The Slippery Slope Argument – Probability, Utility & Category Reappraisal

Adam Corner (corneraj@cardiff.ac.uk)
Department of Psychology, Cardiff University, Tower Building, Park Place,
Cardiff, CF10 3AT, UK

Ulrike Hahn (hahnu@cardiff.ac.uk)
Department of Psychology, Cardiff University, Tower Building, Park Place,
Cardiff, CF10 3AT, UK

Mike Oaksford (m.oaksford@psychology.bbk.ac.uk)
Department of Psychology, Birkbeck, University of London, Malet Street,
London, WC1E 7HX, UK

Abstract

The slippery slope argument (SSA) is generally treated as a fallacy by both traditional and contemporary theories of argumentation, but is frequently used and widely accepted in applied reasoning domains. Experiment 1 tests the hypothesis that SSAs are not perceived as universally weak arguments. The results provide the first empirical demonstration that SSAs vary predictably in their subjective acceptability. Experiment 2 identifies an empirical mechanism on which successful SSAs may be predicated, namely the process of category boundary reappraisal.

Introduction

The "slippery slope" is an intuitive metaphor that is used to refer to a class of arguments with a distinctive form. Classified as a fallacy of reason by most critical thinking textbooks (e.g. Woods, Irvine & Walton, 2000), dismissed as an obstacle to co-operative conversation by an influential socio-linguistic theory of argumentation (van Eemeren & Grootendorst, 2004), yet frequently used and widely accepted in applied domains such as politics (van Der Burg, 1991), law (Lode, 1999) and bioethics (Launis, 2002), the slippery slope argument remains a controversial topic in the field of argumentation. Correspondingly, the argument possesses the somewhat undignified status of "wrong but persuasive", and therefore fits neatly into the category of arguments that argumentation theorists call fallacies.

The slippery slope is a particular breed of *consequentialist* argument (Hahn & Oaksford, in press; for extended analysis of slippery slope arguments see Walton, 1992; for recent experimental work on other forms of consequential conditional see Bonnefon & Hilton, 2004; Thompson, Evans & Handley, 2005). A dissuasive consequentialist argument (or deterrent) warns against a particular course of action on the grounds that it will lead to an undesirable outcome, or consequence. A slippery slope argument (SSA), however, posits not only a negative outcome but the idea that this outcome might in the future be *re-evaluated as positive*, if an initial proposal goes ahead.

A general consequentialist argument, for example, might oppose the legalisation of cannabis because it would lead to an increase in smoking related respiratory problems. A slippery slope argument would oppose legalisation on the grounds that attitudes towards harder drugs might become more positive in the process, and in the future a substance like cocaine might also become legal. The slippery slope has four distinct components:

1. An initial proposal (A)
2. An undesirable outcome (B)
3. The belief that allowing (A) will lead to a re-evaluation of (B) in the future
4. The rejection of (A) based on this belief

The alleged danger lurking on the slippery slope is, therefore, the fear that a presently unacceptable proposal (B) will (by any number of psychological processes – see, e.g. Volokh, 2003) in the future be re-evaluated as acceptable. If we withhold the right of free speech from a neo-Nazi organisation, what will prevent us from censoring legitimate political dissent in the future? The proponent of this argument is inherently appealing to the malleability of public opinion to reject an otherwise appealing course of action. The uncertainty of the future is such that any reasoning about it is at best presumptive. Yet SSAs trade on the uncertainty of the future, and appear to be acceptable in a number of contexts (see, e.g. Volokh, 2003; Lode, 1999). In light of the fact that there has been no empirical investigation of the slippery slope, a pressing task is to examine if, when and how SSAs are successfully employed.

