

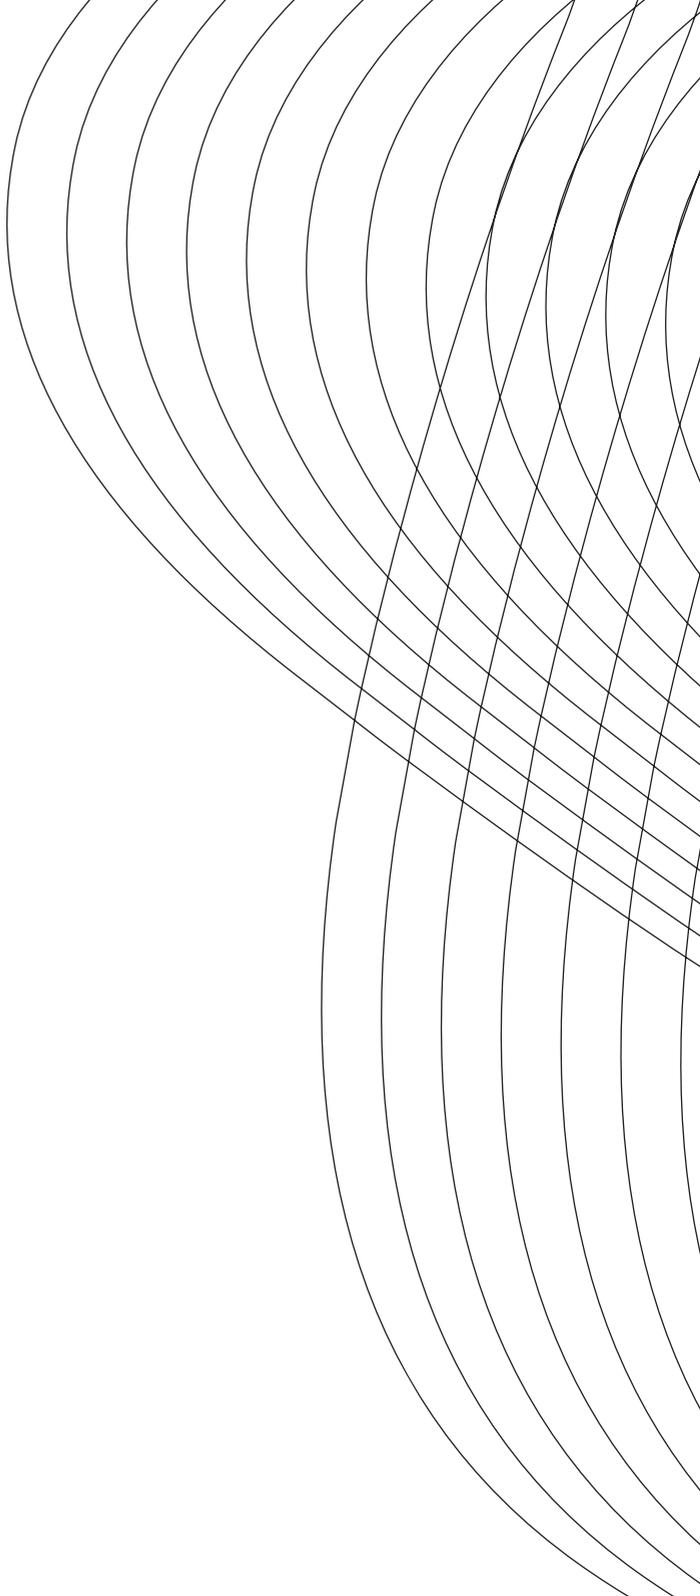


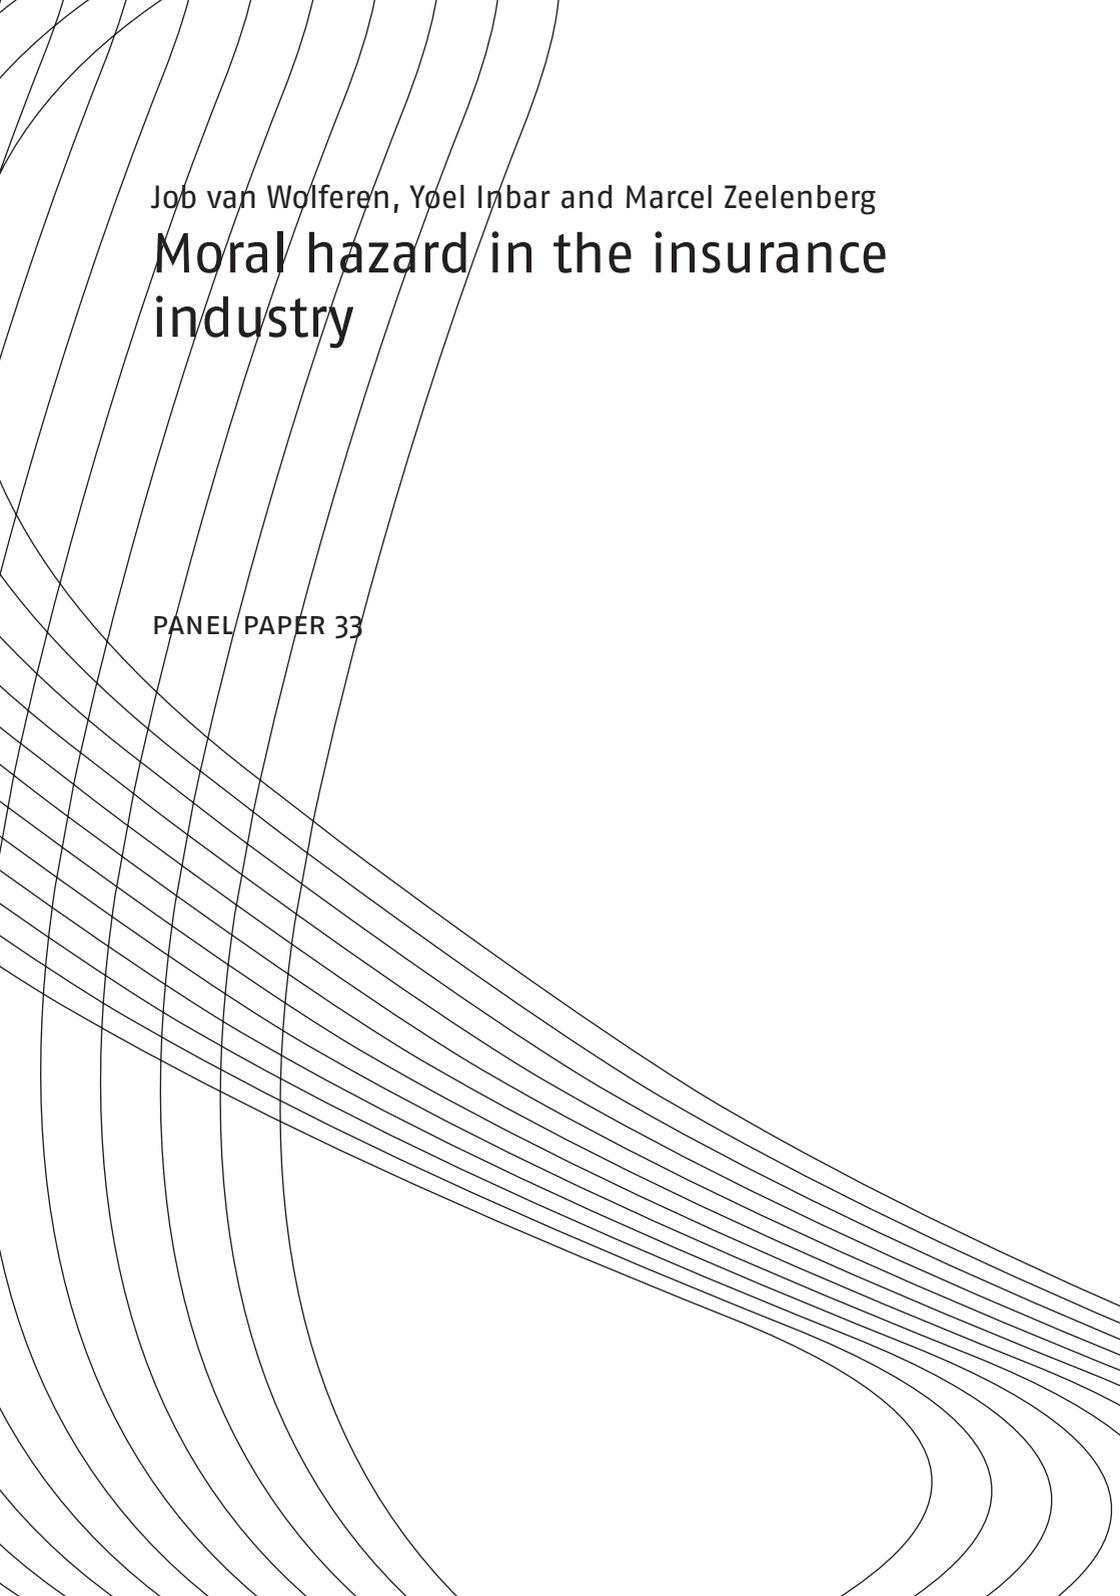
Network for Studies on Pensions, Aging and Retirement

Netspar PANEL PAPERS

*Job van Wolferen, Yoel Inbar
and Marcel Zeelenberg*

Moral hazard in the insurance industry



The background of the page is composed of numerous thin, black, curved lines that sweep across the frame from the top left towards the bottom right. These lines vary in curvature and spacing, creating a sense of movement and depth. The lines are most dense on the left side and become more sparse towards the right.

Job van Wolferen, Yoel Inbar and Marcel Zeelenberg

Moral hazard in the insurance industry

PANEL PAPER 33



Netspar

Network for Studies on Pensions, Aging and Retirement

Colophon

Panel Papers is a publication of Netspar
March 2013

Editorial Board

Roel Beetsma (Chairman) – University of Amsterdam
Bart Boon – Ministry of Finance
Eddy van Doorslaer – Erasmus University Rotterdam
Thomas van Galen – Cardano Risk Management
Kees Goudswaard – Leiden University
Winfried Hallerbach – Robeco Netherlands
Martijn Hoogeweegen – Nationale Nederlanden
Arjen Hussem – PGGM
Frank de Jong – Tilburg University
Alwin Oerlemans – APG
Marine Regnault-Stoel – AEGON Nederland
Maarten van Rooij – De Nederlandsche Bank
Peter Schotman – Maastricht University
Lou Spoor – Achmea
Peter Wijn – APG

Design

B-more Design
Bladvulling, Tilburg

Printing

Printing Office Tilburg University

Editorial address

Netspar, Tilburg University
PO Box 90153, 5000 LE Tilburg
info@netspar.nl

No reproduction of any part of this publication may take place without permission of the authors.

CONTENTS

<i>Preface</i>	7
<i>Overview of policy recommendations</i>	11
<i>Abstract</i>	13
<i>1. Introduction</i>	14
<i>2. What is moral hazard?</i>	15
<i>3. Empirical tests of moral hazard theory</i>	22
<i>4. Moral hazard in pensions</i>	61
<i>5. Discussion and conclusion</i>	67
<i>References</i>	70

PREFACE

Netspar stimulates debate and fundamental research in the field of pensions, aging and retirement. The aging of the population is front-page news, as many baby boomers are now moving into retirement. More generally, people live longer and in better health while at the same time families choose to have fewer children. Although the aging of the population often gets negative attention, with bleak pictures painted of the doubling of the ratio of the number of people aged 65 and older to the number of the working population during the next decades, it must, at the same time, be a boon to society that so many people are living longer and healthier lives. Can the falling number of working young afford to pay the pensions for a growing number of pensioners? Do people have to work a longer working week and postpone retirement? Or should the pensions be cut or the premiums paid by the working population be raised to afford social security for a growing group of pensioners? Should people be encouraged to take more responsibility for their own pension? What is the changing role of employers associations and trade unions in the organization of pensions? Can and are people prepared to undertake investment for their own pension, or are they happy to leave this to the pension funds? Who takes responsibility for the pension funds? How can a transparent and level playing field for pension funds and insurance companies be ensured? How should an acceptable trade-off be struck between social goals such as solidarity between young and old, or rich and poor, and

individual freedom? But most important of all: how can the benefits of living longer and healthier be harnessed for a happier and more prosperous society?

The Netspar Panel Papers aim to meet the demand for understanding the ever-expanding academic literature on the consequences of aging populations. They also aim to help give a better scientific underpinning of policy advice. They attempt to provide a survey of the latest and most relevant research, try to explain this in a non-technical manner and outline the implications for policy questions faced by Netspar's partners. Let there be no mistake. In many ways, formulating such a position paper is a tougher task than writing an academic paper or an op-ed piece. The authors have benefitted from the comments of the Editorial Board on various drafts and also from the discussions during the presentation of their paper at a Netspar Panel Meeting.

I hope the result helps reaching Netspar's aim to stimulate social innovation in addressing the challenges and opportunities raised by aging in an efficient and equitable manner and in an international setting.

