Beneficial Risk Increases

Till Grüne-Yanoff and Holger Rosencrantz

Royal Institute of Technology

Abstract

We construct a model of choices under risk with biased risk perception. On its basis, we argue that sometimes, the regulator should raise the population's risk exposure. In particular: Individuals rationally adjust their behavior to perceived risks. But their risk perception is biased: low risks are underestimated. Individuals incur a welfare loss from choosing on the basis of biased perceived risks while being exposed to real risks. By raising the real risk, actions become riskier, but perceived risk is closer to real risk. Increase in real risk can therefore lead to a welfare improvement. Regulators choose risk exposure by optimizing the population's utility. We present four aggregation models in the light of population heterogeneity, and discuss their legitimacy with respect to standard normative principles concerning risk.

Introduction

Driving in his car through Skåne last summer, one of the authors encountered a very narrow and unmanageable traffic roundabout. This seemed surprising, given the generally good quality of Swedish traffic regulation and concern for road safety. The surprise grew bigger when it became clear that the roundabout had been recently *reduced* in size, and hence *made* more narrow and less manageable. What reason could there be for that?

In this paper we try to answer this question. We argue that under specific conditions, the traffic regulator is justified, or even obliged, to change the existing traffic infrastructure in a way that will be less manageable, and hence more risky, for the traffic participants.

Our argument starts with a simple model of how drivers choose their optimal risk exposure (section 1). We then propose that drivers are biased in the way they perceive risks (section 2). This bias, we argue, prevents drivers

from realizing their optimal risk exposure in their actions. But perception bias is not constant across all risks. Rather, it is large for mid-level risks, and small for low- and high-level risks. By shifting the risk level of traffic regulations away from this intermediate level, a regulator can make drivers choose actions closer to their optimal risk exposure (section 3). Increasing risks in this way can therefore be beneficial for individual drivers. In section 4 we discuss the conditions under which such risk-increasing policies can also be beneficial for a population of heterogeneous drivers. Section 5 concludes.

1 Rational Choice of Accident Rate

Traffic accidents have persistently been one of the most common causes of accidental deaths. Traveling in a car is an important occupational risk of everyday life. This fact is known and acknowledged by most drivers and passengers. They accept a certain risk for the benefits that car travel offers in return. With slight idealization, one may therefore attribute to the agent a rationally accepted accident rate. This rate may differ for different agents and different kinds of accidents. It specifies the risk of a specific accident an agent is rationally accepting when choosing to travel in a car, driving with a certain 'style', and reacting to specific traffic situations in a certain manner. A regulatory policy with the goal of reducing car accidents or traffic fatalities (of car drivers or passengers, as opposed to pedestrians) will commonly have to respect such an acceptance rate. Unless external parties are affected, it is usually not justified to regulate travel risks an agent rationally and voluntarily accepts.

To discuss the properties of the driver's choice in more depth, we construct an idealized choice model. The agent chooses how risky she wants to drive. The option she chooses has one of three consequences. With probability p, she will be harmed in an accident (H). With probability q, she will obtain some advantage A (e.g. speedy arrival, pleasure of driving, etc.). We abstract from differences in harms (for life, limb or property) and differences in advantages. With probability 1-p-q, she will neither obtain A nor H (i.e. $\neg(A \vee H)$). In accord with the interpretation of these consequences, we assume that the agent prefers A to $\neg(A \vee H)$, and $\neg(A \vee H)$ to H.

The agent's driving style options can be described as a lottery $\{p, H; q, A; 1\text{-}p\text{-}q, \neg(A \vee H)\}$. Graphically, it can be represented by a probability triangle, as depicted in Figure 1. By convention, the horizontal axis mea-

sures the probability q of the worst consequence H, increasing from left to right; the vertical axis measures the probability p of the best consequence A, increasing from bottom to top. Hence $\neg(A \vee H)$, the intermediate consequence, is located at the bottom left corner of the triangle. Not every locus of the triangle represents a possible choice option for the agent. Her choices are restricted by the following considerations. First, the agent can choose freely the chance of being harmed (e.g. she can impact the nearest solid object). But she cannot choose freely the chance of obtaining the advantageous consequence. This depends on the environment in which she drives, and her driving skills. The chance of obtaining the advantageous consequence also depends on how much risk the agent chooses to take. The agent's available options are therefore bound by a possibility frontier, which in Figure 1 is characterized by function $f(p)$. This function is specified for an agent's individual skills and a given environment.

