subjects: random inspections and intermittent renewable energy sources. The reason for this split can by traced back to my first research project which focused on stochastic control models in electricity markets. My supervisor and I set up a model of an oil monopoly (e.g., OPEC) that is confronted with random demand shocks. After having studied the model and its implications for some time, it became clear that the results were neither surprising nor did they offer a lot of insights. However, I stuck with the main components of the first research project, namely stochastic control theory and energy markets. The path of stochastic control led me to the topic of random inspections and its effects, while the path of electricity markets led me to renewable intermittent energy sources.

Finally, the thesis comprises of three articles: two on random inspections and one on renewable energy sources. In all three articles – like in many managerial problems – stochastics lie at the core of the problem. In the following I give a short overview over the articles and over their current status in the publication process.

0.2 Overview

The thesis is written in a "journal style", as the goal is to publish each of the three articles. The arguments are given in a compact form but more details can be provided upon request.

0.2.1 Random Inspections

The first two articles deal with random inspections in two different settings – quality inspections and doping inspections – and apply two different methods, namely optimal control and classical game-theory. Both articles have one finding in common: they show that in some situation misbehavior is not reduced by an increased inspection rate. Even more, an economic agent who decides on the rate of inspections might prefer a lower inspection rate to a higher one. In both models, the reason is not a classical cost-benefit analysis, i.e., costs of higher inspection rates are not recovered, but other, unwanted implications which are detailed in the corresponding articles.

The first of the two articles on random inspections is titled "Optimal quality provision when reputation is subject to random inspections" and was published in Operations Research Letters

(2014) Vol. 42 (1), p. 64–69.[1] According to the latest VHB-JOURQUAL 2.1 (2011) ranking, Operations Research Letters is ranked B.

The article considers a firm which sets the quality level of its products in order to maximize its profits. The crucial assumption is that the quality is not observable by consumers, at least not before consumption. If the consumers can never observe the quality – as it is the case for so-called credence goods – the firm would have strong incentives to offer the lowest possible quality: this behavior would save the firm money during the production process but would not influence the sales quantity since consumers do not know the true quality. Therefore, a third party is introduced which has the power to conduct quality inspections at random times (e.g., consumerism or a public authority). Its findings are made public and determine the firms reputation which in turn determines the sales quantity. The smiley-system in place in Denmark is an example for such a model. In this system compliance with food regulations are checked in retail food enterprises, and depending on the findings "smileys" are awarded. The results have to be published on the enterprise's website and have to be displayed such that consumers deciding to enter a restaurant or a shop can see them.[2]

The model implies that for credence goods quality is constant over time and increasing in the inspection rate. For so-called experience goods, for which information about the quality can be observed after consumption and shared among consumers, the model exhibits counter-intuitive implications from a regulator's viewpoint: it implies not only that a low inspection rate might be more beneficial than a high one, but even more, depending on the initial reputation of the firm, a regulator might prefer to not conduct any inspections at all.

The second article is titled "May increasing doping sanctions discourage entry to the competition?" and is currently in the second round of revision in the Journal of Sports Economics. The journal is published by the North American Association of Sports Economists and by SAGE publications. It does not appear in the latest VHB-JOURQUAL 2.1 (2011) ranking, but has an impact factor of 0.743 according to 2012 Journal Citation Reports by Thomas Reuters (2013).

The article was inspired by an article by Haugen [3] who considers a simple static game in which two competing athletes have the choice of using a performance-enhancing drug (doping). Haugen, as others before, shows that the athletes are caught in some sort of prisoner's dilemma and that doping is the rational strategy for both athletes. The results led me to take a different viewpoint and ask how these findings influence young talents who consider becoming professional athletes. In the thesis it is not only shown that the usage of doping might detain talents from entering, but that under certain circumstances this effect is even starker for a higher doping inspection frequency.

Overall, it is noteworthy that both models show how a higher inspection rate can be disadvantageous. These counter-intuitive results should be kept in mind when setting up corresponding inspection policies.

[1] For the thesis, a slightly more detailed version of the paper was incorporated.

[2] For more information see the corresponding website of the Danish Veterinary and Food Administration, `http://www.findsmiley.dk` (retrieved on 07.Aug.2014).

[3] Haugen, K. K., "The Performance-Enhancing Drug Game". Journal of Sports Economics (2004) Vol. 5 (1), p. 67–86.

0.2.2 Intermittent Renewable Energy

The third article “The effect of intermittent renewables on the electricity price variance” covers a different topic. In literature, it is widely believed that the intermittent character of wind power and of photovoltaics increases the price variance on the electricity markets. The idea is that the capacity increase in wind power and photovoltaic increases the variance in the amount of produced energy (as the amount of wind and sun varies during a day) which in turn translates into a higher price variance. The article challenges this idea: First, a theoretical framework is formulated which analyzes how the build-up of intermittent renewables influences the electricity prices. The model implies that a high amount of renewable energy indeed increases the price variance, but that a low amount actually decreases price variance. Then, the theory is put to test by using data from the German/Austrian EPEX Market and empirical evidence for the effects derived from the theoretic model is provided.

The article is currently under revision in OR Spectrum which is ranked A according to the latest VHB-JOURQUAL 2.1 (2011) ranking. It is joined work with Prof. David Wozabal and my former PhD colleague Dr. Christoph Graf. I contributed to setting up the theoretical framework, as well as translating the theory into an empirical model. I was responsible for writing the corresponding chapters as well as the literature review.

While working on the topic of renewable energy, another article emerged (but is not included in this thesis). It is titled “Renewable energy and its impact on power markets” and is also joined work with Prof. David Wozabal and Dr. Christoph Graf. The article gives an overview over the influence of renewable energy on power markets, the stochastics and costs of renewable energy production, as well as related regulatory frameworks, and it was published as a book chapter in the Handbook of Risk Management for Energy Production and Trading (2013), (Kovacevic, Pflug, and Vespucci; Eds.), International Series in Operations Research and Management Science, Volume 199, Springer. 283-311.

Stochastic Modelling: Applications to Managerial Problems

First Article

Operations Research Letters (2014) Vol. 42 (1), p. 64–69.

Chapter 1

Optimal Quality Provision When Reputation is Subject to Random Inspections

We consider a firm whose profit is determined by its reputation. The quality of its products is unobservable, at least before consumption, but random inspections reveal the true quality and change the reputation. We obtain closed-form solutions for the provision of quality and show that increasing the inspection rate can be disadvantageous for consumers.

1.1 Introduction

We consider a firm whose reputation is related to the quality of the products it offers. The quality is in general not observable by customers, at least not before purchasing the product. Therefore, the customers' buying decision depends - apart from price - crucially on the reputation of the firm. The main mechanism for determining the reputation is random inspections. An authority conducts inspections which reveal the true quality. If the true quality deviates from the current reputation, then the reputation is immediately changed. The main questions are: how does a firm behave under such a framework? Which level of quality does the firm offer and how does its reputation evolve? And what are the implications for a regulator that wants to set an optimal inspection rate?

In the paper, we focus on two different types of products and quality, respectively: *credence goods* and *experience goods*. For credence goods (see, e.g., Dulleck and Kerschbamer, 2006), quality relates to the production process. At no point can the customer observe the quality of the product, not even after having consumed it. We think, e.g., of products with the Fairtrade certification or organic production. Customers cannot observe the production process. Their only reference point is the reputation of the firm. By random inspections, an authority (e.g. consumerism) can determine the "quality" of the production process.

Conversely, the quality of experience goods (see Nelson, 1970) can be observed by customers *after* having consumed it. Examples include meals in restaurants, hotels and medical services. Again, we assume that reputation is determined by random inspections by authorities, consumerism or reviewers. Additionally, consumers can share their experience by word-of-mouth. This *diffusion of information* influences the reputation continuously.

We find that for credence goods the firm offers a constant level of quality. This level is independent of the firm's initial reputation. The firm raises the quality level if the rate of inspections is increased, but the marginal increase in quality diminishes. With this knowledge, a regulator can trade off inspection costs against an increased level of quality offered by the firm.

For experience goods, the quality level offered still converges to a constant level in the long run. In the short run, however, the firm offers a varying quality level below or above the long-run level. The offered level is inverse to its reputation (high reputation yields low quality and vice versa) and is adjusted continuously. Therefore, reputation may be misleading since it deviates from the true quality level. However, in time, the deviation decreases. From the viewpoint of a regulator, we show that in some cases a lower inspection rate is better than a higher inspection rate, and that this depends in particular on the initial reputation of the firm.

The model is formulated as a stochastic optimal control problem. Our concept of reputation is similar to the stock of goodwill in the Nerlove-Arrow model (Nerlove and Arrow, 1962) which is still widely used, e.g., Grosset et al. (2011). The difference is that in the Nerlove-Arrow model, the capital stock is built up and depreciates continuously and deterministically, while in our model the capital stock can exhibit random discontinuities due to inspections.

The inspections are modeled as a Poisson jump process. The jumps correct the difference between the reputation and the actual quality and, in some sense, reveal the otherwise hidden behavior of the firm. Although optimal control problems with jumps "are a standard tool in economic literature", (see Sennewald, 2007), they are rarely used to repeatedly reveal states that can be changed continuously. One of the few applications is found in Kumar (1980). There, a finite time differential game between a pursuer and an evader is formulated. A noisy signal of the position of both parties is revealed due to random jumps. A vivid example of this is given in Snyder (1975): the parties are interpreted as a normally shaped swarm of fireflies whose members emit a flash of light at random time instants. The solution of this problem requires the solution of an infinite system of coupled differential equations and can hardly be done analytically. Our model has the nice property of allowing for a closed-form solution. In particular, to the best of our knowledge, this is the first time that random inspections are treated in such a framework.

There is a second component in the model dynamics, diffusion of information, that is widely used in economics, see, e.g., Godes and Mayzlin (2004). For an optimal control setting, many different models were proposed; for an early overview see Sethi (1977). The actual type of dynamics and the analysis of how information diffuses do not lie at the center of our attention. We assume that the true quality of the product is revealed continuously in time. Our simple but very intuitive dynamics allow for a closed-form solution.

The structure of the paper is as follows. In Section 1.2, we state the model and the assumptions. In Section 1.3, we formulate deterministic variations of the model without inspections, which serve as benchmarks. In Section 1.4, we consider only random inspections which correspond to credence goods. In Section 1.5, diffusion of information is added. This model corresponds to experience goods. Section 1.6 concludes and presents ideas for future research.

1.2 Model

Consider a firm and its *reputation*, which is expressed in terms of the publicly *perceived* level of quality Q. The revenue of the firm $\hat{\Pi}(Q,p)$ depends upon reputation and price p. By assuming that the price for each level Q is set accordingly as to maximize revenues, we can focus on reputation only. Hence, the revenue function of the firm is $\Pi(Q) = \max_p \hat{\Pi}(Q,p)$. Independent of its reputation, the firm may choose to offer a different level of quality q. This *offered* or *true* quality incurs costs $\hat{c}(q)$ on the firm. In general, offered quality is not observable by customers. It is revealed by stochastic, Poisson-distributed inspections and possibly by diffusion of information. Offering a lower level of quality saves the firm money instantaneously but, if revealed, lowers the reputation Q. We denote the costs of offering the quality level q by $\hat{c}(q)$. Let r be the discount rate. The firm is risk-neutral and maximizes its expected profit:

$$\max_q \mathbb{E}\left[\int_0^\infty (\Pi(Q) - \hat{c}(q))e^{-rt}\ dt\right], \tag{1.1}$$

subject to the dynamics

$$dQ = \delta(q - Q)\ dt + (q - Q)\ dJ. \tag{1.2}$$

As indicated above, the dynamics allow for two possibilities of revelation of the true level of quality q. The deterministic term corresponds to diffusion of information. The difference between the true level q and reputation Q is observed after the consumption of the good and the information is spread at rate δ. The stochastic term, dJ, corresponds to a random inspection and denotes the increments of a Poisson process with arrival rate λ. At random times, the Poisson process jumps and the increment equals 1. This immediately corrects the state Q for the difference between the true level q and the perceived level Q. In between jumps, the stochastic term does not contribute to the dynamics.

We assume that the revenue function is quadratic and concave in state Q, $\Pi(Q) = -p_2/2\, Q^2 + p_1 Q$. Due to the form of the revenue function we place an upper bound upon the control $q \leq \overline{q} = p_1/p_2$. This is done to avoid an increase in reputation causing a decrease in revenue: as every control can become the actual state, $q > p_1/p_2$ would imply $\Pi(q)' > 0$. Costs are assumed to be linear and increasing in the control q, $\hat{c}(q) = cq$. These assumptions allow for a closed-form solution.

1.3 Deterministic Cases

The following deterministic models without inspections serve as benchmark models. For the models to have a solution, we have to restrict the feasible controls, $q \in [\underline{q}, \overline{q}]$.

1.3.1 Static Case

Proposition 1. *If the true quality level is observable, then the optimal strategy is to offer a constant quality level*

$$q^S = \frac{p_1 - c}{p_2}. \tag{1.3}$$

If the true quality level is never observable, then the optimal strategy is to offer the lowest possible level $q = \underline{q}$.

Proof. Since we are in a deterministic setting, we can omit the expectation in the maximization problem (1.1). If the true quality level is observable, then the problem is equivalent to the following maximization problem:

$$\max_{q\in[\underline{q},\overline{q}]} \int_0^\infty \left(\frac{-p_2}{2}q^2 + p_1 q - cq\right) e^{-rt}\, dt,$$

where q is the level of quality offered. Pointwise maximization yields the claimed result.

On the other hand, if the true quality level is not observable at all, it is clearly optimal to provide the lowest possible quality level $q = \underline{q}$. This minimizes the costs and maximizes the profit. □

We refer to the solution in (1.3) as the *complete information solution.*

1.3.2 Dynamic Case - Diffusion of Information

The proposition above describes the behavior of the firm if offered quality is always observable or if it is never observable. In some cases, the true quality level q is not observable at first but can be experienced after consumption. The experience is shared by word-of-mouth and the information diffuses. The maximization problem is (1.1) - omitting the expectation - and the dynamics of diffusion are given by the deterministic part of (1.2).

Proposition 2 (Diffusion-Only)**.** *Assume that the information about the difference between reputation Q and the true quality level q diffuses at rate δ. Then the optimal strategy is a bang-bang type strategy that approaches a constant level as fast as possible and remains at this level for all times:*

$$q^*(Q) = \begin{cases} \underline{q}, & \text{for } Q > q^\delta \\ q^\delta, & \text{for } Q = q^\delta \\ \overline{q}, & \text{for } Q < q^\delta \end{cases}$$

with

$$q^\delta = \frac{p_1 - (1+\frac{r}{\delta})c}{p_2}.$$

Proof. For a proof apply the standard arguments, see, e.g., Section 3.3. in Feichtinger and Hartl (1986). □

The economic interpretation of the solution is that for a high initial level of perceived quality the firm can exploit its reputation. It saves costs by offering the lowest possible quality for some time while reputation decreases. This is done until reputation reaches the level q^δ. Conversely, starting with a low initial level, the firm wants to increase reputation. This is costly, and the increase in reputation is sluggish. Reaching the complete information solution is too costly for the firm, $q^\delta < q^S$ for all $\delta \geq 0$, although $q^\delta \to q^S$, for $\delta \to \infty$.

Figure 1.1 shows an example of an optimal control and state path. The firm starts with reputation Q_0. It offers the lowest possible level $\underline{q} = 0$ which decreases reputation. When reputation reaches the state Q^δ, the firm switches and offers level Q^δ; perceived and actual quality coincide from now on.

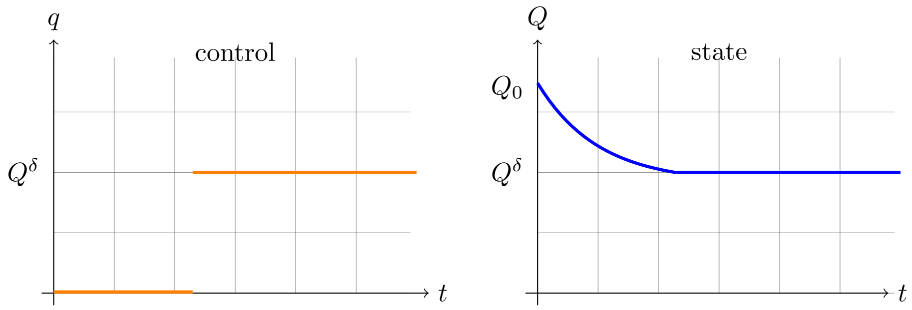

Figure 1.1: Optimal quality path (control) and corresponding reputation (state) under diffusion of information (without random inspections)

1.4 Random Inspections Only

Now we turn to random inspections and credence goods. We assume that an authority has the power to inspect the firm and reveal the difference between the reputation and the true behavior. As an example we consider a restaurant and the quality of the processed food: for customers the origin of the ingredients (e.g., claimed organic origin) is not observable. Only an inspection by authorities reveals the true quality of the ingredients. The same model describes firms which claim to comply with certain production standards (e.g., Fairtrade companies).

Proposition 3 (Inspection-Only)**.** *Assume that the true quality level q is revealed by Poisson-distributed inspections with rate λ. Then the optimal strategy is to offer a constant quality level*

$$q(Q)^* = q^\lambda := \frac{p_1 - (1 + \frac{r}{\lambda})c}{p_2}.$$

For the proof, we guess a candidate for the value function and show that it satisfies the Hamilton-Jacobi-Bellman (HJB) equation. (In Chang (2004) an entire chapter is devoted to guessing value functions). Verification of optimality of the guessed value function is done as shown in Sennewald (2007); Sennewald and Wälde (2006).

Proof of Proposition 3. The firm is confronted with the following maximization problem:

$$\max_{q \le \overline{q}} \mathbb{E}\left[\int_0^\infty \left(-\frac{p_2}{2}Q^2 + p_1 Q - cq\right) e^{-rt}\, dt\right]$$

subject to the dynamics

$$dQ = (q - Q)\, dJ.$$

The corresponding HJB equation reads (see Sennewald, 2007),

$$(r + \lambda)V(Q) = \max_q \{-\frac{p_2}{2}Q^2 + p_1 Q - cq + \lambda V(q)\}. \tag{1.4}$$

We guess a quadratic value function $V(Q) = v_2 Q^2 + v_1 Q + v_0$. Plugging this into (1.4) and applying the first order condition yields

$$q^* = \frac{c - \lambda v_1}{2\lambda v_2}. \tag{1.5}$$

For now, we simply assume that $q^* \le \overline{q}$ and will confirm it later. Plugging q^* and the guessed value function back into the HJB equation, the comparison of coefficients yields

$$
\begin{aligned}
v_2 &= -\frac{p_2}{2(r+\lambda)}, \\
v_1 &= \frac{p_1}{(r+\lambda)}, \\
v_0 &= \frac{1}{2r}\frac{\lambda}{\lambda+r}\frac{(p_1 - \frac{r+\lambda}{\lambda}c)^2}{p_2}.
\end{aligned}
$$

Therefore, the optimal control in (1.5) reads

$$
q^* = \frac{p_1 - (1+\frac{r}{\lambda})c}{p_2}.
$$

First, note that the concavity of $\Pi(Q)$ implies the concavity of $V(q)$. Therefore, the right-hand side of the HJB (1.4) is concave and the assumption of an inner maximum was justified. It is immediately clear that $q^* \leq \overline{q} = p_1/p_2$. Hence, we have guessed a candidate for the value function and retrieved the optimal control corresponding to this guess.

It still remains to verify that $V(Q)$ is indeed a value function. This is done by applying (Sennewald, 2007, Theorem 4) which can be found in the appendix, Theorem 6. First, note that Equations (1.28) and (1.29) are satisfied by construction of $V(Q)$. Then, note that the two limiting conditions also hold: The first condition, (1.30), states that along the optimal state process, the function V does not increase or decrease at a rate faster than $\exp(-rt)$. This clearly holds true, since $V(Q)$ is a polynomial. The second condition, (1.31), holds for the same reason. □

1.4.1 Properties of the Solution and Policy Implications

First, note that the level of offered quality is constant and independent of the reputation of the firm. This is similar to the complete information solution in (1.3) - only the level is lower now, $q^\lambda < q^S$. Random inspections do not yield the complete information solution but ensure a constant level of quality at all times. The possibility of being inspected keeps the firm from exploiting its good reputation by offering low quality.

For an example of an optimal control path and the corresponding state path see Figure 1.2. The firm has an initial reputation Q_0 and offers quality level Q^λ from the beginning. After the first inspection, the control Q^λ is revealed and the reputation changes. From now on, control and state coincide and are equal to Q^λ.

Consider the implications for a regulator which wants to determine the optimal inspection rate λ. The firm starts with an initial reputation Q_0 but offers quality level q^λ. After the first inspection, the control q^λ is revealed and state and control coincide from now at q^λ. With a higher inspection rate λ, the quality leakage between reputation and offered quality is revealed faster. Furthermore, the long-run level of offered quality increases and converges to the complete information solution, $q^\lambda \to q^S$, for $\lambda \to \infty$. A regulator can readily balance the costs of inspection rate λ against the benefits of the quality level in the long run and earlier revelation of the quality leakage. The findings imply that a higher inspection rate is always beneficial.