Roel Beetsma

Chairman of the Netspar Editorial Board

Affiliations

Job van Wolferen – Tilburg University

Yoel Inbar – Tilburg University

Marcel Zeelenberg – Tilburg University

MORAL HAZARD IN THE INSURANCE INDUSTRY

Overview of policy recommendations

1. Insurance companies need to collect more data to obtain insight into whether, when, why, and how people respond to insurance-induced incentives. Specifically, it is unclear which preventive health care behaviors are sensitive to ex ante moral hazard; the demand elasticity for health care is surely positive, but estimates vary considerably; and the true extent of insurance fraud is unknown. Collecting more detailed data is the first step that must be taken, but insurance companies are also encouraged to conduct simple experiments.
2. Ex ante moral hazard seems especially likely for preventive behaviors that are hard to maintain. Insurers should implement systems that aid policyholders in keeping up those behaviors – examples include encouraging exercising, enforcing chronic medication intake and using simple reminders (e-mail or text) to encourage patients to stick to a diet.
3. Policymakers are advised to use a demand elasticity of -0.20 when forecasting the effects of changes in insurance policies. They should be aware, however, that different levels of elasticity apply to different types of care and that all estimates have a wide margin of error. In addition, it is wrong to assume that cutting health care expenditures through cost-sharing leads to less demand for inappropriate care only, with equal or greater demand for appropriate care. Increased cost-sharing

decreases the demand for both appropriate and inappropriate care.

4. Policyholders should be educated about the nature, goals, and workings of insurance. Insurance companies are seen as profiting from an unfair relationship with the policyholder, where the policyholder pays premiums but gets nothing in return. Specifically, insurers need to communicate that premiums are the price of risk-reduction, not advance payments for care. The money collected through premiums is mostly used for paying for (other) people's care, and thus does not disappear into the CEO's pockets. Additionally, the perceived likelihood of successful audits can be increased through publicizing detected cases of fraud.

Abstract

This Panel Paper reviews recent evidence on moral hazard in the insurance industry. We discuss three types of moral hazard and detail how each is an asymmetric information problem. For each of the types, we summarize the empirical evidence and discuss the policy implications that follow from it. The evidence for ex ante moral hazard (i.e., insurance-induced increases in risk-taking) is rather weak but suggests that people engage less in preventive behaviors that are costly and hard to maintain when they obtain insurance. The evidence for ex post moral hazard (i.e., insurance-induced increases in usage of insured services) overwhelmingly indicates that it exists. However, the exact size of this effect in the Dutch health care system remains to be empirically determined. The numbers on insurance fraud indicate that it poses a significant problem. Furthermore, insurance fraud is deemed acceptable and common by many policyholders, and most seem unaware of the nature of insurance. Despite the large amount of data that has been gathered and analyzed, policyholders often lack the knowledge required for accurately predicting what effect policy changes will have. An important implication that follows from this is that insurance companies need to collect more data and run policy-experiments to gain the knowledge they need.

1. Introduction

In the past five years, health care spending in the Netherlands has continued its upward trend. Going from €70.7 billion (13.1% of GDP) in 2006 to €90 billion (14.9% of GDP) in 2011, the demand for health care places an increasing burden on the Dutch economy and its citizens (CBS, 2012). A considerable share of the increased demand is likely due to the aging of the Dutch population, and could have been foreseen. The combination of this trend, however, with the current financial climate, has led to an unprecedentedly loud call for budget cuts in the health care industry.

The budget cuts are also part of the Dutch political discussion, as health insurance coverage is mandatory under the Dutch law. Politicians thus decide which treatments should (at least) be covered under the universal policy, and which should not. Although it is clear that political parties disagree *which* budgets should be cut, they all agree that spending should be reduced. One important factor in deciding which budgets to cut is to determine where money is spent unnecessarily. While much of the discussion centers on content – how valuable given treatments for certain diseases are, how much of care for the elderly should be covered – we focus on a more universal problem that leads to overspending. We discuss three different types of moral hazard and explain how each of them contributes to the increasing burden on the Dutch health care system.

For each of the three types, we start with an explanation of how that type of moral hazard affects the behavior and/or demand of a given insured person. Then, we review the available empirical evidence and discuss what those findings imply for policy and policymakers.

2. What is moral hazard?

Moral hazard broadly refers to the increase in people's use of a service when it is covered by insurance, compared to when it is not. We will use doctor-patient situations and related health risk examples at different points in this paper to illustrate what specific theories or predictions entail. A general example would be that when people have health care insurance they are more likely to visit a doctor compared to when they do not have health care insurance.

Moral hazard theory describes what happens when one party is partly insulated from risk because another party agrees to wholly or partly indemnify losses that the first party might suffer. In a health insurance context, for example, everyone is at risk of needing medical care at some point. This care, of course, costs money, and the insurer agrees to pay for it provided that the insured pays a monthly premium. The main point of moral hazard theory is that insurance removes all or part of the incentive to restrict the use of insured services. With full insurance (i.e., 100% of the costs are covered) there is no financial reason for someone with insurance to not visit the doctor.

The presence of asymmetric information underlies the moral hazard problem (Arrow, 1963). The asymmetry occurs because the insurer has less information about the health status and reasons for health care usage of the insured than the insured themselves do. The insurer cannot check whether a physician visit would not have been made without insurance (when the visit would cost the insured more), or because of some avoidable risks that the insured took. The insured may even have faked his doctor visit to obtain money from the insurance company. In each of these examples, the insurer is unable to perfectly observe the risks that the insured

takes and adjust the contract accordingly. It is important to note that these risks need to be under the control of the insured. If they are, moral hazard may occur: The insurer prefers that the insured take precautions to avoid risks, but cannot enforce this wish. The insured is able to take risks without this being priced into the insurance contract and this asymmetry may lead to greater usage of insured services.

Large theoretical frameworks elaborate on the market failures that may occur as a result of asymmetric information between insurer and insured (Akerlof, 1970; Arrow, 1963; Ehrlich & Becker, 1972; Pauly, 1968; Rothschild & Stiglitz, 1976). These market failures include the complete eradication of insurance but also more subtle forms where asymmetric information leads to sub-optimal levels of insurance. Empirical tests of moral hazard theory in the insurance industry have become more prevalent recently, and three types of moral hazard dominate the discussion.

2.1 Different types of moral hazard

The different health care examples described above hint at the fact that moral hazard may occur for different reasons. This section describes *ex ante* moral hazard, *ex post* moral hazard, and insurance fraud. These are theoretically, psychologically, and behaviorally distinct ways in which health insurance might increase the use of medical services. It should be noted that these types of moral hazard also apply to different domains.¹ We focus on health insurance specifically, as this is a domain in which moral hazard theory has spurred much discussion (Savedoff, 2004) and where it may have far-reaching practical consequences.

1 For example, workers compensation (Dionne & Michel, 1991), auto insurance (Chiappori & Salanie, 2002; Richaudeau, 1999) and even pensions (Keating, 1991).

To illustrate why the first type of moral hazard is predicted, we present a formula on the basis of which people are predicted to determine their actions:

$$\text{Expected loss} = p \times L$$

where p represents the probability or likelihood of the loss whereas L represents the size of that loss. Running the risk of over-explaining: if the risk of getting sick during a given period is 5% (p) and it costs 100 Euros to visit the general practitioner (L), the expected loss is $.05 \times 100 = 5$ Euros. Economic theory assumes that people are rational utility maximizers. This means that they will try to find the optimal solution to this formula. Let us stick with an expected loss of 5 Euros. For the purposes of this example, we will not specify a full model but will simply assume that the expected loss of 5 Euros is an optimal solution for a given insured person named Jeff. Assuming Jeff is a rational utility maximizer, he should change his behavior when p or L changes to return to the – optimal – expected loss of 5 Euros.

Now look at what happens when health insurance is introduced. By compensating for the costs of visits to the general practitioner (GP), insurance basically reduces L . If insurance covers 90% of the costs of the visit, L becomes $.10 \times 100 = 10$ Euros – the amount that Jeff will need to pay the doctor. If insurance by itself does not change the likelihood that Jeff gets sick (p) the expected loss becomes $.05 \times 10 = 50$ Eurocents. Jeff – the rational utility maximizer – does not like this suboptimal expected utility and will change his behavior to return to the optimal level of expected utility.

Behaviors that are referred to as *ex ante moral hazard* change p and take place only before (hence, *ex ante*) an incident takes place. By taking more risk, people increase their expected loss.

For example, Jeff used to take vitamin pills and therefore only had a 5% chance of getting sick (the incident). Now that he is insured, he stops taking his vitamin pills and thereby increases his chances of being in a position to need a doctor. Without vitamin pills, the likelihood that Jeff needs to see a doctor is 50% (they were effective pills, indeed): $.50 \times 10 = 5$ Euros. So by refraining from using preventive care, Jeff returns to his optimal level of an expected loss of 5 Euros and thereby maximizes his utility. This negative effect that insurance has on preventive health behaviors is predicted to hold across behaviors that positively or negatively affect p , including smoking, exercising and alcohol consumption.²

The second type of moral hazard – *ex post moral hazard* – relates mainly to behaviors that change L and take place only after (hence, *ex post*) an incident has taken place. By using more of an insured service, people again try to maximize the utility they experience. Considering Jeff again, imagine that he is sick. Remember that without insurance Jeff pays 100 Euros for a GP visit. Given that Jeff has a limited amount of money to spend, he has to decide whether visiting the doctor is worth 100 Euros to him. If Jeff were seriously ill, he would probably think the visit is easily worth the price. If he just had a cold, however, he might decide to wait and see if the cold passes without visiting a doctor. In both cases, the rationality assumption implies that Jeff determines whether the utility he (thinks he) gets from a GP visit outweighs the costs.

2 A notable theoretical exception is the prediction by Ehrlich and Becker (1972). They mathematically examine the case where an insurance company observes risk-taking and prevention efforts perfectly, and prices insurance contracts accordingly. In this case, policyholders face an incentive to be careful when insured, as carefulness lowers their premium. This is such an unlikely assumption that we do not discuss its practical implications further.

Now let's look again at what happens when Jeff has insurance. Suddenly, visiting the doctor costs 10 Euros. To Jeff – who thinks that 100 Euros is too much to cough up for a cold-related GP visit – 10 Euros might seem like a fair price. With insurance, Jeff thus visits the doctor when he has a cold. Without insurance, he does not. By reducing the price of care, insurance is thus hypothesized to increase the demand for care.

In a related yet different scenario, imagine that when Jeff has a cold, he is willing to spend 100 Euros on a visit to the physician. When insurance reduces the price of the visit to 10 Euros Jeff has 90 Euros 'left' to spend on care. In order to maximize his utility, Jeff might demand additional tests from his doctor or make additional appointments that he does not necessarily need.

To someone who is not trained in economics it might seem strange to describe people's choices as a function of some utility calculations. Indeed, why would someone who is insured stop taking vitamin pills or see the doctor more often when sick? Modeling human behavior as a calculation-based enterprise is central to economic theory. Throughout the paper, we discuss the benefits and limits to economic theory and how its specific rationality assumptions sometimes prove hard to maintain when describing the behavior of irrational people.

The final phenomenon that resides under the moral hazard umbrella is insurance fraud. Unlike the other two types of moral hazard, fraud does not increase the demand for care. However, given that insurance fraud wouldn't be possible without asymmetric information – and cheating on insurance companies is deemed immoral – it is referred to as a moral hazard. The asymmetry of information comes into play because the insurer is unable to directly observe why the policyholder files a claim – so

whether the claim reflects genuine or faked use of insured services.

The simplest theory of insurance fraud states that a policyholder commits fraud when the benefits of committing fraud outweigh the expected loss. The expected loss is determined as the (subjective) probability of being caught, times the punishment that follows being caught (e.g., fines, exclusions, jail time). Often, rather than modeling the decision to commit insurance fraud, the propensity to commit fraud among policyholders is *assumed* to be non-negative. Much of the scientific literature on insurance fraud focuses on how fraud can be detected or what types of incentives should be built into insurance contracts to prevent insurance fraud.

Returning to the calculative Jeff: when he thinks he can get away with claiming reimbursements for a treatment he never received, he will do so if that is profitable enough. Similarly, when signing up for insurance, Jeff presumably knows his risks better than the insurer can observe them (e.g., family history of diseases, smoking and alcohol consumption, how often he exercises etc.). Jeff might hide some information in order to get a lower premium or greater coverage – which would have been denied him, had he shared his knowledge with the insurer. In both cases, Jeff uses the informational advantage he has over the insurer to his own benefit.

2.2 A short note on the morality of moral hazard

Before we consider the empirical evidence that has been gathered to test moral hazard theory, we briefly discuss the morality of moral hazard. Historically, economists and the insurance industry have disagreed on whether increased usage of services as a

result of moral hazard is immoral.³ Economists regard insurance-induced changes in behavior as rational and do not have a strong opinion about these changes – if anything, they consider rational behavior to be good because it leads to ‘optimal’ outcomes. In contrast, insurers, unsurprisingly, have considered increased use of insured services as immoral. Their reasoning is that people should not use something that they apparently do not consider to be worth its cost, simply because it’s free. We consider insurance fraud to be obviously immoral but is hard to condemn ex ante and ex post moral hazard similarly strongly. We will see that some instances of moral hazard are desirable, while others are clearly not. We thus refrain from taking a clear stand in this discussion and continue with describing the extent and strength of the evidence for each type of moral hazard.

3 We refer interested readers to (Baker, 1996) and (Rowell & Connelly, 2012) for in-depth analyses of how moral hazard got its name and how theorizing has changed over time.

