 doi:10.20944/preprints202106.0676.v1

Article

How to discourage adversaries from affecting decision outcomes of a repeated patent application decision–making process

Willem van der Sluis [1,2]

[1] DMO, Utrecht, The Netherlands
[2] The work was conducted when the author was an external doctoral candidate at the Faculty of Sciences, Department of Mathematics, Vrije Universiteit, Amsterdam, The Netherlands, with a workplace at the Centrum Wiskunde & Informatica, Amsterdam

Abstract: Outcomes of repeated decision–making processes may be affected by adversarial actors, without being noticed. Adversaries may try to gain knowledge about a particular decision–making process, identify its decision–makers, and guess which underlying decision support model is used. Then they can simulate the process, and craft different scenarios to affect its decision outcomes. Therefore, designers of decision support systems need to incorporate this in the decision modeling phase. The purpose of this study is to demonstrate this for the repeated decision–making in a patent application process. In this process, two sequential decision outcomes can be affected by adversarial actors: a company's decision to which type of patent office to send a patent request to, and the decision of a specialized patent officer to grant an application, or not. It is motivated that the company's decision–maker is *bounded* rational. A theory for information–theoretic bounded rational decision–making under uncertainty proposed by Ortega et al. is adopted to model this type of decision–maker. A framework is provided to simulate a number of scenarios that adversaries may deploy to affect decision outcomes of a repeated patent application decision–making process. The framework is also utilized for statistically testing the presence of the scenarios, and to demonstrate how to discourage adversaries from deploying them.

Keywords: Adversarial risk analysis and decision analysis; information–theoretic bounded rational decision–making; simulation

1. Introduction

Nowadays companies and organizations are increasingly using mathematical models to support their decision–making processes. However, in most cases, the actual decisions are still taken by humans. Adversaries may try to gain knowledge about a decision–making process and the used decision criteria, may try to guess the supporting mathematical model, may seek ways to degrade the performance of this model, and may try to influence decision–makers by presenting them wrong insights from data. To make a supporting decision model less vulnerable to such adversarial influencing, a solution would be manual and ad hoc reconstruction of the decision support within the parameters of the used decision model, and adapt the model to the adversary's evolving manipulations [1]. An area of research that focuses on the subfields risk analysis and decision analysis is *adversarial risk analysis* (ARA). In ARA, one asserts that analysts should use Bayesian thinking to describe their beliefs about an opponent's goals, resources, optimism and type of strategic calculation, while placing subjective probability distributions on all unknown quantities. This in order to enable analysts to maximize their expected utilities [2]. Not all decision–making processes, however, are suitable for applying Bayesian thinking, such as the patent application decision–making process.

A *patent application* is a request pending at a patent office for the grant of a patent for an invention, being described in a patent specification and a set of one ore more claims stated in a formal document, including necessary official forms and related correspondence [3]. Companies and organizations normally use tools with patentability criteria, like a patenting

Citation: . *Preprints* **2021**, *1*, 0. https://doi.org/

Received:
Accepted:
Published:

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

decision tree and an additional machine learning system, to assist in determining which technological inventions should be patented [4]. Once the decision is made to apply for a patent, a company's intellectual property department evaluates the proposed technology (novelty) in relation to its patenting strategy. Depending on the outcome, a decision–maker X in a company's intellectual property department decides which geographic coverage a patent must have (country, or region), and to which patent office a patent application has to be filed to (decision A: country office or regional office). Lastly, the relevant documentation is provided to and examined by a responsible patent officer at the chosen patent office (i.e. decision–maker Y), who in a process of negotiating or arguing decides whether to grant a patent or not (decision B). So, a binary decision A made by decision–maker X is followed by X observing the outcome of a subsequent binary decision B taken by decision maker Y. In X's deliberation process there is interaction with the environment in that he/she selects the choice alternative a according to some optimized probability distribution $P_{\mathcal{A}=\{regional,country\}}(a)$. This has a stochastic effect on the environment according to the probability distribution $P(o|a)$, where o is X's observation of the outcome of decision B. Decision maker X's resources to extensively evaluate all choice aspects of decision A are limited, and this limitation reduces X to a state of bounded rationality, a term coined by [5]. To model decision–maker X's decision–making, the theory for bounded rational decision–making under uncertainty developed by [6,7] is adopted. Put in a repeated patent application context, the resulting repeated bounded rational decision–making model (here, referred to as model M_1) requires to choose the values of the so–called boundedness parameter and the value of a utility parameter.

In a patent application process, both of the binary decisions A and B are vulnerable to adversarial influencing. For example, decision–maker X may be an adversarial actor, or he/she may be influenced by an adversarial co–worker of the intellectual property department who is presenting wrong insights from data. On the patent office–side of the process, decision maker Y may be an adversary, or he/she may be manipulated by an adversarial co–worker. There is even the possibility that adversarial actors on both sides of the patent application process are closely cooperating. In the present study, a simulation framework is proposed to generate for both binary decisions A and B a sample of decision outcomes in complete absence of adversarial influencing, and an equally sized sample in case some adversarial influencing scenario has been active in the same time window. By pairing corresponding samples, a two samples *(paired) proportion test* can be conducted to test whether there is a significant difference between two proportions of the same decision outcomes, or not. To be precise, an asymptotic McNemar-test without continuity correction [8].