In keeping with the probabilistic approach to understanding human reasoning and rationality (see Oaksford & Chater, 2001), we seek to explain argument strength by invoking Bayesian principles (Howson & Urbach, 1996) that interpret reasoning patterns as probabilistic changes in subjective degrees of belief. An argument's strength is a function of an individual's initial level of belief in the claim, the availability and observation of confirmatory (or disconfirmatory) evidence, and the existence and perceived strength of competing hypotheses (Oaksford & Hahn, 2004; Hahn, Oaksford & Corner, 2005; Hahn & Oaksford, in press; see also Korb, 2004). An individual's belief in an argumentative claim can vary from 0 (no conviction) – 1 (total conviction), and is constantly being updated by (relevant) incoming information. Reformulated in Bayesian terms, a fallacy is simply a

probabilistically weak inductive argument (Oaksford & Hahn, 2004; Hahn & Oaksford, in press). Furthermore, there is no expectation that all instantiations of an argument type should be treated as equally strong by a Bayesian analysis. The context, but more importantly the content of the argument and the beliefs of the arguer themselves will dictate the probabilistic value of an argumentative claim (Hahn & Oaksford, in press). Essentially, Bayes' Theorem provides a consistent and rational model of the incorporation of new information to an existing knowledge state following argumentation, and offers a normative model for assessing argument strength.

Using a Bayesian model of argument strength, an SSA is convincing to the extent that its consequences seem probable given the available evidence. In one sense, therefore, an SSA can be analysed as a simple conditional probability – i.e. what is the chance of (B) occurring given (A)?Consequently, we should expect SSAs whereby the initial proposal is likely to bring about the feared outcome to be stronger than ones where that probability is low.

An account of SSA strength would be incomplete, however, if the *utilities* of the components involved were ignored. In particular, philosophers interested in applied domains such as law or bio-ethics where SSAs are popular have implicitly recognised that probabilistic *and* utilitarian concerns are crucial determinants of consequential and slippery slope argument acceptability (e.g. Holtug, 1993; Lode, 1999). This distinguishes SSAs from most other fallacies of argumentation (for overviews of the traditional catalogue of fallacies see e.g., Woods et al. 2004). This distinction stems from the fact that SSAs advocate *decisions* and as such are not just arguments about factual claims as most (so-called) fallacies are. Bayesian decision theory provides a normative framework for decision-making in situations where outcomes are uncertain, based on the probabilities and utilities involved; it can provide the normative framework that philosophical analyses of SSAs within the argumentation literature have lacked. According to decision theory, agents should maximize subjective expected utility in their choice (see, e.g. Savage, 1954; Keeney & Raiffa, 1976). One consequence of this is that we might even choose to avoid an action with only a low probability outcome, as long as that outcome is catastrophic enough. Whilst the ability of a normative account of utility to accurately model the intricacies of decision making is not without its detractors, it should minimally be the case that "the more probable the causal connection is, and the more we want to avoid (B), the stronger the argument" (Holtug, 1993, p404). Our first experiment seeks to show that these intuitions voiced by some philosophical analysts of slippery slope arguments, and encapsulated in Bayesian decision theory, are indeed widely shared.

Experiment 1

Participants in Experiment 1 were required to read several short scenarios containing slippery slope arguments, and provide a rating of argument strength as illustrated in Figure

1. The experiment was designed to demonstrate experimentally that slippery slope arguments vary in convincingness, and that this variation can be predicted by manipulating (i) the conditional probability, and (ii) the utility of the predicted future outcome.

Regarding (i), an argument where the probability of the outcome (B) given the initial proposal (A) is high should be more convincing than an argument where P(B|A) is low. In the present experiment, the conditional probabilities presented to participants were varied by describing either a probable or an improbable mechanism by which the proposed outcome of the argument could occur.

Regarding (ii), a predicted outcome is a necessary component of slippery slope argumentation, but predicted outcomes that are perceived to have a moderately negative expected utility will not be "feared" or avoided as much as outcomes with very negative expected utility. Predicted outcomes with very negative utilities will provide a stronger argument against the proposed course of action. In the present experiment, the outcome utilities of the arguments presented to participants were set as either moderately negative or very negative. Figure 1 shows an example scenario as seen by participants. In this example, the probability of the outcome (B) given the initial proposal (A) is high (because of the alleged difficulty of formulating clear medical guidelines), whilst the utility of the predicted outcome is very negative (in the form of an increase in involuntary euthanasia). The information in parenthesis relates to the opposing low probability and moderately negative utility experimental condition for this topic.