3. Empirical tests of moral hazard theory

This section discusses the available studies according to the type of moral hazard they test. We start with ex ante moral hazard, continue with ex post moral hazard, and conclude by discussing insurance fraud. Each section opens with an explanation of what the specific type of moral hazard might look like in a health care context, followed by a discussion of the available empirical evidence, and closes with a discussion of policy implications. However, before a start can be made, we have to explain the biggest problem in empirical tests of moral hazard in real-world data: disentangling moral hazard from adverse selection.

3.1 Adverse selection

Asymmetric information is conducive to both moral hazard and adverse selection. Adverse selection means that people with high risk are more likely to opt for insurance than people with low risk. In health care, people who have a family history of sickness and disease are more likely to buy health insurance than people whose family members have lived long and healthy lives.⁴ Insurers cannot perfectly observe a person's health and are thus unable to adjust the insurance premium according to the expected health care expenses.

Although adverse selection has to do with the demand for health insurance, while moral hazard concerns behavior caused by insurance, both lead to the same outcome. When unhealthy people are more likely to buy insurance, insured people will turn

4 Note that adverse selection is different from insurance fraud. The former is merely the result of a correlation between health status and preference for insurance. The latter is unrelated to the health status of the prospective policyholder and implies that he is consciously shielding information from the insurer to obtain personal gain.

out to be unhealthier than uninsured people. As a result, the demand for health care will be greater for insured people than for uninsured people.

If insurance induces moral hazard and causes the insured to excessively use health care, the insured also demand more health care than the uninsured. If we wish to understand how insurance affects people's health decisions it is thus not sufficient to merely test whether the insured demand more health care than the uninsured. Such a comparison would leave it unclear whether the insured demand more health care because they were sicker than the uninsured in the first place, or because insurance lowered the cost of health care and thereby increased demand.

The scientific community has developed several ways to estimate the size of each effect independently (e.g., instrumental variables, natural experiments, longitudinal studies, or response to changing incentives within insurance plans; for an intriguing example, see Abbring, Chiappori & Pinquet, 2003). However, none of these are as good as randomized control trials – and we will see that different estimation strategies sometimes lead to different estimates. In the following sections it will become clear that uncertainty remains with respect to the size of the ex post moral hazard effect and the prevalence of insurance fraud. This highlights the need for randomized control trials in the insurance industry.

3.2 Ex ante moral hazard

This section takes a look at the empirical tests of ex ante moral hazard to find out whether insurance leads people to take more risk and exert less prevention effort. In the context of health care insurance, the studies mentioned here thus test whether Jeff –

upon signing an insurance contract – would indeed stop taking his vitamin pills, start drinking and smoking, or stop exercising.

3.2.1 Empirical evidence for ex ante moral hazard

Most work on ex ante moral hazard has been theoretical and there have been only a few empirical investigations into the matter (Kenkel, 2000; Zweifel & Manning, 2000). We review the articles that report empirical investigations of the ex ante moral hazard effect of health insurance.

The first study that we discuss here (Stanciole, 2007) analyzes data on health behaviors between 1999 and 2003 from the U.S. Panel Study of Income Dynamics (PSID). The PSID surveys close to 8,000 families, collecting data on a range of topics – including whether or not people have health care insurance, how healthy they are, and to what extent they engage in a healthy lifestyle. Stanciole (2007) specifically focuses on heavy smoking, heavy drinking, sedentarism (i.e., never engaging in light physical activity), and obesity.

A simple comparison of the lifestyle of insured (approximately 93% of the sample) and uninsured people would lead one to conclude that the uninsured live unhealthier lives than the insured – a case of reversed moral hazard! However, we have seen that a simple comparison between the two groups is not sufficient. The richness of the data provided by PSID allows the researcher to try and filter out many factors that affect people's lifestyle choices and insurance decisions (e.g., income, education, health, area of residence, and so forth). Such procedures make it more likely that the true effect of insurance on lifestyle choices is extracted. Using an elaborate econometric model, Stanciole (2007) finds that being insured is associated with a greater likelihood of heavy smoking, sedentarism, and obesity, but less heavy drinking

– controlling for the many factors mentioned above. The author then concludes that insurance causes policyholders to engage in unhealthy behaviors. However, the finding is a simultaneity rather than causality. Even though 93% of the sample was insured, and despite the attempt to filter out any selection effects, the presence of adverse selection cannot be ruled out.

A similar study uses data from 6,000 households (approximately 15,000 individuals) stored in the British Household Panel Survey on smoking behavior and the answer to a yes/no question 'do you play sports, go walking or swimming at least once a year' (Courbage & de Coulon, 2004). Approximately 10% of the population opts for private insurance coverage in addition to the publicly funded national health care system. The effect of this additional coverage on 'preventive' health behaviors is estimated. Note, however, that the measures were taken from a dataset that was not specifically designed to test for ex ante moral hazard, which means that the measures of preventive behaviors are therefore rather crude.

The data do not support the ex ante moral hazard prediction. In fact, people with private insurance coverage are more likely to exercise and are less likely to smoke. The authors carry out an additional test, estimating the effect of private insurance on the likelihood that women undergo a breast cancer screening or seek cervical smear testing. Both are preventive behaviors and the ex ante moral hazard hypothesis thus predicts that insurance decreases the likelihood of these screenings. The complication, however, is that these preventive behaviors are covered by the insurance policy. The ex post moral hazard would thus predict an increased use of these services among the privately insured – and indeed a small positive effect for breast cancer screening is observed (Courbage & de Coulon, 2004).

A limitation to this study is that it controls for far fewer factors than the previous one (Stanciole, 2007) did, and it suffers from a similar selection problem. That is, the authors cannot exclude the possibility that people opted for insurance because they were unhealthier than the people who did not opt for insurance. With observational datasets such as these, causal effects may be estimated using 'instrumental variables (IV)'. We will not go into the details of this technique, but it is useful when randomized control trials are impractical or impossible. However, the IV estimate also shows that private insurance does not lead to less exercising or more smoking. In sum, this study thus contradicts the ex ante moral hazard prediction.

A very specific and interesting test of the ex ante moral hazard effect was conducted after different U.S. states – at different times – adopted mandates that required basic insurance policies to cover “supplies, services, medications, and equipment for treating diabetes [...] without charging higher premiums” (Klick & Stratmann, 2007). If insurance coverage of diabetes treatments causes diabetics to reduce their prevention efforts (i.e., watching their diet) then their Body Mass Index (BMI) should increase. The researchers use BMI data from the Behavioral Risk Factor Surveillance System accumulated between 1996 and 2000, and tested whether BMI went up for diabetic individuals after the state in which they lived adopted an insurance mandate. Astonishingly, they find that when states adopt such mandates the average BMI of diabetics increased by 1.7 points, or 6 percent! Importantly, the BMI of non-diabetics did not respond to the mandates. It thus seems that diabetics trade prevention for treatment, possibly because prevention is so costly. Namely, prevention of diabetes worsening is strenuous work – it is hard to avoid sugary foods and to stick to a healthy diet. Furthermore, the treatment of diabetes–

related problems is effective. It is thus unclear to what extent this finding generalizes to other specific diseases and treatments for which prevention efforts are less costly or treatment is less effective.

Recently, another highly specific test of ex ante moral hazard was conducted on data obtained in a malaria-prone area in Ghana (Debebe, van Kempen & de Hoop, 2012). There, insurance for malaria treatments was introduced and the authors tested whether this led to a reduction in six different malaria prevention activities (e.g., repellent use and expenditure, presence of mosquito-proof windows, use of mosquito nets). Interestingly, insurance for malaria treatments only reduced the use of mosquito nets that have to be maintained by the owner – a demanding and time-consuming task that involves treating the net with insecticide. For 'factory-nets', which do not require treatment with insecticide by the owner, no reduction in use was observed. This finding, together with the Klick and Stratmann (2007) study, indicates that people reduce their preventive efforts only to the extent that it is hard to engage in them.

By far the best and least confounded test of whether insurance leads to less healthy behavior was conducted by Dave and Kaestner (2009). In the U.S., Medicare provides the opportunity to everyone over who is 65 years of age or older to enroll in an insurance plan. Compared to private insurance plans, Medicare is very generous and cheap. In fact, part A (hospital insurance) is free and the optional part B (medical insurance) costs very little per month. Part C involves the option to receive Medicare benefits through their private insurance plans and part D deals with prescription drug plans (both are optional). Unsurprisingly, the participation rate in plan A and B among people eligible for Medicare was 97.13% in 2008 (New Geography, 2008; State Health

Facts, 2012). Dave and Kaestner (2009) describe the provision of Medicare as causing an 'exogenous' shock in health care insurance. Since virtually all adults over 65 opt for Medicare there is no selection problem in testing the effect of insurance on health behavior. That is, people do not choose to (not) enroll in insurance; they are 'assigned to insurance'. The researchers exploit this situation to estimate the pure effect of insurance on health behavior using data of 3,396 people from the Health and Retirement Study, collected in eight measurement waves. These data are specifically collected for testing hypotheses such as moral hazard and thus include a vast amount of information at the individual level.

The authors note that physicians generally encourage healthy lifestyle choices and discourage unhealthy alternatives. Confirming the ex post moral hazard hypothesis, obtaining Medicare increases the number of visits to the physician – thereby leading Medicare enrollees to live healthier lives. The authors then cleverly control for the positive effect of increased physician contact on health behavior to estimate the true effect of insurance on preventive behaviors.

Even though this study avoids the selection problem, it has to deal with the fact that many people have insurance before they are eligible for Medicare. This means that they can only test the effect of insurance on people who have never had insurance prior to obtaining Medicare. This narrows the population under investigation drastically and introduces a confounding variable. People who have never been able or willing to buy insurance are likely to differ from people who have had insurance in many ways. The authors attempt to control for these differences by filtering out the effects of age, sex, race, employment status, wealth, and so forth, but it remains hard to determine how well they capture

pre-existing differences between the enrollees that have and those that have not had insurance.

Controlling for the positive effect of increased contact with physicians, an *ex ante* moral hazard effect is found for men (who had never been insured prior to obtaining Medicare). Most notably, when those men enroll in Medicare, they become almost 40% less likely to exercise, 18% less likely to stop smoking, 15.8% more likely to smoke daily, and 14.8% more likely to drink daily. No such effects were observed for women who enrolled in Medicare. More importantly, the positive effect of increased contact with physicians on these behaviors fully counteracted the negative effects of insurance. The net effect of insurance on health was therefore predicted to be zero. In conclusion, *ex ante* moral hazard is of little concern in the Medicare system – the *ex ante* moral hazard effect that is found occurs for only a small subset of the insured, and it is not clear that the finding generalizes to the entire Medicare population.

Although the Dutch health insurance system is rather different from some of the systems in which the studies above have been conducted, we now elaborate on how the described findings could inform Dutch policymaking.

3.2.2 Policy implications of the findings on ex ante moral hazard

Unfortunately, the empirical evidence thus far does not paint a clear picture of whether, when, and why *ex ante* moral hazard occurs. One obvious 'policy' implication is that insurers should work together with academics to come up with new ideas to test the *ex ante moral hazard* hypothesis, while using data from insurance companies. This will advance both theoretical and practical insights into whether, when, and how people's risky

decisions are affected by insurance. We highlight the importance of small-scale experiments again later in this section.

Despite considerable variability in estimates of the size of the ex ante moral hazard effect at least two conclusions can be drawn. First, insurance does not seem to cause excessive risk-taking. Rather, ex ante moral hazard effects are only found for preventive behaviors. The findings would thus allow us to predict that Jeff is less likely to take his vitamin pills when he obtains insurance – but not that he would start to smoke or use drugs. Namely, even the effects in studies that test for perverse effects of insurance on drinking behavior and exercising might reasonably be construed as reductions in preventive behaviors: Upon enrolling in Medicare, elderly men were less inclined to stop or reduce their smoking and alcohol intake.

Second, the empirical evidence suggests that ex ante moral hazard is most likely to occur for diseases where people have to take many medications, follow a strict diet, or exercise a lot (Debebe et al., 2012; Klick & Stratmann, 2007). These behaviors are very 'costly', in that they require a lot of time, effort, and devotion on the part of the insured. To give some concrete examples, asthma patients are required to use their asthma inhalers twice a day, every day; people with heart problems have to exercise regularly to stay fit; and people with high blood sugar have to follow restrictive diets. Such preventive efforts are hard to maintain, and the reviewed evidence suggests that these behaviors are especially susceptible to ex ante moral hazard effects. Insurers that cover the treatments for these diseases would therefore be well-advised to look into procedures that enforce people's adherence to these preventive measures. The reasonable thing to do for companies that are trying to promote these healthy behaviors is thus to eliminate or at least try to

reduce some of the hassle associated with preventing the disease. In the case of asthma, companies should allow their clients to sign up for a daily e-mail or texting service that reminds them to take their medication. Similarly, people with heart problems could be encouraged to define an exercising schedule or to develop a routine. Adherence to the schedule could be enforced in a similarly simple fashion. These are relatively cheap and likely effective measures to mitigate ex ante moral hazard.

The studies described in this section of the paper also provide guidelines that help in determining how much policymakers should worry about ex ante moral hazard when new treatments become available. In deciding whether or not to cover the treatment, and to what extent incentives to counter ex ante moral hazard should be provided, it would help to determine how much trouble it is for people to engage in preventive efforts. If taking precautions against a specific disease is easy, it is unlikely that providing coverage for the treatment of that disease diminishes the extent to which people engage in them. In contrast, if people have to exert a lot of effort to prevent a given disease, they are more likely to be responsive to the incentive effects of insurance.