In the present study, six adversarial influencing scenarios have been defined and implemented: three different basic scenarios and three combinations of these scenarios. A measure has been introduced to express the *attractiveness* of an influencing scenario from the perspective of an adversarial actor. The measure is based on the observed average number of times the presence of a scenario can statistically be proven, and the observed average Cohen effect size [9] of the proven presences. A multi–objective optimization model (referred to as model M_2) is formulated to minimize the set of object functions corresponding to the six considered scenarios, with regard to the to be chosen boundedness parameter and utility parameter of model M_1. It has been made plausible that solving model M_2 for a time window yields the most favorable parameter value pair of model M_1 in this time window, and implementing this parameter pair will make it less attractive for adversaries to deploy the six considered influencing scenarios in the time window.

2. Results

This section provides the results of the performed simulation study, and the conclusions that have been drawn. As stated in the introduction, the purpose of the study is to

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

demonstrate how a repeated patent application decision–making process, on average, can be made less vulnerable to adversaries trying to affect its decision outcomes by deploying six considered influencing scenarios in some time window W. Three time windows are considered (1 year, 2 years, and 3 years). Mathematical details about these scenarios, the statistical test(s) conducted to test for their presence, the modeling of the repeated patent application decision–making process (i.e. model M_1), and the used simulation framework can be found in Section 3. Two parameters $0.30 < \beta^d < 0.60$ and $1.1 < U^d(1,R) < 5.0$ of model M_1 remain to be specified (see Subsection 3.2). A second model M_2 (see Subsection 3.8) is developed to determine the most favorable parameter pair $(\beta_W^{d,*}, U_W^{d,*}(1,R))$ for a time window W. In model M_2, a set of objective functions is to be minimized with regard to the parameters β_W^d and $U_W^d(1,R)$, where each objective function corresponds to a considered influencing scenario $\bullet$. An objective function represents the *attractiveness* of the corresponding influencing scenario from the perspective of an adversarial actor, and requires two statistical quantities as input: the sample mean of positive test results $\mu_{ptr^{\bullet}_{(\cdot,W)}}$ (with $\cdot = X$ or Y) and the associated sample mean Cohen distance $\mu_{\Delta^{Cohen,\bullet}_{(\cdot,W)}}$ (see Subsection 3.7). Here X and Y correspond to the statistical test conducted for the decision outcomes of decision–maker X and Y, respectively (see Figure 9 below). To obtain these statistical quantities, 50 simulation runs with 50 sub–runs per simulation run were performed for each time window (see Subsection 3.6).

Subsection 2.1 illustrates how the attractiveness of each of the three basic influencing scenarios $\bullet = 1, 2$ and 3 in a time window depends on the behavior of the two statistical quantities $\mu_{ptr^{\bullet}_{(\cdot,W)}}$ and $\mu_{\Delta^{Cohen,\bullet}_{(\cdot,W)}}$ in the $(\beta_W^d, U_W^d(1,R))$–landscape. All plots shown in this subsection were generated on a parameter grid in which β_W^d is ranging from 0.32 to 0.59 with steps of 0.02, and $U_W^d(1,R)$ is ranging from 0.11 to 4.9 with steps of 0.1. The main contribution of the present study is to propose a mathematical model (M_1) for repeated patent application decision–making that inherently includes a second mathematical model (M_2) that takes into account that adversaries may guess the structure of model M_1 and its parametrization, and deploy six different crafted scenarios to affect its decision outcomes. Moreover, model M_2 provides a parametrization for model M_1 that makes it less attractive for adversaries to deploy the six scenarios. This is the subject of Subsection 2.2.

In total six adversarial influencing scenarios have been implemented in the proposed simulation framework, for a time window W of 1 year, 2 years, and 3 years. Three basic scenarios, denoted by $\bullet = 1$, $\bullet = 2$ and $\bullet = 3$, and three combinations of these scenarios, denoted by $\bullet = 2 + 3$, $\bullet = 1 + 2$ and $\bullet = 1 + 3$.

2.1. First inspection of the attractiveness of the three basic influencing scenarios

Figure 1 and Figure 2 below show surface plots of the sample mean *ptr*–scores and sample mean Cohen distances on the z-axis for the basic influencing scenario $\bullet = 1$ and the time windows $W = 1$ and $W = 2$, respectively. In the bumpy surface plots of the mean Cohen distance, the heights are almost similar for both time windows (between a medium effect 0.5 and a large effect 0.8). And the surface plots for the sample mean *ptr*–score show a relatively smooth landscape surface. The surfaces of both statistical quantities rise with increasing values of the grid parameters, where the heights in the upper right warm colored area of the sample mean *ptr*–score surface for time window $W = 1$ are considerably higher than those for time window $W = 2$. Hence, it does not seem to be attractive for an adversarial actor to deploy this scenario for a period longer than 1 year.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