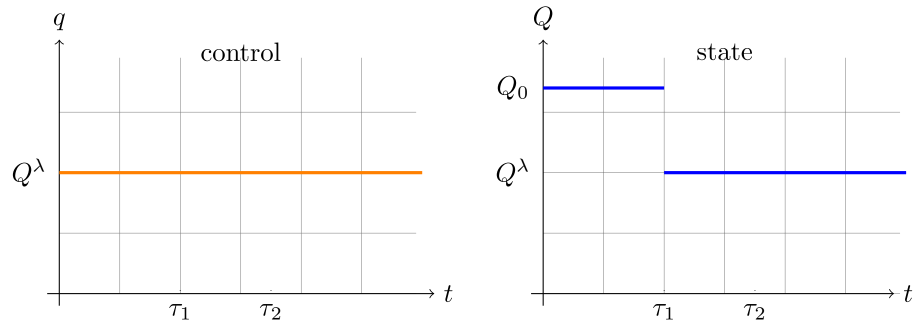

Figure 1.2: Optimal quality path (control) and corresponding reputation (state) with random inspections (without diffusion of information)

1.5 Random Inspections and Diffusion of Information

In this section, we combine random inspections and diffusion of information to model experience goods. As before, Poisson-distributed inspections occur at rate λ. Additionally, consumers observe the quality of the product after having consumed it. Their experience, or more precisely, their knowledge about the quality diffuses at rate δ. Again, a restaurant serves as an example. The quality of the food and service is observable by every customer after consumption. Additionally, reviewers from magazines or organizations (e.g., Gault & Millau) visit the restaurant and evaluate it. We assume that their reviews are accessible by every customer and that their reviews change the reputation instantaneously.

We derive the optimal control depending on the current state and then formulate a time-dependent version of the solution. With these results we obtain policy implications for a regulator who wants to determine an optimal inspection rate.

Proposition 4 (Diffusion and inspection)**.** *Assume that the true quality level q is revealed by Poisson-distributed inspections with rate λ. Between inspections, the information about the difference between reputation Q and the true quality level q diffuses at rate δ. Then the optimal strategy is to offer the quality level*

$$q(Q)^* = Q^F + \frac{\delta}{\lambda}(Q^F - Q), \tag{1.6}$$

where

$$Q^F = \frac{p_1 - (1 + \frac{r}{\lambda+\delta})c}{p_2}$$

is the fixpoint of $q(Q)^*$.

Proof. The maximization problem is given by (1.1) subject to (1.2). By assuming a quadratic value function, the parameters are determined analogously to the proof of Proposition 3. The proof can be found in the appendix. □

The optimal control equals Q^F and a "distortion term" $(Q^F - Q)$. Loosely speaking, for reputation $Q < Q^F$ a level of quality higher than Q^F is offered, while for reputation $Q > Q^F$ quality below Q^F is supplied. This implies that, except for the state Q^F, reputation and offered quality never coincide. The actual deviation of the offered quality from Q^F depends on the fraction δ/λ. The higher λ in relation to δ, the closer is the offered quality to Q^F. For $\lambda \to \infty$, the optimal control converges toward the constant value Q^F, whereas $\delta \to \infty$ yields values far from Q^F. The direct influence of the parameters of the revenue function p_1, p_2 is not as strong. They only change the value Q^F. The higher the curvature p_2 of the revenue function, the lower Q^F, whereas a higher p_1 increases Q^F.

1.5.1 Time-Dependent Formulation of the Solution

The solution of Proposition 4 is state-dependent. For control problems that involve, e.g., Brownian motion, this is the only way to write down the optimal control explicitly. Here, we can do a little more.

We note that in the model, stochastics arise only due to inspections. In between inspections, the model evolves in a deterministic way. Therefore, the optimal control and the optimal state path can be calculated explicitly in between inspections and the pieces can be pasted accordingly at the random inspection times.

Corollary 5. *Denote by $0 \leq \tau_1 \leq \ldots \leq \tau_n \leq t$ the inspection times until time t. Under the assumptions of Proposition 4, the optimal state path is given by*

$$Q^*(t) = Q^F + (Q(0) - Q^F)e^{-\frac{\delta}{\lambda}(\lambda+\delta)t} \left(-\frac{\delta}{\lambda}\right)^{|\{\tau_1,\ldots,\tau_n\}|}, \tag{1.7}$$

and the corresponding optimal control is given by

$$q^*(t) = Q^F + (Q(0) - Q^F)\ e^{-\frac{\delta}{\lambda}(\lambda+\delta)t}\ \left(-\frac{\delta}{\lambda}\right)^{1+|\{\tau_1,\ldots,\tau_n\}|}. \tag{1.8}$$

Here, $|.|$ denotes the cardinality of the set (i.e. the number of inspections which occurred until time t).

Proof. In between inspections the model is deterministic. The deterministic part of the dynamics is given by $dQ = \delta(q - Q)\ dt$. Plugging in the optimal control $q^*(Q)$ from (1.6) yields the following ordinary differential equation (ODE):

$$\begin{aligned} Q'(t) = \frac{dQ}{dt} &= \delta(q^* - Q) \\ &= \frac{\delta}{\lambda}(\delta + \lambda)\left(Q^F - Q(t)\right). \end{aligned}$$

By differentiation, the following solution of the ODE is easily verified:

$$Q^*(t) = Q^F + (I - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)t}, \tag{1.9}$$

where I is the initial reputation at time $t = 0$, $Q(0) = I$. This equation gives the optimal state path until the occurrence of the first inspection. Plugging it back into the formula for the state dependent control (1.8), we find the corresponding optimal control until the first inspection:

$$q^*(t) = q^*(Q^*(t)) = Q^F - \frac{\delta}{\lambda}(I - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)t}. \tag{1.10}$$

Assume that the first inspection occurs at time $t = \tau_1$. Then, the state at τ_1 is given by $Q(\tau_1) = q(\tau_1)$. This is because an inspection reveals the current state. From now on, the problem is deterministic again, but the initial value has changed to $I = Q(\tau_1)$.

We now use an induction type argument that applies the logic from above to prove (1.7) and (1.8). Denote by $0 = \tau_0 < \tau_1 \leq \ldots \leq \tau_n \leq t$ the inspection times until time t (we set $\tau_0 = 0$ for convenience of notation). From (1.9) and (1.10) it is clear that (1.7) and (1.8) hold true for $t \in [\tau_0, \tau_1)$. Now, assume that for $t \in [\tau_{i-1}, \tau_i)$, the control is given by (1.8). We show that the same equations hold true for the interval $[\tau_i, \tau_{i+1})$. For $t \in [\tau_i, \tau_{i+1})$ and $\hat{t} = t - \tau_i$, we have

$$\begin{aligned}
Q^*(t) &= Q^*(\tau_i + \hat{t}) \\
&\overset{(1.9)}{=} Q^F + (Q^*(\tau_i) - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)\hat{t}} \\
&= Q^F + (q^*(\tau_{i-}) - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)\hat{t}} \\
&\overset{(1.8)}{=} Q^F + e^{-\frac{\delta}{\lambda}(\delta+\lambda)\hat{t}} \times \\
&\quad \left((Q(0) - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)\tau_i} \left(-\frac{\delta}{\lambda}\right)^{1+|\{\tau_1,\ldots,\tau_{i-1}\}|} \right) \\
&= Q^F + (Q(0) - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)t} \left(-\frac{\delta}{\lambda}\right)^{|\{\tau_1,\ldots,\tau_i\}|},
\end{aligned} \tag{1.11}$$

which proves (1.7). Furthermore, recalling the state dependent formulation of $q^*(Q)$ in (1.6), we have

$$\begin{aligned}
q^*(t) &= q^*(Q^*(t)) \\
&\overset{(1.6)}{=} -\frac{\delta}{\lambda}Q^*(t) + \frac{\delta+\lambda}{\lambda}Q^F \\
&\overset{(1.11)}{=} Q^F + (Q(0) - Q^F)e^{-\frac{\delta}{\lambda}(\delta+\lambda)t} \left(-\frac{\delta}{\lambda}\right)^{1+|\{\tau_1,\ldots,\tau_i\}|}.
\end{aligned}$$

This proves (1.8). $\square$

1.5.2 Properties of the Solution

The structure of the time-dependent solution q^* in (1.8) is given as follows:

$$q^*(t) = Q^F + \underbrace{(Q(0) - Q^F)}_{\text{distance}} \underbrace{e^{-\frac{\delta}{\lambda}(\lambda+\delta)t}}_{\text{diffusion}} \underbrace{\left(-\frac{\delta}{\lambda}\right)^{1+|\{\tau_1,\ldots,\tau_n\}|}}_{\text{inspection}}.$$

It consists of the fixpoint Q^F and a term related to the distance between this fixpoint and the initial reputation, $Q(0) - Q^F$. The latter term has two more components: an exponential term which relates to continuous adjustments in quality (labeled "diffusion") and a fraction which relates to the sudden changes of an inspection (labeled "inspection").

First, consider the diffusion term. As time passes, the distance between $Q(0)$ and Q^F decreases due to this exponential term. The solution is drawn continuously toward the fixpoint Q^F. Still, Q^F is never attained since the exponential term never vanishes. Rewriting the exponent as $-(\delta + \delta^2/\lambda)t$, we see that an increase in the diffusion rate δ increases the speed of convergence. Conversely, an increase in the inspection rate λ decreases the speed of convergence in between inspections.

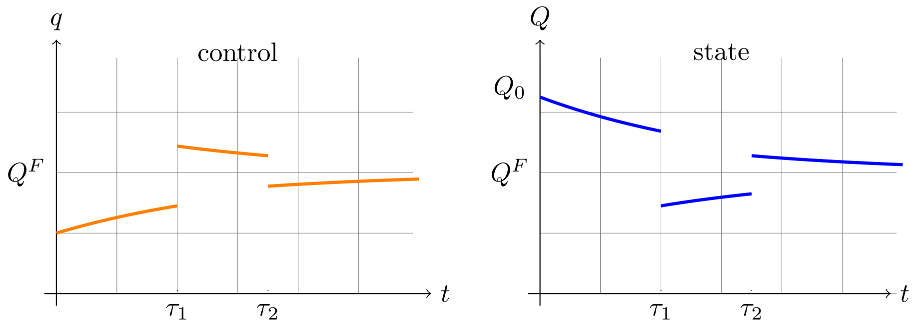

Figure 1.3: Optimal quality path (control) and corresponding reputation (state) under random inspections with diffusion of information.

Now, turn to the inspection term. The expression $|\{\tau_1, \ldots, \tau_n\}|$ counts the number of inspections that have occurred so far. After an inspection, the expression is increased by one and the sign of the term changes. This implies a sudden change in the quality level: if the offered quality q^* was higher than Q^F before, it is lower than Q^F after an inspection and vice versa.

The absolute size of the fraction $(-\delta/\lambda)$ determines by how much the offered quality level is higher or lower than before. If $\delta/\lambda < 1$, an inspection decreases the distance between the offered quality and the long-run optimum Q^F. If $\delta/\lambda > 1$, then an inspection pushes the optimally offered quality away from the long-run optimum. In expectation, the state still converges toward Q^F: the exponential term draws the control closer toward the long-run optimum than an inspection pushes it away. This can be concluded from (1.7). Still, it can happen that inspections occur earlier. This implies that at some point the control or the state could grow above the boundary $\bar{q}$. (In this area, increasing reputation yields lower revenues). For the sake of the closed-form solution, we have to ignore the positive probability of this happening. Also, the probability can be made arbitrarily small by narrowing an interval of feasible initial reputations around Q^F.

Summarizing, the optimal offered quality level and reputation converge toward Q^F. This convergence is alternating in the sense that a quality level below Q^F is offered while reputation is above Q^F, or vice versa. The alternation is triggered by the random inspections. Figure 1.3 shows an example of an optimal quality path and the corresponding reputation. The firm starts with initial reputation Q_0. It offers a low level of quality which it increases continuously. At the same time, the reputation decreases. At time τ_1, an inspection occurs and the reputation changes immediately. Now, as reputation is low, the firm offers a high level of quality, but lowers it continuously. At time τ_2, another inspection occurs and changes reputation immediately, again. Both, offered quality and reputation, converge toward the long-run optimum Q^F.

1.5.3 Policy Implications

Assume that a regulator wants to set an optimal inspection rate λ. First, one needs to define what "optimal" means in this context. At first sight, it may seem optimal to maximize the offered quality level $q^*(t)$ over time. This would ignore the possibility of the reputation Q being higher than the offered quality q^*. A regulator would want to avoid customers receiving a lower quality than they expect. On the opposite, if customers receive quality above the expected level

Q, the regulator would not mind. Therefore, we assume that the regulator solves the following minimization problem:

$$\min_{\lambda} \mathbb{E}\left[\int_0^{\infty} \max\{Q^*(t) - q^*(t), 0\}\ e^{-Rt}\ dt\right], \tag{1.12}$$

where R is the discounting rate of the regulator. We refer to this as "minimizing the *cumulative quality gap*".

We were not able to find a closed-form solution for this problem, but we investigate its solution numerically. To do so, we fix a time horizon $T < \infty$. As the difference between offered and perceived quality vanishes over time, this can be considered an adequate approximation, even for small values of R. Denote by $0 = \tau_0 < \ldots < \tau_n < T$ the inspection times until time T. Using the time-dependent formulation, part of the integrand in (1.12) can be rewritten:

$$Q^*(t) - q^*(t) = (Q(0) - Q^F)\left(\frac{\delta+\lambda}{\lambda}\right) \times \sum_{i=0}^{n}\left(-\frac{\delta}{\lambda}\right)^i \frac{e^{-(\frac{\delta}{\lambda}(\lambda+\delta)+R)\tau_i} - e^{-(\frac{\delta}{\lambda}(\lambda+\delta)+R)\tau_{i+1}}}{\frac{\delta}{\lambda}(\lambda+\delta)+R}. \tag{1.13}$$

Using this, we can apply a Monte Carlo simulation. We fix λ, simulate several thousand scenarios of inspection times $\tau_1, \ldots \tau_n$, and average over the outcomes. The procedure is then repeated for a range of inspection rates.

For the simulation, we choose a yearly time scale and time horizon $T = 100$. We assume that the regulator's discount rate and the firm's discount rate are $R = r = 0.1$. We set costs $c = 1$, diffusion rate $\delta = 0.2$, and finally choose the revenue function $\Pi(Q) = -0.5Q^2 + 10Q$. Recall that the parameters of the revenue function influence rather the long-run level of the solution than the actual dynamics. We simulate the parameters $0 \leq \lambda \leq 3$, which corresponds to a maximum of three inspections a year. This is sufficient to grasp the qualitative behavior of the solution. In each step we increase λ by 0.1, and run 50.000 simulations. (In the interval $[0, 1]$ we increase λ by 0.05 in each step as the solution changes quickly in this area).

To make sure that a fast closing of the quality gap does not affect consumers negatively over time, we have controlled for the offered quality discounted from the viewpoint of the regulator: $\int q(t)exp(-Rt)\ dt$, over the time interval $[0, T]$. The variation in this discounted quality is negligible for different inspection rates and therefore omitted in the presentation of the results.

The central result of the simulations is that the qualitative behavior of the solution depends crucially on the initial reputation of the firm. It is even the case that for a range of initial reputations a regulator prefers not to introduce inspections at all.

We can characterize three different types of solutions with respect to the initial reputation. We speak of (i) *low* initial reputation if $Q_0 < q^\delta$, (ii) *medium* initial reputation if $Q_0 \in [q^\delta, q^S)$, and (iii) *high* initial reputation if $Q_0 \geq q^S$. The choice will be clear from the following discussion. It relates to the fact that $Q^F \in [q^\delta, q^S)$ for all $\lambda \geq 0$, since $Q^F \to q^\delta$ for $\lambda \to 0$, and $Q^F \to q^S$ for $\lambda \to \infty$.

In (i), the initial reputation Q_0 is lower than q^δ. Proposition 3 implies that for $\lambda = 0$, the firm first raises its reputation by offering the highest possible quality. During this time, offered quality is always higher than the firm's reputation. This is beneficial for the consumers and the quality gap equals zero. Once the reputation reaches q^δ, the firm sets offered quality

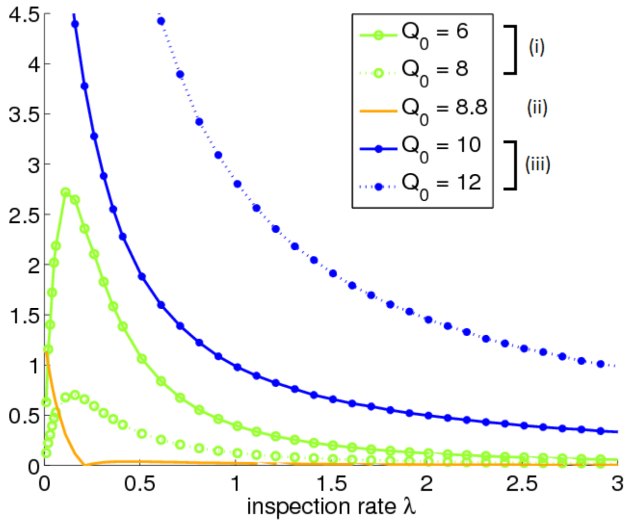

Figure 1.4: Cumulative quality gap for (i) low, (ii) medium, (iii) high initial reputation.

equal to reputation and both stay constant from now on. Hence, the quality gap remains zero for all times if $\lambda = 0$. For $\lambda > 0$, reputation and offered quality converge to the long-run optimum alternating, i.e. if reputation is high, offered quality is low and vice versa. Hence, the cumulative quality gap is greater 0 under the presence of inspections. In this case, omitting inspections is optimal from the viewpoint of a regulator.

For (ii), the initial reputation is higher than q^δ. If $\lambda = 0$, the firm chooses the lowest possible quality in the beginning and exploits its reputation according to Proposition 3. This results in a cumulative quality gap greater than 0. But now, there exists a rate λ^* for which $Q^F = Q_0$ (recall that Q^F depends, among others, on λ). Under λ^*, constant quality Q^F is offered for all times since Q^F is a fixpoint; see Proposition 4. Hence, the cumulative quality gap is minimized to zero. Deviating λ in either direction from λ^* is disadvantageous since it causes offered quality and reputation to alternate.

For (iii), there exists no $\lambda \geq 0$ for which Q^F and Q_0 coincide, since Q_0 is always greater than Q^F. Numerical simulations suggest that in this case the cumulative quality gap decreases if λ is increased. This is plausible, since (1.6) implies that a higher rate λ stabilizes the offered quality around Q^F. This reduces deviations from the long-run level. Hence, this is the only case in which a higher inspection rate is always beneficial.

The three scenarios are represented in Figure 1.4. The parameters of the simulation imply $q^\delta = 8.5$ and $q^S = 9$. The graphs correspond to the cases (i)-(iii). The reason why we depicted only one solution for case (ii) of medium initial reputation is that the variations in the cumulative quality gap are rather small and difficult to identify on this graph. Still, from Proposition 4 it is clear that for (ii) the minimizing λ^* increases as the initial reputation increases.

In Figure 1.4, the spikes for low initial reputation are striking. Their presence can be explained by recalling Corollary 5. For a low λ, convergence toward Q^F is fast, and it takes long before an inspection occurs. Until the first inspection, the quality gap equals 0 since offered quality is higher than reputation. If λ is increased, convergence slows down and the first inspection until which the quality gap is non-zero arrives faster. On the other hand, a very high λ pushes the solution stronger toward Q^F. The "pushing" effect seems to prevail over the decreasing speed of convergence if λ is sufficiently increased, and the quality gap is reduced more quickly.

The numerical studies were conducted for $\delta = 0.2$ which can be related to durable consumer goods, since the information about quality spreads slowly. If the diffusion rate δ is increased, the goods can be seen as fast-moving consumer goods. For these, the qualitative behavior depicted in Figure 1.4 remains unchanged. The only differences are that the curves are lower and flatter and that the spikes for (i) move to the right. On one hand, this implies that the effect of inspections is weaker for fast-moving consumer goods than for durable goods. On the other hand, the move of the spikes implies that the effect of an increased diffusion rate δ on the cumulative quality gap is not clear: for a fixed λ, the effect can either be positive or negative, depending on the parameters and on the initial reputation.

With respect to variation of the remaining parameters, the following is noted. Changing the discounting rate R does not change the qualitative behavior of the solution. Varying the parameters p_1, p_2, the long-run level Q^F changes. As the three scenarios described above depend on the initial reputation, raising p_1 can turn a high initial reputation scenario in a low one. The same holds true for decreasing the curvature p_2. Therefore, a parameter study on p_1, p_2 can be related to the parameter study for initial reputations, as in Figure 1.4.