Looking at the current literature, the findings suggest that even if there is an ex ante moral hazard effect, it is probably rather small. In addition, it is likely that the net effect of insurance on preventive efforts is zero. This follows from Dave and Kaestner (2009) who find that insurance also leads to healthier lifestyle choices through increased contact with physicians. It is important to note that this does not mean that there is no reason to try to prevent ex ante moral hazard. Namely, the net effect of insurance might turn out to be positive if insurers could prevent the decrease in healthy lifestyle choices and/or preventive efforts

while maintaining the positive effect of increased contact with physicians.

We end this policy recommendation section with a word of caution and a call for experiments. We were unable to find tests of ex ante moral hazard in the Dutch health care system and it is unclear whether the findings in the studies we described are directly applicable to the Dutch system. In contrast to some of the systems that have been studied, health care insurance is legally required and thus nearly universal in the Netherlands. Even if there were a test of ex ante moral hazard in the Dutch system, there would be no uninsured group to compare the Dutch behavior to. This would cause difficulty with estimating the ex ante moral hazard effect.

One possibility is to look at how different *levels* of coverage affect people's health decisions. Then, however, the problem of disentangling adverse selection from moral hazard arises. This highlights the need for insurers to engage in small-scale experiments in which they randomly vary the coverage for certain treatments for some (randomly selected) group of policyholders. Randomized control trials are the golden standard in scientific research and are the best method to obtain reliable estimates of the effect under study. In the case of ex ante moral hazard, an example would be to increase diabetes treatment coverage for a randomly selected group of diabetics and non-diabetics, and then measure their BMI and how well they stick to their diet. Our guess is that few policyholders would decline greater coverage at the same price and there are nearly limitless variations on this experimental design that could be applied to test predictions in specific domains of interest. Running experiments is the best way to gather the knowledge that is needed to inform policy.

3.3 Ex post moral hazard

This section reviews the empirical tests of *ex post* moral hazard. Holding preventative behavior constant, does insurance lead to an increased demand for care? Again, we specifically focus on health insurance, and these papers thus test whether demand for care increases when insurance lowers the price of care. In contrast to *ex ante* moral hazard – which is mainly perceived as malignant – *ex post* moral hazard may be construed as both malignant and benign. The malignancy argument holds that insurance induces people to use care that they otherwise would not have bought for the same price. It thus leads to a welfare loss because the consumption of care costs more than its real worth to the consumer. The benign argument holds that being uninsured leads people to forego even basic and necessary care, which makes some increased use of care under insurance appropriate (Nyman, 2004). In addition, especially if care is cost-effective and the treatment does not have negative side effects, more moral hazard is better (Pauly & Held, 1990).

Most *ex post* moral hazard tests use existing datasets and exploit some sort of exogenous variation in the price of care – for example, when co-payment rates for physician services rise in one group of people but not in the other (Chiappori, Durand & Geoffard, 1998), when turning 65 makes U.S. citizens eligible for Medicare (Dave & Kaestner, 2009), or when co-insurance rates change nationwide (Cockx & Brasseur, 2003). It is then determined how much of the change in demand for care is attributable to these changes in price. Unfortunately, many of these studies are unable to directly estimate the causal effect of the price change on the change in demand for care. This might be due to the absence of an appropriate control group, the cross-sectional design of the study, or the inability to observe the ‘unobservables’ that interact

with the price change and the demand for care. Consequently, a review that includes multiple studies that approach the ex post moral hazard test differently shows that the estimates of the ex post moral hazard effect are somewhat inconsistent. However, most studies conclude that demand for care goes down when the price of care goes up – but only by a bit (Skriabikova, Pavlova & Groot, 2010).

Another problem that troubles research on ex post moral hazard is the limited availability of data. Selecting and finding the right information is frequently a difficult undertaking. Most insurance companies do not keep track of the health behaviors of their beneficiaries and they often have limited information on their backgrounds. In addition, data may be inaccessible for researchers because of privacy laws. Many articles use data from nationwide databases or very 'local' databases such as data from one particular insurer. These have the problem that they may not contain all the information the researcher needs to answer his questions (e.g., only 'proxies' for demand and health behavior) or the generalizability of the sample under study may be questioned.

3.3.1 Empirical evidence for ex post moral hazard

On July 1st, 1993, a law was passed in France that increased the co-payment percentage in the public health care plan. Most for-profit insurers that provided supplementary coverage responded to this law by also introducing cost-sharing in their plans. Importantly, some of these plans were paid for by firms on behalf of their employees (ruling out adverse selection effects) and some firms decided not to implement the cost-sharing. This situation provided an interesting natural experiment where the employees of one firm kept their insurance plan without cost-sharing, while

employees of a similar firm suddenly faced 10.4% cost-sharing for physician services (Chiappori et al., 1998).

In testing the ex post moral hazard hypothesis the researchers distinguish between office visits, home visits (i.e., the physician goes to the home of the sick), and specialist visits. A comparison was made between the amount of visits per person in the year before the implementation of the law and in the year after it. The number of office and specialist visits remained relatively constant in both groups. However, the average number of home visits dropped by 18.75% only among employees of the firm that introduced cost-sharing.

Chiappori et al. (1998) conclude that small monetary costs may not induce large changes in the demand for care because of significant non-monetary costs associated with office and specialist visits (waiting time, etc.). Specifically, they state that the relative size of the monetary cost is not large enough to cause changes in these visits. An alternative explanation could be that 'luxury' care (i.e., having the doctor visit you, rather than the other way around) is more sensitive to price changes than basic care. In addition, the beneficiaries were not randomly assigned to their insurance plans. Therefore, it may be questioned whether the employees of firms that introduced cost-sharing are directly comparable to those of the firm that did not.

A study conducted in the U.S. also distinguishes between different types of care and also finds differences in the extent to which ex post moral hazard applies to them. The data come from the Medical Expenditure Panel Survey and all doctor visits in the year 2002 are coded to be either a general physician visit or a specialist visit. In addition, for every visit it was determined whether it was related to an acute, chronic, or no disease (Koç, 2010). The specialist visits may include episodes related to

anything ranging from physiotherapy (no disease) to respiratory or heart problems (acute or chronic). Elaborate statistical techniques are used to address the endogeneity problem of insurance – people are not randomly assigned but rather choose to be publicly, privately, or not insured at all. Therefore, it remains troublesome to estimate the true causal effect of insurance on the demand for care.

Analyses on the 17,419 observations (of people between the ages of 18 and 64) show the highest moral hazard estimate for specialist visits that are unrelated to diseases. Demand for these visits for those with private insurance is 231% that of the uninsured. In comparison, this number is 'only' 139% for specialist visits related to chronic diseases. There is much less variability in the moral hazard estimate for general practitioner visits across the different types of visits (104% to 151%). This finding makes intuitive sense, as the visits with highest moral hazard are the ones most easily deferred (or perhaps even unnecessary), while the visits with lowest moral hazard are hardly deferrable and most likely necessary. If cost-sharing reduces moral hazard, a policy implication that follows from this (and the Chiappori et al. (1998) finding) is that the level of cost-sharing should be different for different visits. However, while Chiappori et al. (1998) did not find an ex post moral hazard effect for specialist visits, Koç (2010) did.

In the ex ante moral hazard section we described the study that exploited the exogenous variation in health care due to Medicare eligibility (Dave & Kaestner, 2009). Both males and females who were uninsured prior to Medicare had more doctor visits, as well as an increased likelihood of visiting a doctor once and having a hospital stay upon receiving Medicare. That finding is in line with the ex post moral hazard hypothesis.

Recall that *ex ante* moral hazard predicts that people engage less in prevention when they have insurance. *Ex post* moral hazard predicts that people demand more care. So what happens when the prevention of a given disease is a form of care? Consider preventive screenings such as breast or prostate examinations, vaccinations, and other prophylactic measures. These treatments might reasonably be considered to be preventive efforts and thus, following from research on *ex ante* moral hazard, one would expect lower demand for these treatments (i.e., less prevention) among people who have insurance compared to those who do not. In contrast, it is equally defensible to consider these treatments as instances of care and therefore predict that, consistent with *ex post* moral hazard, the demand for these treatments would be higher for people who have insurance compared to those who do not.

The empirical evidence is in favor of the latter hypothesis – that is, greater use of preventive care by the insured. This evidence comes from a study conducted using data from 12,100 Mexicans over the age of 50 in the Mexican Health and Aging Study (MHAS) to test for differences between the insured and uninsured in the extent to which they use preventive services (Pagan, Puig & Soldo, 2007). This study (unlike Dave & Kaestner, 2009 and Stanciole, 2008) does not find differences in smoking and drinking behaviors. However, a positive effect of insurance on the use of insured services is found. The uninsured in the sample are less likely to use preventive screenings for hypertension, high cholesterol, diabetes, and three types of cancer. This finding is not very strong, as only the effect of insurance on cholesterol and diabetes screening remains significant after controlling for other health risk factors. Nevertheless, this study definitely provides

no evidence for reduced use of preventive medical care, whereas there is some evidence for increased use.

It is clear from the studies described above that the ex post moral hazard hypothesis has truth to it. It is fair to say that the demand for care is responsive to the price of care. When insurance reduces the price of care, demand goes up. However, many studies point out that it is hard to draw firm conclusions from the data they present because their nature remains correlational. Next, we discuss a study that circumvents this problem by randomly assigning families to different insurance plans. However, before we do so, we explain a common metric to measure the size of ex post moral hazard: elasticity of demand. Readers with a background in economics may feel free to skip this section.

3.3.1.1 Demand elasticity

Although not every paper on ex post moral hazard estimates the price elasticity of the demand for health care, it is a very important number in policymaking. Demand elasticity gives insight in the extent to which demand for health care changes when the price goes up. The price elasticity of demand for a certain good is calculated using the following formula:

$$\text{Demand elasticity} = \frac{\% \text{ change in quantity}}{\% \text{ change in price}}$$

To illustrate what this formula means in the context of health care we return to Jeff. We know from the previous example that with insurance, Jeff has to pay 10 Euros to visit the doctor. For the purpose of this example, let us say Jeff visits the GP five times a year. The next year, the insurer decides to raise the price of physician visits to 15 Euros (a 50% increase), and Jeff only visits the GP four times (a 20% decrease). This means that the elasticity of Jeff's demand for health care is $-0.20 / 0.50 = -0.40$. This

number aids policymaking, as it helps predicting the changes in demand if deductibles or coinsurance rates would be increased or decreased. An intuitive way of interpreting the elasticity in this example is to look at it this way: If the price of a physician visit is doubled (100% increase), demand goes down by 40%.

Of course, policymakers are not interested in Jeff per se; rather, they wish to know what a population of Jeffs would do. So let us assume a demand elasticity of -0.40 and a population of a million Jeffs. With insurance, this population has five million physician visits every year (i.e., each member visits the physician five times a year). Increasing the price of a physician visit from 10 to 12 Euros (20%) would lead to $-0.40 \times 0.20 = 8\%$ fewer visits. At the population level, this thus means $5,000,000 \times 8\% = 400,000$ fewer physician visits.

Often, articles report *arc* demand elasticities where changes in both quantity and price are measured as percentages of the average of the old and new price/quantity. The elasticity we refer to above is a point estimate and reflects the elasticity of demand for health care given the current situation. This number is thus important for determining the effect of changes in a given health care system. Arc elasticities, however, refer to the elasticity of demand along a 'range' and might be seen as the best guess of the mean elasticity between two points on the demand curve. It is calculated as follows:

$$\text{arc elasticity} = \frac{\left(\frac{\text{new demand} - \text{old demand}}{(\text{new demand} + \text{old demand}) / 2} \right)}{\left(\frac{\text{new price} - \text{old price}}{(\text{new price} + \text{old price}) / 2} \right)}$$

This number is of interest when designing a completely new health care system, or when contemplating large changes in an existing health care system. Depending on the specifics of the

situation, arc elasticities may be different from regular elasticities (e.g., arc elasticity of Jeff's demand for health care would be -0.55 instead of -0.40). However, the interpretation for both types of elasticity is similar: both give an estimate of how much demand decreases when price increases.

3.3.1.2 The RAND experiment

The RAND Health Insurance Experiment (Newhouse, 1996) is the pinnacle of experimental tests of moral hazard in the insurance industry and gives unique insights into the price responsiveness of the demand for health care. The Experiment was conducted between 1974 and 1982 and involved the random allocation of people in four large cities and two rural areas in the U.S. to different insurance plans. Approximately 2,750 families composed of 7,700 individuals participated for three to five years. Families were randomly assigned to different insurance plans so that researchers would be able to estimate the effects of insurance on the demand for medical care, without having to worry about selection problems. More specifically, each family was assigned to one of four fee-for-service plans with different levels of coinsurance (0%, 25%, 50%, and 95%; this is the percentage of the medical care costs that the insured pays for himself) or an HMO-style plan in which participants faced no cost-sharing and always needed a referral from their primary care physician to visit a specialist. The researchers also varied the maximum amount of out-of-pocket health care expenditures (5%, 10%, or 15% of annual income; the insurance covered all expenditures over this limit). The study tested how the different insurance plans affected the demand for hospital care, outpatient medical care, care related to acute and chronic episodes of sickness, dental care, and well-care (a label used for deferrable care such as immunizations,

examinations, and birth control – deferrable does not mean unnecessary). The incredibly large amount of data gathered in this experiment has been described in more than 300 scientific articles and numerous reports. The interested reader is referred to a detailed book by the lead researchers in which they discuss the policy implications of the RAND experiment (Newhouse, 1996). In this section we highlight the most important findings with respect to how different levels of (co)insurance affect the demand for health care.