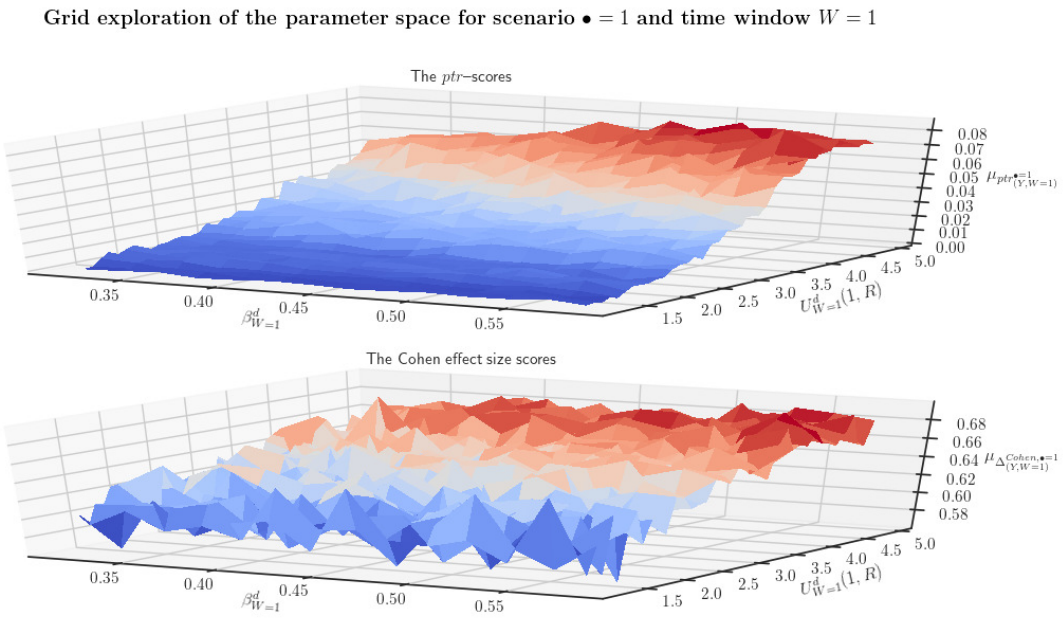

Figure 1. Surface plots of $\mu_{ptr^{\bullet=1}_{(Y,W=1)}}$ and $\mu_{\Delta^{Cohen,\bullet=1}_{(Y,W=1)}}$ for scenario $\bullet = 1$ and time window $W = 1$.

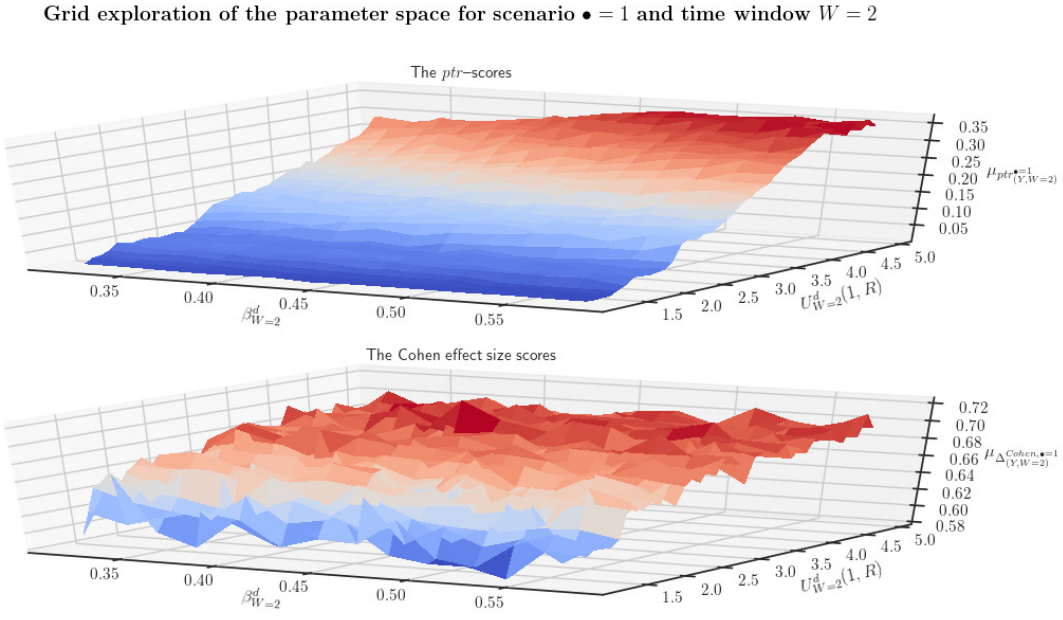

Figure 2. Surface plots of $\mu_{ptr^{\bullet=1}_{(Y,W=2)}}$ and $\mu_{\Delta^{Cohen,\bullet=1}_{(Y,W=2)}}$ for scenario $\bullet = 1$ and time window $W = 2$.

The computed mean power of the conducted statistical tests (i.e. McNemar tests) for the time windows are $\beta = 0.96 \pm 0.02$ for $W = 1$ and $\beta = 0.95 \pm 0.02$ for $W = 2$, so the power of the conducted McNemar tests is sufficient for security analysts of a patent applying company.