Finally, we give a short remark on how the results change if the regulator minimizes $|q - Q|$ instead of $\max\{q - Q, 0\}$. This means that the regulator is reluctant to any quality deviations, even if offered quality is higher than reputation. For low initial reputation the shapes are similar to Figure 1.4; only, the quality gap does not approach 0 for $\lambda \to 0$, but a positive value. This is because now every deviation is penalized, even if offered quality exceeds reputation. Still, there exists λ that maximizes the difference. For high initial reputation the solution is similar to the just described solution for low initial reputation. This is clear because the absolute value puts the same weight on quality deviations, independent of over- or understating reputation. For medium initial reputation, the shape of the curve changes, but the minimizing rate λ^* which results in a cumulative quality gap of 0 remains the same.

1.6 Conclusion and Further Research

We have investigated how a firm manages its reputation when facing random inspections. The inspections are Poisson-distributed and reveal the quality the firm offers, which consequently influences the firm's reputation. In case of inspections-only, the firm chooses to offer a constant level of quality at all times. After introducing diffusion of information, the optimal control changes. The long-run optimum is approached by an unsteady path, as the offered quality is changed continuously in time and abruptly after an inspection. This finding helps a regulator to decide on an optimal inspection rate. The optimal rate minimizes the quality gap, that is, the time in which customers receive lower quality than they expect by looking at the firm's reputation. The results show that a higher inspection rate is not always beneficial. In particular, if the firm's initial reputation is not too high, an increased inspection rate can increase the quality gap. This result is quite counter-intuitive. It implies that policy makers have to take more than costs into account when choosing an inspection rate. In particular, if a firm has very low initial reputation and information is spread by word-of-mouth, introducing inspections might not be necessary at all.

One could think about modifying the framework by adding additional uncertainty. In between inspections, reputation is subject to small random shocks and rumors. This can be modeled by a Brownian motion (see, e.g., Chang (2004) for optimal control problems involving Brownian motion). The optimal control is similar to the one obtained here. The difference is that the (long-run) optimum is lowered by a multiplicative constant depending on the variance of the small random shocks. The drawbacks are that optimality is more difficult to prove, see Øksendal and Sulem (2004), and that the appealing time-dependent reformulation is not possible.

An extension could be to incorporate different distributions for the random inspections. In the paper, a Poisson distribution is assumed. One could make the inspection rate contingent on a second state, which captures the outcome of past inspections, or incorporate time-dependency. Unfortunately it will be difficult to derive an analytical solution even under restrictive assumptions.

Acknowledgment

The author wants to thank Franz Wirl and Klaus Wälde for discussion, and an anonymous referee for his remarks which helped improving the paper. This research was supported by the Vienna Science and Technology Fund (MA09-019).

Bibliography

Chang, F., 2004. Stochastic Optimization in Continuous Time. Cambridge University Press.

Dulleck, U., Kerschbamer, R., Mar. 2006. On doctors, mechanics, and computer specialists: The economics of credence goods. J. Econ. Literature 44 (1), 5–42.

Feichtinger, G., Hartl, R., 1986. Optimale Kontrolle ökonomischer Prozesse. de Gruyter, Berlin; New York.

Godes, D., Mayzlin, D., Fall 2004. Using online conversations to study word-of-mouth communication. Mark. Sci. 23 (4), 545–560.

Grosset, L., Roberti, P., Viscolani, B., 2011. A goodwill model with predatory advertising. Operations Res. Lett. 39, 419–422.

Kumar, P., 1980. A differential game with jump process observations. J. Optim. Theory and Appl. 31 (2), 219–231.

Nelson, P., Mar. - Apr. 1970. Information and consumer behavior. J. Political Econ. 78 (2), 311–329.

Nerlove, M., Arrow, K., 1962. Optimal advertising policy under dynamic conditions. Econometrica 29, 129–142.

Øksendal, B., Sulem, A., 2004. Applied Stochastic Control and Jump Diffusion. Springer.

Sennewald, K., 2007. Controlled stochastic differential equations under poisson uncertainty and with unbounded utility. J. Econ. Dyn. & Control 31, 1106–1131.

Sennewald, K., Wälde, K., Oct. 2006. Itô's lemma and the bellman equation for poisson processes: An applied view. J. Econ. 89 (1), 1–36.

Sethi, S., 1977. Dynamic optimal control models in advertising: a survey. SIAM Rev. 19 (4), 658–725.

Snyder, D., 1975. How to track a swarm of fireflies by observing their flashes. IEEE Trans. Inf. Theory IT-21, 692–695.

1.7 Appendix

1.7.1 Proof of Proposition 4

Proof. The firm faces the following maximization problem:

$$\max_{q\in[0,\overline{q}]} \mathbb{E}\left[\int_0^\infty \left(-\frac{p_2}{2}Q^2 + p_1 Q - cq\right) e^{-rt}\, dt\right], \tag{1.14}$$

subject to the dynamics

$$dQ = \delta(q - Q)\, dt + (q - Q)\, dJ. \tag{1.15}$$

The corresponding HJB Equation reads (see Sennewald (2007)):

$$(r + \lambda)V(Q) = \max_q\{-\frac{p_2}{2}Q^2 + p_1 Q - cq + \delta(q - Q)V'(Q) + \lambda V(q)\}. \tag{1.16}$$

We assume a quadratic value function $V(Q) = v_2 Q^2 + v_1 Q + v_0$ and use the FOC.

$$0 = -c + 2\delta v_2 Q + \delta v_1 + 2\lambda v_2 q + \lambda v_1 \tag{1.17}$$

$$q^* = \frac{c - (\lambda + \delta)v_1}{2\lambda v_2} - \frac{\delta}{\lambda}Q \tag{1.18}$$

Plugging this into the HJB - Equation (1.16), we compare coefficients and find the following:

$$v_2 = -\frac{\lambda p_2}{2\left(\delta^2 + 2\delta\lambda + \lambda^2 + \lambda r\right)}, \tag{1.19}$$

$$v_1 = \frac{\delta c + \lambda p_1}{\delta^2 + 2\delta\lambda + \lambda^2 + \lambda r}, \tag{1.20}$$

$$v_0 = \frac{(\delta c + c\lambda - \delta p_1 - \lambda p_1 + cr)^2}{2p_2 r\left(\delta^2 + 2\delta\lambda + \lambda^2 + \lambda r\right)}. \tag{1.21}$$

Therefore, q^* in (1.18) becomes

$$q^*(Q) = -\frac{\delta}{\lambda}Q + \frac{-\delta c + \delta p_1 + \lambda p_1 - c(\lambda + r)}{\lambda p_2}. \tag{1.22}$$

We reformulate the expression above. First, we define

$$Q^F = \frac{p_1 - (1 + \frac{r}{\lambda+\delta})c}{p_2}. \tag{1.23}$$

Then we note that the following term can be rewritten:

$$\frac{-\delta c + \delta p_1 + \lambda p_1 - c(\lambda + r)}{\lambda p_2} = \frac{\delta + \lambda}{\lambda}\,\frac{p_1 - (1 + \frac{r}{\lambda+\delta})c}{p_2} = \frac{\delta + \lambda}{\lambda}Q^F.$$

Therefore, the optimal control can be written as

$$q^*(Q) = -\frac{\delta}{\lambda}Q + \frac{\delta + \lambda}{\lambda}Q^F. \tag{1.24}$$

It is easy to check that $q^*(Q^F) = Q^F$. The value Q^F is the fixpoint to which the optimal control converges (see Section 1.5 of the paper). $\square$

1.7.2 Verification Theorem

The following verification theorem is found in (Sennewald, 2007, Theorem 4). Formulation and notation are adapted to our needs. Consider a stochastic optimal control problem given by

$$\max_{q\in\Theta} \mathbb{E}\left[\int_0^\infty u(Q,q)e^{-rt}\ dt\right], \quad \textbf{s.t.} \tag{1.25}$$

$$dQ = \alpha(Q,q)\ dt + \beta(Q,q)\ dJ. \tag{1.26}$$

We denote by Θ the set of feasible states and the set of feasible controls (they coincide in our model). Define for a continuous differentiable function $f(Q) : \mathbb{R} \to \mathbb{R}$ the operator

$$D^q(f(Q)) := \alpha(Q,q)f_Q(Q) + \lambda[f(Q+\beta(Q,q)) - f(Q)]. \tag{1.27}$$

Denote by Q^q the state process corresponding to the policy q.

Theorem 6. *(Sennewald, 2007, Theorem 4) Let a continuous differentiable function $V : \Theta \to \mathbb{R}$ satisfy*

$$rV(Q) \geq u(Q,q) + D^qV(Q), \quad \forall Q \in \Theta, \forall q \in \Theta, \tag{1.28}$$

and suppose in addition that there exists a policy $q^* \in \Theta$ such that

$$rV(Q) = u(Q,q^*) + D^{q^*}V(Q), \quad \forall Q \in \Theta. \tag{1.29}$$

If furthermore for all $Q \in \Theta$ the limiting condition on the expected value

$$\lim_{t\to\infty} \mathbb{E}[e^{-rt}V(Q^{q^*})] = 0, \tag{1.30}$$

and the limiting inequality

$$\lim_{t\to\infty} \mathbb{E}[e^{-rt}V(Q^q)] \geq 0, \quad \forall q(Q) \in \Theta, \tag{1.31}$$

are satisfied, then V is the value function of the optimal control problem and the policy q^* is optimal.

Proof. See (Sennewald, 2007, Theorem 4). ☐

Stochastic Modelling: Applications to Managerial Problems
Second Article
submitted to an international journal.

Chapter 2

May Increasing Doping Sanctions Discourage Entry to the Competition?

This paper studies the effect that the existence and usage of doping has upon talents who consider becoming high-performance athletes. To this end, a two-player doping model is extended by introducing an additional agent who decides whether to enter the competition or not. It is shown that the agent's willingness to invest into entry decreases if doping is present. An increase in doping sanctions can even pronounce this effect when the initial sanction level is low. It is also shown that the results can be generalized to an n player framework.

2.1 Introduction

In literature it has been shown that doping constitutes a type of prisoner's dilemma for athletes, see, e.g., Bird and Wagner (1997); Haugen (2004). Each athlete finds it optimal to dope to gain an advantage over her opponents, and eventually all athletes end up doped. Hence, nobody has gained an advantage but the additional risk of being caught makes all athletes worse off than without doping. Apart from seeing this as a mechanism that explains the existence of doping,[1] one can interpret the results in the following way: it can also be in the interest of the athletes to have policies in place that prevent doping, since they are better off if nobody dopes.

This relates to one of three reasons Buechel et al. (2013) mention, why it is socially desirable to reduce the extent of doping: first, doping can cause health problems for athletes. Second, athletes serve as role models which interferes with being doped and breaking the rules of the game. And third, closely related to the former argument, sport may become uninteresting if athletes systematically violate the rules.

In this paper, another negative aspect of doping will be highlighted: doping can constitute an entry barrier for future athletes. This is already noted in other articles, e.g., Preston and Szymanski (2003), but there costs are the key argument: athletes from rich countries can afford

[1] There exists a stream of literature that aims at understanding and explaining the phenomenon of doping by game-theoretic approaches; good literature overviews are given in Dilger et al. (2007); Buechel et al. (2013).

doping more easily, while athletes from poorer countries are disadvantaged due to missing money. We will show that even if pecuniary costs of pursuing a performance-enhancing drug are neglected, the existence and usage of doping may still keep future athletes from entering the sport. This indicates that anti-doping policies do not only affect current athletes, but also young talents that consider becoming high-performance athletes. [2]

To prove our point, we extend the performance-enhancing drug game of Haugen (2004).[3] While Haugen considers a simple and intuitive two-player game to illustrate the prisoner's dilemma, we add a third agent to the model who considers entering the competition. Entry - or, equivalent, participation - is costly since it involves training, equipment and opportunity costs, but costs are such that Agent 3 would join a fair and clean competition. Now if doping sanctions are too low, we show that the doping decision constitutes either a prisoner's dilemma or a chicken game for the competitors. Under these circumstances, the expected payoff for Agent 3 is reduced, either (i) due to an increased probability of loosing if she does not dope while her opponents do, or (ii) due to the fine when doping and being caught. If the decreased expected payoff is lower than the participation costs, the negative payoff will detain her from entering the competition.

The remainder of the article is organized as follows. In Section 2.2, we recall the two-player performance-enhancing drug game of Haugen (2004). In Section 2.3, the entry game is presented and analyzed while Section 2.4 discusses the results. Section 2.5 shows that the qualitative results also hold true for the more general case of n athletes, and presents further insights on the performance-enhancing drug game for n players. Section 2.6 concludes, and the deferred calculations are found in the appendix.

2.2 The Performance-Enhancing Drug Game

We quickly review the performance-enhancing drug game in Haugen (2004) before extending it. In this game, two athletes compete against each other for a prize a. They are equally likely to win in the first place, but both have access to a performance-enhancing drug: taking this drug - referred to as doping - ensures victory if the competitor does not take it. If both athletes dope, then their winning probabilities are equal again. There is also a risk involved in taking the drug: with probability r, a doped athlete is caught and fined an amount c.

The competition is a single shot game, that is, there is only one round played. The risk-neutral athletes decide simultaneously on whether to dope or not. Furthermore, perfect information is assumed (the agents know everything mentioned above). Denoting the strategies by **D** for doping, and **ND** for not doping, the normal form of the game looks as follows:

		Agent 2	
		D	**ND**
Agent 1	**D**	$(\frac{1}{2}a - rc, \frac{1}{2}a - rc)$	$(a - rc, 0)$
	ND	$(0, a - rc)$	$(\frac{1}{2}a, \frac{1}{2}a)$

The outcome of the game depends on the parameters, in particular on the expected sanctions rc relative to the prize a.

[2]We avoid the term "professional athlete" as there are no-media Olympic sports that are still inhabited by "amateurs" which, nevertheless, are subject to the same dynamics as professionals.

[3]This model has already been extended by Eber (2008) to incorporate fair-play norms.

- If $rc > \frac{1}{2}a$, the Nash equilibrium are the strategies $(\mathbf{ND}, \mathbf{ND})$. With high sanctions, the agents do not dope and engage in a clean competition.

- If the sanctions are low, $rc < \frac{1}{2}a$, then the Nash Equilibrium is the strategy in which both players dope, $(\mathbf{D}, \mathbf{D})$. In this case, the athletes are caught in a sort of prisoner's dilemma. Their utility is lower as if both athletes would not dope, but still, the strategic best decision is to dope.

The assumption that athletes dope if the expected benefits exceed the costs is standard in literature (e.g., Maennig (2002); Berentsen and Lengwiler (2004); Eber (2008); Frenger et al. (2011); Buechel et al. (2013); Kirstein (2014)). A recent article challenges the assumption that payoffs influence the doping behavior in the often claimed straight-forward manner: Frenger et al. (2012) provide evidence that the absolute value of prize money influences the doping affinity, but that the unequal payouts due to the final ranking (so-called spreads) have no significant influence on the doping decision. However, in a simplified game-theoretic framework it is not clear that the payoff a should be interpreted as the height of the prize spread (compared to the zero payoff in case of loosing) instead of interpreting it as the level of absolute prize money. Furthermore, the prize a does not necessarily refer to prize money only. It may be interpreted as a sum of material benefits (e.g., prize money and endorsements) and immaterial benefits (e.g., fame and prestige), which may trigger different behavior than winnings only. We therefore assume that the model at hand is a valid model for doping behavior, keeping in mind that there is research potential for more sophisticated formulations.

2.3 Entry to the Performance-Enhancing Drug Game

In the model described above, two athletes compete against each other. Now think of the two athletes as incumbents and assume that a third agent considers entering the competition. Participation in the competition is costly and incurs costs p on the entrant. This includes financial costs (travelling, buying equipment), training effort, finding sponsors, participating in preliminary competitions as well as opportunity costs.[4] Clearly, Agent's 3 decision to enter the competition depends on the actual height of these costs.

In principle, the incumbents also face participation costs that may not differ a lot from the entrant's cost. They also have to invest into training and also bear opportunity costs. We can assume their costs to be equal to p, too. However, the assumption that the incumbents have already taken their decision to participate in the competition implies that the actual height of costs does not matter. We will see that the participation costs influence the agent's payoff but do not alter their doping strategies. The participation costs are only relevant for the agent that takes the entry decision.

The model aims at approximating the viewpoint of a young talent: she is confronted with a group of high-performance athletes that compete against each other independent of her entry. She has to decide whether to join or abstain from becoming an athlete herself by reflecting on the environment she will find herself in after entry.

[4] Note that besides participation *costs*, there can also be participation *benefits*, i.e., an agent may receive pleasure just from being an athlete, independent of her performance in the competition. Without loss of generality, one can assume that these benefits are included in the costs p, respectively that participation benefits are netted against participation costs.

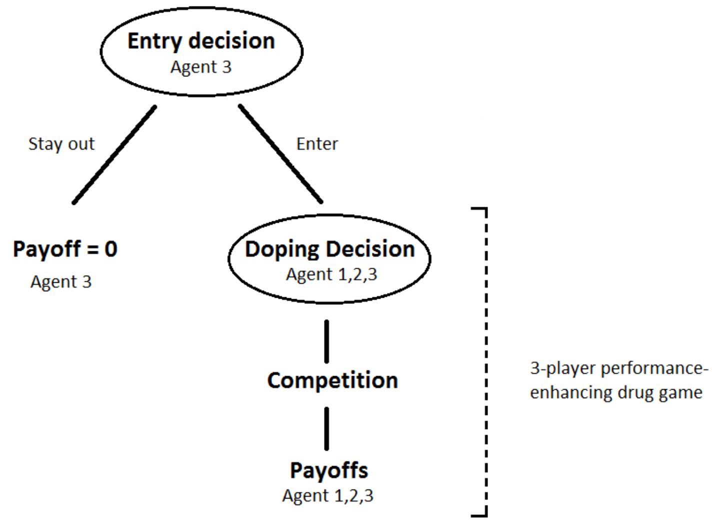

Figure 2.1: Entry game from the viewpoint of Agent 3.

As a benchmark, assume for a moment that no performance-enhancing drug exists. To take the entry decision, Agent 3 compares the expected payoff from entering the competition to the payoff when staying out. If she enters the competition, she pays participation costs p, but wins the prize a with probability $\frac{1}{3}$. Her expected payoff is given by $EP = \frac{1}{3}a - p$. If she decides to stay out, she receives an alternative payoff which we normalize to 0. Therefore, she enters the competition only if $EP > 0$, otherwise she abstains from entering. [5] This implies that Agent 3 enters the competition if the participation costs satisfy the following:

$$p < \frac{1}{3}a. \tag{2.1}$$

Now, we introduce a performance-enhancing drug. The assumptions of the two-player game in Section 2.2 are easily generalized to three athletes: using the drug ensures victory as long as no other athlete has taken it; otherwise, the winning chances are equal for the doped athletes.

The timing of the entry game is as follows: First, Agent 3 decides whether to enter the game at costs p or not. If she enters, all three athletes simultaneously decide to dope or to stay clean and then compete against each other. If Agent 3 decides not to enter, she receives payoff 0 and only the two incumbents compete. The competition between the incumbents has no effect on Agent 3 anymore (but note that it coincides with the two-person game reviewed in Section 2.2). The timing and decision structure from the viewpoint of Agent 3 are shown in Figure 2.1.

To solve the game, we look for a subgame perfect Nash equilibrium. To this end, we simply examine the outcomes of the two branches in Figure 2.1. We identify the Nash Equilibria in the three-player performance-enhancing drug game (the right branch in Figure 2.1) and compare it to the zero payoff Agent 3 receives when staying out of the game (the left branch in Figure 2.1). Agent 3 then chooses the higher payoff and takes her entry decision accordingly.

[5] More precisely, for $EP = 0$ she would be indifferent between entering the competition and staying out. We abstain from an explicit analysis of limit cases as we are more interested in general and robust results which hold for a range of parameters. However, note that these limit cases do not interfere with our results.

2.3.1 The Three-Player Performance-Enhancing Drug Game

Assume that all three athletes compete against each other while having access to a performance-enhancing drug. It is a generalization of the two-player game, and the payoff matrix is bigger now: the additional rows indicate the choice of Agent 3.

		Agent 2: **D**	Agent 2: **ND**
Agent 1:**D**	Agent 3:**D**	$(\frac{1}{3}a - rc, \frac{1}{3}a - rc, \frac{1}{3}a - rc)$	$(\frac{1}{2}a - rc, 0, \frac{1}{2}a - rc)$
	Agent 3:**ND**	$(\frac{1}{2}a - rc, \frac{1}{2}a - rc, 0)$	$(a - rc, 0, 0)$
Agent 1:**ND**	Agent 3:**D**	$(0, \frac{1}{2}a - rc, \frac{1}{2}a - rc)$	$(0, 0, a - rc)$
	Agent 3:**ND**	$(0, a - rc, 0)$	$(\frac{1}{3}a, \frac{1}{3}a, \frac{1}{3}a)$

In the table above, the participation costs p are omitted in the payoffs. They lower the actual payoffs for all agents, but as they are constant and independent of the scenario, they do not influence the choice of the strategy in the drug game. (On the other hand, they will be relevant when it comes to deciding whether to enter the competition or not).