The main conclusion that may be drawn from the RAND experiment is that the demand for health care is sensitive to changes in price. The overall elasticity of demand for care health care is approximately -0.20 . Yet, co-insurance and deductibles are crude and unrefined in the way they cut back health care expenditures. In addition, although the demand for health care decreases as the co-insurance rate increases, the largest decrease in demand is observed when co-insurance rises from 0% to 25%. This finding holds across almost all types of care, ranging from treatment of chronic diseases, to in- and outpatient care, dental care, and even to acute care such as emergency room visits.

There are some findings in the RAND experiment that require some elaboration. First, the researchers have gone to great lengths to test whether co-insurance increases *only* the demand for effective health care (something standard economic theory would predict). To this end, they had a panel of doctors go through a large portion of the medical files in the experiment to determine whether – given the disease or symptoms a particular patient had – the obtained care was appropriate or not. They found that co-insurance (across the different levels) decreases the demand for inappropriate care as well as for appropriate care. One might use this finding as an argument for free health care for all, to

ensure that everyone gets the (appropriate) care they need. However, the authors are quick to point out that this finding may well be framed in opposite manner. Free care increases the demand for both inappropriate and appropriate care. Inappropriate care not only causes precious health care resources to go to waste; it may actually have negative health effects (side-effects of unnecessary drugs, complications that occur during superfluous operations, and so forth). The decreased use of health care is thus not unequivocally positive or negative.

Second, not all types of health care were equally sensitive to changes in price. More specifically, treatments that were deemed most easily deferrable – such as dental care and well care – showed the highest sensitivity to price (arc elasticities were -0.30 and -0.43 respectively). Historically, the economic argument has been to have greater coverage for treatments that are relatively less price elastic because they are less prone to perverse (ex post moral hazard) effects. The assumption thus is that the larger the elasticity, the higher the proportion of inappropriate treatments. Looking at the price elasticities in the RAND experiment, one would then propose that dental and well care should have great levels of co-insurance, as their price elasticity is high. However, the relatively high level of price elasticity in these two categories is not due to a larger number of inappropriate treatments. Rather, people simply defer these treatments because they can. Again, the policy implications that follow directly from basic economic theory require some care.

Third, besides differences in elasticity between types of care, there is also considerable variance in estimates of elasticity for the same type of care along the range of cost-sharing possibilities. For example, between 0 and 25% co-insurance, the arc elasticity of demand for outpatient care is -0.17 while it is -0.31 between 25%

and 95% co-insurance (Newhouse, 1993). This dissimilarity makes it difficult to use the estimates of the RAND experiment in health care systems that already have 45% cost-sharing, for example.

An important finding in the RAND experiment is that cost-sharing barely affects health outcomes for most people in the study. This combats the argument that people forego so much appropriate care under cost-sharing that it increases costs in the long run – through an increased need for more complicated and expensive care when untreated health problems deteriorate. Put differently, the free plan did also not markedly improve the health of its beneficiaries. This is likely due to the increase of inappropriate care on the free plan that – through side effects and complications – counteracted the beneficial effects of the additional appropriate care received. One important caveat is that those who were sick and poor at the beginning of the experiment did experience negative consequences from the cost-sharing plan. They exited the experiment with worse health than did the sick and poor on the free plan. This finding is corroborated by other studies that find that price elasticity of demand for medical care is highest for low-income people (Skriabikova et al., 2010).

Importantly, the long-term effects of being on the free plan or cost-sharing plan could not be tested because the maximum period of time people were on one of the experimental plans was five years. Indeed, a recent review of studies on the causal effect of health insurance on the demand for care and health outcomes (in U.S. samples) finds that insurance does improve health (Freeman, Kadiyala, Bell & Martin, 2008).

A final note on the RAND Experiment is that it did not test the ex post moral hazard hypothesis in the traditional sense: it did not compare the behavior of the insured to that of the uninsured. While the data are valuable for estimating the size of the effect of

different levels of cost sharing, it is not clear what the demand for care of the uninsured would look like. However, it is clear from the data that there is no linear relationship between the level of cost sharing and demand for care. Therefore, extrapolating the RAND Experiment finding to no insurance is difficult, and estimating the net effect of some health insurance compared to no health insurance based on the RAND estimates remains tricky (but see Keeler, 1988, for an attempt to do so).

A short note on demand elasticity. In the introduction of the section on ex post moral hazard we touched briefly upon the relative disagreement between different articles on the size of demand elasticity. A review of studies that estimate the elasticity of demand for physician services in U.S. finds that estimates are most often between -0.10 and -0.50 (Skriabikova et al., 2010). However, estimates range from smaller than -0.10 to greater than 1. A review that includes estimates from natural experiments as well as cross-sectional studies observes a similar disparity in the estimates and concludes that “the demand for health care falls with increases in out-of-pocket costs [...] the magnitude of the estimated response varies widely, however” (Zweifel & Manning, 2000, p. 436). It is likely that the difference in the estimates between countries is due to differences between the specific health care systems. The best estimate of demand elasticity for a given country is thus one that is estimated using data from that same country. The next section reviews the (meager) empirical evidence on ex post moral hazard and the elasticity of demand for health care in the Dutch system.

3.3.1.3 Findings for the Dutch health insurance system

As noted, it is unclear to what extent the exact estimates of the elasticity of demand of one country or setting are applicable to

another. Alarming, there are instances of considerably different estimates within the same country (see Cockx & Brasseur, 2003; and for a Belgian example, Van de Voorde, Van Doorslaer & Schokkaert, 2001). To get a sense of whether the Dutch system is subject to the same (moderate) ex post moral hazard effect that is found in other countries, we searched for ex post moral hazard tests in the Dutch system.

We could find only two empirical tests of ex post moral hazard in the Dutch health care system and both were conducted by the same author (Van Vliet, 2001, 2004). In the first paper, data on health care utilization, health insurance status, and background variables from 6,039 families (13,362 persons) who were privately insured between 1990 and 1994 were obtained from the Dutch Central Bureau of Statistics (CBS; Van Vliet, 2001). The author cleverly determined the average payment rate (similar to marginal cost of care; dependent on expectations of exceeding the deductible and health status and the size of the deductible) to estimate that the demand elasticity for health care on average is -0.08 . There is almost no elasticity for hospital visits (-0.007) and relatively much for physiotherapy (-0.12). However, even the highest estimate is well below that of the RAND experiment. The author notes that this might be due to differences in the health care system between the U.S. and the Netherlands, and the relative size of health care spending in terms of GDP (9% in the Netherlands versus 14% in the U.S.; a higher share leads to higher price elasticity). Note that the sample under study is privately insured, and it is unclear whether the same results would hold for a publicly insured sample.

In a second article, Van Vliet (2004) estimates the elasticity of six types of care (general practitioner visits, physiotherapy, pharmaceutical care, medical specialist visits, hospital care and

'other') using data from a private health insurance company. For each of the 100,487 enrollees their expenditures are predicted if their deductible would be zero. Using the size of the deductible and the difference between the predicted and actual health care expenditures, the elasticity of the demand for each of the six types of care is estimated. Note that this procedure only allows estimating the elasticity of demand for the 48.7% of enrollees that have a deductible. The highest elasticity was found for GP visits (-0.40), followed by physiotherapy (-0.32) and the lowest for hospital visits (-0.04).

The same researcher – using two different datasets and two different methods – thus finds a rather low (-0.08 in 2001) as well as a more moderate (-0.14 in 2004) elasticity of demand for health care. In the 2004 paper the price elasticity for demand for visits to the general practitioner even is -0.4 while it is -0.09 in 2001. In both papers there is considerable variance in the demand elasticities across different types of care. The rank order of the estimates, however, is quite similar across the two papers.

These estimates leave the true demand elasticity of demand for care unknown. Yet, the findings would suggest that when policymakers aim to reduce health care costs, they should increase the cost of GP visits because in both papers the demand elasticity is highest for this type of care. However, physicians serve an important gatekeeper role in the Dutch health care system. Increasing the price of visits may lead people to circumvent the physician and visit a specialist right away. Specialist visits are generally more expensive, and reducing GP visits through increased cost sharing – a measure intended to cut costs – might thus lead ultimately to increased costs.

3.3.2 Policy implications of the findings on ex post moral hazard

It is clear from the findings reviewed above that ex post moral hazard exists. When the price of care goes up, demand generally goes down. However, we have seen that not all types of care are equally responsive to changes in price, and there are conflicting findings even for the same type of care (e.g., specialist visits). Worse, even within a health care system estimates of demand elasticity may differ wildly. While the difference seems small (the difference between estimates of elasticity for GP visits in the previous section is only -0.31), the effects on the population level are of huge importance. Namely, to decrease demand by a certain desired amount, low price elasticity would lead policymakers to increase deductibles or coinsurance rates substantially. In contrast, higher price elasticity would render a smaller increase sufficient. In a more extreme situation, low price elasticity may deter policymakers from introducing price increases at all, because they think that the additional burden placed on the insured does not outweigh the benefit of the slight decrease in demand.

Even though it is unclear whether the RAND experiment estimate is applicable to the Dutch situation, and whether it applies to the current (much more advanced) health care system, it seems that -0.20 is the least confounded estimate of the elasticity of the demand for care. We hasten to add, though, that the estimates for the elasticity of demand differ across the types of treatments. The treatments that were most easily deferred (without regard for their appropriateness) showed the highest price elasticity. In predicting the impact of certain policy changes, the RAND experiment estimates could thus be used, but with caution. At present, there is simply not enough evidence to determine with certainty the exact price elasticity of demand for care in the Dutch health care system.

A clear implication of the reviewed findings is that different levels of co-insurance may be desirable for different types of treatments. Again, however, determining the exact level of coinsurance remains tricky, as the available data do not render consistent estimates. In addition, a word of caution is warranted. Even if the true elasticity of demand were revealed, policymakers would still face tough decisions. The RAND experiment clearly demonstrated that increasing co-insurance or raising deductibles does not lead beneficiaries to select more appropriate care or cut inappropriate care. Co-insurance cuts the demand for care 'across the board'. Ideally, insurers would distinguish between appropriate care and inappropriate care, and reimburse only the appropriate care while having 100% coinsurance for inappropriate care. However, the same treatment might be appropriate in one case but inappropriate in the other, and checking the appropriateness of each claimed treatment is simply not an option. The insurance company thus needs data on which types of treatments are generally effective in which situation, in order to be able to adjust its coinsurance rates accordingly. This would require insurance companies to keep track of large amounts of data (e.g., health status, specifics of disease episode, treatment details, etc.), but this might prove worthwhile in the long run.

If insurance companies are interested in determining the exact elasticity of the demand for care, they should experiment a little bit. We have seen that tests of ex post moral hazard rely on datasets that have not specifically been created for the type of question that is asked. Often, proxies for health status (such as drug use, or number of physician visits in the previous year) have to be used and the distribution of insurance or different levels of coinsurance is endogenous. Insurance companies could boldly attempt a replication of the RAND Experiment, but this seems both

expensive and impractical. Smaller scale experiments, however, seem more plausible and might prove very valuable. For example, with recent and upcoming increases in the mandatory deductible (Rijksoverheid, 2012) it is simple to offer a randomly selected (representative) sample of beneficiaries a lower deductible for free. The rejection rate for such an offer would be low and this set-up allows the experimenter to compare the demand for health care in the treatment group to a (randomly selected representative) control group that is not offered the same discount on its deductible. Monitoring these two groups for one year – measuring their health status and demand for care (both insured and uninsured) – could yield reliable estimates of demand elasticity of care in the Dutch health care system. Similarly simple designs could be created for specific treatments by making them cheaper or free for certain groups of randomly selected people. Controlled, randomized experiments like those just described could drastically improve knowledge of – and thereby predictions based on – the elasticity of demand for health care.

3.4 Insurance fraud

Checking every claim that is filed for truthfulness is highly cost-inefficient, and therefore many insurers approve claims having to trust that policyholders file claims *only* when they are legally entitled to do so. Economic theory predicts that the insured take the opportunity to betray the insurer's trust when they believe the benefit of filing an illegitimate claim outweighs the probability and costs of getting caught. Insurance fraud thus does not lead to greater use of insured services but it does lead to more or larger claims, thereby increasing the burden on the insurance system. We have already established that fraud is an asymmetric information problem where the policyholder abuses the trust

that the insurer is obliged to place in him or her. The ways in which policyholders can exploit their informational advantage are classified as either "hard" or "soft" fraud. Hard fraud includes deliberately staging accidents or faking claims to obtain a payment from the insurance company. Soft fraud includes 'claim padding' and opportunistically misrepresenting the facts on a claim to obtain a higher payment or a lower monthly premium (Insurance Information Institute, 2010). In Jeff's case, hard fraud means he fakes his medical bills, while soft fraud means he has a medical bill and changes some numbers to obtain a higher payment. The distinction between the two is not always as clear as in the previous example, but it seems that soft fraud is harder to detect than hard fraud.

Importantly, the available estimates of the prevalence of fraud are usually based on 'proven' instances of fraud (although it is not always clear how the estimates are derived). Therefore, it is doubtful that available estimates of insurance fraud capture all of the fraud that is committed. Nevertheless, the reports below indicate that insurance fraud is a significant problem.