For the basic scenarios $\bullet = 2$ and $\bullet = 3$, only surface plots for the scenario option *COW* are shown (see Subsection 3.4.2 and Subsection 3.4.3). In this scenario option, an adversarial co–worker in a patent applying company's intellectual property department tries to affect decision outcomes of the company's decision–maker X (who is unaware of any adversarial influencing). Figure 3 and Figure 4 below show the surface plots for scenario option $S^{\bullet=2}_{COW}$ and time window $W = 1$ and $W = 2$, respectively.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

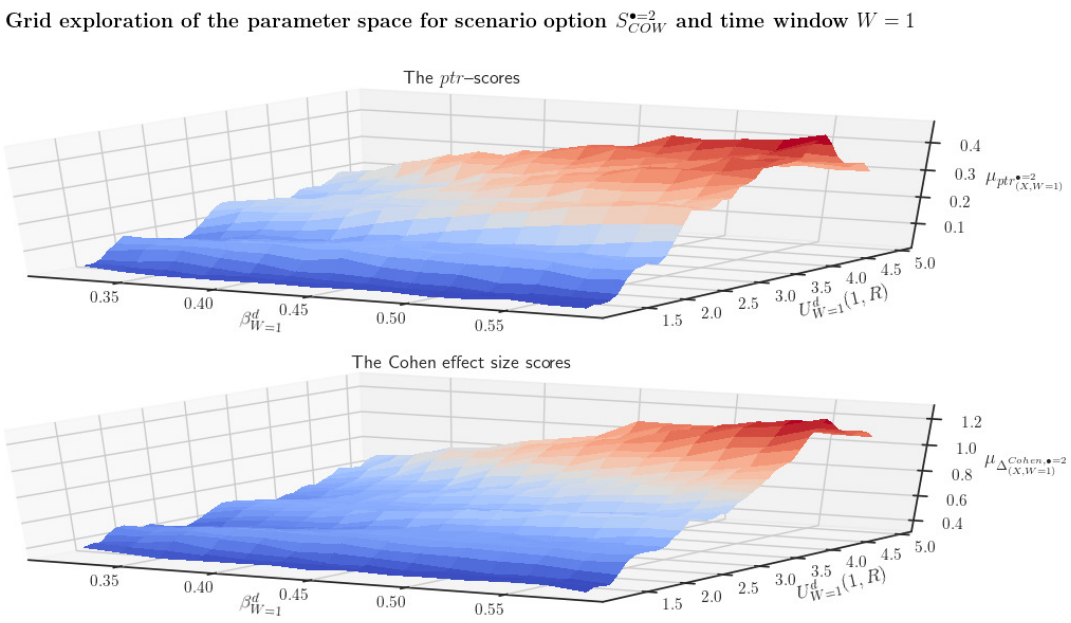

Figure 3. Surface plots of $\mu_{ptr^{\bullet=2}_{(X,W=1)}}$ and $\mu_{\Delta^{Cohen,\bullet=2}_{(X,W=1)}}$ for scenario option $S^{\bullet=2}_{COW}$ and $W = 1$.

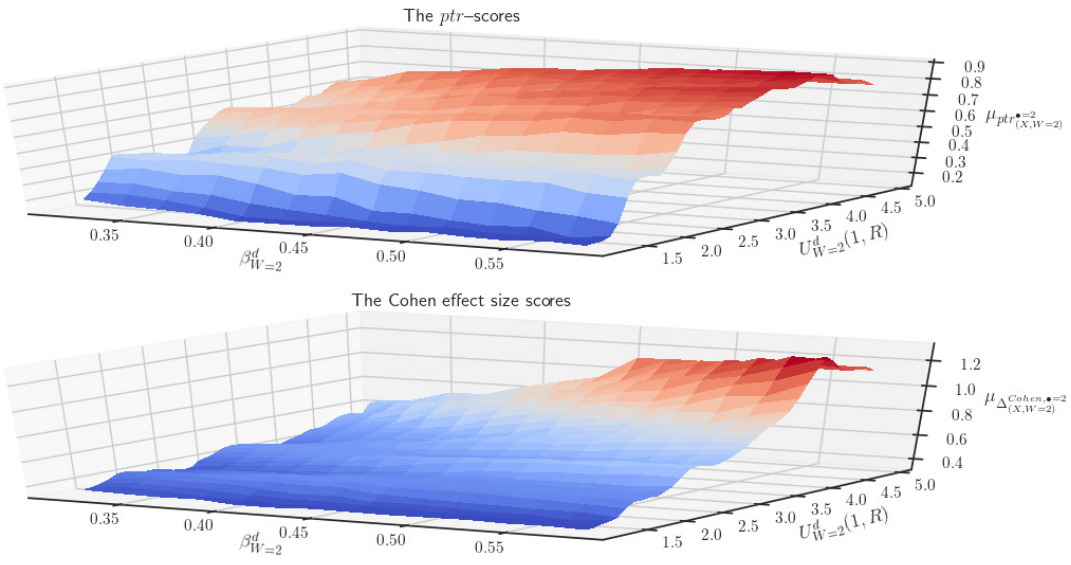

Figure 4. Surface plots of $\mu_{ptr^{\bullet=2}_{(X,W=2)}}$ and $\mu_{\Delta^{Cohen,\bullet=2}_{(X,W=2)}}$ for scenario option $S^{\bullet=2}_{COW}$ and $W = 2$.