For technical reasons, we assume that f is continuous and continuously differentiable for the domain of the probability triangle (i.e. for $f(p) \leq p-1$). We justify the form of this function as follows. First, we abstract from any local non-monotonicity and assume that the possibility frontier is a monotonically increasing function of p. Without this assumption, the choice would be trivial. Rational agents will consider increasing p only if such an increase also increases q. Second, we assume the function to be concave. This implies that the possibility set that f demarcates is convex. If the possibility space was not convex, agents would not be able to obtain mixtures of two options that they can reach. In a probabilistic setting, this is implausible. Therefore, f must be concave. Third, while an incremental increase of a small risk of harm contributes greatly to the chance of gaining an advantageous consequence, an incremental increase of a large risk contributes comparatively little. We therefore assume that the function has a monotonically decreasing slope. Fourth, increasing very low risks will lead to very high increases in the probability of advantageous consequences. hence the slope of f should be very steep for p close to 0. Last, even though the slope of f will be flat for all levels of skills when p is high, difference in skills will have significant impact on the risk/advantage trade-off in this region. I.e. skill and environment influence at what height f crosses the diagonal of the probability triangle. From the class of continuous, monotonically decreasing functions with monotonically decreasing slope we therefore choose the function of the form

$$f(p) = \sqrt{\alpha \times p} \tag{1}$$

α is the parameter that captures changes in skill or environment. We therefore write the possibility frontier as f_α. Because of the last two reasons detailed above, this functional form is more suitable than functions like $f(p) = \sqrt[\alpha]{p}$ or $f(p) = \log_\alpha p$. The form of f_α and the possibility space it delineates are shown in Figure 1.

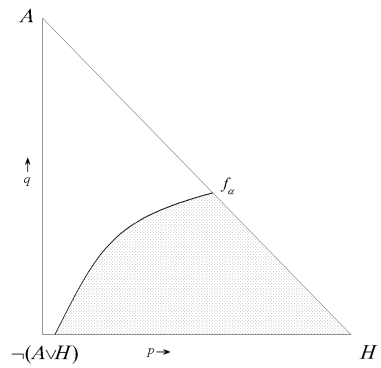

Figure 1: The Possibility Frontier

Given the standard assumptions of ordering plus continuity, the agent's preferences over different driving styles can be represented by a set of indifference curves. The additional standard assumption of independence of expected utility restricts the set of indifference curves to being upward sloping, linear and parallel. An agent's risk preferences that satisfy these assumptions express the agent's evaluation of risk as the slope of these indifference curves, across the whole area of the probability triangle. Figure 2 depicts this situation. The point at which the possibility frontier is tangential to the indifference curves is the optimal risk exposure p^*. Because possibility sets bounded by f_α are convex, the optimum is unique. A rational agent chooses that behavior which realizes the optimal risk exposure. Figure 3 depicts the utility curve of the possibility frontier f_α for all p.

If all agents are rational, a non-paternalistic government has no justification to regulate driving behavior. All agents know their driving skills in a given environment, and therefore know how likely they will obtain the advantageous consequence for a given level of risk exposure. They will choose

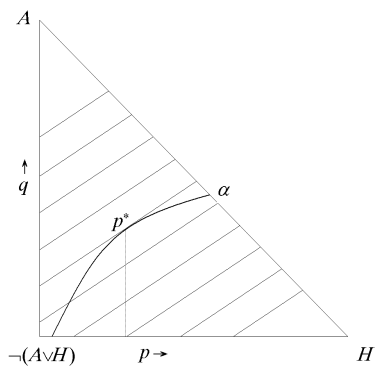

Figure 2: Indifference Curves

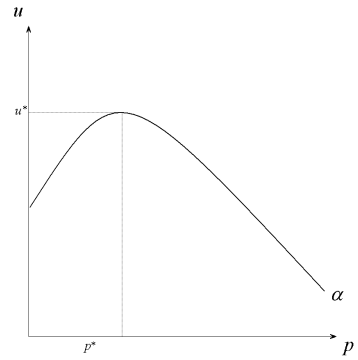

Figure 3: Utility Function of α

how much risk they are willing to dare according to their risk preferences. They will choose the accident rate that they are rationally willing to accept.