The Insurance Information Institute (2010) estimates that 3 to 10 percent of U.S. health care expenditures are due to fraud. This translates to a loss of between \$77 billion and \$259 billion – obviously, a rather wide margin of error. The Coalition Against Insurance Fraud (CAIF) states on its website that \$80 billion is lost annually to insurance fraud (across industries). However, they sum up health care fraud statistics from other sources that report estimates of 'at least 3 percent' of health care expenditures (\$68 billion) as well as '\$125 to \$175 billion annually' (CAIF, 2012). The Dutch association of insurers (Verbond van Verzekeraars, 2011) reports the number of detected fraudulent claims (313 in 2010) instead of a percentage estimate. The 'Convenant aanpak

verzekeringsfraude' (2011) reports that €7.5 million worth of fraud was detected in 2009. The Verbond van Verzekeraars (2011) reports a total volume of €84,524 million spent by health insurers. 0.000089% of health care expenditures would thus be due to fraud. These wildly differing assessments of fraud may reflect the effectiveness of the Dutch health care insurers (relative to U.S. insurers) in deterring fraud, a motivation on the part of U.S. insurers to make fraud seem as bad as possible, or a general lack of accurate information. Given that not all types of fraud are equally easily observable and – by its very nature – insurance fraud is concealed, we suspect that the latter is true.

While the literature on insurance fraud is growing, much of it focuses on the detection of insurance fraud and the deterrent effect of auditing (Derrig, 2002). Few articles examine the causes of insurance fraud or assume that asymmetric information and self-interest are the underlying causes. As a result, many articles are focused at particularities with respect to specific numbers or patterns within claims (e.g., Major & Riedinger, 2002) or present theoretical analyses of the effects of auditing (for an overview, see Picard, 2000).⁵ We are not aware of any studies that systematically test the effects of punishment. However, a relatively straightforward prediction would be that greater punishment is associated with less fraud.

We believe that one implication of the fraud statistics reported above is that it might prove worthwhile to look at the causes

5 Auditing means that the insurer randomly checks the truthfulness of a subset of claims. This deters fraud because the insured cannot anticipate whether or not he will be audited. This economic analysis of the effects of fraud assumes, however, that people like Jeff decide to commit fraud solely based on the consciously calculated odds that their claim is audited – an unlikely proposition. Until 2012, there had not been a test where auditing probability was randomly varied in the field, and the accuracy of the mathematically predicted optimal auditing models is often not tested.

of insurance fraud. A greater understanding of the motivation to commit fraud will lead to a better grip on how fraud can be prevented and might be more effective than punishing offenders after the fact – especially when it is so unclear how often fraud is actually committed. Therefore, we discuss the empirical literature on why people commit insurance fraud and discuss its policy implications.

3.4.1 Empirical evidence on reasons to commit insurance fraud

It is clear that the main reason for people to commit insurance fraud is that they want money. However, a focus on the financial gain could not explain why people commit insurance fraud per se; they might as well cheat on their taxes, steal a pair of shoes, or rig a boxing match they're betting on. Therefore, the evidence we discuss below centers around why people might feel that insurance fraud is justifiable. This way, we hope to shed light on the specific causes of insurance fraud to inform prevention strategies aimed at reducing fraud.

Before we discuss the scientific literature, we highlight a noteworthy investigation by the insurance industry. The CAIF (1997) held interviews in focus groups and surveyed a fairly representative sample of 602 U.S. respondents. Participants in the focus groups indicated that the most important reasons to consider filing a fraudulent claim were (1) saving money (2) getting unaffordable things done (3) getting back at insurance companies. The survey data revealed that 61% of respondents agreed that fraud is a way to seek “a fair return” on premiums. Furthermore, 56% assumed that insurance companies already charge a bit extra to cover the fraudulent expenses. These findings suggest that policyholders feel like insurance companies are making too much money and that fraud is a way of balancing the scales. Perhaps

unsurprisingly, the report also finds that those with a more positive attitude towards the insurer were less tolerant of fraud.

The same report finds that people think the chances of being discovered are rather small, and approximately 60 percent of respondents indicated they thought that different types of fraud (e.g., claim-padding, misrepresenting an accident, or falsifying receipts) were "fairly" or "very" common (CAIF, 1997). Furthermore, almost one-third of the surveyed sample said they knew someone who had cheated an insurance company but only 17 percent of those people had reported the cheater.

Interestingly, the survey also asked respondents to indicate what measures they thought would be effective in curtailing fraud. More than 90 percent indicated that claims should be verified more often – an expensive suggestion. Another much less expensive suggestion was to inform the public about the costs of fraud – we will return to this measure in the section on policy implications.

A survey commissioned by the Insurance Research Council focused on the relationship between the perception of insurance companies and perceived acceptability of insurance fraud (Tennyson, 1997). Specifically, a representative sample of 1,987 adults was asked how acceptable it is to increase a claim "to make up for the deductible" and "to make up for past premiums". People who found paying the premium "very burdensome" thought insurance fraud was more acceptable than those who had less trouble paying for insurance. Also, policyholders who felt "very confident" about the financial stability of their insurer thought fraud was less acceptable than those who were less confident. These results were interpreted as evidence for a positive relationship between perceptions of insurance companies and perceived acceptability of fraud.

Another interesting finding from this article is that the tolerance of insurance fraud was strongly predicted by the average tolerance for fraud in the state where the respondent lived. This suggests that social norms are a strong predictor of how acceptable people think it is to cheat on the insurance company. Building on previous research, Tennyson (2002) analyzes another part of the data collected in the CAIF study reported above. The focus in this study is on how experience with insurance companies affects attitudes towards fraud. Specifically, those respondents who had filed an insurance claim in the two years before the survey were less accepting of fraud. Arguably, policyholders who get their claim approved have a more positive attitude towards the insurer. This finding is thus consistent with the claim that perceptions of insurance companies are correlated with attitude towards fraud.

The evidence discussed so far illustrates that policyholders often feel that it is not unethical to file a fraudulent claim and that the chances of getting caught are slim. However, these studies were run in U.S. samples and it is unclear whether the findings are applicable to the Dutch population. Luckily, the Dutch Coalition of Insurers – in collaboration with Leiden University – commissioned a survey that explored the determinants of insurance fraud among 50 (Dutch) policyholders who had committed fraud, and 51 who had not (Verschuur, 1992). The power of this survey is in the fact that the data are linked to information in the “Centraal Informatie Systeem” (CIS; a system that stores all Dutch insurance claims and uses algorithms to detect possibly fraudulent ones). This allowed the author to compare the answers of policyholders that are registered to have filed a fraudulent claim to those of policyholders who have never committed fraud.

The results of the survey are consistent with those of the U.S.-based studies. The two most important reasons to commit fraud

were (1) financial need (2) making up for past premiums. Also, most respondents thought that the likelihood that a fraudulent claim would be discovered was small and the measure that was expected to lead to the greatest reduction in fraud was increasing the frequency with which claims are verified – again, an expensive undertaking for insurance companies.

The report by Verschuur (1992) also aims to get insight into whether fraudsters differ from non-fraudsters psychologically. On average, fraudsters are more sensitive to self-interest, and less to social norms and collective interests. They think insurance fraud is more acceptable (not surprising), are less satisfied with their socio-economic status, and are more competitive. Although this is interesting in itself, we think this finding does not help in fighting insurance fraud. It is not feasible to have every policyholder fill out questionnaires to determine a psychological profile – and consequently check the claims of those with a fraudster's profile more thoroughly. Additionally, it is likely that the people who would consider committing fraud, would also be likely to fake answers on such questionnaires.

Taken together, the empirical evidence indicates that policyholders think insurance fraud is common, acceptable, and easy to get away with. Importantly, it seems that most people feel the insurer-policyholder relationship is unfair, and that insurance fraud is a way to get back at evil insurance companies. In this light, the underlying cause of insurance fraud may become apparent.

It seems that policyholders do not construe insurance as a mechanism through which they share risk. Also, they do not consider the effects that fraud would have on the other policyholders of their insurance company. Rather, they feel like a big insurance company is undeservingly making money and

that they do not get enough of a return on their premiums. This conceptualization of insurance is conducive to fraud because it makes filing illegitimate claims look like a victimless crime and a justifiable way to compensate oneself for the 'unfairness' of insurance contracts.

Indeed, Clarke (1990, p. 3) already mentioned that policyholders are likely to perceive insurance companies as "fair game". Policyholders are unaware of the fact that fraud eventually leads premiums to rise and that therefore, other policyholders bear the cost of the crime. Another point, also emphasized by Viaene and Dedene (2004), is that insurers should be careful not to assume that policyholders' knowledge of the workings and goals of insurance is correct. Given the responses in the surveys we reviewed, it is unlikely that people think of their insurance contracts as a means through which they transfer their risk of health care expenditures to the insurance company. Similarly, they seem unaware of the fact that the transfer is possible largely because the insurer pools the risks of many policyholders. In fact, it is ironic that people think fraud is a way to compensate for the money they 'lost' on premiums – not having to file an insurance claim obviously means that nothing bad has happened to them!

3.4.2 Policy implications of the findings on insurance fraud

Although the estimates of how much insurance fraud actually costs are only tentative, the evidence we reviewed revealed that insurance fraud is overwhelmingly regarded as common and acceptable. We think insurance fraud is the most wasteful form of moral hazard that we have described and it is worth fighting. We describe two lessons that can be learned from the reviewed literature and suggest ways to improve the way in which insurance fraud is combated.

First, policyholders should be educated about the nature of insurance. It is clear that most people lack an understanding of the essence of insurance – sharing risks – and this seems to lead to a somewhat tolerant attitude towards insurance fraud.⁶ Opportunities for education arise when new clients sign up or when policyholders file a claim. These are situations where the insurer is in close contact with its clients and allows for the dissemination of relevant information. The Inshared (Achmea) commercials are a good example of how policyholders may be informed. These advertisements show visually how a group of policyholders pools money that is redistributed to those who suffer misfortune. Such graphic illustrations of how insurance works are exactly the type of information that is needed. They may reduce the likelihood that people perceive insurance as an individual contract on which they are losing money. Instead, we expect, it leads them to realize that they are not 'alone' in the insurance contract and they get a sense of what happens with the money they spend on insurance.

Relatedly, it is important that insurers make sure policyholders realize that what they get for paying premiums is a reduction in risk. Knowledge that premiums are essentially the price of risk reduction will reduce the feeling that one's payments are in vain or that the policyholder is not getting his money's worth. Initiatives that give the policyholder insight into what 'his' money is used for can serve the same goal. When people are aware that their payments have been used to pay for someone else's

6 We are currently researching the extent to which people are aware of the nature of insurance and how the way that people perceive insurance relates to their attitude towards fraud. Interested readers should contact Job van Wolferen (j.vanwolferen@tilburguniversity.edu) if they would like more information on our project – and especially if they have clever ideas on how to test this relationship.

care – or even better, that their own treatments are partly paid for by others – they will learn how insurance redistributes money to those that need it. Ultimately, such greater understanding should lead people to see the payments they receive from insurance companies as bonuses rather than a reduction of the amount the insurer still ‘owes’ them.

The second lesson is related to the widespread belief that insurance fraud often goes undetected. Although the articles we reviewed do not report a relationship between perceived auditing probability and attitude towards fraud, it is clear that the perceived frequency of successful auditing is extremely low. Most respondents therefore also indicated that they thought fraud would be most successfully reduced if insurers would check claims more often and more precisely. This is in line with theories that predict that greater likelihood of punishment deters people from filing illegitimate claims (for an overview, see Picard, 2000). The low perceived probability of getting caught would also reduce the deterrent effect of penalties for fraud. If the perceived probability could be increased, the size of punishment could become an important factor in the decision whether or not to commit fraud – high penalties would have a greater deterrent effect than low penalties. However, we were unable to find empirical tests of this hypothesis.

The problem with checking claims is that it is expensive and the costs of the audit might be greater than the gain that results from rejecting a fraudulent claim. A relatively cheap way to increase the perceived likelihood of auditing among claimants is to explicitly mention that the claim they are filing gets checked using the CIS database. In addition, it might be wise to publicize cases (e.g., through press releases) in which fraudsters have successfully been caught – especially if the case is sensational. A long line

of research in psychology teaches us that the things that come to mind easily are perceived to be more likely (e.g., Schwarz et al., 1991), and news items on detected insurance fraud will likely increase how often people think fraud is detected. Note that these suggestions can be implemented without having to actually check claims more often – and are thus relatively cheap!

Both implications – increasing awareness of the nature of insurance and increasing the perceived likelihood of fraud being detected – can be implemented credibly and successfully only when the message is perceived to be authentic. When the perception of insurance companies is negative, messages such as the ones we suggest may backfire. Namely, if policyholders would get the impression that the industries' attempt at reducing fraud is aimed at making more profit, they might become even more tolerant towards fraud. It is thus extremely important for insurers to make sure the image their policyholders have of them is positive.