For this scenario option, all surface plots show a smooth landscape. For both considered time windows, the landscape of the mean Cohen distance shows a warm colored area in the upper right corner, with a maximum that is even hot colored (i.e. values ≥ 0.8). For the time window $W = 1$, the warm colored area more or less coincides with the warm colored area of the sample mean *ptr*–scores, whereas the warm colored area of the sample mean *ptr*–scores for the time window $W = 2$ is much broader than is the case for $W = 1$. In addition, the values of the sample mean *ptr*–scores for the time window $W = 2$ are considerably higher than those for the time window $W = 1$. In the cooler areas of the surface plot for time window W=1, however, there are areas with medium sample mean Cohen distances and relatively low sample mean *ptr*–scores. Hence, this area might be attractive for an adversarial actor. The computed mean power of the conducted McNemar tests are $\beta = 0.65 \pm 0.01$ for the time window $W = 1$ and $\beta = 0.72 \pm 0.01$ for the time

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

window $W = 2$, so the conducted McNemar tests lack some power. Overall, this scenario option does not seem to be attractive for an adversarial actor to deploy for a period longer than 1 year.

Figure 5 and Figure 6 below show the surface plots for the scenario option $S^{\bullet=3}_{COW}$ and time windows $W = 1$ and $W = 2$, respectively.

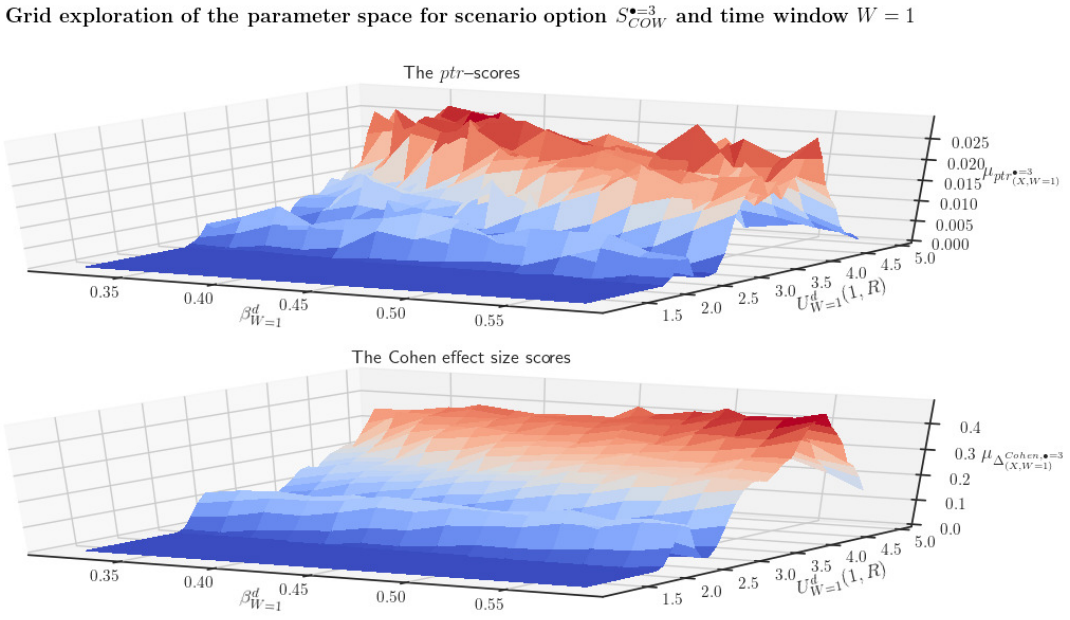

Figure 5. Surface plots of $\mu_{ptr^{\bullet=3}_{(X,W=1)}}$ and $\mu_{\Delta^{Cohen,\bullet=3}_{(X,W=1)}}$ for scenario option $S^{\bullet=3}_{COW}$ and $W = 1$.