2 Risk Perception Bias

Psychological research into risk perception has shown that people judge the risk of a certain action according to certain mental strategies (also called heuristics), which are valid in some circumstances but lead to large and persistent biases in others (Kahneman *et al.* 1982). Relevant for driving behavior in particular is that people tend to underestimate the risk of their actions when this risk-likelihood is below a certain threshold. We distinguish between an agent's perceived risk $\hat{p}$, and the real risk p, associated with an action.[1]

We assume that agents underestimate real risk, if below a threshold τ. The function that maps real risk onto perceived risks therefore assumes a 'bulge', as depicted in Figure 4.[2]

There are at least three arguments for such the shape of this function. First, most drivers consider themselves to be better than the average driver. (Sivak *et al.* 1989, Cauzard and Wittink 1998). Studies for various countries showed that subjects gave similar, strongly asymmetrical responses concerning how relaxed, wise and considerate they were as drivers. This result indicates that drivers widely underestimate the real risks when it comes to their own driving behavior. Second, it has generally been observed that drivers adjust their behavior to maintain a constant level of perceived risk in a changed

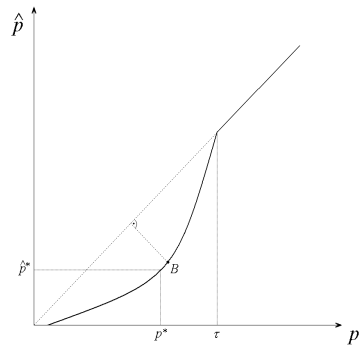

Figure 4: The Bias Function

environment (Adams 1985, Wilde 1989). But some study of this risk compensation effect show that, when the environment is perceptibly made 'safer', drivers overcompensate and get involved in more accidents than before. A good example of this is Kallberg (1993), who studied the effects of providing better nighttime lane delineation. A common-sense prediction would be that better visibility through better delineation could have only positive effects on safety. However, that is not what Kallberg found. Kallberg monitored speeds and accidents on sections of roads that were either equipped or not equipped with post-mounted delineators (retro-reflectors raised substantially above the road surface) for increased visibility. The results showed that on lower-speed, curvy roadways, increased visibility led to faster speeds and more accidents – a net negative effect. Third, introspecting our own risk perceptions in traffic situations, we find that we neglect very small risks, are not sensitive to commonly encountered risks, and only react strongly to risks that go beyond what is commonly encountered. When encountering the latter, we feel that we received a 'wake-up call' to be more alert and cautious, but soon settle again into the standard routine of underestimating low and standard risks.

We therefore assume that perceived risk be a continuous and monotonously increasing function of real risk. Figure 4 depicts a function whose form is determined by the depth of the 'bulge' (the distance of point B from the 45° line) and the threshold τ below which risk perception is biased. A simple algebraic expression of such a functional form is:

$$b(p) = \begin{cases} \tau \times (\frac{p}{\tau})^{\chi} & : \quad p < \tau \\ p & : \quad p \geq \tau \end{cases} \tag{2}$$

Using this function, we abstract from asymmetric biases in which the largest distance between p and $b(p)$ is not in the middle between 0 and τ.

Risk-perception bias has a negative welfare effect for the individual driver. Perception bias will lead an agent to associate a lower risk to his driving behavior, given the environment and his skills, than what is really associated with it. When choosing an action to obtain the optimal risk exposure p^*, the agent will choose an action that yields his perceived risk - that is, $\hat{p}^*$. Because of the existing bias, the action results in a real risk exposure of $p = b^{-1}(\hat{p}^*)$, which will differ from p^* more or less. Figure 5 illustrates this bias effect graphically. Because of the bulge-form of the bias function, drivers will generally choose behavior that results in higher than wanted risk exposures (as long as the optimal risk level is below the threshold τ).

Figure 6 shows the welfare consequences of the risk perception bias. The agent's biased risk perception leads him to choose an action that realizes risk $p = b^{-1}(\hat{p}^*)$, when aiming for optimal risk exposure. The utility of this action is $u(b^{-1}(\hat{p}^*))$. The perception bias thus results in a welfare loss of $u(p^*) - u(b^{-1}(\hat{p}^*))$.[3]

3 Regulative Intervention

Automotive risks commonly have a low likelihood, but have grave consequences. The policymaker is therefore interested in regulating traffic behavior in order to mitigate these consequences. The government can intervene with two different motivations. First, it may simply be interested in reducing the risk people take when driving. Lowered risk-taking will lead to a fall of accidents in the accident statistic. This may be a goal in itself. However, such a goal may be problematic, because it may interfere with people's preferences. As shown in Figure 2, a driver may rationally choose, given a specific environment and her individual skills, a risk level p^* - say driving 130km/h on a highway. If the government sets the speed limit at 110km/h for highways, it prohibits the agents to realize her optimum risk exposure, and forces her to choose actions that are less likely to yield an advantageous result. As long as such regulations affect the driver alone, they create a welfare loss: the driver would be better off, according to her own wants, without the regulation.[4]