A final comment on the fight against insurance fraud is that we need more data on actual fraud. While the surveys that measure tolerance of fraud are extremely valuable, the data on actual fraud are rather limited. It seems that the Dutch Coalition of Insurers has already started gathering more data as they distinguish between acceptance and claim fraud as of 2009. However, for each claim the CIS currently only saves name and address information, the branch in which the claim is filed, the amount claimed, and whether the claim is fraudulent or not. That renders it very hard to run informative analyses on the predictors of insurance fraud. We recommend that the CIS at least save what type of product or service is claimed and through which insurance company the claim is filed, as well as details with respect to the policy (such as level of deductible and how long the customer had been

paying premiums). This would allow insurers to determine, for example, whether awareness-raising commercials such as those by InShared indeed lead to fewer instances of fraud. A comparison of amount claimed for the same type of product or service at two different insurance companies would also be informative. Namely – without having to look at cases where fraud is actually determined – higher average claims in one company could be indicative of claim padding (soft fraud). Similarly, comparisons of the extent to which insurance fraud occurs at different companies could teach insurers something about which fraud deterrence activities are successful and which are not.

4. Moral hazard in pensions

This section reviews the evidence on moral hazard problems in the pension industry. Specifically, for each of the three types of moral hazard we discuss to what extent they are applicable to pensions. First, we provide a short description of the Dutch pension system. We return to Jeff, who could build up his pension through the three pillars that comprise the Dutch pension system. The first pillar guarantees that Jeff has a bit of income after he retires and is referred to as AOW (Algemene OuderdomsWet). Every taxpayer contributes to the AOW in a pay-as-you-go system where the current working population pays for the current retired population. The second pillar comprises an income-dependent contribution to a pension fund that is in charge of making sure Jeff's money is invested or saved, and returned to him after he retires. Participation in this second pillar is mandatory for many people who have a job (which is important for reasons we discuss later) and ensures that Jeff does not have to give up too much of his income once he retires. Jeff could also opt to add to his pension by using a third-pillar option such as investing in real estate or creating a stock portfolio.

To understand what ex ante moral hazard would look like in the Dutch pension system it is helpful to determine what risks people face and what preventive efforts they could engage in. Recall that ex ante moral hazard occurs when people take excessive risks or fail to engage in preventive efforts as a result of being insured. The risk is to be left with too little money to survive and live decently after retirement. The action to prevent this from happening is to save up enough money to pay for retirement. Without mandatory national pension schemes such as the Dutch system, people are completely responsible for their own pensions. In essence, by

imposing the pay-as-you-go system, the Dutch government removes much of the responsibility for the pensions from individuals. Following the *ex ante* moral hazard argument, this should make people more careless about saving for retirement. The theory thus predicts that people would become more careless in planning their financial futures. The expectation that the government will take care of them once they retire might thus inhibit appropriate investing in private plans to supplement the public pension.

Although there are virtually no data that address this hypothesis, it is unlikely that people become less likely to save up for the future if the government imposes a mandatory pension system. Namely, it turns out that in situations where pension planning is not mandatory, participation is very low (e.g., Thaler & Bernatzi, 2004). A national public safety-net pension that guarantees a minimal standard of living could thus very well be enough to increase participation.

Following the *ex ante* moral hazard argument, however, the expectation that the government provides a pension would still decrease people's willingness to participate in the second pillar. After all, they know that a minimal standard of living is guaranteed – the responsibility for their financial future is partly transferred to the government – which would lead to a decrease in additional pension savings. As a result, retirees might have to give up a significant part of their living standard and possibly end up in financial trouble because they can no longer afford their mortgage.

Although it is analytically challenging to test for the kind of *ex ante* moral hazard suggested here, some successful attempts have been made. Most notably, Feldstein (1996) uses U.S. data on the generosity of the Social Security program and national saving rates

to estimate to what extent public pension provision substitutes for private savings. Specifically, exploiting the variation in those numbers between 1930 and 1992, Feldstein (1996) finds that every dollar provided as Social Security wealth reduces private savings by two to three cents.

Returning to the Dutch pension system, it becomes apparent why strong encouragement of participation in the second pillar is important. Without mandatory participation in this pillar, the provision of AOW could decrease the likelihood that employees save up additional money for retirement. It thus seems that the Dutch system is well equipped to combat ex ante moral hazard. One might argue that the government should make participation in the second pillar mandatory to ensure that everyone is able to enjoy a comfortable retirement, but this question requires a political answer. In addition, as long as the current pension system provides a satisfactory solution, changes do not seem desirable.

To determine the extent to which ex post moral hazard poses a problem, it is again useful to determine what factors in the Dutch pensions system might be involved. Recall that ex post moral hazard occurs when a given service is used more often when it is insured than when it is not insured. The 'use' of retirement is non-flexible for everyone who lives long enough to exceed the retirement age. However, the possibility of early retirement does allow for flexibility in the age at which people retire. If the public pension provision is generous it might increase the prevalence of early retirement. So, without public pension provision Jeff might consider retiring at 65, or maybe even at 66 or 67, to obtain a little more financial room during his retirement. With the public pension provision, Jeff might not feel the need to save up a little

extra money because he knows the government has saved up retirement money for him.

A review by Juurikalla (2008) summarizes the evidence that speaks to the question whether people retire sooner when the retirement benefits are more generous. All of the evidence points in the same direction – the generosity of public pension schemes and the generosity of their early retirement clauses positively predict early retirement rates. Put simply, Jeff is more likely to retire early if it means he gets generous retirement benefits, and less likely if the benefits are less generous. It should be noted, though, that this form of ex post moral hazard is not necessarily problematic in the same way ex post moral hazard in health insurance might be. In both cases people use the insured service more when it costs less. However, there is a case to be made that for health care decisions, there is an objective standard by which one could determine whether health care is needed or not. For health care decisions it is thus possible to label some treatments as 'inappropriate'. This is not (or is much less so) the case for early retirement. The fact that there is an option to retire early makes it appropriate to retire early. So, although people do seem to respond to the incentive effects of early retirement clauses, we are not sure that this behavior should be subsumed under the 'undesired' forms of moral hazard.

The final moral hazard – fraud – seems inapplicable to what most people think of when they think of pensions. Specifically, it is hard to fake retirement or to fake one's age to obtain retirement benefits earlier than legally allowed. However, an intriguing feature of many pension systems allows for some exploitation by its subscribers. Specifically, pension systems often have an optional disability insurance clause. Specifically, an enrollee in the pension fund can opt to pay an additional premium so that when

he becomes (partially) disabled, the pension fund compensates the associated loss of income. Like any other insurance, this system may suffer from asymmetric information problems.

There is one specific disability insurance situation that goes to the core of asymmetric information: musculoskeletal disorders. Common examples include back, neck, and shoulder pains, as well as pain in joints and nerves. Such disorders are hard to diagnose, and it is difficult to objectively evaluate the seriousness of these problems, which allows for opportunistic behavior. So, imagine Jeff opted for the disability insurance in his pension plan and feels strain in his back. He goes to the doctor who determines that the back pain is severe enough to declare Jeff 40% disabled: because of the back pain he can only do 60% of the work that he used to be able to do. Consequently, the disability insurance clause in Jeff's pensions plan goes into effect and he starts receiving disability benefits to compensate for his loss of income. Usually, the doctor periodically checks the severity of the disorder that causes the partial disability. Under the disability insurance plan, Jeff has the incentive to overstate the pain he is feeling during the check-ups in order to keep receiving the disability benefits.

It turns out that exactly those diseases that are hard to diagnose – the musculoskeletal disorders – are responsive to the incentive effects in disability insurance. Variation in the generosity of two Canadian pension plans that provided disability insurance, and the rigor with which they checked claims based on hard-to-diagnose-disorders, has been linked to the prevalence of such claims (Campolieti, 2002). Specifically, the Quebec disability insurance plan was consistently tough on musculoskeletal-related claims between 1971 and 1999. However, the Canada plan increased the generosity of its musculoskeletal-related payments by \$150 in 1987 and only adopted tougher medical screenings

in 1995. Interestingly, the Quebec plan received only slightly increasing amounts of musculoskeletal-related claims over the entire period. In contrast, the Canada plan saw a much larger increase of musculoskeletal-related claims after 1987 and a sharp decrease after 1995. Notably, the claims related to diagnoses that are more easily diagnosed did not show the same systematic variation over time in either of the plans. These findings suggest that some claimants exploited the presence of asymmetric information, possibly by overstating the seriousness of their musculoskeletal disorders.

4.1 Moral hazard in pensions: conclusions

We have briefly discussed to what extent each of the three types of moral hazard applies to the Dutch pension system. In sum, the national public pension (AOW) could cause ex ante moral hazard by decreasing participation in pension systems that supplement the AOW. However, this could only happen if it weren't for the strong encouragement to enroll in these plans (participation is often mandatory).

The possibility for ex post moral hazard arises in the flexibility of the age at which people are allowed to retire. When early retirement benefits are generous, they encourage early retirement. To the extent that early retirement causes a loss of welfare (e.g., profitable labor is lost), it is worth considering whether early retirement benefits should be altered to discourage early retirement.

Pension providers usually also provide disability insurance plans. This makes them susceptible to fraud in that people may fake or exaggerate their disability to obtain disability benefits. Stringent medical screenings seem to effectively counteract this problem.

5. Discussion and conclusion

The empirical evidence for each of the three types of moral hazard is not equally strong. We have seen that there is very little evidence for ex ante moral hazard, although it may exist in cases where prevention of a disease is particularly hard and costly. The evidence for ex post moral hazard is much more consistent – almost all studies find that demand for care goes up when its price goes down. However, the estimates of the extent to which the demand for health care is sensitive to price are much less consistent. There is variance in the estimates of price elasticity for different types of care, but also for estimates of elasticity within types of care. The numbers on insurance fraud are by definition noisy estimates, as fraudulent actions are intended to go unnoticed, but the available indices suggest that it is a significant problem. The main reason that people commit insurance fraud seems to be that they are unaware of the nature and goals of insurance.

To determine the extent to which moral hazard poses a welfare loss, it would be worthwhile to determine by how much moral hazard drives up insurance premiums. It would be even better to estimate how many people refrain from buying insurance who would have bought insurance, had the premium not been increased by moral hazard. In other words, does moral hazard lead to suboptimal levels of insurance purchased? We were unable to find literature that empirically estimates the size of this welfare loss, and we discuss below two reasons why we think it is unlikely that it will be precisely and reliably estimated in the near future. First, the empirical challenge is substantial. Ideally, a researcher finds or builds a society and randomly offers one half of the people insurance while denying it to the other half. The researcher

should control the way health care is provided and distributed to keep the true costs of health care constant. Also, the study should run for an extended period of time to allow both societies reach a stable level of demand for health care. This clearly is a daunting task; one that would dwarf the RAND experiment, in which people were 'merely' randomly assigned to different levels of insurance.

A second reason why it is unlikely that the welfare question will soon be answered is more theoretical. Specifically, it deals with whether economically optimal demand for care and optimal demand for insurance are equivalent to the desirable levels of demand for those services. We have seen that in the RAND experiment, the people with 95% co-insurance – which amounts to virtually no insurance – postponed or refrained from getting appropriate care. Although some might argue that it is economically optimal for people to deny health care if they do not value it at its cost, it is questionable whether such a situation is desirable. We already briefly mentioned that some ex post moral hazard might be construed as benign because insurance increases the likelihood that people get care when they need it. One could thus define the optimal level of demand for care as the level of demand in a society that does not have insurance, or as the level of demand – insurance-induced or not – that leads to the healthiest society. An answer to the question whether an insurance system converges to an optimal level of demand thus has a different answer depending on how the optimal level is defined. Even if the empirical challenge of estimating these optimal levels could be overcome, we are certain that researchers would disagree on which definition should be used.

A similar argument can be made with respect to the demand for insurance. People might decide not to buy insurance if it is too expensive, which could be considered economically optimal.

However, again, one could also define the optimal level based on which level of insurance leads to the healthiest society. The estimated welfare loss – which would be estimated by comparing both observed levels of demand to their respective optimal levels of demand – would seem drastically different depending on which definition is used. In sum, although we think that the welfare question deserves an answer, the current empirical literature is unable to provide one.

We conclude this paper by reiterating a point that came up in all three discussions of policy implications: more data are needed. Currently, the data are not detailed enough to get a good understanding of how each type of moral hazard affects the Dutch health care system. Therefore, insurers are encouraged to collect more data on claim, health, and risk behaviors, as well as demographic information. In addition, insurers are advised to consider running experiments to get reliable estimates of how much of a problem each type of moral hazard is. The current climate – with budget cuts and changes in the health care system – is perfectly suited for randomly offering the possibility to some policyholders of keeping their more generous coverage while the coverage in the universal policy becomes less generous.