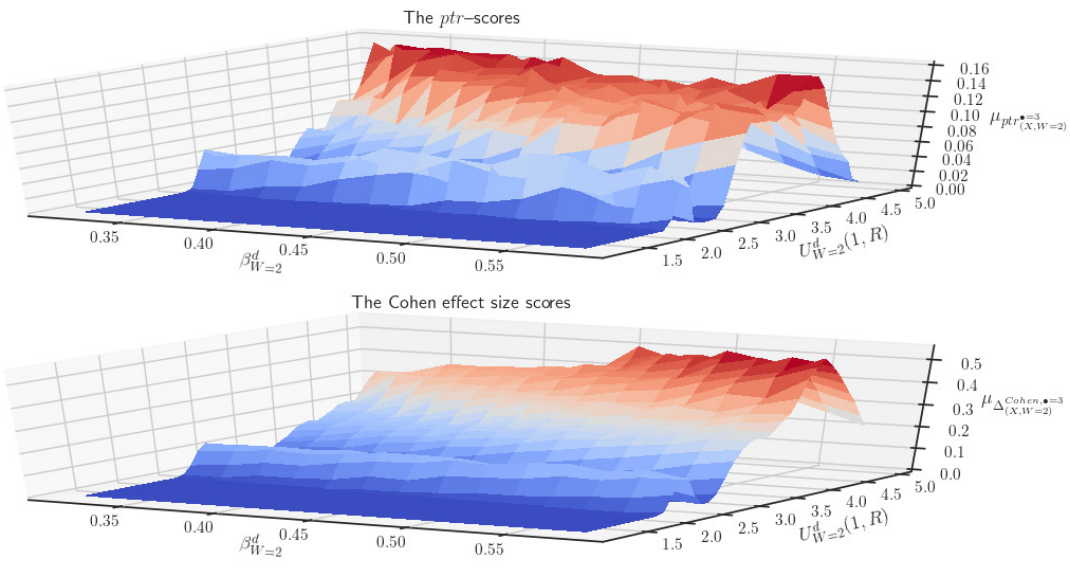

Figure 6. Surface plots of $\mu_{ptr^{\bullet=3}_{(X,W=2)}}$ and $\mu_{\Delta^{Cohen,\bullet=3}_{(X,W=2)}}$ for scenario option $S^{\bullet=3}_{COW}$ and $W = 2$.

Of the three basic influencing scenarios, scenario option $S^{\bullet=3}_{COW}$ seems to be the most attractive for an adversarial actor to deploy, especially in the time window $W = 1$. This is mainly due to the low sample mean ptr–scores at this time window, even in the warm colored area of the surface. Unlike the other two basic scenarios, the sample mean Cohen distances in the warm area of this scenario are small (ranging from below 0.1 to below 0.4).

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

The computed mean power of the conducted McNemar tests are $\beta = 0.62 \pm 0.01$ for the time window $W = 1$ and $\beta = 0.69 \pm 0.01$ for the time window $W = 2$, so the conducted McNemar tests lack some power.

The above inspection makes clear that adversarial actors in their own simulation and analysis somehow need to make a trade–off between the sample mean *ptr*–score (i.e. the likelihood of occurrence of positive test results) and the sample mean Cohen distance (i.e. the expected effect of the scenario), in order to determine the attractiveness of an influencing scenario in a time window. Thereby taking into account that the values of both statistical quantities for each influencing scenario and considered time windows W strongly depend on the positions of the model M_1's parameter pair $(\beta_W^d, U_W^d(1,R))$ in the surface landscape. From an adversarial risk analysis (ARA) perspective, of importance to company security analysts is to find the most favorable model parameter pair $(\beta_W^{d,*}, U_W^{d,*}(1,R))$ for each time window that makes it on average the least attractive for adversaries to deploy either of the six considered scenarios. For company security analysts, as well as for adversaries, it is also of concern to find out whether combining basic scenarios simply implies addition of their sample mean Cohen distances, or that combining may cause some form of cancelling out of sample mean Cohen distances, due to the mathematical structure of model M_1. In other words, what is the most favorable parameter value pair of model M_1 for each time window for the set of six adversarial influencing scenarios? This is the subject of Subsection 2.2.

2.2. The most favorable M_1 model parameter pair for each time window

In model M_2, a multi–objective optimization problem with a set of six attractiveness objective functions is formulated, in order to find the most favorable parameter value pair $(\beta_W^{d,*}, U_W^{d,*}(1,R))$ of model M_1 for a time window W (see Subsection 3.8). The NSGA–II evolutionary optimization method [11,12] is used to solve the optimization problem. This method yields a set of favorable model parameter pairs for a time window W, denoted by $\{(\beta_W^{(i),d,*}, U_W^{(i),d,*}(1,R))\}_{i=1}^{N_{pop}}$, where N_{pop} is the population size used in the NSGA–II method. In the simulations $N_{pop} = 50$, meaning that the method yields 50 favorable model parameter pairs. Associated with each such pair is an attractiveness value $A_{\bullet,W}^{(i)}$ for the corresponding scenario/scenario option. Due to the definition of $A_{\bullet,W}^{(i)}$ (see Subsection 3.7), some favorable model pairs in the above set may correspond with (very) high attractiveness values. By putting some threshold value on the attractiveness value, undesirable favorable model parameters pairs will be dropped, and this yields the reduced set $\{(\beta_W^{(i),d,*}, U_W^{(i),d,*}(1,R))\}_{i=1}^{N_{sel}}$ of selected model parameter pairs, with $1 < N_{sel} \leq N_{pop} = 50$. Figure 7 below shows an example of the frequency distributions of the attractiveness values corresponding to the reduced parameter pair set, for the scenario option *COW*. The figure reveals that the scenario $\bullet = 1$, the scenario option $S_{COW}^{\bullet=3}$ and the combined scenario $\bullet = 1 + S_{COW}^{\bullet=3}$ can potentially do more harm to the decision outcomes of the patent applying repeated decision–making process than scenario $\bullet = 2$ can do, especially in the time window $W = 1$. The figure also reveals that the longer each scenario is deployed, the less harmful it is, and the less spreaded are the corresponding attractiveness values (i.e. smaller bin sizes). This is due to getting more reliable statistics with growing numbers of patent requests in a time window. Figure 7 also reveals that the combination of scenario $\bullet = 1$ with the scenario option $S_{COW}^{\bullet=2}$ can potentially do considerably less harm than scenario $\bullet = 1$ can do on its own. Something similar is the case when combining the scenario options $S_{COW}^{\bullet=2}$ and $S_{COW}^{\bullet=3}$. The harm that the combination of scenario $\bullet = 1$ with the scenario option $S_{COW}^{\bullet=3}$ can do is confuse.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