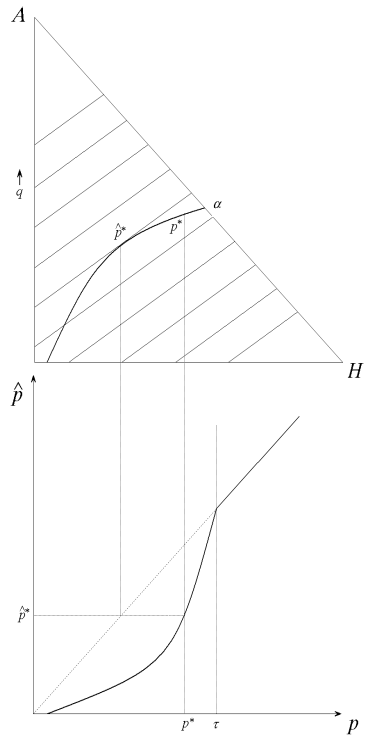

Figure 5: Bias Shifts Choice on Possibility Frontier

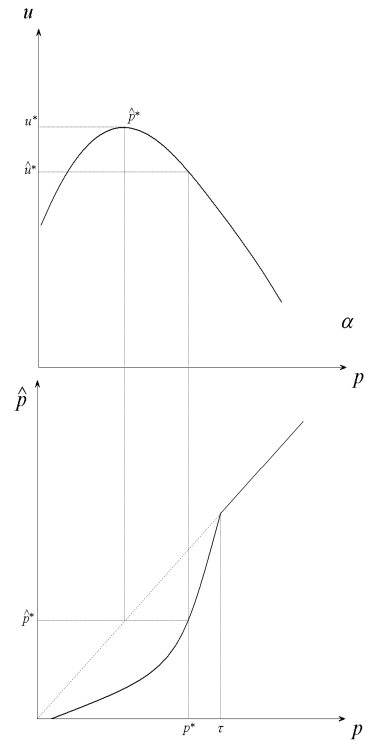

Figure 6: Bias Shifts Utility Result of Choice

The government could justify such regulations paternalistically: that it is not concerned with what the driver wants, but with her welfare, and that her welfare is improved if she is exposed to a lower accident rate than what she wants.

Alternatively, the government may regulate *vicariously*. In that case, it takes the drivers' preferences as the basis of its policy, and tries to ensure that drivers' behavior will be optimal according to their own preferences. That biased heuristics make people underestimate the risks of their own behavior is a reason for the government to intervene vicariously. The argument goes as follows. People rationally adjust their behavior to the perceived risks of that behavior. But people often underestimate the risks involved in car driving. They therefore do not rationally adjust their behavior to the real risks. Supporting people in rationally adjusting their behavior to real risks can therefore benefit all agents and improve upon traffic safety.

Two kinds of regulatory policies can provide such a support. First, the real risk of certain actions can be presented in information campaigns, employing both factual and affective information. In order to offset biases of risk perception, these campaigns may even resort to factual exaggerations. Campaigns of these sorts may raise ethical concerns with respect to the obscenity of affective information, or with respect to deceptive strategies. In any case, they may not be very effective, as people may trust their own risk perceptions more than the information provided in those campaigns.

A second strategy is to increase the real risk to those levels where the heuristics work more effectively, and where perception therefore approximates the real risks. The economist Armen Alchian once suggested as much, when he proposed to fit each car's steering wheel with a spear directed at the heart of the driver, so as to make her acutely aware of the dangers involved in driving (quoted in Landsberg 1993, 5). More to the point, Swedish transport authorities, at least on a local level, seem to follow such a strategy. Toni Rnnow-Rasmussen communicated a case in the vicinity of Lund, where the authorities decreased the size of a roundabout in order to make drivers slow down (inadvertently increasing the risk of an accident). Similarly, the results presented in Kallberg (1993) may lead to a reconsideration of installing post-mounted delineators for increased visibility, even though such a measure will increase the risk for drivers.

But what exactly does increase of real risk mean in this context? Certainly, the government cannot determine the risk exposure of the individual drivers. After all, the drivers choose their risk exposure according to rational criteria. All the government can do is to change the environment in such a fashion that it is rational for the agent to choose that risk exposure which is the government's target. As discussed in section 1, the possibility frontier of an individual driver depends on her skills and the environment in which she exerts these skills. The government can influence that environment. By installing a spike on the steering wheel, by decreasing the size of roundabout or by removing visibility aids, the government reduces the maximally feasible probability of an advantageous consequence for a given level of risk exposure. In other words, after the government increased the risk, the agent has to dare a higher level of risk exposure to get the same benefit as before.