To analyze the game, we distinguish 3 different scenarios and identify the Nash equilibria. We refer to the scenarios as follows:

1. low sanctions, if $rc < \frac{1}{3}a$,
2. medium sanctions, if $\frac{1}{3}a < rc < \frac{2}{3}a$,
3. high sanctions, if $rc > \frac{2}{3}a$.

In his paper Haugen (2004) argues that $rc << \frac{1}{2}a$ by providing examples from the U.K. premier league and the U.S National Football League as well as anecdotal evidence. The examples from team sports might not be appropriate, as in team sports different dynamics can come into action, e.g., free riding on doping team members (see Daumann (2011)). However, Daumann (2011) also provides examples from tennis, golf, cycling and marathon running and contrasts them with costs related to doping. According to his reasoning it still seems safe to assume $rc < \frac{1}{3}a$, which implies that currently low sanctions are in place. Nevertheless, the analysis is done for all scenarios to understand how different inspection rates affect the agents' behavior.

Low sanctions It is easily seen that there exists a unique Nash equilibrium in which all three agents dope, (**D**,**D**,**D**). This corresponds to a three-players prisoner's dilemma in which all athletes would be better off if nobody doped, but, still, the best strategy for every agent is to use the performance-enhancing drug.

Medium sanctions Under medium sanctions the nature of the game changes. It does not constitute a prisoner's dilemma any more. If all athletes dope, then all of them receive a negative expected payoff. Instead, there exist three pure Nash equilibria, in each of which two agents choose the same strategy. This is shown in the appendix.

Furthermore, their exist equilibria in mixed strategies. With a mixed strategy, the agent randomizes her doping behavior, e.g. she tosses a coin with certain probabilities for doping and not doping. The calculation of the probabilities is more or less straightforward and shown in the appendix. There exist four mixed equilibria. In three of them, two agents mix while the third one chooses a pure strategy. However, we want to emphasize the equilibrium in which all three players mix their strategies. The probabilities are given by

$$\text{Probability}_{\sigma}(\mathbf{D}) = 2 - 3\frac{rc}{a}, \tag{2.2}$$

$$\text{Probability}_{\sigma}(\mathbf{ND}) = 3\frac{rc}{a} - 1. \tag{2.3}$$

It is easily seen that these are indeed probabilities: the assumptions on sanctions rc in this scenario imply that (2.2) lies in the interval $[0, 1]$, furthermore, (2.2) and (2.3) add up to 1.

We denote the mixed equilibrium by ($\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$), where the strategy $\boldsymbol{\sigma}$has assigned the probabilities from (2.2) and (2.3).

High sanctions There exists a unique Nash equilibrium in which all three agents decline taking performance-enhancing drugs, (**ND**,**ND**,**ND**). This is equivalent to the scenario where no performance-enhancing drugs are available at all. We refer to a competition without any doping as *clean competition*.

These are all the Nash equilibria for the three different scenarios. In the following subsection we focus again on the multiple equilibria that arise under medium sanctions.

Fierce competition and equilibrium selection for medium sanctions

We have seen that for medium sanctions there exist as many as seven Nash equilibria. Three pure ones, three equilibria in which two athletes mix and one in which all athletes mix. The existence of multiple Nash equilibria disturbs the analysis of the game. One would prefer to have a single equilibrium at hand to understand the effects of doping on the entry decision. There are theoretical approaches on how to find the most reasonable equilibrium, see Harsanyi and Selten (1988), as well as experimental approaches to study which equilibrium is obtained. We quickly review some of the experimental works and then use the assumption of a fierce competition to rule out all but one equilibrium.

For medium sanctions, the structure of the game is a three-player generalization of the "chicken game", see, e.g., Fink et al. (1988). The chicken game constitutes a classical anti-coordination game in which all players prefer a minimum number of players to cooperate, but for each player it is optimal to not cooperate. Here, cooperating corresponds to **ND** while non-cooperation (also referred to as defection) corresponds to **D**. Concerning experimental evidence, de Heus et al. (2010) state that there is fairly little research done for chicken games (at least when comparing it to other classical games like the prisoner's dilemma). Especially for chicken games involving more than two players, literature is scarce.

Some examples of empirical findings include the importance of risk attitude and framing studied by de Heus et al. (2010). The results show that players cooperate more often if the outcome of defection is presented as additional gain, but cooperate less often if the outcome of defection is presented as loss. Holm (2000) show that in the chicken game, people, independent of their own sex, tend to behave more "hawkish" if they know that their counterpart is female. A more qualitative approach is taken by Butler et al. (2011), who conduct an experiment with several games, among them the chicken game, and afterwards interview the participants about their motives and feelings.

The question which of the equilibria is attained in the three-player drug game is closely related to the possibilities of (mis-)coordination among athletes. In experiments, there is often a coordination mechanism implemented, e.g. the possibility of one- or two-way message sending, see, e.g., Bornstein and Gilula (2003). Let us analyze if coordination for a pure Nash equilibrium is possible in the performance-enhancing drug game, and if it is desired by the athletes.

First, note that it is very likely that the athletes can communicate in some way with each other. Athletes will be familiar with their (possible) opponents, and are able to deliver at least one-way messages via the media. Clearly, none of the athletes would publicly announce that she will dope. One could only announce to never use performance-enhancing drugs at all. The question is whether this is particularly credible. It is rather the case that all athletes claim to be clean. And even if the announcement is credible, then it is not very useful for the athlete: the chicken game structure causes the opponents to exploit that knowledge immediately. All multiple equilibria consisting of pure strategies imply that at least one of the other athletes dopes. Hence, the athlete that stays clean and says so right from the beginning will surely loose the competition because one of the opponents will be doped. [6]

The same arguments hold true for mixed equilibria in which only two athletes mix their strategy. One athlete is always either better off or worse off than the mixing two athletes (at least in expectation).

Instead, the following seems more reasonable. We assume a fierce competition among athletes. Each of them is not only willing, but eager to win. Their attitude is spread among competitors and via the media; granting an opponent higher chances for winning is not an option for the athletes. The assumption excludes coordination in pure strategies as argued above, and implies that an equilibrium is attained in which all players mix. [7]

2.3.2 Entry Decision

Now that we have found the Nash equilibria of the three-player game, we can examine the entry decision of Agent 3. Her decision depends on the expected payoff from the game and on the participation costs. Staying out of the competition gives her payoff 0. Consequently, she enters the competition if she expects an amount greater than 0 from doing so. In particular, the highest participation cost Agent 3 is willing to pay equals the expected payoff from the three-player game.

Recall that Agent 3 enters a clean competition as long as the participation cost $p < \frac{1}{3}a$. We show that the presence of the performance-enhancing drug decreases her payoff from the game. Therefore, the costs she is willing to pay decrease accordingly, since she would end up with a negative outcome, otherwise. Although Agent 3 would join a clean competition at costs $\frac{1}{3}a$, she defers entry if doping sanctions are too low. We consider the decrease in acceptable participation costs as an increase in the entry barrier. In the following, we examine the entry decision for all scenarios in detail.

[6] This argument is also used in Haugen (2004). Furthermore, this is in line with the literature on inspection games in which there is no pure Nash equilibrium, see Buechel et al. (2013). If there was a pure Nash equilibrium, then there would be no uncertainty and inspections would be purely a matter of duty.

[7] One could argue against the assumption of fierce competition, claiming that there might be an outsider coordinating the actions. It may be thinkable that by oppression or coercion agents are forced into a pure equilibrium. Imaging for example the influence of the laboratory producing the performance-enhancing drug or a single official or trainer who is in charge of all three agents. On the other side, this is a very strong assumption: if it was true, then understanding the strategic decision of single athletes would by far not be the most relevant part in understanding the doping phenomenon and coining anti-doping policies.

Low sanctions In the unique Nash equilibrium $(\mathbf{D}, \mathbf{D}, \mathbf{D})$ the expected payoff for Agent 3 when joining the competition is given by

$$\frac{1}{3}a - rc - p. \tag{2.4}$$

This implies that Agent 3 enters as long as

$$p < \frac{1}{3}a - rc, \tag{2.5}$$

which ensures a positive payoff in expectation. This amount is lower than the participation costs she would accept for a clean competition. Therefore, the existence of doping increases the entry barrier.

Furthermore, the following is noteworthy. Increasing sanctions rc, the acceptable participation costs for Agent 3 decrease further, which is counterintuitive. One would expect that increasing sanctions may decrease usage of doping and encourage entry to the competition. In this scenario, this will be true for a stark increase in doping sanctions. For a mild increase of low sanctions rather the opposite can be the case: Agent 3 still has to dope to have a chance of winning but higher sanctions decrease her expected payoff. Hence increasing the sanctions only mildly can result in an increase in the entry barrier. This observation will be important when discussing the implications of the model.

Medium sanctions In the mixed equilibrium $(\boldsymbol{\sigma},\boldsymbol{\sigma},\boldsymbol{\sigma})$, the expected payoff of Agent 3 is given by

$$\frac{3(\frac{1}{3}a - rc)^2}{a} - p. \tag{2.6}$$

She would enter the competition only if the participation costs are

$$p < \frac{3(\frac{1}{3}a - rc)^2}{a}. \tag{2.7}$$

The parameter-dependent range of acceptable participation costs lies in the interval $[0, \frac{1}{3}a)$. This is lower than the participation costs of $\frac{1}{3}a$ which she would accept for entering a clean competition. In this scenario, increasing the sanctions decreases the entry barrier for new athletes.

High sanctions In this case, the unique Nash equilibrium implies a clean competition, (**ND**,**ND**,**ND**). We have already seen that Agent 3 enters the competition as long as $p < \frac{1}{3}a$. With sufficiently high sanctions, doping is prevented even if a third athlete enters the competition and the existence of drugs does not change the entry barrier.

2.4 Results and Discussion

Putting the findings together, the maximal participation costs Agent 3 is willing to pay are depicted in Figure 2.2(a). The maximal acceptable participation costs (which are equal to the expected payoff after having joined the competition) are plotted against the sanction parameters rc. For participation costs below the solid line, Agent 3 enters the competition while she abstains from entry for costs above the line. For a fair competition, Agent 3 would pay as much as $\frac{1}{3}a$,

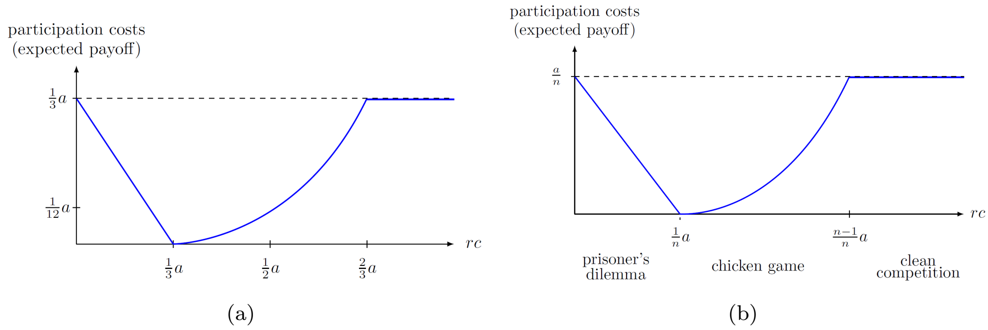

Figure 2.2: (a) Maximal participation costs Agent 3 is willing to pay to enter the competition if performance-enhancing drugs are available. The x-axis indicate sanctions, the y-axis indicate participation costs in terms of prize a. For participation costs above the solid line Agent 3 defers entry, for participation costs below the solid line Agent 3 enters the competition. The dashed line indicates the participation costs Agent 3 accepts for entering a clean competition. (b) Analog graph for $n-1$ incumbents and an n-th athlete who considers entry. The different types of games are indicated on the x-axis.

indicated by the dashed line. In the performance-enhancing drug game, the expected payoff is below $\frac{1}{3}a$ most of the time, and so is her willingness to pay.

Recall that for low sanctions, her expected profit decreases if rc increases. This is because in equilibrium all athletes dope, and higher sanctions for doping reduce the expected payoff. For medium sanctions, the mixed strategies come into play. Under these, Agent's 3 expected payoff rises when sanctions are increased, since there is less doping and hence less sanctioning to expect.

For high sanctions, Agent 3 accepts the highest costs $\frac{1}{3}a$ to enter, since no performance-enhancing drugs will be used.

This implies that the entry barrier is lowest if either sanctions are so high that any doping is prevented or so low that everybody dopes but expects no consequences. Anything in between may defer entry of new athletes. In particular, having low sanctions in place and increasing them only mildly with the intention to lower the level of doping can be counterproductive. It might not change the doping behavior of the athletes at all, but discourage agents even more from entering the competition.

2.5 Extension to n Agents

We have seen how doping influences the entry barrier in the case of two players. In particular, under the assumption of fierce competition, the barrier is always increased if doping is present. In this section we generalize the result to n players. To this end, we analyze the performance-enhancing drug game for n players. This extends the work of Haugen (2004) in the following way. We have already seen that with three-players the structure of the game changes. Whereas for two players the only outcomes are a clean competition and the prisoner's dilemma, there exists a range of parameters for which the game with three players constitutes a prisoner's dilemma. We show that the range of parameters for which the game constitutes a chicken game increases in the number of players, whereas the range of the other two scenarios gets restricted. Then, we show that the presence of doping increases the entry barrier also for n players.

2.5.1 The n-Player Performance-Enhancing Drug Game

The extension of the setting to n players is straight-forward (compare to the three-player extension in Section 2.3). Displaying the payoff matrix for n players is difficult but as the game is symmetric, the payoff structure is quite simple. Let n be the number of athletes. The expected payoff for player i, payoff$_i(\delta)$, depends on her strategy δ and on the total number of athletes doping:

# agents doping	0	1	2	...	n-1	n
payoff$_i$(**D**)		$a - rc$	$\frac{a}{2} - rc$	...	$\frac{a}{n-1} - rc$	$\frac{a}{n} - rc$
payoff$_i$(**ND**)	$\frac{a}{n}$	0	0	...	0	

Note that we have omitted the participation costs in this stage of the game (by the same reasoning as before). From the payoff structure we will deduce again three qualitatively different types of competition, depending on the sanctions rc, the prize a, and on the number of competitors n:

- a Prisoner's dilemma if $rc < a\frac{1}{n}$,
- a clean competition, if $rc > a\frac{n-1}{n}$, and
- a chicken game for the parameters in between.

Prisoner's dilemma: Consider the outcome in which all athletes dope, $(\mathbf{D}, \ldots, \mathbf{D})$. The corresponding payoff for each agent is $\frac{a}{n} - rc$. As long as $\frac{a}{n} - rc > 0$, no agent has an incentive to deviate from the outcome, since choosing the strategy **ND** would leave him with payoff 0. On the other side, starting from any other outcome, each player has an incentive to dope to increase her payoff. Therefore, $(\mathbf{D}, \ldots, \mathbf{D})$ is the only Nash Equilibrium if $rc < a\frac{1}{n}$. The expected payoff for each agent is given by

$$\text{payoff}_i = \frac{a}{n} - rc. \tag{2.8}$$

Clean competition: For a clean competition, no agent must have an incentive to deviate from the outcome $(\mathbf{ND}, \ldots, \mathbf{ND})$. The payoff for a deviating agent has to be smaller than in the non-doping outcome: $a - rc < \frac{a}{n}$ which is equal to $rc > a\frac{n-1}{n}$.

Chicken game: For the remaining parameters the resulting game constitutes an n-player chicken game. Assume an outcome in which k agents choose **D** and $n-k$ agents choose **ND**. As long as $\frac{a}{k+1} - rc > 0$, a non-doping agent has an incentive to deviate from her strategy: with **ND** her payoff is 0, while switching to **D** yields a positive payoff. Hence, the number of agents doping in equilibrium is given by k, with $0 < k < n$, where k satisfies the following:

$$\frac{a}{k} - rc > 0 \quad \text{and} \quad \frac{a}{k+1} - rc < 0. \tag{2.9}$$

Clearly, any k athletes can be the doping ones. The number of possible combinations is given by the binomial coefficient $\binom{n}{k}$, defined by

$$\binom{n}{k} := \frac{n!}{k!(n-k)!}. \tag{2.10}$$

Here, $n! = n \cdot (n-1) \ldots 2 \cdot 1$, denotes the factorial. Hence, there exist $\binom{n}{k}$ Nash equilibria in pure strategies.

These are the pure Nash equilibria. As argued in the three-player game, a solution in mixed equilibria is more realistic (under the assumption of fierce competition). We show that there always exists an equilibrium in which all players mix their strategy. A requirement for such a mixed equilibrium is that each agent is indifferent between **D** and **ND**. Let D be the probability of an agent to choose **D** (as the game is symmetric all agents will choose the same mixed strategy). Then, for an agent to be indifferent, the expected payoff from doping has to be equal to the expected payoff from not doping:

$$\sum_{k=0}^{n-1} \binom{n-1}{k} D^k (1-D)^{n-1-k} \left(\frac{a}{k+1} - rc \right) = (1-D)^{n-1} \frac{a}{n}. \tag{2.11}$$

The left-hand side sums the scenarios in which k of the remaining $n-1$ competitors dope: $\binom{n-1}{k}$ is the number of the possible outcomes, $D^k(1-D)^{n-1-k}$ is the probability of each outcome occurring, and $\frac{a}{k+1} - rc$ is the payoff for the agent in each outcome. The right-hand side yields the payoff for a non-doping agent: if the remaining $n-1$ agents do not dope, her payoff is $\frac{a}{n}$, otherwise her payoff is zero.

Equation (2.11) can be rewritten as

$$\frac{rc}{a} = \frac{1}{nD} \left(1 - (1-D)^{n-1}\right). \tag{2.12}$$

This is shown in the appendix. The right-hand side yields a polynomial in D of degree $n-1$ and there does not exist a closed-form solution for D. However, one can show that the right-hand side is downwards sloping for $D \in (0,1]$, and that the values range from $a\frac{n-1}{n}$ to $a\frac{1}{n}$. Therefore, there exists a unique solution for D (and hence a unique mixed equilibrium) if $a\frac{1}{n} \leq \frac{rc}{a} \leq a\frac{n-1}{n}$.

As in the mixed equilibrium the expected payoffs for doping and for not doping coincide, we can read the payoff for the mixed strategy outcome directly off from Equation 2.11 (i.e. by looking at the right-hand side):

$$\text{payoff}_i = \frac{a}{n}(1-D)^{n-1}. \tag{2.13}$$

2.5.2 Entry Decision

Now assume that there are $n-1$ incumbents in the competition and a n-th agent considers entering. The entry decision depends on the actual participation costs p, as well as on the scenario defined above. If no performance-enhancing drug exists, then the agent can expect to win with probability $\frac{1}{n}$, hence she would be willing to enter the competition as long as $E < \frac{a}{n}$. This is our reference scenario. If a performance-enhancing drug exists, we can group the results according to the scenarios defined before:

Prisoner's dilemma: The expected payoff for the n-th player for the outcome $(\mathbf{D}, \ldots, \mathbf{D})$ is given by

$$\text{payoff}_n = \frac{a}{n} - rc - p. \tag{2.14}$$

Hence, the agent enters only if $p < \frac{a}{n} - rc$. This is clearly lower than the participation costs she accepts for a clean competition.

Clean competition As the name already indicates, the outcome is $(\mathbf{ND}, \dots, \mathbf{ND})$ and hence the agent's payoff is

$$\text{payoff}_n = \frac{a}{n} - p. \tag{2.15}$$

Under high sanctions, the payoff – and therefore the entry barrier – remains unchanged even with the existence of drugs.

Chicken game In this case, the expected payoff for Agent n is given by

$$\text{payoff}_n = \frac{a}{n}(1 - D)^{n-1} - p. \tag{2.16}$$

The payoff is lower than in a clean competition since $(1 - D) < 1$, but the actual difference depends on the doping probability D. From (2.12) it follows that for an increasing sanction rc the probability D decreases which in turn increases the payoff.

2.5.3 Results and Discussion

This section shows that the findings derived for the three-player game generalize to the n-player game. The results are displayed in Figure 2.2(b). Depending on sanctions rc, the types of games are indicated. Again, the solid line shows the expected payoff for the entering agent. For participation costs below the line, the Agent enters the competition, and stays out for costs above the line. The acceptable participation costs for a clean competition are indicated by the dashed line. It is seen that doping increases the entry barrier, as long as sanctions are not high enough to ensure a clean competition. Remarkably, for very low sanctions, a modest increase in fines or inspection rates increases the entry barrier, as the expect payoff is decreased due to sanctions. Only above a certain threshold, higher sanctions lower the entry barrier again.