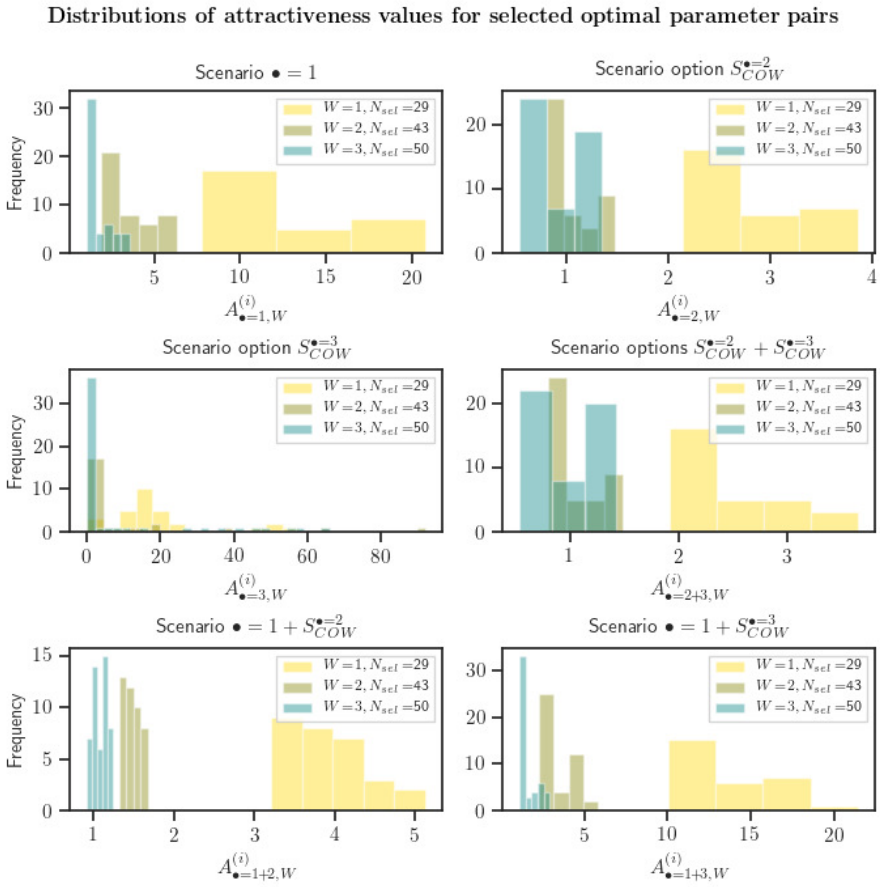

Figure 7. Frequency distributions of the attractiveness values $A^{(i)}_{\bullet,W}$ corresponding to the reduced set of favorable M_1 parameter value pairs $\{(\beta^{(i),d,*}_W, U^{(i),d,*}_W(1,R))\}_{i=1}^{N_{sel}}$ for the three time windows and six considered scenarios, for the scenario option *COW*.

From the perspective of an adversarial actor, scenario $\bullet = 1$ and scenario option $S^{\bullet=3}_{COW}$ are potentially attractive on their own, even when considering favorable parameter pairs, whereas combining each of them with another scenario option reduces attractiveness. This is especially true for the time window $W = 1$. Based on all of the above findings, company security analysts have to apply some *selection procedure* on the reduced set $\{(\beta^{(i),d,*}_W, U^{(i),d,*}_W(1,R))\}_{i=1}^{N_{sel}}$ of M_1 model parameter pairs, in order to arrive at the single most favorable model parameter $(\beta^{d,*}_W, U^{d,*}_W(1,R))$ for a time window (see Subsection 3.8). This selection procedure is company specific and may therefore be hard for adversaries to guess. Table 1 below shows an example of the selected most favorable M_1 model parameter pair for each time time window. The coordinates of the selected most favorable parameter pairs are in agreement with the inspection results described in Subsection 2.1, based on inspecting surface plots.

Table 1. Selected most favorable M_1 model parameter pair for each of the three time windows, and the set of six considered influencing scenarios, for scenario option *COW*.

	COW	
W	$\beta^{d,*}_W$	$U^{d,*}_W(1,R)$
1	0.571	4.997
2	0.574	3.468
3	0.560	3.334

Figure 8 below is similar to Figure 7 above, except that the scenario option X is

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

considered instead of *COW*. In scenario option X, the company decision maker X is the adversarial actor who negatively influences the outcomes of decision A (see Subsection 3.4.2 and Subsection 3.4.3).