In our model, a government's risk increase is represented by an downward shift of the possibility frontier. In terms of function (1), this is realized by a decrease of the parameter α. Three functions with different parameter values are depicted in Figure 7. Every function of form (1) has a continually

decreasing slope. An decrease in α decreases the function's slope at every point. Therefore, a rational agent with unchanged standard preferences will choose a lower risk exposure when the government increases the risk.[5] For agents who are free of perception bias, such an increase of risk yields a welfare loss, as depicted in Figure 8. Adjusting to a lower 'yield of risk exposure, the agent cannot reach as high a probability for advantageous consequences anymore, given his risk preferences. By increasing risk, the government forces the driver to rationally adjust and decrease his risk exposure, and it reduces the maximal utility the driver can reap by his driving behavior. For drivers without perception bias, the risk increase policy therefore has a purely paternalistic effect: it reduces drivers' risk exposure at the cost of their welfare.

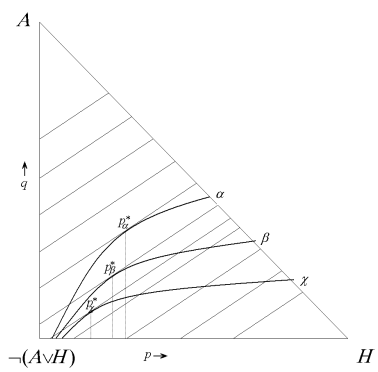

Figure 7: Shifting the Possibility Frontier

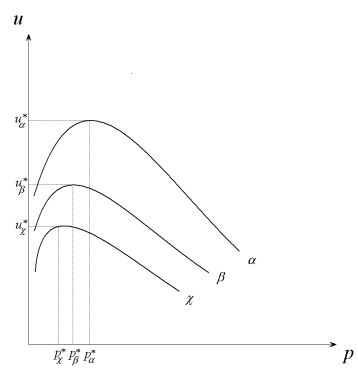

Figure 8: Effects of Shifting the Possibility Frontier on teh Utility Function

The situation changes when drivers are subject to perception biases. As we argued in section 2, perception bias causes the agent to choose a risk exposure higher than the optimal level, which leads to a welfare loss. Because of the 'bulge'-shape of the bias function, a change in the driver's choice of perceived risk exposure may bring his perception closer to the optimal real risk exposure. Choosing accordingly, such an adjustment may result in a net welfare gain. We illustrate this case in Figure 10. The agent's optimal risk exposure in environment α is p^*_α, which would give him utility u^*_a. Because p^*_α is located close to the maximum perception bias B, the agent will choose an action that realizes risk level $\hat{p}^*_\alpha$, instead, with utility level $\hat{u}^*_\alpha$. Now the government increases the risk by changing the environment to β, as shown in Figure 9. Here, the agent's optimal risk exposure is p^*_β. Because of the

perception bias, the agent will choose an action that realizes risk level $\hat{p}^*_\beta$, with utility level $\hat{u}^*_\beta$. p^*_β is located further away from the maximum perception bias B than p^*_α, hence $\hat{p}^*_\beta$ is located closer to p^*_β. Because $\hat{u}^*_\beta > \hat{u}^*_\alpha$, the government's risk increase yields a net welfare benefit for the driver, even though it generally shrunk his utility function.

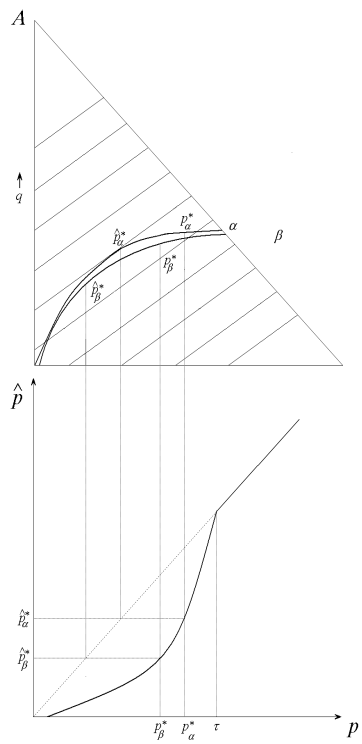

Figure 9: Bias Shifts Choice to a lower indifference curve in environment α than in β