From the analysis above one can also gain insights into the performance-enhancing drug game of Haugen (2004) when extended to n players:

- For more than two players, the type of game does not necessarily constitute a prisoner's dilemma or a clean competition any more, but possibly a chicken game. Even more, the higher the number of athletes n, the smaller the range of parameters that constitute a prisoner's dilemma.
- Since a competition is only clean if $a\frac{n-1}{n} < rc$, we see that the minimum sanctions for a clean competition increase in the number of competitors. This is in line with the findings of Ryvkin (2013). It implies that doping can occur in an initially clean competition not only if sanctions are lowered but also if either more athletes participate in the competition or if the prize a is raised.
- On the other hand, if the expected sanctions exceed the prize, that is $a - rc < 0$, then doping would never occur, independent of the number of players. However, it is doubtable that such sanctions can be implemented if there is limited liability of athletes.
- In the mixed equilibriums, the doping probability increases with prize a and decreases with sanctions rc: this follows immediately from Equation (2.12) and the considerations below. The change for an increasing number of competitors is not clear: for high sanctions, an increase in the number n decreases the individual doping probability, whereas for low sanctions, the individual doping probability is increased if more players compete.

2.6 Conclusion

The aim of this paper is to study the effect that doping policies have upon "future" athletes: does the existence of doping defer entry of new athletes to the sport? Whereas the two-player game of Haugen (2004) implies that the competing athletes can be worse off due to doping if sanctions are too low, this paper indicates that with low sanctions, future athletes are also negatively affected: if doping is present, they have to dope as well to keep their chance of winning, but the risk of being caught decreases their payoff and therefore their willingness to invest into a career in sports. The same mechanism implies that if the expected punishment increases but is still too low to change the doping behavior of the agents, then the readiness to become a high-performance athlete decreases even further. This should be kept in mind when coining anti-doping policies. In particular, the model relates doping policies to taking responsibility towards a future generation of athletes. Additionally, if talented athletes decide to stay out of sports, excitement for fans may be reduced due to a decreased number of competitors.

Acknowledgments

The author wants to thank Wieland Müller for discussion as well as two anonymous referees for their valueable comments.

Bibliography

Berentsen, A., Lengwiler, Y., 2004. Fraudulent accounting and other doping games. Journal of Institutional and Theoretical Economics 160 (3), 402–415.

Bird, E., Wagner, G., 1997. Sport as a common property ressource - a solution to the dilemmas of doping. Journal of Conflict Resolution 41 (6), 749–766.

Bornstein, G., Gilula, Z., 2003. Between-group communication and conflict resolution in assurance and chicken games. Journal of Conflict Resolution 47 (3), 326–339.

Buechel, B., Emrich, E., Pohlkamp, S., 2013. Nobody's innocent: the role of customers in the doping dilemma. Tech. rep., MPRA (Munich Personal RePEc Archive), no. 44627, posted 27. February 2013 12:57 UTC, available at: `http://mpra.ub.uni-muenchen.de/44627/`.

Butler, D., Burbank, V., Chisholm, J., 2011. The frames behind the games: Players perceptions of prisoners dilemma, chicken, dictator, and ultimatum games. The Journal of Socio-Economics 40 (2), 103–114.

Daumann, F., 2011. Doping in elite sport from the economic perspective. In: Sport and Doping. The Analysis of an Antagonistic Symbiosis. Lang, Peter Frankfurt, pp. 55–70.

de Heus, P., Hoogervorst, N., van Dijk, E., 2010. Framing prisoners and chickens: Valence effects in the prisoners dilemma and the chicken game. Journal of Experimental Social Psychology 46 (5), 736–742.

Dilger, A., Frick, B., Tolsdorf, F., 2007. Are athletes doped? Some theoretical arguments and empirical evidence. Contemporary Economic Policy 25 (4), 604–615.

Eber, N., 2008. The Performance-Enhancing Drug Game Reconsidered - A Fair Play Approach. Journal of Sports Economics 9 (3), 318–327.

Fink, E., Gates, S., Humes, B., 1988. Game Theory Topics. Incomplete Information, Repeated Games and N-Player Games. Quantitative Applications in the Social Sciences. Sage Publications, Inc.

Frenger, M., Pitsch, W., Emrich, E., 2011. The impact of anti-doping legislation in europe - outlines for the development of model-based hypotheses. In: Sport and Doping. The Analysis of an Antagonistic Symbiosis. Lang, Peter Frankfurt, pp. 71–100.

Frenger, M., Pitsch, W., Emrich, E., 2012. Erfolg(+)reich und verdorben. Eine empirische Überprüfung weit verbreiteter Vorurteile im Sport. Sportwissenschaft 42 (3), 188–201.

Harsanyi, J., Selten, R., 1988. A General Theory of Equilibrium Selection in Games. MIT Press.

Haugen, K. K., 2004. The Performance-Enhancing Drug Game. Journal of Sports Economics 5 (1), 67–86.

Holm, H., 2000. Gender-based focal points. Games and Economic Behavior 32 (2), 292–314.

Kirstein, R., 2014. Doping, the inspection game, and bayesian enforcement. Journal of Sports Economics 15 (4), 385–409.

Maennig, W., 2002. On the economics of doping and corruption in international sports. Journal of Sports Economics 3 (1), 61–89.

Preston, I., Szymanski, S., 2003. Cheating in contests. Oxford Review of Economic Policy 19 (4), 612–624.

Ryvkin, D., 2013. Contests with Doping. Journal of Sports Economics 14 (3), 253–275.

parameters	type	equilibrium
$\frac{1}{3}a < rc < \frac{1}{2}a$	pure	($\mathbf{D}$,$\mathbf{D}$,$\mathbf{ND}$) a.p.
	2 players mix	($\boldsymbol{\delta}_0$,$\boldsymbol{\delta}_0$,$\mathbf{D}$) a.p.
	3 players mix	($\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$)
$\frac{1}{2}a < rc < \frac{2}{3}a$	pure	($\mathbf{D}$,$\mathbf{ND}$,$\mathbf{ND}$) a.p.
	2 players mix	($\boldsymbol{\delta}_1$,$\boldsymbol{\delta}_1$,$\mathbf{ND}$) a.p.
	3 players mix	($\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$,$\boldsymbol{\sigma}$)

Table 2.1: Nash Equilibria for the scenario of medium sanctions. Greek letters denote mixed strategies. Permutation of strategies are indicated by "a.p.", for example ($\mathbf{D}$,$\mathbf{D}$,$\mathbf{ND}$) a.p. includes all permutations, namely ($\mathbf{D}$,$\mathbf{ND}$,$\mathbf{D}$) and ($\mathbf{ND}$,$\mathbf{D}$,$\mathbf{D}$).

2.7 Appendix

2.7.1 Nash Equilibria for Three Players under Medium Sanctions

The pure equilibria follow readily from the table in Section 2.3.1. For the scenario of medium sanctions, there are additional equilibria in which two or three agents mix their strategies. The equilibria depend on the actual parameters and are shown in Table 2.1. Recall that the game is symmetric, and therefore all permutation of strategies are equilibria. $\boldsymbol{\delta}_0$,$\boldsymbol{\delta}_1$,$\boldsymbol{\sigma}$ denote mixed strategies; the corresponding probabilities of choosing $\mathbf{D}$ under strategy δ are denoted by $p_\delta(\mathbf{D})$ and are given by:

$$p_{\boldsymbol{\delta}_0}(\mathbf{D}) = 3 - 6\frac{rc}{a}, \tag{2.17}$$

$$p_{\boldsymbol{\delta}_1}(\mathbf{D}) = 4 - 6\frac{rc}{a}, \tag{2.18}$$

$$p_{\boldsymbol{\sigma}}(\mathbf{D}) = 2 - 3\frac{rc}{a}. \tag{2.19}$$

In the following, all possible outcomes are examined and the results shown in Table 2.1 are derived.

Two agents mix, remaining agents plays fair

We show that there exists an equilibrium in which two agents choose the strategy $\boldsymbol{\delta}_1$, while the remaining agent plays fair. Assume without loss of generality that Agent 3 plays fair, while Agents 1,2 mix their strategies. As the game is symmetric, the mixing probability is the same for Agents 1 and 2. Let D denote the probability of doping for a mixing agent. In equilibrium, a mixing agent has to be indifferent between doping and playing fair. Hence, for $i = 1, 2$ the following equation has to hold true:

$$\text{payoff}_i(\mathbf{D}) = \text{payoff}_i(\mathbf{ND}) \tag{2.20}$$

$$\left(\frac{1}{2}a - rc\right) D + (a - rc)(1 - D) = \frac{1}{3}a(1 - D) \tag{2.21}$$

$$-\frac{1}{2}D + (a - rc) = \frac{1}{3}a - \frac{1}{3}aD \tag{2.22}$$

$$-\frac{1}{6}D = -\frac{2}{3}a + rc \tag{2.23}$$

$$D = 4 - 6\frac{rc}{a}. \tag{2.24}$$

In order for D to be a probability, the following has to hold:

$$0 \leq 4 - 6\frac{rc}{a} \leq 1. \tag{2.25}$$

This is equivalent to

$$3 \leq 6\frac{rc}{a} \leq 4 \tag{2.26}$$

$$\frac{1}{2}a \leq rc \leq \frac{2}{3}a. \tag{2.27}$$

Hence, this equilibrium is only possible under the parameter restrictions of (2.27).

Let us check that this is really an equilibrium, i.e., that no agent has an incentive to deviate. The mixing agents have no reason to change their strategy, since their payoff for playing **D** or **ND** is the same by construction (see (2.20)). We only have to check that Agent 3 cannot do better by choosing **D** instead of **ND**. When choosing **ND**, her expected payoff is given by

$$(1-D)^2\frac{a}{3} = \frac{3(a-2\mathrm{rc})^2}{a}. \tag{2.28}$$

If she dopes, her expected payoff is:

$$D^2\left(\frac{a}{3} - rc\right) + 2D(1-D)\left(\frac{a}{2} - rc\right) + (1-D)^2(a - rc) = \frac{7a}{3} - 11\mathrm{rc} + \frac{12\mathrm{rc}^2}{a}. \tag{2.29}$$

Subtracting (2.28) from (2.29) yields

$$\frac{3(a-2\mathrm{rc})^2}{a} - \left(\frac{7a}{3} - 11\mathrm{rc} + \frac{12\mathrm{rc}^2}{a}\right) = \frac{2a}{3} - \mathrm{rc} \geq 0, \tag{2.30}$$

as long as (2.27) holds.

We conclude that no agent has an incentive to deviate and that the outcome is an equilibrium, as claimed. Furthermore, the analysis implies that for the parameter restrictions (2.27) this is the only equilibrium in which two agents mix their strategies.

Two agents mix, remaining agent dopes

Similarly, we show that two agents choose $\boldsymbol{\delta}_0$while the remaining agent dopes. Assume without loss of generality that Agent 3 plays the pure strategy **D**, while Agents 1,2 mix their strategies. In the mixed strategy, D denotes again the probability of doping. Then for $i = 1, 2$ it follows:

$$\mathrm{payoff}_i(\mathbf{D}) = \mathrm{payoff}_i(\mathbf{ND}) \tag{2.31}$$

$$\left(\frac{1}{3}a - rc\right)D + \left(\frac{1}{2}a - rc\right)(1-D) = 0 \tag{2.32}$$

$$-\frac{1}{6}D = \frac{1}{2}a - rc \tag{2.33}$$

$$D = 3 - 6\frac{rc}{a}. \tag{2.34}$$

Again, we need to check under which conditions D yields a probability.

For D to be a probability, the following has to hold:

$$0 \leq 3 - 6\frac{rc}{a} \leq 1 \tag{2.35}$$
$$\frac{1}{3}a \leq rc \leq \frac{1}{2}a. \tag{2.36}$$

As in Section 2.7.1, we show again that no agent has an incentive to deviate. We only need to check that the agent playing **D** cannot do better by playing **ND**. If she dopes, her payoff is given by:

$$D^2\left(\frac{a}{3} - rc\right) + 2D(1-D)\left(\frac{a}{2} - rc\right) + (1-D)^2(a - rc) = a - 7\text{rc} + \frac{12\text{rc}^2}{a}. \tag{2.37}$$

If she plays fair, her payoff is given by:

$$(1-D)^2\frac{a}{3} = \frac{4(a - 3\text{rc})^2}{3a}. \tag{2.38}$$

Subtracting the two yields:

$$a - 7\text{rc} + \frac{12\text{rc}^2}{a} - a - 7\text{rc} + \frac{12\text{rc}^2}{a} = rc - \frac{a}{3} > 0, \tag{2.39}$$

if (2.36) holds true. This shows that the outcome indeed constitutes an equilibrium. This is the only equilibrium in which two agents mix under the parameter restrictions (2.36).

All three agents mix

In an equilibrium where all three agents mix, all of them have to be indifferent between doping and playing fair. Denote by D the probablity that a agent dopes in the mixed strategy. Then payoff$_i$(**D**) has to be equal to payoff$_i$(**ND**) for $i = 1, 2, 3$:

$$\left(\frac{1}{3}a - rc\right)D^2 + 2\left(\frac{1}{2}a - rc\right)(1-D)D + (a - rc)(1-D)(1-D) =$$
$$-2D(1-D) + \left(\frac{1}{3}a\right)(1-D)(1-D). \tag{2.40}$$

Solving for D yields,

$$D = 2 - \frac{3rc}{a}. \tag{2.41}$$

For D to act as probabilities the following has to hold:

$$0 \leq 2 - 3\frac{rc}{a} \leq 1 \tag{2.42}$$
$$1 \leq 3\frac{rc}{a} \leq 2 \tag{2.43}$$
$$\frac{1}{3}a \leq rc \leq \frac{2}{3}a. \tag{2.44}$$

Since by construction no athlete has an incentive to deviate, this is indeed a Nash Equilibrium. Furthermore, it is the only equilibrium in which all three agents mix their strategies.

2.7.2 Mixed Nash Equilibrium for $n \geq 3$ Players

To obtain (2.12) from (2.11), first, note that the left-hand side of (2.11) can be rewritten:

$$a\left[\sum_{k=0}^{n-1}\binom{n-1}{k}D^k(1-D)^{n-1-k}\frac{1}{k+1}\right]-rc. \tag{2.45}$$

The coefficients of the series are given by:

$$\begin{aligned}
&\binom{n-1}{k}\frac{1}{k+1} && (2.46)\\
&=\frac{(n-1)!}{k!(n-1-k))!}\frac{1}{k+1} && (2.47)\\
&=\frac{(n-1)!}{(k+1)!(n-1-k))!} && (2.48)\\
&=\frac{1}{n}\binom{n}{k+1}. && (2.49)
\end{aligned}$$

Therefore, we can rewrite the sum as

$$\begin{aligned}
\left[\sum_{k=0}^{n-1}\binom{n-1}{k}D^k(1-D)^{n-1-k}\frac{1}{k+1}\right] &= \frac{1}{n}\sum_{k=0}^{n-1}\binom{n}{k+1}D^k(1-D)^{n-1-k} && (2.50)\\
&=\frac{1}{n}\sum_{k=1}^{n}\binom{n}{k}D^{k-1}(1-D)^{n-k} && (2.51)\\
&=\frac{1}{nD}\sum_{k=1}^{n}\binom{n}{k}D^{k}(1-D)^{n-k} && (2.52)\\
&=\frac{1}{nD}\left[\sum_{k=0}^{n}\binom{n}{k}D^{n-k}(1-D)^{k}-(1-D)^n\right] && (2.53)\\
&=\frac{1}{nD}\left[1-(1-D)^n\right]. && (2.54)
\end{aligned}$$

For the last equality we have used the fact that $\sum_{k=0}^{n}\binom{n}{k}D^{n-k}(1-D)^k=(D+(1-D))^n=1$. Now plugging this into (2.11) we obtain:

$$\frac{a}{nD}\left[1-(1-D)^n\right]-rc=\frac{a}{n}(1-D)^{n-1}, \tag{2.55}$$

which can be manipulated further:

$$\begin{aligned}
\frac{a}{nD}\left[1-(1-D)^n\right]-\frac{a}{n}(1-D)^{n-1} &= rc && (2.56)\\
\frac{a}{n}\left[\frac{1}{D}-\frac{(1-D)^n}{D}-(1-D)^{n-1}\right] &= rc && (2.57)\\
\frac{a}{n}\left[\frac{1}{D}-(1-D)^{n-1}\left(\frac{1-D}{D}+1\right)\right] &= rc && (2.58)\\
\frac{a}{n}\left[\frac{1}{D}-(1-D)^{n-1}\frac{1}{D}\right] &= rc && (2.59)\\
\frac{a}{n}\frac{1}{D}\left[1-(1-D)^{n-1}\right] &= rc, && (2.60)
\end{aligned}$$

which yields (2.12) as claimed.

2.7.3 Unique Solution for D

Consider Equation (2.12). For $D \to 0$, the right-hand side converges to $\frac{n-1}{n}$ while for $D = 1$ the right-hand side equals $\frac{1}{n}$. Since the right-hand side is continuous, a solution in these interval exists. Is it unique? To this end, we show that the function is downward sloping by proving that its derivative is negative (for $D \in (0, 1]$. Consider the derivative of the right-hand side.

$$\frac{d}{dD}\left(\frac{1}{nD}\left[1-(1-D)^{n-1}\right]\right) = \frac{-1+(1-D)^{-1+n}+(1-D)^{-2+n}D(-1+n)}{D^2 n}. \tag{2.61}$$

We want to show that the derivative is negative for $D \in (0, 1]$, since then the Equation (2.12) has only one solution which implies that there exists a unique mixed strategy. For the term to be negative, the nominator has to be smaller than 0.

$$-1+(1-D)^{-1+n}+(1-D)^{-2+n}D(-1+n) < 0 \tag{2.62}$$

$$(1-D)^{-1+n}+(1-D)^{-2+n}D(-1+n) < 1 \tag{2.63}$$

This is indeed the case, since for the left-hand side of (2.63) it holds true that it tends to 1 as $D \to 0$ and tends to 0 as $D \to 1$. In between, the function is decreasing for $n \geq 3$, which is seen by considering the derivative:

$$-(1-D)^{-3+n}D(-2+n)(-1+n) < 0, \tag{2.64}$$

for $0 < D < 1$. Therefore, we have shown that a unique mixed equilibrium exists.

Stochastic Modelling: Applications to Managerial Problems
Third Article
submitted to an international journal.
with D.Wozabal and C.Graf

Chapter 3

The Effect of Intermittent Renewables on the Electricity Price Variance

The dominating view in the literature is that renewable electricity production increases the price variance on spot markets for electricity. In this paper, we critically review this hypothesis. Using a static market model, we identify the variance of the infeed from intermittent electricity sources (IES) and the shape of the industry supply curve as two pivotal factors influencing the electricity price variance. The model predicts that the overall effect of IES infeed depends on the produced amount: while small to moderate quantities of IES tend to decrease the price variance, large quantities have the opposite effect. In the second part of the paper, we test these predictions using data from Germany, where investments in IES have been massive in the recent years. The results of this econometric analysis largely conform to the predictions from the theoretical model. Our findings have wide-ranging implications for policy makers coordinating subsidy schemes for renewables, flexible production capacities, and electricity storages.

3.1 Introduction

In the last decades electricity markets were liberalized in many countries, most of which established electricity exchanges where electricity is traded on an hourly basis. Since the possibilities of storing electricity are limited and demand is highly inelastic, electricity prices exhibit strong daily, weekly, and yearly seasonalities which are mainly driven by fluctuating demand. This peculiar property of electricity is shared by no other commodity and prompted active research investigating the competitiveness of the markets, the statistical properties of the price process, and the determinants of price formation.

In the last years the debate centered on the transformation of electricity production driven by the goal to decrease the carbon footprint of the industry. For the understanding of the dynamics of electricity prices, the most interesting aspect of this transition is the interplay of large intermittent production capacities with the above mentioned characteristics of electricity demand and insufficient storage capacities.

The impact of intermittent electricity sources (IES) on electricity prices is investigated in many papers. Most of them focus on the price level and it is widely agreed on that an increase in IES reduces the average spot prices (see, e.g., Sensfuß et al., 2008; Green and Vasilakos, 2010; Jacobsen and Zvingilaite, 2010; Woo et al., 2011; Paraschiv et al., 2014). In this study, we focus on another aspect: we assess the impact of intermittent production on the variance of electricity prices.

The price variance is an important quantity for several reasons. Its evolution is a major determinant when investing in plants, for which the return is directly connected to the variation of prices such as electricity storage plants. These plants earn money by buying electricity when prices are high and selling it when prices are low. The higher the price variance, and therefore the difference between high and low prices, the more profitable this strategy is. Hence, high price variance directly translates to higher earnings of storages and therefore higher incentives to invest in this type of plants (see Muche, 2009; Dunne and Mu, 2010). Likewise the returns on investments in smart grid infrastructure depend heavily on the variance of prices. The higher the price variance, the higher the incentive to invest in smart grid and metering technology that smoothens the demand. Since these technologies are necessary to deal with the challenges induced by an increased share of IES production, understanding the mechanisms that drive price variance on electricity exchanges is of importance to academics, policy makers, and investors alike.