Whilst flicking through a copy of Ethics magazine, you come across an article on the thorny issue of euthanasia. Despite almost unanimous agreement (from both the medical profession and terminally ill individuals) on the acceptability of helping some patients to end their suffering, opponents claim that **the legalisation of voluntary euthanasia will lead to an increase in cases of involuntary euthanasia – or "medical murder" (other hospital patients feeling that their lives are not as worthwhile).** The British Medical Association has warned that once voluntary euthanasia is permitted in some cases, **it will be difficult to formulate clear guidelines about when doctors can euthanize patients (the British Medical Association has indicated that most hospital patients are unconcerned by the thought of voluntary euthanasia in hospitals)** . The article ends with the view of the author about the future of euthanasia legislation;

"We should oppose the legalisation of euthanasia in the UK, as it will lead to an increase in the number of instances of 'medical murder'".
Please indicate below how convincing the author's view is:

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
Unconvincing   Very Convincing
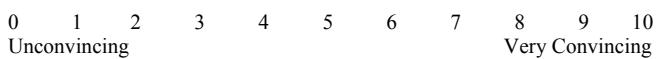
Figure 1: An example scenario from Experiment 1. Probability and utility information is highlighted in bold.

Method

Participants

Sixty undergraduate psychology students from Cardiff University participated in Exp. 1 for course credit.

Design

Both experimental variables were manipulated at two levels (probable/improbable mechanism, moderately/very negative utility) across four different argument topics, creating a total of sixteen distinct arguments and a 2 x 2 factorial design. All participants were presented with four slippery slope arguments, each concerning a different topic, and were required to provide a rating of argument strength for each argument. Each participant saw only one argument from each topic, and participated once in each experimental condition (i.e. a *Latin Square Confounded* design - Kirk, 1995). The topics of the arguments and the order they were presented in were randomised for each participant.

Materials & Procedure

Each participant received an experimental booklet containing four slippery slope arguments on different topics (see Figure 1 for an example scenario). The topics were (i) the introduction of I.D. Cards, (ii) the distribution of newspapers to a small General Store, (iii) the legalisation of voluntary euthanasia, and (iv) the cessation of postal deliveries to houses where vicious dogs live.

Results & Discussion

Figure 2 shows the mean ratings of argument strength by condition.

To statistically analyse data from Latin Square Confounded designs, subject effects within the ratings are factored out and the analyses are conducted on the residuals (Kirk, 1995). An ANOVA revealed that the probability manipulation had a significant effect on ratings of argument strength, $F(1,239) = 22.26$, $p < .001$, as did the utility manipulation, $F(1,239) = 28.71$, $p < .001$. These main effects suggest that both probability and utility are important factors in determining SSA strength, as predicted by Bayesian decision theory. The interaction between these two factors (which is also predicted by Bayesian decision theory) failed to reach statistical significance. Having found evidence for the existence of such an interaction in subsequent work, we suspect that lack of power is to blame.

Two planned comparisons indicated that the ordinal predictions made by Bayesian decision theory about argument strength were clearly supported. Ratings of argument strength were highest in the High Probability/Very Negative Utility condition, and lowest in the Low Probability/Moderately Negative Utility condition. This difference was confirmed as significant with a t-test, $t(119) = 7.48$, $p < .001$. As expected, ratings of argument strength in the two mixed conditions, where the effects of the variables were expected to work in opposition, were not significantly different from each other.
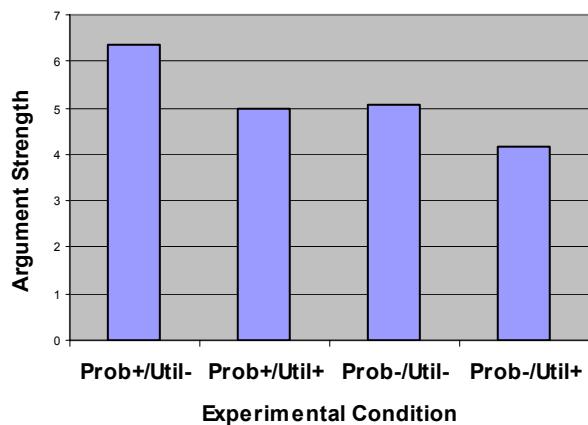

Figure 2: Mean ratings of argument strength from Exp. 1. Prob+ refers to the condition where the outcome is made to seem probable, Prob- refers to the condition where it is less likely. Util- indicates that the consequence is very undesirable, Util+ that it is less so.