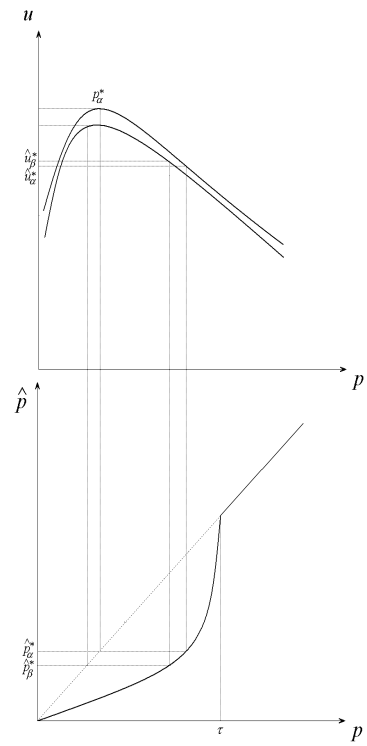

Figure 10: Due to Bias, $\hat{u}^*_\alpha$ is lower than $\hat{u}^*_\beta$

The possibility of such a welfare increase are limited by properties of the original environment, and the driver's reaction to it. If the driver's optimal risk exposure p^* lies close to the point of maximal perception bias B, a change in the environment, as shown in Figure 10, can improve the driver's welfare. If p^* lies far to the left of B, however, the welfare increase from mitigating the bias will likely not offset the loss from the loss incurred by the governmental risk increase. Worst, if p^* lies to the right of B, a governmental risk increase will increase the bias, and hence create a double negative welfare effect.

Under the right conditions, however, a government's risk increase will make the driver better off according to his own risk preferences (even though, being unaware of his perception bias, he will not want such a risk increase). A vicarious government is therefore justified, or even obliged, to choose such a risk-increasing policy. The question is whether a government can justify such a risk increase for a collective of drivers, when risk preferences, driving skills and perception biases are heterogeneously distributed in the population.

4 Risk Increases for a Population

A vicarious government is not concerned with the preferences of a single citizen, but with the preferences of all citizens. Traffic policy, even on the municipal level, involves many people with different capacities, characteristics and attitudes. In section 3 we showed that under certain conditions, a single driver may benefit from a risk-increasing policy. We now investigate the rationale for such a policy in a heterogeneous population.

Our model allows drivers to vary in three dimensions. First, their risk preferences may vary. Honoring standard preference assumptions, this means that individual's indifference curves may have different slopes. Second, drivers may exhibit different driving skills in one and the same environment. For the model, this means that different drivers come with different αs for the function (1). The government's impact on these bases when increasing risks may be a uniform reduction of the bases, or some reduction that takes prior bases as weights. Third, drivers may vary in their perception biases. More specifically, they may vary in the threshold τ below which their perceptions are biased, and in the depth of the 'bulge', which is described by χ in function (2).

If drivers are homogenous in all three parameters, the results of section 3 apply directly. If the driver's optimal risk exposure p^* lies close to the left of B, and if their utility functions are 'shallow enough, then a risk increase will improve all drivers welfare. Maximizing agent's utility with respect to α under these circumstances is a Pareto-efficient policy.

If drivers are heterogeneous in one or more of the parameters, a risk increase have different consequences for a population. We distinguish two cases. In the first, everybody's lot improves, and hence the policy is a Pareto-improvements. In the second, it does not improve everybody's lot, and hence it is not a Pareto-improvement.

If drivers vary only with respect to risk preferences, and the most risk-prone ones have optimal risk exposures close to the left of B, a risk increase will improve the welfare of the risk-proner most. Depending on how widely risk preferences are distributed, and how 'deep' utility function are, it still may be the case that the policy improves all drivers' welfare, and hence is Pareto efficient.

If drivers vary only with respect to skills, those with the highest skills have optimal risk exposures close to the left of B, and a risk-increase policy affects all drivers equally, a risk increase will improve the welfare of the higher skilled most. Depending on how widely skills are distributed, and how 'deep' utility function are, it still may be the case that the policy improves all drivers' welfare, and hence is Pareto efficient. It may be argued that such a policy affects the higher-skilled comparatively more than the lower-skilled (who are quite cautious already). In this case, the welfare effect will be increased.

If drivers vary only in their risk perception, those with B close to the right of p^* will gain most. Whether all will obtain welfare increases will depend on how 'deep' their bias functions' bulges are, whether their Bs are located to the left or the right of p^*, and how much the risk increase pushes their p^* to the left.

If drivers vary in two or three of the parameters, interrelations between these parameters must be investigated. For example, it seems plausible to assume that risk preferences vary with skills, and that risk preferences also vary with perception biases. Such correlations may have special impacts on the welfare properties of a risk increase.

In those cases where some risk increase leads to a Pareto-improvement, it is likely that many non-identical risk-increasing policies also lead to a Pareto-improvement. What further criteria can be used to rank these policies?

First, a utilitarian principle will choose that policy that maximizes aggregated utility, irrespective of individual welfare effects. Applying the utilitarian principle only to the set of Pareto-improving policies avoids the usual problems. It even gets around the exemption problem (Hansson 2002, 25), which requires explicit justification if people's right not to be exposed to risk of negative impact is overridden. Because everybody benefits to some degree from any of the assessed policy measures, it may be justified to choose that one which maximizes aggregate utility.