The extant literature is not clear on the issue of price variance for two reasons. First, many studies refer to price variability, a vague term for which many possible definitions are used, e.g., the range of observable prices, the frequency of price spikes, the price variance, or the volatility. Second, findings are contradictory. Many authors predict an increase in price variability due to IES. However, some empirical works find a reduced price variability of power prices as the amount of intermittent production increases. In this paper, we introduce a theoretical framework that reconciles these seemingly contradictory findings and is consistent with the data observed on the German electricity market.

Among the theoretical works predicting an increased price variability is Milstein and Tishler (2011) who consider a two stage game for electricity producers. In the first stage, energy producers decide which share of their production portfolio will be intermittent. In the second stage, energy is produced accordingly, but with some probability the IES cannot be exploited which causes price spikes due to inelastic demand. This implies that build-up IES can increase price volatility. Green and Vasilakos (2010) simulate the impact of wind generation on the spot market prices in the UK for the year 2020. They rely on real world wind data, information on existing and planned wind farms, and demand data. Assuming a competitive market, they simulate monthly price distributions. The results imply that volatility of prices increases in future and that the magnitude of the increase depends on the competitiveness of the market. In Chao (2011) a central planer chooses a production portfolio and sets prices as to maximize expected social welfare. Stochastic demand depends on prices, and supply is subject to random outages or intermittence. For two different pricing schemes, scenarios with and without IES are compared. While average prices decrease due to IES, the load-weighted standard deviations increase.

Other papers focus on historical data and find that price variance has increased due to IES. Woo et al. (2011) examine the effect of wind generation on spot prices in Texas by means of a (partial-adjusted) linear regression model. By extrapolation of their models they argue that price variance tends to increase, if wind generation is raised. More recently, Ketterer (2012) analyzed the German electricity market by using a GARCH model. She finds that an increased production from wind reduces the mean price but increases the variance while an increased load shows the opposite effects.

The conclusions concerning price variability are ambiguous in Jónsson et al. (2010) who study the effects of wind energy forecasts on prices in Denmark. They consider different levels of wind penetration, that is, the percentage of electricity produced by wind turbines. Their findings suggest that higher predicted wind penetration lowers the mean spot price and reduces the intra-day variance of the prices. Still, they conclude that price volatility will increase in future, due to the varying price patterns induced by different wind penetration scenarios. Jacobsen and Zvingilaite (2010) also study the current market situation in Denmark on a broad level and focus on the large share of IES. With respect to prices, they conclude that prices fluctuate more but that the frequency of peak prices is reduced.

Tveten et al. (2013) report that photovoltaic (PV) infeeds in Germany have decreased price variance substantially from 2009–2011. They compare days with similar electricity demand and different levels of PV production and find that on days with high PV infeeds the price variance is lower. This is attributed to the merit order effect: under high PV power penetration, the production mix is changed, i.e., peak power plants with high short-run marginal costs are used less. Therefore, prices stay at a lower level and do not jump up to peak-load prices, which reduces fluctuations.

The merit order effect also forms the base of this paper, but we study it more systematically and in a broader context for wind and PV energy. We contribute to the discussion by showing that IES can either increase or decrease price variance. For our theoretical investigations, we use a static market model with random residual demand influenced by IES infeed. We identify two factors driving the price variance: (1) the curvature of the supply function (merit order), and (2) the distribution of the random IES supply. To test our theoretical findings, we perform statistical analysis on daily price variances from Germany and are able to verify the existence of the two effects. Furthermore, our statistical analysis provides evidence that the merit order in the German market is "inverse S-shaped."

Our choice of the variance as a measure of dispersion of electricity prices is motivated by its nice theoretical properties: both the market model as well as the decomposition of the price variance into its components in the empirical part of the paper hinge on that choice. Other measures of variation such as the mean absolute deviation or the standard deviation do not permit a similar analysis. We therefore choose the variance over its alternatives, although it presents us with certain challenges in the econometric analysis due to its sensitivity to outliers.

The remainder of this paper is structured as follows. In Section 2, we detail the theoretical approach, identifying the effects which influence the price variance. In Section 3, we analyze the effects of increased intermittent electricity production on the variance of electricity prices in an econometric model. Section 4 is devoted to the policy implications of our findings, while Section 5 concludes the paper.

3.2 Theory

In the following, we provide a theoretical framework explaining how IES infeeds affect the price variance. We first explain the market model and then give an intuitive idea of two opposing effects that intermittent production has on the distribution of the electricity prices. We conclude with a formal analysis, which identifies the relevant factors influencing the price and show that an increased IES infeed does not necessarily increase the price variance, but may even lower it.

Our investigations are based on a standard static market model with inelastic, random aggregate demand for electricity Q. As common in the literature, we start with the assumption of a competitive convex market supply curve S, which equals the aggregate industry cost function (see e.g., Barlow, 2002; Kim and Knittel, 2006; Redl et al., 2009; Arocena et al., 2012). The assumption of convexity will be relaxed later on to implicitly incorporate dynamic aspects of the plant dispatch problem of electricity producers. The market price is obtained as $S(Q)$, i.e., by intersecting the inelastic demand with the supply curve. We therefore implicitly assume that no player is able to exercise market power.

In the setting outlined above, the most natural way to include intermittent production would be to change the supply function. However, for the purpose of our model, we find it more convenient to subtract the intermittent production I from the demand. In this way we obtain the so called *residual demand* $Q - I$, which is the demand that has to be covered by conventional plants. The obvious advantage of this kind of modeling is that all the stochastic quantities, i.e., the demand and the intermittent infeed, enter the residual demand while the supply curve remains deterministic. Incorporating IES via a change in the supply function or via the residual demand is equivalent, if intermittent production is always fed-in regardless of the price, i.e., if it merely shifts the demand curve to the left. Indeed, typically intermittent sources of electricity have marginal costs close to zero and therefore supply all the energy they produce into the grid as long as prices are positive. Additionally, many countries subsidize intermittent renewable production by guaranteed feed-in tariffs which also include preferential access to the market. Therefore it is quite standard to model renewable production via residual demand (see e.g., Sensfuß et al., 2008; Nicolosi, 2010; Graf and Wozabal, 2013; Hirth, 2013; Wagner, 2014).

Figure 3.1 intuitively explains how an increase in IES infeed affects the market price in our model. The x-axis indicates the quantity of electricity and the y-axis the price. The densities on the x-axis represent the stochastically varying residual demands $X = Q - I$, which are mapped to the stochastic prices $P = S(X)$ whose densities are drawn over the y-axis. In Figure 3.1(a), the increased amount of IES shifts the residual demand to the left to a region where the supply curve is flatter. If the shape of the residual demand distribution remains unchanged, the variance of the prices *decreases*. If, however, an increased share of intermittent production changes the shape of the distribution of X, then the situation is ambiguous. For example, if the density of the residual demands broadens by an increase in intermittent production, then the variance could increase, even though the residual demand is pushed into the flat area of the supply function as is illustrated in Figure 3.1(b). Hence, the overall effect of intermittent production depends crucially on the shape of the distribution of the residual demand and on the slope of the supply function. These considerations imply that the build-up of IES can either increase or decrease price variance, depending on which effect has a stronger influence.

We state this in a more rigorous way: Let the supply function $S : \mathbb{R} \to \mathbb{R}$ be increasing,

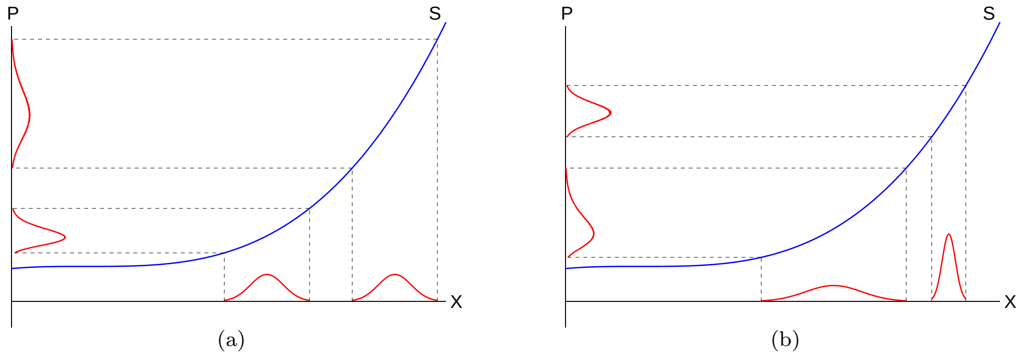

Figure 3.1: Panel (a): Shift of the residual demand with otherwise constant distribution. Panel (b): Shift of the residual demand with broadened distribution.

convex, and differentiable. We compare two scenarios with different shares of intermittent energy sources. In the first scenario the residual demand is given by X with price $S(X)$. In the second scenario, we assume a larger share of IES and denote the additional random IES infeed by $I' \geq 0$ and the resulting residual demand by $Y = X - I'$. We start by investigating the case where the distribution of Y is not broader than the distribution of X, which we formalize as

$$F_Y^{-1}(\beta) - F_Y^{-1}(\alpha) \leq F_X^{-1}(\beta) - F_X^{-1}(\alpha), \quad \forall \alpha, \beta : 0 \leq \alpha < \beta \leq 1, \tag{3.1}$$

i.e., that the interquantile ranges of X are as least as wide as that of Y. It is shown in Lemma 3 (see Appendix 3.6) that (3.1) implies $\operatorname{Var}(Y) \leq \operatorname{Var}(X)$ (while the opposite is not necessarily true). We modify a theorem of See and Chen (2008) and obtain the following result, which we prove in the appendix.

Proposition 1. *Let X and Y be non-negative random variables with inverse distribution functions F_X^{-1}, F_Y^{-1} respectively, and let $S : \mathbb{R}^+ \to \mathbb{R}$ be a convex increasing function which is twice differentiable. If* (3.1) *holds and $Y \leq X$, then*

$$\operatorname{Var}(S(Y)) \leq \operatorname{Var}(S(X)). \tag{3.2}$$

The proposition states under which conditions *more* IES leads to *less* price variance. Looking at the two conditions, we see that $Y \leq X$ holds by definition, but Condition (3.1) need not hold in practice. It is fulfilled for the special case depicted in Figure 3.1(a) where $Y = X - c$ for some constant $c > 0$, but is does not hold in general. A thorough study of the proof of the above proposition shows that, depending on the interaction between the shape of the supply function and the distributions of the residual demand, the variance may still decrease, even if (3.1) is violated. However, this analysis involves tedious algebraic manipulation and leads to conditions that do not lend themselves to economic interpretation.

For a practical study, we disentangle the effect of the shape of the supply function and the distribution of residual demand on the price variance. We employ a first-order Taylor approximation of the supply function S around $\mathbb{E}(X) = \mu$, which reads

$$S(X) \approx S(\mu) + S'(\mu)(X - \mu). \tag{3.3}$$

Taking variances on both sides using the fact that $\operatorname{Var}(aX + b) = a^2 \operatorname{Var}(X)$ and $\operatorname{Var}(b) = 0$ for

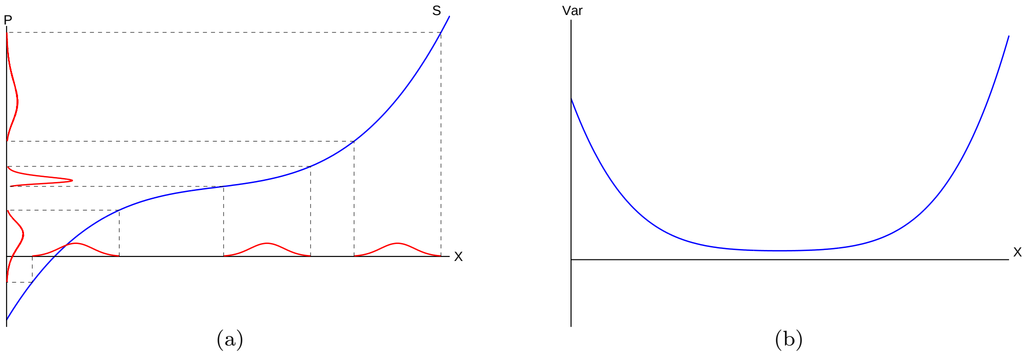

Figure 3.2: Panel (a): Price effects with a concave-convex supply function. Panel (b): Effects of a shift in the mean demand on the variance assuming a concave-convex supply function.

$a, b \in \mathbb{R}$, we find

$$\operatorname{Var}(S(X)) \approx S'(\mu)^2 \operatorname{Var}(X). \tag{3.4}$$

Hence, the variance of the prices is approximated by the squared slope of S evaluated at the mean of the residual demand times the variance of the residual demand. Analogously, the first-order approximation for $\operatorname{Var}(S(Y))$ can be found and we see that $\operatorname{Var}(S(Y)) \leq \operatorname{Var}(S(X))$ is *approximately* true if

$$S'(\mathbb{E}(Y))^2 \operatorname{Var}(Y) \leq S'(\mathbb{E}(X))^2 \operatorname{Var}(X). \tag{3.5}$$

We can use the above criterion to explain what happens in Figure 3.1(b): From $Y \leq X$, it follows that $\mathbb{E}(Y) \leq \mathbb{E}(X)$ and therefore, since S is convex, that $S'(\mathbb{E}(Y)) \leq S'(\mathbb{E}(X))$. Hence, to outweigh the effect that Y is shifted to the flatter part of the supply curve, the variance of Y has to be correspondingly bigger so that (3.5) is violated and the variance of $S(Y)$ is larger than the variance of $S(X)$. Conversely, if Y is shifted far enough to the left, the convex shape of the supply function can outweigh a possible positive effect of $\operatorname{Var}(Y)$ on the price variance.

Condition (3.5) is no strict criterion, since it is based on an approximation which tends to get worse as the values of the random variables move away from their respective means. However, it formalizes the intuition underlying Figure 3.1.

Until now we have assumed a monotonically increasing, convex supply curve in a static model. In reality, the dispatch problem in one period is coupled with the dispatch decision in the neighboring periods, i.e., the dispatch is a dynamic decision. Nevertheless, the decisions of the players in the market result in an aggregate supply function for each hour, which can be used to study the effects of IES integration. Some authors argue that inflexible, slow base-load producers cause a dramatic decrease in prices near to the left end of the supply function. Such plants would rather accept producing at a loss for a few hours than switching off, which would induce longer down times and considerable fuel costs for restarting the plant (see e.g., Fanone et al., 2013). This corresponds to the observations made in Germany in times of low demand and high intermittent infeed, where prices frequently are close to zero and sometimes even drop to negative values. He et al. (2013) adopt a dual-exponential model to capture this effect. In their paper, they show that a model with a supply curve consisting of a concave part followed by a convex part fits the merit order curve on the German spot market better than a purely convex supply function. We call a function of this type *inverse S-shaped*. This modification of the model thus incorporates the most important aspect of the dynamics of the dispatch decision, namely plant inflexibility and the resulting load changing costs.

The arguments in this section can be adapted to this setting as is demonstrated in Figure 3.2(a): If the demand is shifted into the concave part of the supply function then the variance tends to increase because of the steeper slope of the function in this region. Hence, in this setting, the variance as a function of the location of the mean of the residual demand, keeping all other characteristics of the distribution fixed, is U-shaped, i.e., high for very high and very low demand and smaller in the middle section as is depicted in Figure 3.2(b). In our empirical analysis, we show that the U-shaped dependence of the variance on the residual demand is present in German market data.

3.3 Empirical Analysis

In Section 3.2, we identified the shape of the supply function and the distribution of the residual demand as the two pivotal influences on the electricity wholesale price variance. The findings imply that IES can either decrease or increase price variance, depending on which of the two influences prevail. In this section, we put these theoretical considerations to the test. For this purpose, we conduct a statistical investigation of the price variances currently observed in the German day-ahead spot market of electricity. More specifically, we set up OLS regression models, aiming to explain daily spot price variance by the variances and means of demand and renewable infeed, while controlling for other factors that also potentially affect electricity price variation. We are able to capture the variance increasing as well as the variance reducing effect in the regression model.

In this study, we define renewables as wind power and PV and do not consider other technologies. We do this for two reasons: First, because the bulk of intermittent production comes from these two technologies and second, because infeed from other relevant renewables, as for example biomass or hydro plants, can be controlled and therefore does not contribute much to the variance of the residual demand.

We start with an overview of the German market design with a focus on the day-ahead spot market and the marketing of renewable electricity and then discuss the variables we use in the regressions. In the main part of this section, we discuss the setup of the regressions as well as the results of our study of the price data from 2007 until 2013.

3.3.1 The German Day-Ahead Market

We pick Germany as a showcase because investment in renewable sources has been considerable in the last decade. Figure 3.3(a) documents this, showing the installed capacity of wind power and photovoltaics: installed capacity of wind plants was growing steadily between 2001 and 2012 with an average compound annual growth rate of about 12% and installed capacity of PV more than tripled between 2009 and 2012 (see also Wozabal et al., 2013).

For Germany the relevant market places to trade electricity are the European Energy Exchange (EEX) and the European Power Exchange (EPEX). While on the EEX options and futures on power can be traded, the EPEX hosts the day-ahead and intra-day market. The hourly prices on the day-ahead market are determined in sealed bid, uniform price auctions at 12 p.m. for every hour of the next day. On the intra-day market electricity can be traded

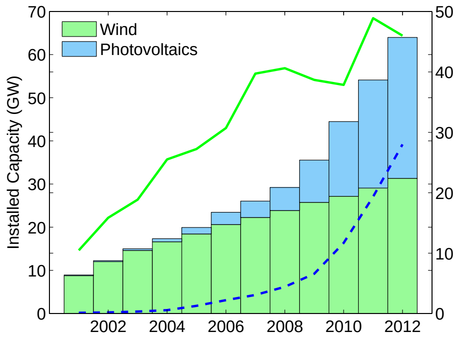

(a)

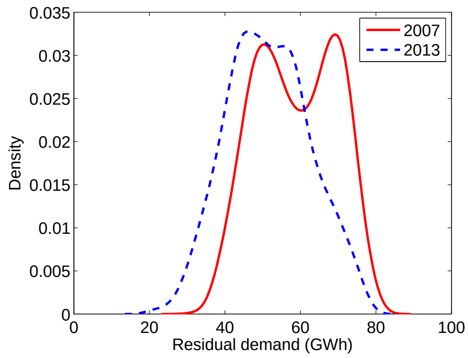

(b)

Figure 3.3: Panel (a): Installed capacity (bars) and production (lines) of wind power (solid) and PV (dotted) in Germany. Source: Wozabal et al. (2013); Panel (b): Residual demand distributions (load reduced by forecast wind power and PV production).

continuously for delivery in single hours on the same or the following day until 75 minutes before delivery begins.

The legal framework for subsidizing renewable power production in Germany is regulated in the German Renewable Energy Sources Act (EEG, see Federal Ministry for the Environment, Nature Conservation, and Nuclear Safety, 2013) and its amendments, which requires the grid operators to buy all the power produced by renewable sources in their grid area for a regulated tariff and sell it further on the day-ahead and intra-day market. The grid operators have to publish forecasts of renewable production for the next day. The forecast quantities are sold on the day-ahead market and the differences between the day-ahead forecasts and actual volume are traded on the intra-day market.

The difference of the revenues on the spot markets and the payments to the producers of renewable energy are typically negative and are charged along with the grid fees to the final consumers. Hence, although producers receive a fixed feed-in tariff, renewable power is traded on the wholesale markets as if producers themselves would be active in these markets. In a recent modification of the EEG, producers of renewables can also opt for a market premium system instead of a feed-in tariff. In the market premium regime, producers directly sell their quantities without any support from the grid operators (see e.g., Gawel and Purkus, 2013). The modification of the EEG is in place since 2012 and aims to incentivize market-orientated behavior of the suppliers of renewables. However, since most of the renewable production is sold on the day-ahead market, we focus on the day-ahead prices at the EPEX, where Germany and Austria are treated as one market zone.[1]

To guarantee transparency, the EEG requires grid operators to provide data on renewable production to the public. In particular, data about production and day-ahead forecasts of renewables is available in an hourly resolution. The good data availability adds to the fact that the EPEX is an ideal showcase. Furthermore, Graf and Wozabal (2013) showed that the EPEX works efficiently, which means that the day-ahead market complies with the theoretical

[1] In principle, suppliers of renewables under the market premium regime are free in selecting a marketplace. However, the day-ahead spot market is way bigger than the intra-day market although the latter is growing: the traded volume in 2013 at the day-ahead market was 245.6 TWh (245.3 TWh in 2012) versus 19.7 TWh on the intra-day market (15.8 TWh in 2012), see EPEX (2014). Therefore, we assume that most of the produced quantities are sold on the day-ahead market.

Variable	Description	Source	Unit
Price	Hourly day-ahead spot price for Germany/Austria	EPEX (2014)	EUR
Demand	Hourly load for Germany/Austria	ENTSO-E (2014)	GWh
Wind	Hourly day-ahead wind forecast for Germany/Austria	TSOs (2014)	GWh
Photovoltaic (PV)	Hourly day-ahead PV forecast for Germany	TSOs (2014)	GWh
Temperature	Daily temperature index	Mathematica	Celsius
Oil	Daily Europe Brent Spot Price FOB	EIA (2014)	EUR
Daylight minutes	Daily daylight minutes		

Table 3.1: Overview of used variables and data sources.

framework in Section 3.2.