Experiment 2

The results of Experiment 1 indicate that SSAs are differentially convincing. The higher the probability of the consequent event, and the more negative the utility associated with it, the more acceptable the arguments appear. Given, furthermore, that Bayesian decision theory provides a normative framework whereby it is rational to incorporate these factors into decisions about actions, it cannot be maintained that SSAs are universally weak arguments. People seem to find some SSAs acceptable and there are circumstances, according to Bayesian decision theory, whereby such a view is *subjectively rational*.

However, the ultimate rejection of the "wrong but persuasive" tag would be provided by a demonstration that the differential convincingness of SSAs has some *objective*, empirical basis. In addition to identifying and manipulating the factors that dictate subjective SSA strength, therefore, it is equally important to ask whether people have good reason to be persuaded by at least some slippery slope arguments. Is there reason to believe that 'slippage' occurs in the real world?

It is often claimed by those authors that have been positive about SSAs that conceptual vagueness (e.g. the difficulty of providing a precise definition of "terminally ill") and a fear of constructing arbitrary distinctions (e.g. deciding which terminally ill patients' lives are "worthwhile") provides the rationale for many SSAs (e.g. Lode, 1999, p1499). Govier (1982) suggests that the process of psychological assimilation acts as a catalyst for slippery slope arguments, and the ancient philosophical paradox of *Sorites* provides an example:

It is morally wrong to kill a sentient being, and a foetus at the time of birth (T) is a sentient being.
A foetus at one second (T-1) before the time of birth is also a sentient being, as the addition or subtraction of one second cannot affect a being's sentience.
Therefore, a foetus at (T-2) is also a sentient being.
Therefore, a foetus at (T-n) is also a sentient being; a foetus at the moment of conception is a sentient being.

The S*orites* argument plays on the vagueness of the predicate "sentient", and the inevitability of *modus ponens* to achieve paradox, but the idea that certain SSAs might be based on some kind of category boundary reappraisal mechanism has been articulated implicitly by many authors – indeed, the notion that a slippery slope might exist between an ostensibly positive proposal and a negative outcome directly implies an extension process of some kind.

The message to unwary reasoners is that the majority of the concepts that pervade our everyday argumentation are indeterminate. When advances in gene therapy are discussed, therefore, the spectre of Nazi Eugenics is raised because the concept of pro-social genetic engineering is vague (Holtug, 1993), and membership of the category "acceptable practice" is a dynamic and fluctuating process.

Because our everyday concepts lack necessary and sufficient features and do not, as a consequence, have clear cut boundaries (for references see e.g., Pothos & Hahn, 2000), classification is heavily dependent on the set of instances to which the category label has been applied. Though very different accounts of the nature of conceptual structure exist, theorists are agreed that there is a systematic relationship between the items that have been classified as belonging to a category and subsequent classification behavior. It is a fundamental of a wide range of current theories of conceptual structure, that encountering instances of the category at the category boundary will extend that boundary for subsequent classifications. Furthermore, there is a range of empirical evidence which is consistent with these assumptions. In particular, there are numerous experimental demonstrations of so-called exemplar effects, that is, effects of exposure to particular instances and their consequences for subsequent classification behavior (e.g., Nosofsky, 1986, 1988a, 1988b; Lamberts, 1995). For example, observing that a dog that weighs 10kg is considered underweight invites the conclusion that a dog that weighs 10.5kg is also underweight. With only the information that a 5kg dog is underweight, and a 15kg dog is overweight, however, one might not be so compelled to draw this conclusion.