Alternatively, Pareto-improving policies may be ranked according to an egalitarian principle. Because policy interventions improve drivers' situations by reducing their risk perception biases, the appropriate metric for an

equality assessment is the individual reduction of perception bias. The best policy (from the set of Pareto-improving policies) according to the egalitarian principle is that which brings individual perception biases closest to a common level.

It may be the case that the population is so heterogeneous that no policy is an Pareto improvement. Clearly, a policy that makes all worse off should not be considered. Additionally, we do not see a good justification of the utilitarian principle in such a situation. It would inevitably expose some people to higher risks for the sole benefit of making the aggregate, and hence some other, and possibly undeserving, person, better off. Still, one may judge some risk increasing policies to be positive, according to the following principles.

For egalitarian reasons, a policy may be acceptable if it converges the risk biases of the population towards a common level, even if this means that those with lower biases are made a little worse off through the policy. Similarly, a prioritarian principle may justify policies that bring about a decrease of the most biased perceptions, while making those with smaller biases a little worse off.

Conclusion

We have shown under which conditions drivers are better off when the traffic regulator deteriorates elements of the traffic infrastructure. The motivation for such regulatory intervention is to decrease biases in the drivers' perceptions and hence to allow them to choose more in accord with their rational judgment than it was possible for them before the intervention. By giving a rationale for our introducing case of the Skåne roundabout, we have given an example how non-paternalistic regulation is possible. The regulator in our model regulates to improve individual's adherence to their own values, instead of imposing its own.

Notes

[1]This view of risk perception, taken as a general rule or even as a reasonable assumption, is far from uncontroversial (e.g., Sjberg, xxxx). Nevertheless, it may serve as a starting point for interesting discussions. In this paper, we do not want to depend on any particular theory of behavior. Rather, we wish to base our discussion on relatively simple and weak assumptions. When we say that people underestimate the probability of accident when this probability is low, this should be understood simply as a basic idea and we do not want the discussion to stand or fall with any particular theory.

[2]This function must be distinguished from the probability weighting function popular in recent literature. The probability weighting function measures risk attitudes and in particular over- and undervaluation of risks. It assumes an inverted S-shape, because observations reveal 'risk-aversion for small probability losses and risk-seeking for large probability losses' (Tversky and Kahneman 1992, 306). Our function, in contrast, only presents the biased estimation of real risk. The two functions therefore do not contradict each other.

[3]Before proceeding, some remarks should be made on the notion of 'correct' versus 'incorrect' choices. On first glance, it may seem that these notions could be replaced by 'rational' and 'irrational'. Since correct choice is closely linked to what the agent rationally accepts, such an interpretation appears adequate. However, this may be misleading. According to the Bayesian account, rational choice under risk is defined as relative to subjective probability - or, in this context, perceived risk (REFERENCE?). Therefore, even if it is the case that an agent chooses to take a risk that she would not rationally accept, choices should be referred to as 'incorrect' rather than 'irrational'.

[4]Because accidents often involve other people, and in particular people who are more vulnerable than driver (e.g. pedestrians), real situations are more complicated to evaluate. We abstract from these complications for the moment.

[5]It is noteworthy that risk homeostasis cannot be modeled with standard preferences. Risk homeostasis claims that drivers adjust their behavior to a changing environment to maintain a constant level of perceived risk. To model a constant p* under a changing possibility frontier, the indifference curves must not be parallel. One possible modification that would allow modeling risk homeostasis is the 'fanning out' hypothesis. In the context of

the probability triangle, this means that choices between prospects located in the bottom right-hand corner appear more risk prone than should be expected given preferences revealed for choices located leftwards or upwards in the triangle. The fanning out hypothesis has been extensively discussed in the literature (For an overview, see Starmer, (2000)) We will not use it further in this paper.

References

Adams, J. (1985). Smeed's law, seat belts, and the emperor's new clothes. In L. Evans, and R. C. Schwing (Eds.), Human behavior and traffic safety. New York: Plenum.

Cauzard, J.-P., and Wittink, R. (Eds.) (1998). The attitude and behaviour of European car drivers to road safety (SARTRE 2 reports, Part 1: Report on principal results). SWOV Institute for Safety Research, Leidschendam, The Netherlands.

Hansson, Sven Ove (2004) 'Philosophical Perspectives on Risk, Techne 8(1).