3.3.2 Data

In this section, we discuss the variables used in our econometric analysis. An overview of these variables along with their descriptions and sources is provided in Table 3.1.

For our analysis, we use daily price variances, i.e., 24 hourly observations to compute one variance per day. The time frame to calculate the variance is chosen for the following reasons: On the one hand, it should be long enough such that the renewables can cause price variance, e.g., due to changed weather conditions. Thus, a higher resolution, e.g., half-daily, would be meaningless. On the other hand, the time frame should not be too long, since in this case one would loose important information due to aggregation and furthermore, we would end up with too few observations for a sound empirical analysis.

In order to estimate the effect of residual demand characteristics on the price variance, we require hourly data on prices, on electricity demand, and on infeed from renewable sources of power. The data on prices can be obtained directly from the EPEX website, the data on demand from the European Network of Transmission System Operators for Electricity (ENTSO-E) and the data on renewable infeed from the main grid operators (TSOs) in Germany/Austria.[2]

Considering the mechanisms of the day-ahead market, we use the day-ahead forecasts for renewable production instead of the actual infeeds, if available. For Germany, wind forecasts are available for the entire time period from 2007–2013. German PV data has not been published prior to 2010.[3] Austrian PV production is negligible and forecasts are not available, while Austrian wind production forecasts are published since 2010.

For the demand, day-ahead forecasts are not available in general, and we therefore substitute them by the actual load. The effect of this substitution should be marginal, since demand can be very well predicted. The mean average percentage error (MAPE)[4] of the demand forecast is 2% for state of the art methods (see McSharry and Taylor, 2007; Taylor, 2010), as compared to the MAPE for wind power and PV production forecasts which range from 12% to 18%, (Wozabal et al., 2013).

The characteristics of demand and infeed of renewable power are not the only factors driving price variance. To control for influences that originate from the supply side of the market, we include oil prices and three month lagged oil prices to capture the influence of the price of

[2]Germany: EnBW Transportnetze, Tennet TSO, Amprion, and 50Hertz Transmission; Austria: Austrian Power Grid.

[3]For 2010, only one of the german grid operators provided PV forecasts (Tennet TSO). Therefore, we extrapolated PV forecasts for the other grid operators by scaling up and down the Tennet forecasts according to the average ratios between the Tennet forecasts and the forecasts of the corresponding operator in 2011.

[4]The MAPE is defined as $n^{-1}\sum_{i=1}^{n}\left|\frac{y_i-\hat{y}_i}{y_i}\right|$ where $(\hat{y}_i)_{i=1}^{n}$ are the forecasts, $(y_i)_{i=1}^{n}$ are the actual values and n is the total numbers of hours in the underlying period.

	2007	2008	2009	2010	2011	2012	2013	Total
Price Avg	35.85	64.83	38.93	44.51	51.50	42.51	37.60	45.11
Price Var	430.04	672.55	252.16	174.68	158.01	202.90	235.11	391.62
Wind infeed Avg	4.25	4.48	4.16	4.53	5.29	5.69	5.86	4.90
Wind infeed Var	13.06	13.79	11.60	13.00	18.65	18.89	22.72	16.39
PV infeed Avg				1.11	2.27	3.18	3.40	1.43
PV infeed Var				3.25	11.08	22.86	27.93	11.31
Renewables infeed Avg	4.25	4.48	4.16	5.64	7.56	8.87	9.27	6.32
Renewables infeed Var	13.06	13.79	11.60	14.88	25.98	36.57	43.26	26.97
Demand Avg	63.15	63.03	58.92	62.44	62.78	60.12	60.75	61.60
Demand Var	117.50	109.15	122.93	109.44	125.34	132.11	121.96	122.14
Residual demand Avg	58.90	58.55	54.75	56.80	55.22	51.26	51.49	55.28
Residual demand Var	113.14	103.24	119.03	103.88	112.98	127.36	120.87	122.49
# Observations	8,729	8,753	8,730	8,760	8,759	8,784	8,760	61,275
# Outliers$^+$	114	80	56	35	17	66	52	420
# Outliers$^-$	0	5	28	46	89	25	20	213

Table 3.2: Descriptive statistics. Prices are reported in EUR/MWh while quantities are in GWh.

primary energy. Furthermore, we include the number of daylight minutes to capture seasonal trends, and an index of daily average temperatures. To construct the daily temperature index for Austria and Germany we use a population weighted mean of all the Austrian and German cities with more than 10,000 inhabitants. The index is constructed using the Mathematica 7.0 functions CityData and WeatherData (see Graf and Wozabal, 2013). For the oil price, we use the daily Europe Brent spot price provided by EIA (2014).

The reason we use the variance as the measure of dispersion is its nice theoretical properties. It allows for a neat decomposition of the residual demand in the regression models. This, in turn, permits to study in detail the two pivotal effects driving the price variance; other measures, e.g., standard deviation would not allow for such a decomposition. However, the variance is a quadratic expression and we are performing OLS regressions, minimizing the sum of squared residuals. The combination of the two renders our analysis extremely sensitive to outliers. Therefore, we remove hourly observations which deviate from the average price for every year by more than three standard deviations (see Mugele et al., 2005; Ketterer, 2012, for a similar approach). The corresponding numbers are shown at the bottom of Table 3.2.[5] Moreover, we remove days where we have less than 13 price observations left, and we end up with 2553 daily observations (out of 2557 days in the considered time period).

3.3.3 Descriptive Statistics

Table 3.2 displays the descriptive statistics of our sample. Having a first glimpse on the raw data reveals several peculiarities: The yearly price variances of the years 2009 to 2013 are significantly lower than in the years 2007 and 2008, although the yearly average production of renewables increased drastically in the same time. The yearly empirical variance of the renewables has increased with the massive build-up of photovoltaics (2011 – 2013). Wind production obviously exhibits long-term variability, since e.g., average yearly production is lower in 2009 than in 2007 although installed capacity is significantly higher in 2009. We also note that the number

[5] For a few hours, data points are missing in the publicly available data, therefore the number of observations varies slightly over the years.

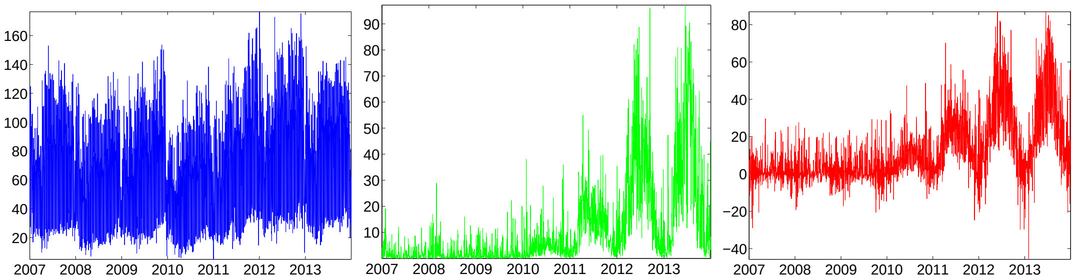

(a) (b) (c)

Figure 3.4: Panel (a): Variance of demand; Panel (b): Variance of renewable infeed; Panel (c): Covariance between demand and renewable infeed.

of negative outliers in the price peaks in 2011 while the number of positive outliers decreased from 2007 to 2011.

While the average renewable production increased considerably between 2007 and 2013, the yearly average demand decreased slightly, leading to a changed residual demand. Figure 3.3(b) shows the results of a kernel density estimation of residual demand for the first and the last year in our sample. Comparing the two years yields the following observations: (i) the mean decreased between 2007 and 2013, (ii) the shape of the distribution changed, in particular the two peaks are less pronounced in 2013 than in 2007, (iii) the maximum residual demand stayed almost constant while the minimum residual demand reduced drastically. The changing shapes of the distributions can be explained by PV production – a technology which supplies most electricity during the midday peak hours. Therefore, an increase in PV production eliminates the traditional noon peak which has an effect on the shape of the residual demand distribution (see e.g., Gawel and Purkus, 2013; Tveten et al., 2013).

When studying the influence of IES on the variance of residual demand, the covariances between demand and renewable infeed obviously play an important rule too. In Figure 3.4, we depict the daily development of the variances of demand and renewables as well as their covariances. While the daily demand variance changes only slightly between the years (Panel a), the introduction of PV in 2010 has a significant impact on the intra-day variance of renewable infeed (Panel b). PV also affected the covariance between demand and renewables (Panel c), since this technology supplies the most power during the noon hours, a time of the day which usually coincides with high demand.

3.3.4 Results

To test the influence of residual demand, wind power, and PV infeed on the variance of German wholesale electricity prices, we estimate the following OLS regression models with the price variance as dependent variable: A model that uses residual demand as an explanatory variable (model 1), a model in which the components that make up residual demand, i.e., demand and IES infeed, are treated as separate regressors (model 2), and a third model where we further disentangle the infeed from renewable sources into their two constituting elements, i.e., wind power and PV (model 3).

All models include the other factors discussed in the last section in daily resolution. We perform all regressions using heteroskedasticity and autocorrelation robust standard errors (Andrews, 1991; Zeileis, 2006), since both the Durbin-Watson test and the Breusch-Pagan test reject the hypothesis of i.i.d. residuals for our models. For the statistical analysis we used the program R (R Core Team, 2013).

We first discuss the results from model 1, where we use the distribution of the residual demand to explain price variance. In particular, we are interested in the effects of shifts in the residual demand and in changes in the shape of the distribution. In the regressions, we capture the former by average daily residual demand and model the latter by daily variances of demand. To include the possibility of an inverse S-shaped supply function, we also include the squared average daily demand as a regressor. It turns out that both the AIC and the BIC favor the corresponding model over the model without the squared term.

The results of the regression are depicted in Table 3.3 in the column marked with (1). The average residual demand has a convex quadratic influence on the price variance, which leads to higher variances for very low and very high levels of residual demand. This can be seen as an indication of an inverse S-shaped supply function. The variance of the residual demand has a strong positive influence on the variance of the prices. Both empirical findings support the theoretical considerations in Section 3.2, which state that changes in the price variance are driven by shifts into areas with different slope of the supply curve and by changes in the dispersion of the residual demand distribution.

To capture the influence of renewables on the variance, we estimate model 2, i.e., we break up the residual demand and repeat the analysis. Motivated by the results about the quadratic relationship of the price variance on the average residual demand, we also include all the terms of the expansion

$$X^2 = (Q - I)^2 = Q^2 - 2QI + I^2,$$

i.e., the squared demand, the squared amount of intermittent production, and the interaction of the two. Following a similar logic, we split up the variance of the residual demand as

$$\mathrm{Var}(X) = \mathrm{Var}(Q - I) = \mathrm{Var}(Q) - 2\,\mathrm{Cov}(Q, I) + \mathrm{Var}(I)$$

and include all the terms on the right hand side of the above equation as regressors.

Table 3.3, Column 2, shows the results of the regression for model 2. The coefficients of the covariance between demand and renewables and the interaction term of these two variables have the anticipated negative signs. Figure 3.5(a) illustrates how the predicted variance is influenced by the average amount of renewable production. The results suggest that a higher demand generally leads to a higher price variance than a lower demand, except if the low demand occurs in combination with high infeed from renewables. The co-occurrence of low demand and high infeed of renewables pushes the residual demand into the steep area of the inverse S-shaped supply curve. This results in extremely low prices and a higher price variance. It also relates to the phenomenon of negative prices for electricity on the day-ahead spot market that are sometimes observed in the German market (see e.g., Fanone et al., 2013).

With regard to the variance of demand and renewables, the coefficients of model 2 show that the price variance increases with an increase in the variance of demand and renewables while the covariance between those has a negative effect. Figure 3.5(b) depicts these effects. The upper plane shows the situation for the lowest observed covariance between demand and renewable infeed in our data while the lower plane represents the situation with the highest observed covariance. The interpretation of this effect is the following: an increased variance of

	(1)	(2)	(3)
Intercept	981.73	152.55	4.67
	(0.001)	(0.729)	(0.990)
Residual Demand	−40.06		
	(0.000)		
Residual Demand^2	0.43		
	(0.000)		
Demand		−15.59	−16.19
		(0.256)	(0.186)
Demand^2		0.23	0.23
		(0.0067)	(0.036)
Renewables		18.64	
		(0.096)	
Renewables^2		1.07	
		(0.000)	
Renewables * Demand		−0.69	
		(0.002)	
Wind			29.74
			(0.011)
Wind^2			0.50
			(0.007)
PV			−53.71
			(0.023)
PV^2			9.93
			(0.000)
Wind * Demand			−0.74
			(0.001)
PV * Demand			−1.13
			(0.023)
PV * Wind			2.52
			(0.004)
Residual Demand Var	1.70		
	(0.000)		
Demand Var		1.15	1.02
		(0.000)	(0.000)
Renewables Var		1.88	2.75
		(0.000)	(0.000)
Demand, Renewables Cov		−5.92	−3.59
		(0.000)	(0.000)
Daylight Hours	-0.06	0.03	0.11
	(0.510)	(0.720)	(0.222)
Temperature	3.76	4.42	3.95
	(0.019)	(0.004)	(0.004)
Oil Price	−2.66	−2.34	−1.07
	(0.136)	(0.122)	(0.420)
Oil Price 3M	2.71	4.05	5.21
	(0.126)	(0.012)	(0.000)
Adjusted R^2	0.25	0.29	0.36
N	2,553	2,553	2,553

Table 3.3: Coefficients (p-values below) of the OLS regressions for the years 2007 to 2013 (heteroskedasticity and autocorrelation consistent estimation of the variance-covariance matrix). All variables represent daily averages unless indicated otherwise.

demand and renewables leads to a higher price variance. However, this effect is less pronounced for a high covariance between these two variables than for a low covariance. This is in particular interesting because Figure 3.4(c) shows fluctuating level of covariance over the years.

Finally, we discuss the results of model 3, where we split up renewable production into its parts, i.e., wind power and PV. However, due to the almost perfect correlation (0.95) between the variance of PV and its mean, we do not include both terms in the regression. The high correlation is caused by the peculiarity of PV, i.e., it supplies only during daylight hours and

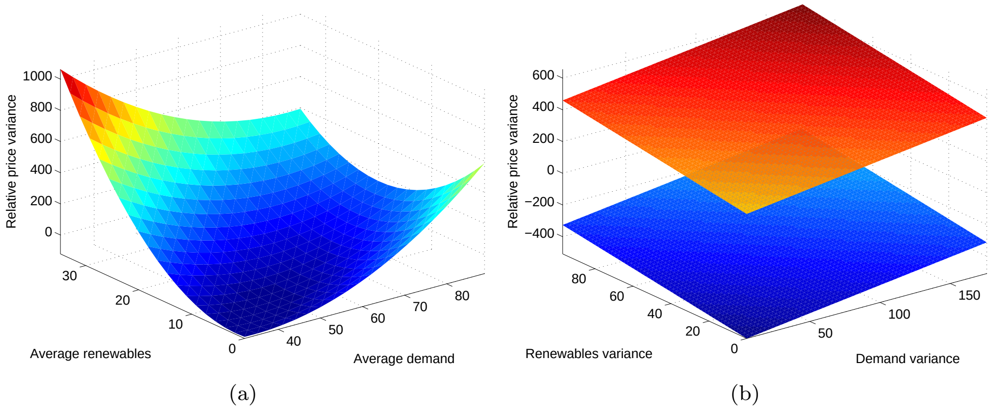

Figure 3.5: Panel (a): Predicted effects of average renewable infeed and demand (in GW) on relative price variance; Panel (b): Predicted effects of renewable infeed and demand variance on relative price variance. All other regressors are set to zero.

a higher variance is observed when more electricity from PV is supplied. Thus, a full decomposition into the sample mean and variance of wind and PV is not possible. Instead, we only decompose the mean of the renewable production. We present the results in Table 3.3, Column 3. The effects of the average wind and PV production on the relative price variance are depicted in Figure 3.6. Notice that the relative price variance is evaluated at the average demand of 61.6 GW (see Table 3.2). The estimated coefficients of PV infeed and the coefficient of the interaction effect between PV and demand are reflected in the sharp decrease of price variance when average PV infeed increases. Wind infeed also reduces the price variance, but not as drastically. Although both coefficients of wind infeed are positive, the price variance is reduced because of the negative interaction term between wind infeed and demand. The average marginal effects of PV and wind infeed (per GWh) on the price variance evaluated at the observations yield an average reduction of 83 for PV compared to 7 in the case of wind. Therefore, the impact of average PV infeed on the price variance is about twelve times larger than that of wind.

The coefficients of the variables which are not related to the residual demand are of similar magnitude in all three models: Price variance is increasing with average temperature, leading to higher price variance in summer than in winter.[6] As expected, the average three month lagged oil price which serves as a proxy for the fossil fuel prices has a positive influence on the price variance. The coefficients of the average daylight minutes and the oil price are not significant. The explained variation ranges from 25 to 36 percent.

As an additional exercise, we performed yearly regressions in order to identify anomalies between the different years. With one exception all relevant coefficients have the same sign as in the full sample. The results of these regressions are available on request.

3.4 Policy Implications

In this paper, we argue that the effect of IES infeed on the spot price variance depends crucially on the amount and the type of installed capacity. First, the higher the correlation of the infeed with stochastic demand the higher the variance dampening effect brought about by a decrease

[6]Since Germany and Austria are traditionally winter-peaking countries this seems counterintuitive at first glance. An explanation is that high PV production in summer coupled with lower demand pushes the residual demand in the concave area of the merit order resulting in a higher variance.

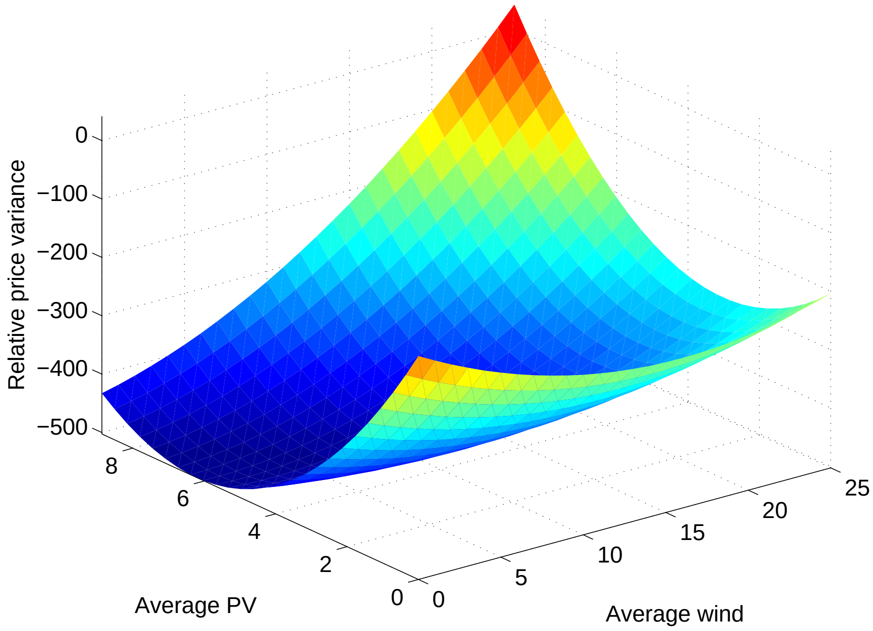

Figure 3.6: Predicted effects of average PV and wind infeed (in GW) on relative price variance evaluated at the average demand level of 62.6 GW. All other regressors are set to zero.

of the variance of residual demand. In this sense, the installation of PV capacities tends to decrease price variance more than the installation of wind turbines, since the infeed from the former has a pronounced positive correlation with demand, while the infeed from the latter shows no such clear correlation pattern. Political assessments of the desirability of wind power versus PV should take this aspect into account.

Second, according to our analysis in markets with inverse S-shaped supply functions the introduction of a small to moderate amount of IES is likely to have a variance damping effect while a larger increase may lead to an increased price variance on the electricity exchanges, as the residual demand is pushed into the concave part of the merit order.

In fact, we can quantify the tipping point for the German market using the models' estimates from Section 3.3. Given the estimated coefficients of model 2, we can find the daily demand and the renewable infeed that minimize the predicted variance in prices. Carrying out this exercise yields a demand of 42.8 GW and a renewables infeed of about 5.7 GW. Fixing the demand at its average of about 61.6 GW, the minimal variance is reached at a level of renewable infeed of 12 GW. Notice that these numbers are daily mean and variance levels, i.e., the daily price variance is minimized given a daily infeed of renewables and a fixed demand. Although these calculations give only a rough picture, it becomes clear that in the first years of our analysis, the IES infeed will have had a variance dampening effect on most days, while in the later years, with increased capacity and a large number of high infeed days, IES will have increased the variance on many days.

For policy makers this has wide ranging implications. Clearly, the structural changes in electricity markets, triggered by the build-up of intermittent renewable electricity production, require increased flexibility of production and consumption. This flexibility may either be provided by electricity storage plants, by a demand response via for example smart grids, or by flexible production technologies such as gas turbines (see Denholm and Hand, 2011).