There is then a feedback loop inherent in the classification of new data into an existing category, whereby that classification alters the category itself (see also e.g. Lakoff, 1987). In appropriate circumstances this extends the category boundary in a way that could naturally give rise to slippery slope arguments (Hahn & Oaksford, in press).

Experiment 2 was designed to demonstrate experimentally the link between category boundary reappraisal and slippery slope arguments. If SSAs have an objective basis in category expansion driven by exemplar effects, there should be agreement between the perception of an SSA's strength and corresponding categorisation decisions, given identical data to evaluate.

Participants were given a fictitious scenario describing a debate between the Finnish Government and the Finnish Housing Association concerning the allocation of *Outstanding Natural Beauty* status to candidate areas of Finnish land. The Government was allocated the role of preserving as much Finnish countryside as possible, while the Housing Association was portrayed as being primarily concerned with providing affordable housing space. Participants were informed that land was awarded *Outstanding Natural Beauty* status if it contained an "unusually high number of large animal species", and that if this status was conferred, no further housing development was permitted in that location. Using these materials, we sought to demonstrate an exemplar effect that would influence both ratings of SSA strength and categorization decisions.

Method
Participants
Sixty undergraduate psychology students from Cardiff University participated in Exp. 2 for course credit.

Design
A single variable (the numerical value of the exemplar provided for evaluation) was manipulated between the groups of participants, and two experimental measures (either a categorisation decision or a rating of argument strength) were recorded, creating a total of four experimental groups. Figure 3 represents a summary of the information given to participants in the categorisation groups.
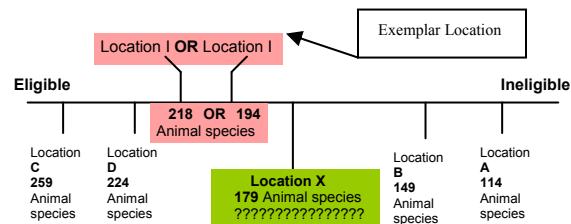

Figure 3: A schematic representation of the information provided to participants in Exp. 2. Depicted are different locations and their eligibility for '*Outstanding Natural Beauty Status*'. An exemplar effect concerning the target location, *Location X*, is induced by manipulating the information associated with *Location I*. One group of participants is told that eligible *Location I* has *218* animal species, the other group is told it has *194*. This should systematically affect both categorization and argument strength ratings focused on *Location X*.

The experimental variable was manipulated by altering the number of large species of animal contained in the

exemplar location (*Location I)*, which participants were informed was eligible for *Outstanding Natural Beauty* status: either *194*, or *218*.

Participants were asked to make a categorisation decision of their own, based on the information they had just read: whether a further location (*Location X*, containing *179* animal species) was eligible for *Outstanding Natural Beauty* status.

The exemplar manipulation for *Location I* was expected to differentially influence categorisation decisions between the two groups. Participants who were told that *Location I* contained *194* animal species should be more likely to categorise *Location X* as eligible for *Outstanding Natural Beauty* status, as they should perceive *Location X* as closest in resemblance to the *Location I* (already designated as eligible for *Outstanding Natural Beauty* status). When *Location I* contained 218 animal species, however, participants should perceive *Location X* as closest in resemblance to the ineligible locations.

Argument strength was assessed by presenting *Location I* as part of a slippery slope argument. Instead of being informed that *Location I* was eligible for *Outstanding Natural Beauty* status, participants were told that while the Finnish Housing Association were not too concerned about *Location I* being awarded *Outstanding Natural Beauty* status, this would lead to a further location (*Location X*) also receiving *Outstanding Natural Beauty* status, which the Finnish Housing Association viewed as problematic. It was predicted that participants who viewed this argument when *Location I* contained *194* animals would provide a higher rating of argument strength, as they should perceive *Location X* as sufficiently close to the existing category boundary to be vulnerable to a slippery slope style reappraisal (mirroring the exemplar effect predicted in the categorisation groups).

Materials & Procedure

Each participant received an experimental booklet containing a brief description of the fictitious scenario and a list of locations that had already been adjudicated, followed by either the categorization decision or the argument strength rating as an experimental task.