References

- Abbring, J.H., Chiappori, P.A. & Pinquet, J. (2003). Moral hazard and dynamic insurance data. *Journal of the European Economic Association*, 1(4), 767–820. doi: 10.1162/154247603322493159
- Akerlof, G.A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- Arrow, K. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5), 941–973. doi: citeulike-article-id:580320
- Baker, T. (1996). On the genealogy of moral hazard. *Texas Law Review*, 75(2), 237–292.
- Centraal Bureau voor de statistiek (2012). *Zorgrekeningen; uitgaven (in lopende en constante prijzen) en financiering*. Retrieved from <http://statline.cbs.nl/StatWeb/publication/?DM=SLNL&PA=71914NE&D1=025,32,37-45&D2=4-13&HDR=G1&STB=T&VW=T>
- Coalition Against Insurance Fraud . (1997). *Four faces: Why some Americans do – and do not – tolerate insurance fraud*. Coalition Against Insurance Fraud. Retrieved from <http://www.insurancefraud.org/downloads/Four%20Faces.pdf>
- Coalition Against Insurance Fraud (2012). Go figure: Fraud data. Retrieved from <http://www.insurancefraud.org/healthinsurance.htm>
- Convenant aanpak verzekeringfraude (2011). Retrieved from http://www.verzekeraars.nl/UserFiles/File/download/Convenant_Aanpak_Verzekeringfraude.pdf
- Chiappori, P.A., Durand, F. & Geoffard, P.Y. (1998). Moral hazard and the demand for physician services: First lessons from a French natural experiment. *European Economic Review*, 42, 499–511.
- Chiappori, P.A. & Salanie, B. (2002). *Testing contract theory: A survey of some recent work*. CESifo working paper 738.
- Clarke, M. (1990). The control of insurance fraud. *British Journal of Criminology*, 30(1), 1–23. doi: n/a
- Cockx, B. & Brasseur, C. (2003). The demand for physician services – Evidence from a natural experiment. *Journal of Health Economics*, 22(6), 881–913. doi: 10.1016/S0167-6296(03)00047-X
- Courbage, C. & de Coulon, A. (2004). Prevention and private health insurance in the UK. *Geneva Papers on Risk and Insurance-Issues and Practice*, 29(4), 719–727.
- Dave, D. & Kaestner, R. (2009). Health insurance and ex ante moral hazard: Evidence from Medicare. *International Journal of Health Care Finance & Economics*, 9(4), 367–390. doi: DOI 10.1007/s10754-009-9056-4

- Debebe, Z.Y., van Kempen, L. & de Hoop, T. (2012). A perverse 'net' effect? Health insurance and ex-ante moral hazard in Ghana. *Social Science & Medicine*. doi: 10.1016/j.socscimed.2012.02.035
- Derrig, R.A. (2002). Insurance fraud. *Journal of Risk and Insurance*, 69(3), 271-287. doi: 10.1111/1539-6975.00026
- Dionne, G. & Michel, P. (1991). Workers compensation and moral hazard. *Review of Economics and Statistics*, 73(2), 236-244.
- Ehrlich, I. & Becker, G.S. (1972). Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80, 623-648.
- Feldstein, M. (1996). Social security and saving: New time series evidence. *National Tax Journal*, 49(2), 151-164. doi: n/a
- Freeman, J.D., Kadiyala, S., Bell, J.F. & Martin, D.P. (2008). The causal effect of health insurance on utilization and outcomes in adults: A systematic review of US studies. *Medical Care*, 46(10), 1023-1032. doi: 10.1097/MLR.0b013e318185c913
- Insurance Information Institute. (2010). Insurance handbook: *A guide to insurance - what it does and how it works*. Retrieved from http://www2.iii.org/assets/docs/pdf/Insurance_Handbook_20103.pdf
- Juurikkala, O. (2008). Retire early, save little: The ironic disincentives of public pension systems. In Booth, P., Juurikkala, O. & Silver, N. (eds.) *Pension Provision: Government failure around the world* (pp. 75-94). London: The institute of economic affairs.
- Keating, D. (1991). Pension insurance, bankruptcy and moral hazard. *Wisconsin Law Review*(1), 65-108.
- Keeler, E.B. (1988). *The demand for episodes of medical treatment in the health insurance experiment*. Santa Monica: Rand Corporation.
- Kenkel, D.S. (2000). Prevention. In A.J. Culyer & J.P. Newhouse (eds.), *Handbook of Health Economics* (vol. 1B, pp. 1676-1720). Amsterdam, The Netherlands: Elsevier.
- Klick, J. & Stratmann, T. (2007). Diabetes treatments and moral hazard. *Journal of Law & Economics*, 50(3), 519-538.
- Koç, Ç. (2010). Disease-specific moral hazard and optimal health insurance design for physician services. *Journal of Risk and Insurance*. doi: 10.1111/j.1539-6975.2010.01384.x
- Major, J.A. & Riedinger, D.R. (2002). EFD: A hybrid knowledge/statistical-based system for the detection of fraud. *Journal of Risk and Insurance*, 69(3), 309-324. doi: 10.1111/1539-6975.00025
- Newhouse, J.P. (1996). *Free for all? Lessons from the RAND Health Insurance Experiment*. Cambridge, MA: Harvard University Press.
- Nyman, J.A. (2004). Is 'moral hazard' inefficient? The policy implications of a new theory. *Health Affairs*, 23(5), 194-199. doi: 10.1377/hlthaff.23.5.194

- Pagan, J.A., Puig, A. & Soldo, B.J. (2007). Health insurance coverage and the use of preventive services by Mexican adults. *Health Economics*, 16(12), 1359–1369. doi:10.1002/Hec.1226
- Pauly, M.V. (1968). The economics of moral hazard: comment. *American Economic Review*, 58(3), 531–537.
- Pauly, M.V. & Held, P.J. (1990). Benign moral hazard and the cost-effectiveness analysis of insurance-coverage. *Journal of Health Economics*, 9(4), 447–461.
- Picard, P. (2000). Economic analysis of insurance fraud. In G. Dionne (Ed.), *Handbook of Insurance* (pp. 315–362). Norwell, MA: Kluwer Academic Publishers.
- Richaudeau, D. (1999). Automobile insurance contracts and risk of accident: An empirical test using French individual data. *Geneva Papers on Risk and Insurance Theory*, 24(1), 97–114.
- Rijksoverheid (2012). *Wat is het eigen risico in de zorgverzekering?* Retrieved from <http://www.rijksoverheid.nl/onderwerpen/zorgverzekering/vraag-en-antwoord/wat-is-het-eigen-risico-in-de-zorgverzekering.html>
- Rothschild, M. & Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4), 629–649.
- Rowell, D. & Connelly, L.B. (2012). A history of the term "Moral Hazard". *Journal of Risk and Insurance*, 1–25. doi: 10.1111/j.1539-6975.2011.01448.x
- Savedoff, W.D. (2004). Kenneth Arrow and the birth of health economics. *Bulletin of the World Health Organization*, 82, 139–140.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauerschkatka, H. & Simons, A. (1991). Ease of retrieval as information – Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61(2), 195–202.
- Skriabikova, O., Pavlova, M. & Groot, W. (2010). Empirical models of demand for out-patient physician services and their relevance to the assessment of patient payment policies: A critical review of the literature. *International Journal of Environmental Research and Public Health*, 7(6), 2708–2725. doi: 10.3390/ijerph7062708
- Stanciole, A.E. (2007). *Health insurance and life style choices: Identifying the ex ante moral hazard*. IRISS Working Paper Series. Retrieved from <http://econpapers.repec.org/RePEc:irs:iriswp:2007-10>
- Tennyson, S. (1997). Economic institutions and individual ethics: A study of consumer attitudes toward insurance fraud. *Journal of Economic Behavior & Organization*, 32(2), 247–265. doi: 10.1016/S0167-2681(96)00904-3
- Tennyson, S. (2002). Insurance experience and consumers' attitudes toward insurance fraud. *Journal of Insurance Regulation*, 21(2), 35–55.
- Thaler, R.H. & Benartzi, S. (2004). Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(1), s164–s187. doi: <http://dx.doi.org/10.1086/380085>

- Van de Voorde, C., Van Doorslaer, E. & Schokkaert, E. (2001). Effects of cost sharing on physician utilization under favourable conditions for supplier-induced demand. [Article; Proceedings Paper]. *Health Economics*, 10(5), 457–471. doi: 10.1002/hec.631
- Van Vliet, R.C.J.A. (2001). Effects of price and deductibles on medical care demand, estimated from survey data. *Applied Economics*, 33(12), 1515–1524. doi: 10.1080/10003684001001362
- van Vliet, R.C.J.A. (2004). Deductibles and health care expenditures: Empirical estimates of price sensitivity based on administrative data. *International Journal of Health Care Finance and Economics*, 4(4), 283–305. doi: 10.1023/B:IHFE.0000043759.93644.e0
- Verbond van Verzekeraars (2011). Verzekerd van Cijfers 2011. Retrieved from http://www.verzekeraars.nl/UserFiles/File/cijfers/VvC2011_totaal.pdf
- Verschuur, W.L.G. (1992). *Determinanten van verzekeringsfraude*. Leiden: Werkgroep Veiligheid.
- Viaene, S. & Dedene, G. (2004). Insurance fraud: Issues and challenges. *Geneva papers on Risk and Insurance*, 29, 313–333. doi: 10.1111/j.1468-0440.2004.00290.x
- Zweifel, P. & Manning, W.G. (2000). Moral hazard and consumer incentives in health care. In J.C. Anthony & P.N. Joseph (eds.), *Handbook of Health Economics* (vol. 1, pt 1, pp. 409–459). Amsterdam: Elsevier.

PUBLICATIONS IN THE PANEL PAPERS SERIES

1. *Saving and investing over the life cycle and the role of collective pension funds*
Lans bovenberg , Ralph Koijen, Theo Nijman and Coen Teulings
2. *What does behavioural economics mean for policy? Challenges to savings and health policies in the Netherlands*
Peter Kooreman and Henriëtte Prast
3. *Housing wealth and household portfolios in an aging society*
Jan Rouwendal
4. *Birth is the sessenger of death – but policy may help to postpone the bad news*
Gerard van den Berg and Maarten Lindeboom
5. *Phased and partial retirement: preferences and limitations*
Arthur van Soest and Tunga Kantarci
6. *Retirement Patterns in Europe and the U.S.* (2008)
Arie Kapteyn and Tatiana Andreyeva
7. *Compression of morbidity: A promising approach to alleviate the societal consequences of population aging?* (2008)
Johan Mackenbach, Wilma Nusselder, Suzanne Polinder and Anton Kunst
8. *Strategic asset allocation* (2008)
Frank de Jong, Peter Schotman and Bas Werker
9. *Pension Systems, Aging and the Stability and Growth Pact* (2008) Revised version
Roel Beetsma and Heikki Oksanen
10. *Life course changes in income: An exploration of age- and stage effects in a 15-year panel in the Netherlands* (2008)
Matthijs Kalmijn and Rob Alessie
11. *Market-Consistent Valuation of Pension Liabilities* (2009)
Antoon Pelsser and Peter Vlaar
12. *Socioeconomic Differences in Health over the Life Cycle: Evidence and Explanations* (2009)
Eddy van Doorslaer, Hans van Kippersluis, Owen O'Donnell and Tom Van Ourti
13. *Computable Stochastic Equilibrium Models and their Use in Pension- and Ageing Research* (2009)
Hans Fehr
14. *Longevity risk* (2009)
Anja De Waegenaere, Bertrand Melenberg and Ralph Stevens

15. *Population ageing and the international capital market* (2009)
Yvonne Adema, Bas van Groezen and Lex Meijdam
16. *Financial Literacy: Evidence and Implications for Consumer Education* (2009)
Annamaria Lusardi and Maarten van Rooij
17. *Health, Disability and Work: Patterns for the Working-age Population* (2009)
Pilar García-Gómez, Hans-Martin von Gaudecker and Maarten Lindeboom
18. *Retirement Expectations, Preferences, and Decisions* (2010)
Luc Bissonnette, Arthur van Soest
19. *Interactive Online Decision Aids for Complex Consumer Decisions: Opportunities and Challenges for Pension Decision Support* (2010)
Benedict Dellaert
20. *Preferences for Redistribution and Pensions. What Can We Learn from Experiments?* (2010)
Jan Potters, Arno Riedl and Franziska Tausch
21. *Risk Factors in Pension Returns* (2010)
Peter Broer, Thijs Knaap and Ed Westerhout
22. *Determinants of Health Care Expenditure in an Aging Society* (2010)
Marc Koopmanschap, Claudine de Meijer, Bram Wouterse and Johan Polder
23. *Illiquidity: implications for investors and pension funds* (2011)
Frank de Jong and Frans de Roon
24. *Annuity Markets: Welfare, Money's Worth and Policy Implications* (2011)
Edmund Cannon, Ian Tonks
25. *Pricing in incomplete markets* (2011)
Antoon Pelsser
26. *Labor Market Policy and Participation over the Life Cycle* (2012)
Pieter Gautier and Bas van der Klaauw
27. *Pension contract design and free choice: Theory and practice* (2012)
Henk Nijboer and Bart Boon
28. *Measuring and Debiasing Consumer Pension Risk Attitudes* (2012)
Bas Donkers, Carlos Lourenço and Benedict Dellaert
29. *Cognitive Functioning over the Life Cycle* (2012)
Eric Bonsang, Thomas Dohmen, Arnaud Dupuy and Andries de Grip
30. *Risks, Returns and Optimal Holdings of Private Equity: A Survey of Existing Approaches* (2012)
Andrew Ang and Morten Sorensen
31. *How financially literate are women? Some new perspectives on the gender gap* (2012)
Tabea Bucher-Koenen, Annamaria Lusardi, Rob Alessie and Maarten van Rooij
32. *Framing and communication: The role of frames in theory and in practice* (2012)
Gideon Keren
33. *Moral hazard in the insurance industry* (2013)
Job van Wolferen, Yoel Inbar and Marcel Zeelenberg

Moral hazard in the insurance industry

This Panel Paper by Job van Wolferen, Yoel Inbar and Marcel Zeelenberg (all TiU) reviews recent evidence on moral hazard in the insurance industry. They discuss three types of moral hazard: ex ante moral hazard, ex post moral hazard, and insurance fraud, and detail how each is an asymmetric information problem. For each of the types, they summarize the empirical evidence and discuss the policy implications that follow from it.