Kahneman et al 1982 Judgement under Uncertainty, CUP

Landsberg 1993, The Armchair Economist

Sivak, M., Soler, J., and Tränkle, U. (1989). Cross-cultural differences in driver self assessment. Accident Analysis & Prevention, 21, 371-375.

Sivak, Michael (2002) 'How common sense fails us on the road: contribution of bounded rationality to the annual worldwide toll of one million traffic fatalities', Transportation Research Part F 5: 259-269.

Sjberg (xxxx)

Starmer, Chris (2000) 'Develoments in Non-Expected Utility Theory' Journal of Economic Literature 38(2): 332-382.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cu-

mulative representation of uncertainty, Journal of Risk and Uncertainty, 5, 306

Wilde, G. J. S. (1989). Accident countermeasures and behavioural compensation: the position of the risk homeostasis theory. Journal of Occupational Accidents, 10, 267-292.

A Algebraic Treatment of the utility maximization problem

As depicted in Figure 2, the driver's indifference curves are linear. Her risk preferences can therefore be represented by the function

$$q_k(p) = u + k \times p \tag{3}$$

where k is the driver's trade-off rate of k increments of q for one increment of p, and u is the utility level on which this trade-off takes place.

As specified in (1), the possibility frontier is represented by the function

$$f_\alpha(p) = \sqrt{\alpha \times p}$$

The utility an agent can achieve when choosing a specific p is thus given by the intersection of f_α with q_k:

$$\begin{aligned} f_\alpha(p) &= q_k(p) \\ \Leftrightarrow \sqrt{\alpha \times p} &= u + k \times p \\ \Leftrightarrow u_{\alpha,k}(p) &= \sqrt{\alpha \times p} - k \times p \end{aligned}$$

The optimal risk exposure is given by setting the first derivative of $u_{\alpha,k}$ to 0.

$$\begin{aligned}
\frac{\partial u_{\alpha,k}}{\partial p} &= \frac{\sqrt{\alpha}}{2\sqrt{p}} - k = 0 \\
\Leftrightarrow p^*_{\alpha,k} &= \frac{\alpha}{4k^2}
\end{aligned}$$

However, the agent's risk perception is biased. Since the agent chooses an action with optimal *perceived* risk, the real risk chosen under perception bias will be equal to

$$\hat{p} = b^{-1}(p)$$

As specified in (A, the bias function is

$$b(p) = \begin{cases} \tau \times (\frac{p}{\tau})^{\chi} & : \quad p < \tau \\ p & : \quad p \geq \tau \end{cases}$$

The inverse of b is

$$b^{-1}_{\tau,\chi}(\hat{p}) = \begin{cases} \tau \times (\frac{\hat{p}}{\tau})^{\frac{1}{\chi}} & : \quad p < \tau \\ p & : \quad p \geq \tau \end{cases}$$

Thus, when the agent chooses an action that she perceives as yielding the optimal risk exposure $\hat{p}^*$, she really exposes herself to risk $p = b^{-1}_{\tau,\chi}(\hat{p^*})$, and gets utility

$$\begin{aligned}
\hat{u}^*_{\alpha,k}(\hat{p}^*) &= \sqrt{\alpha \times b^{-1}_{\tau,\chi}(\hat{p^*})} - k \times b^{-1}_{\tau,\chi}(\hat{p^*}) \\
&= \sqrt{\alpha \times \tau \times (\frac{\hat{p}^*}{\tau})^{\frac{1}{\chi}}} - \tau \times \ k \times (\frac{\hat{p}^*}{\tau})^{\frac{1}{\chi}} \\
&= \sqrt{\alpha \times \tau \times (\frac{\alpha}{4 \times \tau \times k^2})^{\frac{1}{\chi}}} - \tau \times \ k \times (\frac{\alpha}{4 \times \tau \times k^2})^{\frac{1}{\chi}}
\end{aligned}$$

For any policy that increases risk from β to α (hence $\beta > \alpha$) to be beneficial, it must satisfy the following inequality:

$$\hat{u}^*_{\alpha,k,\tau,\chi}(\hat{p}^*) - \hat{u}^*_{\beta,k,\tau,\chi}(\hat{p}^*) \quad > \quad 0.$$

With appropriate algebraic transformations, this yields:

$$\frac{\alpha^{\frac{\chi+1}{2\chi}} - \beta^{\frac{\chi+1}{2\chi}}}{\alpha^{\frac{1}{\chi}} - \alpha^{\frac{1}{\chi}}} \quad > \quad \tau^{\frac{\chi+1}{2\chi}} \times k^{\frac{\chi+1}{\chi}} 4^{-\frac{1}{2\chi}}$$

B Graphical Treatment of Solutions