Results & Discussion

The yes/no data obtained from the categorisation groups were analysed using a ranked sign test. Participants who had been told that *Location I* contained *194* animals categorised the new location as deserving of *Outstanding Natural Beauty* status on 11 of 15 occasions. Participants who had been told that *Location I* contained *218* animals categorised the new location as deserving of *Outstanding Natural Beauty* status on 0 of 15 occasions. This difference was significant at $p < .01$.

The argument rating data were analysed using a t-test. Participants who had been told that *Location I* contained *194* animals rated the arguments as significantly more convincing (M = 4) than participants who had been told that

Location I contained *218* animals (M = 2.6), $t(28) = 2.1$, $p < .05$.

These results provide empirical support for the philosophical analysis of slippery slope arguments by authors such as Govier (1982) and Volokh (2003) by demonstrating, in a tightly coupled design, how slippery slopes may rest on a category boundary extension process.

General Discussion

The results of Experiment 1 are the first empirical demonstration that SSAs vary predictably in their acceptability, and that this variation is broadly captured by a Bayesian account of argument strength. With regard to argumentation theory and the study of the fallacies in general, this is of interest because variation in strength for arguments of identical structure has typically been problematic for existing theories of fallacy (e.g. van Eemeren & Grootendorst, 2004), but the idea that argument strength is a graded concept is a central tenet of the Bayesian account. In this respect the results mirror those recently obtained for other fallacies such as the 'argument from ignorance' (Oaksford & Hahn, 2004; Hahn, Oaksford & Bayindir, 2005).

With regard to the philosophical debate about slippery slope arguments specifically, our results suggest that the credibility (some) slippery slope arguments possess in applied domains such as law or medical ethics can be justified. The clear implication of the data obtained in Experiment 1 is that SSAs are differentially persuasive. That they are not simply 'persuasive but wrong' follows from the fact that the key variables involved in their evaluation – probability and utility- have a normative basis in Bayesian decision theory. Moreover, there can be an *objective*, non-zero probability that the slippage on which SSAs are predicated can, in fact occur. Experiment 2 demonstrates this by linking experimentally categorization and slippery slope argument acceptability. Exemplar effects provide the kind of empirical mechanism that the fear of outcome re-evaluation inherent in slippery slope arguments requires.

Other mechanisms of real world slippery slopes have also been suggested. The legal philosopher Volokh, for example, points to the possibility of multi-peaked voter preferences that could bring about slippery slopes (Volokh, 2003). Taking the currently debated topic of I.D. card introduction in Britain as an example, it might be the case that the introduction of voluntary I.D. cards will hasten the introduction of compulsory I.D. cards, not because individuals alter their attitudes to them, but because some voters, though preferring no I.D. Cards to compulsory I.D. Cards, would prefer compulsory I.D. to the woolly and costly compromise of semi-compulsory I.D. Cards, thereby making probable the direct move from no I.D. to compulsory I.D. (Volokh, 2003).

Future work will be directed at addressing the existence of other such processes of slope facilitation. In some ways the most important contribution of the present studies is the demonstration that experimental investigation of slippery

slope arguments and their underlying mechanisms is feasible. This adds a new tool for the argumentation theorist, as argumentation studies have been dominated by textual analysis on the one hand, and –to a lesser extent- the development of logics for argumentation and computational argument systems. However, it is not just argumentation theory for which slippery slope arguments are of interest. Cognitive psychology has a long research tradition concerned with how we reason with conditionals (see e.g. Evans & Over, 2004; Oaksford & Chater, 1998). Yet this research has focused on an extremely narrow set of conditional statements and only very recent experimental work has sought to branch out and investigate systematically different kinds of conditional statement (Bonnefon & Hilton, 2004; Thompson et al. 2005). Slippery slope arguments provide another, distinct form of conditional with its own characteristics to add to this developing set. Finally, with their emphasis on values, and their real-world importance, slippery slope arguments also link to social psychological research on attitude and attitude change under the header of expectancy value theory (e.g., Fishbein & Ajzen, 1975). This body of work currently stands separate both from the cognitive psychological work on conditionals and from argumentation theory. Slippery slope arguments would seem an ideal focal point for a much needed interdisciplinary integration.