Policies to promote renewable electricity production typically directly target the building of corresponding production capacities, tacitly assuming that the demand for increased flexibility will result in market conditions that act as incentives to invest in demand side response technologies and storage facilities. However, all corresponding technologies require substantial investments whose returns are mainly driven by price variance. In particular, if price variance goes down in the first phase of renewable expansion, then these projects become less profitable and correspondingly less capacities will be build or projects will be delayed, even if firms correctly anticipate an increase in price variance at a later point in time.

Since construction of the corresponding infrastructure takes time, this in turn might result in a situation where the overall system is left with too little flexibility to handle increased IES infeed which may lead to network problems and, in the worst case, blackouts. Examples of near grid failures in Germany could be observed in the year 2011 (see Bundesnetzagentur, 2012). Since there is no market for the economic costs of blackouts these costs constitute a negative externality which might decrease overall welfare, even if all players act rational on the existing electricity markets.

In conclusion one can say that policies that subsidize the build-up of renewable intermittent capacities should also incentivize the expansion of storage capacity, some means of demand response, or transmission capacities to neighboring grid areas. A large scale demand response by consumers could for example be facilitated by a smart electricity grid in combinations with flexible pricing schemes for electric power, (see e.g., Strbac, 2008). Flexibility in the production and storage capacity can either be directly subsidized or ensured by capacity markets (see e.g., Crampton et al., 2013, for a discussion on the importance of capacity markets). The promotion of electric cars whose combined battery capacity can act as a storage device might be another option if combined with a smart grid (see e.g., Guille and Gross, 2009).

3.5 Conclusion

While the negative effect of renewables on average spot prices is beyond controversy in the literature, research on price variance is still rather scarce and most authors conclude that variance increases with the installation of intermittent capacities. In this paper we showed that the build-up of IES can not only increase but also decrease electricity price variance. The latter effect is expected to occur for rather small amounts of IES capacity, while the former is expected once a significant amount of IES capacity is installed. We argue this point theoretically using a static market model and identify the slope of the supply function as well as the variance of the IES infeed as two crucial factors affecting price variance. Whereas the variance of IES infeed increases price variance, the actual amount of IES infeed can also influence the price variance negatively. In Section 3.3, we test our theoretical considerations and quantify the two effects using data from the German spot market. As we argue in the last section, the findings in this paper have significant policy implications, since they show that market based incentives for investments in flexibility alone might not result in sufficient capacities built in time to ensure network stability. Hence, policies that require producers to provide flexibility, the build-up of infrastructure such as a smart grid, or direct subsidies for flexible generation and storage facilities are called for.

Further investigations about the price variance might want to explicitly incorporate the dynamics in the bidding that determines the market prices, for example in a dispatch model.

Another direction of future research could be the modelling of strategic interactions of the players on the market, an aspect which we did not consider here.

Acknowledgments

We want to express our thanks to Klaus Gugler, Felix Höffler, Franz Wirl, and the participants of the Mannheim Energy Conference 2013 for stimulating discussions and valuable inputs on a draft version of this paper. David Hirschmann was partially supported by the WWTF (Project Number: MA09-019).

Bibliography

Andrews, D. W. K., 1991. Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. Econometrica 59, 817–858.

Arocena, P., Saal, D. S., Coelli, T., 2012. Vertical and Horizontal Scope Economies in the Regulated U.S. Electric Power Industry. The Journal of Industrial Economics 60 (3), 434–467.

Barlow, M., 2002. A diffusion model for electricity prices. Mathematical Finance 12 (4), 287–298.

Bundesnetzagentur, 2012. Report on Energy Supplies in Winter 2011/12. Tech. rep.

Chao, H., 2011. Efficient pricing and investment in electricity markets with intermittent resources. Energy Policy 39 (7), 3945–3953.

Crampton, P., Ockenfels, A., Stoft, S., 2013. Capacity Market Fundamentals. Economics of Energy & Environmental Policy 2 (2), 27–46.

Denholm, P., Hand, M., 2011. Grid flexibility and storage required to achieve very high penetration of variable renewable electricity. Energy Policy 39 (3), 1817–1830.

Dunne, T., Mu, X., 2010. Investment spikes and uncertainty in the petroleum refining industry. The Journal of Industrial Economics 58 (1), 190–213.

EIA, 2014. U. S. Energy Information Administration. Online, available: http://www.eia.gov.

ENTSO-E, 2014. European Network of Transmission System Operators for Electricity. Online, available: http://www.entsoe.eu, http://www.entsoe.net.

EPEX, 2014. European Power Exchange. Online, available: http://www.epexspot.com.

Fanone, E., Gamba, A., Prokopczuk, M., 2013. The case of negative day-ahead electricity prices. Energy Economics 35, 22–34.

Federal Ministry for the Environment, Nature Conservation, and Nuclear Safety, 2013. Renewable Energy Sources Act (EEG). Online, available: http://www.erneuerbare-energien.de.

Gawel, E., Purkus, A., 2013. Promoting the market and system integration of renewable energies through premium schemes—A case study of the German market premium. Energy Policy 61, 599–609.

Graf, C., Wozabal, D., 2013. Measuring competitiveness of the EPEX spot market for electricity. Energy Policy 62, 948–958.

Green, R., Vasilakos, N., 2010. Market behaviour with large amounts of intermittent generation. Energy Policy 38 (7), 3211–3220.

Guille, C., Gross, G., 2009. A conceptual framework for the vehicle-to-grid (v2g) implementation. Energy Policy 37, 4379–4390.

Gurland, J., 1967. An inequality satisfied by the expectation of the reciprocal of a random variable. The American Statistician 21, 24–25.

He, Y., Hildmann, M., Herog, F., Andersson, G., 2013. Modeling the merit order curve of the european energy exchange power market in germany. IEEE Transactions on Power Systems 28 (3), 3155–3164.

Hirth, L., 2013. The market value of variable renewables: The effect of solar wind power variability on their relative price. Energy Economics 38, 218–236.

Jacobsen, H. K., Zvingilaite, E., 2010. Reducing the market impact of large shares of intermittent energy in Denmark. Energy Policy 38 (7), 3403–3413.

Jónsson, T., Pinson, P., Madsen, H., 2010. On the market impact of wind energy forecasts. Energy Economics 32, 313–320.

Ketterer, J. C., 2012. The Impact of Wind Power Generation on the Electricity Price in Germany. Ifo Working Paper Series 143, Ifo Institute for Economic Research at the University of Munich.

Kim, D.-W., Knittel, C. R., 2006. Biases in static oligopoly models? Evidence from the California electricity market. The Journal of Industrial Economics 54 (4), 451–470.

McSharry, P., Taylor, J., 2007. Short-term load forecasting methods: an evaluation based on European data. IEEE Transaction on Power Systems 22 (4), 2213–2219.

Milstein, I., Tishler, A., 2011. Intermittently renewable energy, optimal capacity mix and prices in a deregulated electricity market. Energy Policy 39 (7), 3922–3927.

Muche, T., 2009. A real option-based simulation model to evaluate investments in pump storage plants. Energy Policy 37 (11), 4851–4862.

Mugele, C., Rachev, S. T., Trück, S., 2005. Stable modeling of different european power markets. Investment Management and Financial Innovations 3, 65–85.

Nicolosi, M., 2010. Wind power integration and power system flexibility – An empirical analysis of extreme events in germany under the new negative price regime. Energy Policy 38 (11), 7257–7268.

Paraschiv, F., Erni, D., Pietsch, R., 2014. The impact of renewable energies on eex day-ahead electricity prices. Energy Policy (Article in Press).

R Core Team, 2013. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
URL `http://www.R-project.org/`

Redl, C., Haas, R., Huber, C., Böhm, B., 2009. Price formation in electricity forward markets and the relevance of systematic forecast errors. Energy Economics 31 (3), 356–364.

See, C.-T., Chen, J., 2008. Inequalities on the variances of convex functions of random variables. Journal of inequalities in pure and applied mathematics 9, Article 80.

Sensfuß, F., Ragwitz, M., Genoese, M., 2008. The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in germany. Energy Policy 36, 3086–3094.

Strbac, G., 2008. Demand side management: Benefits and challenges. Energy Policy 36, 4419–4426.

Taylor, J. W., 2010. Triple seasonal methods for short-term electricity demand forecasting. European Journal of Operational Research 204 (1), 139 – 152.

TSOs, 2014. German TSOs: Amprion, 50 Hertz, TenneT, and TransnetBW; Austrian TSO: APG. Online, available: http://www.amprion.net, www.50hertz.com, http://www.tennettso.de, www.transnetbw.de, www.apg.at.

Tveten, Å. G., Bolkesjø, T. F., Martinsen, T., Hvarnes, H., 2013. Solar feed-in tariffs and the merit order effect: A study of the German electricity market. Energy Policy 61, 761–770.

Wagner, A., 2014. Residual Demand Modeling and Application to Electricity Pricing. The Energy Journal 35 (2), 45–73.

Woo, C., Horowitz, I., Moore, J., Pacheco, A., 2011. The impact of wind generation on the electricity spot-market price level and variance: The texas experience. Energy Policy 39 (7), 3939–3944.

Wozabal, D., Graf, C., Hirschmann, D., 2013. Renewable Energy and its Impact on Power Markets. In: Kovacevic, Pflug, Vespucci (Eds.), Handbook of Risk Management for Energy Production and Trading. International Series in Operations Research and Management Science. Springer, pp. 283–311.

Zeileis, A., 2006. Object-oriented Computation of Sandwich Estimators. Journal of Statistical Software 16 (9), 1–16.

3.6 Appendix

The proof of Proposition 1 requires the following proposition:

Proposition 2. *Let X be a positive random variable which takes values on an interval $[a, b]$. Let $h : \mathbb{R}^+ \times \mathbb{R} \to \mathbb{R}$ be a function such that the following hold:*

1. *$h(X, \cdot)$ is twice differentiable*
2. *$h(\cdot, \epsilon)$ is increasing $\forall \epsilon \in [a, b]$, and*
3. *$\frac{\partial}{\partial \epsilon} h(\cdot, \epsilon)$ is decreasing $\forall \epsilon \in [a, b]$.*

Then for $\epsilon_1, \epsilon_2 \in [a, b]$ with $\epsilon_1 < \epsilon_2$ it holds true that $Var[h(X, \epsilon_2)] \leq Var[h(X, \epsilon_1)]$.

Proof. The proof is analogue to the proof of (See and Chen, 2008, Theorem 3.1). It suffices to show that $Var[h(X, \epsilon)]$ is a decreasing function of ϵ. We have

$$\begin{aligned}
\frac{\partial}{\partial \epsilon} \mathrm{Var}(h(X, \epsilon)) &= \frac{\partial}{\partial \epsilon} \mathbb{E}(h(X, \epsilon)^2) - \frac{\partial}{\partial \epsilon} (\mathbb{E}(h(X, \epsilon))^2 \\
&= \mathbb{E}\left[\frac{\partial}{\partial \epsilon} h(X, \epsilon)^2\right] - 2\mathbb{E}(h(X, \epsilon)) \frac{\partial}{\partial \epsilon} (\mathbb{E}(h(X, \epsilon)) \\
&= \mathbb{E}\left[2h(X, \epsilon) \frac{\partial}{\partial \epsilon} h(X, \epsilon)\right] - 2\mathbb{E}(h(X, \epsilon))(\mathbb{E}(\frac{\partial}{\partial \epsilon} h(X, \epsilon)) \\
&\leq 0.
\end{aligned}$$

The last inequality follows from the fact that for a random variable X, and for continuous functions f, g on $\mathbb{R}$, $\mathbb{E}[f(X)g(X)] \leq \mathbb{E}[f(X)]\mathbb{E}[g(X)]$ if f is monotonically increasing and g is monotonically decreasing (see e.g., Gurland (1967)). The interchange of the differential operator and the expectation is clearly justified by the assumptions we have made. □

Proof of Proposition 1. The proof is analogue to (See and Chen, 2008, Theorem 3.3). We sketch the important steps.

Define $h(u, \epsilon) = g\left(F_X^{-1}(u) + \epsilon[F_Y^{-1}(u) - F_X^{-1}(u)]\right)$, with $\epsilon \in [0, 1]$. First, note that $h(\cdot, \epsilon)$ is an increasing function $\forall \epsilon \in [0, 1]$. Secondly, note that $\frac{\partial}{\partial \epsilon} h(\cdot, \epsilon)$ is a decreasing function since

$$\begin{aligned}
\frac{\partial}{\partial u} \frac{\partial}{\partial \epsilon} h(\cdot, \epsilon) &= \frac{\partial}{\partial u} \left\{[F_Y^{-1}(u) - F_X^{-1}(u)]\; g'\left(F_X^{-1}(u) + \epsilon[F_Y^{-1}(u) - F_X^{-1}(u)]\right)\right\} \\
&= [(F_Y^{-1})' - (F_X^{-1})']\; g'\left(F_X^{-1}(u) + \epsilon[F_Y^{-1}(u) - F_X^{-1}(u)]\right) \\
&\quad + \epsilon(F_Y^{-1})'\; [F_Y^{-1} - F_X^{-1}]\; g''\left(F_X^{-1}(u) + \epsilon[F_Y^{-1}(u) - F_X^{-1}(u)]\right) \\
&\quad + (1 - \epsilon)(F_X^{-1})'\; [F_Y^{-1} - F_X^{-1}]\; g''\left(F_X^{-1}(u) + \epsilon[F_Y^{-1}(u) - F_X^{-1}(u)]\right) \\
&\leq 0
\end{aligned}$$

For the last inequality note that the first term is negative due to condition (2) and g being increasing ($g' > 0$), while the second and third term are negative due to condition (1), convexity of g ($g'' > 0$) and properties of distribution functions, $(F_X^{-1})' \geq 0$ and $(F_Y^{-1})' \geq 0$.

Now one can apply Proposition 2 and finds

$$Var[g(Y)] = Var[h(u, 1)] \leq Var[h(u, 0)] = Var[g(X)].$$

□

The following lemma links inequality (3.1) to the variance of the random variables.

Lemma 3. *If for two random variables X and Y*

$$F_Y^{-1}(\alpha) - F_X^{-1}(\alpha)$$

is monotonically decreasing, then $\mathrm{Var}(Y) \leq \mathrm{Var}(X)$.

Proof. Notice that for $\alpha > \alpha'$, we have

$$F_Y^{-1}(\alpha) - F_X^{-1}(\alpha) \leq F_Y^{-1}(\alpha') - F_X^{-1}(\alpha')$$

and therefore

$$F_Y^{-1}(\alpha) - F_Y^{-1}(\alpha') \leq F_X^{-1}(\alpha) - F_X^{-1}(\alpha'),$$

i.e., the inter-quantile range of Y is smaller than that of X for all quantiles α, α'.

Setting $\alpha_1 = F_X(\mathbb{E}(X))$ and $\alpha_2 = F_Y(\mathbb{E}(Y))$, yields

$$\begin{aligned}
\mathrm{Var(X)} &= \int_{-\infty}^{\infty} (x - \mathbb{E}(X))^2 f_X(x) dx = \int_0^1 (F_X^{-1}(\alpha) - F_X^{-1}(\alpha_1))^2 du \\
&\geq \int_0^1 (F_Y^{-1}(\alpha) - F_Y^{-1}(\alpha_1))^2 du \geq \int_0^1 (F_Y^{-1}(\alpha) - F_Y^{-1}(\alpha_2))^2 du \\
&= \mathrm{Var}(Y),
\end{aligned}$$

where the last inequality follows from

$$\mathbb{E}(Y) = F_Y^{-1}(\alpha_2) = \arg\min \left\{ a : \mathbb{E}\left[(Y-a)^2\right] \right\}.$$

$\square$

Summaries & Curriculum Vitae

Summary

This PhD thesis comprises of three articles. Two of the articles investigate the effect of random inspections, whereas the third article concerns itself with the effect of renewable energy sources on the electricity price variance.

The first article considers a firm whose reputation is subject to random quality inspections. It is assumed that consumers can either never observe the quality of the firm's products (credence goods), or can know about the quality only after having consumed the product (experience goods). Therefore, consumers rely on the findings of quality inspections. The article studies the optimal behavior of the firm and derives the optimal inspection rate from the viewpoint of a regulator. One of the main results is that under certain circumstances a lower inspection rate can be more beneficial than a higher one.

The second article studies the effect of doping inspections on the career decision of a young sports talent. It is assumed that a performance-enhancing drug (doping) is available which increases the likelihood of winning the competition at the risk of being caught cheating. The model implies that low doping sanctions can discourage the entry to the competition and that this effect can even be starker if the inspection rate is increased only mildly.

The first two articles show situations in which an increase in inspection rates can be disadvantageous from the viewpoint of a regulator. These results should be kept in mind when taking corresponding policy decisions.

In the third article, the effect of intermittent renewable energy sources on the wholesale electricity price variance is studied. By means of a theoretical model two pivotal drivers for the variance are identified: the distribution of the renewable energy generation and the shape of the industry supply function. Using data from the German/Austrian day-ahead market, empirical evidence for the two effects is provided. The findings imply that large amounts of renewable energy sources increase price variance whereas small amounts can actually decrease the variance. The results have implications for policy makers coordinating subsidy schemes for renewable energy sources, flexible production capacities, and electricity storages.

Zusammenfassung

Diese Dissertation besteht aus drei Artikeln. Zwei der Artikel untersuchen den Effekt von zufälligen Inspektionen, während der dritte Artikel sich mit dem Einfluss von intermittierenden erneuerbaren Energiequellen auf die Varianz von Strompreisen beschäftigt.

Der erste Artikel behandelt eine Firma deren Reputation von Qualitätskontrollen bestimmt ist. Dem Modell liegt die Annahme zu Grunde, dass Konsumenten entweder nie, oder erst nach Konsum, die Qualität des Firmenerzeugnisses erkennen können. Daher verlassen sich die Konsumenten auf die Ergebnisse der Qualitätskontrollen. Der Artikel untersucht das optimale Verhalten der Firma unter diesen Annahmen, und leitet daraus die optimale Inspektionsrate aus Sicht eines Regulators ab. Eines der Hauptresultate der Arbeit ist, dass unter bestimmten Umständen eine niedrigere Inspektionsrate im Sinne des Regulators bessere Ergebnisse liefern kann als eine hohe.

Der zweite Artikel beschäftigt sich mit dem Einfluss von Dopingtests auf die Karriereentscheidung von Nachwuchssportlern. Es wird angenommen, dass ein leistungssteigerndes Mittel (Doping) zur Verfügung steht, dessen Einnahme die Gewinnchance erhöht, aber die Gefahr als Betrüger entdeckt zu werden mit sich bringt. Das Modell zeigt, dass niedrige Dopingsanktionen Nachwuchstalente davon abhalten können Sportler zu werden, und dass eine zu geringe Erhöhung der Sanktionen das Problem sogar verstärken kann.

Die ersten beiden Artikel zeigen, dass es aus Sicht eines Regulators in manchen Situation nachteilig sein kann die Inspektionsrate zu erhöhen. Diese Ergebnise sollten bei der Festlegung von Regeln im entsprechenden Kontext berücksichtigt werden.

Der dritten Artikel untersucht den Effekt von intermittierenden erneuerbaren Energiequellen auf die Strompreisvarianz. Mittels eines theoretischen Modells werden zwei Einflussfaktoren identifiziert: die Verteilung der Einspeisung von erneuerbarer Energie und die Form der aggregierten Angebotskurve. Anhand von Daten des deutschen/österreichischen Strommarkts werden die beiden Effekte empirisch belegt. Die Resultate zeigen, dass eine große Menge an erneuerbarer Energie die Strompreisvarianz erhöht, während eine geringere Menge die Preisvarianz sogar reduzieren kann. Die Resultate haben einen Einfluss auf das Zusammenspiel von Subventionen für erneuerbare Energiequellen, flexiblen Produktionskapazitäten und Energiespeichern.

Curriculum Vitae

Name	David Hirschmann
Education	since September 2010: PhD Management Program University of Vienna, Austria September 2002 - September 2009: Diploma in Mathematics (M.Sc. equivalent) University of Vienna, Austria 2006 / 2007: Studies of Mathematics University Autónoma de Madrid, Spain
Publications	D. Hirschmann, 2014. Optimal Quality Provision when Reputation is subject to Random Inspections. Operations Research Letters, Vol. 42 (1), p. 64–69. Wozabal, D., Graf, C., Hirschmann, D., 2013. Renewable Energy and its Impact on Power Markets. In: Handbook of Risk Management for Energy Production and Trading (Kovacevic, Pflug, and Vespucci; Eds.), International Series in Operations Research and Management Science, Volume 199, Springer. 283–311.
Submitted	D. Hirschmann, May increasing doping sanctions discourage entry to competition? (under review for Journal of Sports Economics). Wozabal, D., Graf, C., Hirschmann, D., The effect of intermittent renewables on the electricity price variance. (under review for OR Spectrum).