Acknowledgements

References

Bonnefon, J.F. & Hilton, D.J. (2004). Consequential Conditionals: Invited and Suppressed Inferences From Valued Outcomes. *Journal of Experimental Psychology; Learning, Memory & Cognition 30 (1) 28-39.*

Der Burg, W. van. (1991). The Slippery Slope Argument. *Ethics 102, 42-45.*

Eemeren, F.H. van. & Grootendorst, R. (2004). *A systematic theory of argumentation – the pragma-dialectical approach.* Cambridge: Cambridge University Press.

Evans, J. St. B.T. & Over, D. (2004). *If.* Oxford: Oxford University Press.

Fishbein, M. & Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research.* Addison-Wesley: Reading, MASS.

Govier, T. (1982). What's wrong with slippery slope arguments. *Canadian Journal of Philosophy 12, 303-316.*

Hahn, U. & Oaksford, M. (*in press*). A Bayesian Approach to Informal Reasoning Fallacies. *Synthese.*

Hahn, U, Oaksford, M. & Bayindir, H. (2005). How convinced should we be by negative evidence? *Proceedings of the 27th Annual Meeting of the Cognitive Science Society.*

Hahn, U, Oaksford, M. & Corner, A. (2005) Circular arguments, begging the question and the formalization of

argument strength. *Proceedings of AMKCL05 –Adaptive Knowledge Representation and Reasoning.* Helsinki.

Holtug, N. (1993). Human Gene Therapy: Down the Slippery Slope. *Bioethics 7, 402-419.*

Howson, C. & Urbach, P. (1996). *Scientific Reasoning: The Bayesian Approach (2nd Ed).* Chicago: Open Court.

Keeney, R.L. & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs.* New York: Wiley.

Kirk, R.E. (1995). *Experimental Design – Procedures for the Behavioural Sciences.* London: Brooks/Cole.

Korb, K. (2004). Bayesian Informal Logic and Fallacy. *Informal Logic, 24, 41-70.*

Lakoff, G. (1987). Cognitive Models and Prototype Theory. In, Neisser, U. (Ed). *Concepts and conceptual development: Ecological and Intellectual Factors.* Cambridge: Cambridge University Press.

Lamberts, K. (1995). Categorization under time pressure. *Journal of Experimental Psychology: General, 124, 161-180.*

Launis, V. (2002). Human Gene Therapy & the Slippery Slope Argument. *Medicine, Health Care and Philosophy 5, 169-179.*

Lode, E. (1999). Slippery Slope Arguments and legal Reasoning. *California Law Review 87, 1468-1543.*

Nosofsky, R.M. (1986) Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115, 39-57.*

Nosofsky, R.M. (1988a). Exemplar-based accounts of the relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory and Cognition, 14, 700-708.*

Nosofsky, R.M. (1988b). Similarity, frequency and category representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 54-65.*

Oaksford, M. & Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Sciences 5 (8), 349-357.*

Oaksford, M., & Hahn, U. (2004). A Bayesian approach to the argument from ignorance. *Canadian Journal of Experimental Psychology*, *58*, 75-85.

Oaksford, M. & Chater, N. (1998). *Rationality in an uncertain world : essays on the cognitive science of human reasoning.* Hove: Psychology Press.

Pothos, E. & Hahn, U. (2000). So concepts aren't definitions, but do they have necessary or sufficient features? *British Journal of Psychology, 91, 439-450.*

Savage, L. J. (1954). *The foundations of statistics.* John Wiley: New York.

Thompson, V.A, Evans J. St. B.T. & Handley, S.J. (2005). Persuading and dissuading by conditional argument. *Journal of Memory & Language 53, 238-257.*

Volokh, E. (2003). The Mechanisms of The Slippery Slope. *Harvard Law Review 116, 1026-1137.*

Walton, D. (1992). *Slippery Slope Arguments.* Oxford: Clarendon Press.

Woods, J, Irvine, A. & Walton, D. (2004). *Critical Thinking, Logic & The Fallacies.* Toronto: Prentice Hall.