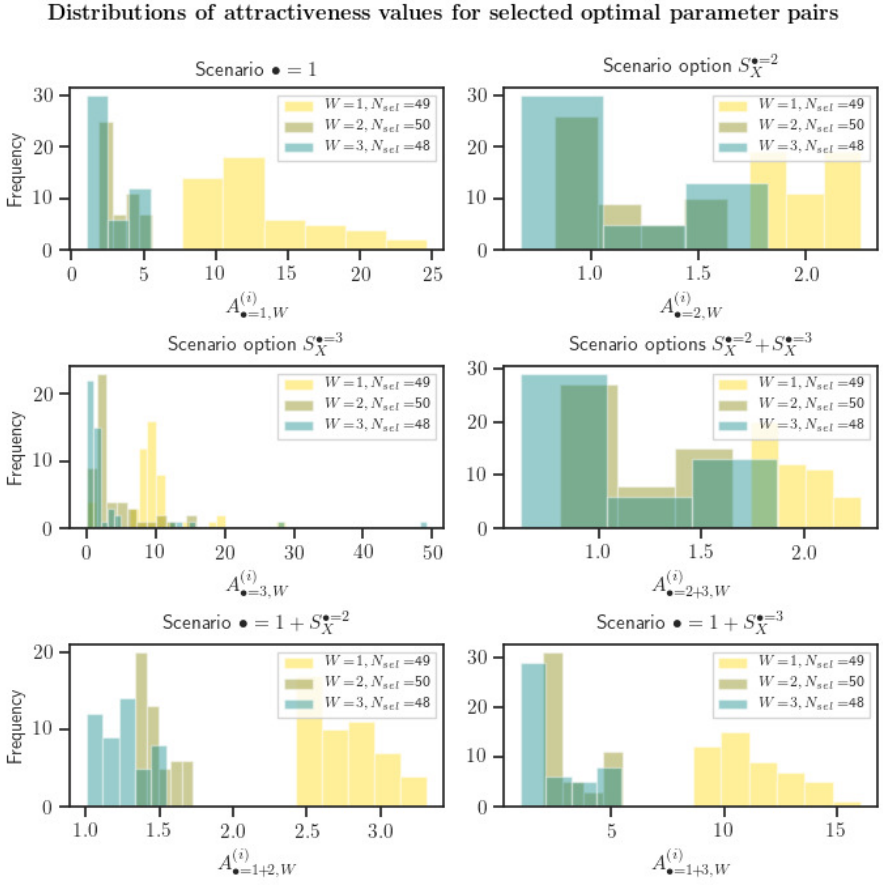

Figure 8. Frequency distributions of the attractiveness values $A^{(i)}_{\bullet,W}$ corresponding to the reduced set of favorable M_1 parameter value pairs $\{(\beta_W^{(i),d,*}, U_W^{(i),d,*}(1,R))\}_{i=1}^{N_{sel}}$ for the three time windows and six considered scenarios, for the scenario option X.

Figure 8 reveals that the frequency distributions of the attractiveness values for scenario option X resemble those of scenario option *COW*, except that they are less spreaded and most frequency mass has shifted to the left. The frequency distribution for the scenario $\bullet = 1$ is slightly more spreaded than is the case for scenario option *COW*, within the normal statistical variation. As is the case for scenario option *COW*, the scenario $\bullet = 1$ and scenario option $S_X^{\bullet=3}$ are potentially more harmful than scenario option $S_X^{\bullet=2}$, and combinations of them with another scenario option reduce their attractiveness. Table 2 below shows an example of the selected most favorable M_1 model parameter pairs for each time time window for scenario option X.

Table 2. Selected most favorable M_1 model parameter pairs for each of the three time windows, and the set of six considered influencing scenarios, for scenario option X.

	COW	
W	$\beta_W^{d,*}$	$U_W^{d,*}(1,R)$
1	0.471	4.784
2	0.333	4.887
3	0.336	4.246

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

3. Materials and Methods

As stated in the introduction, the probabilistic scenario of the patent application decision–making process is not Bayesian, in that the company decision–maker X selects a choice alternative (decision A: region office or country office) before observing the outcome of decision B (patent request granted or not granted). Moreover, decision–maker X is in fact a bounded rational decision–maker. Subsection 3.1 briefly formalizes the theory of Ortega et al. that is used to model a bounded rational decision–maker for the above probabilistic scenario, and provides an example too

3.1. Formalization of the theory of Ortega et al. for the probabilistic scenario of the patent application decision–making process

In real–world decision problems, a decision–maker does not always have enough resources to exhaustively evaluate all aspects of each choice alternative of a decision. Ortega et al. have shown that this limitation changes a decision problem in a fundamental way. Their theory first requires defining a finite outcome space $\mathcal{X}$, be defined as:

$$\begin{aligned} \mathcal{X} &= \mathcal{A} \times \mathcal{O}, \\ &= \underbrace{\{a_1, \cdots, a_M\}}_{\text{Choice alternatives}} \times \underbrace{\{o_1, \cdots, o_N\}}_{\text{Observations}}, \end{aligned}$$

where $0 < N, M \in \mathbb{N}$, $\mathcal{A}$ is a finite space of choice alternatives and $\mathcal{O}$ is a finite space of possible observations. Furthermore, in the probabilistic scenario in which the decision–maker first selects a choice alternative a and then observes the stochastic state of the world o, the theory conceptualizes a decision–maker's deliberation and planning process as follows. The decision–maker first chooses a (what Ortega et al. call) prior decision policy, i.e. a probability distribution $P_{0,\mathcal{X}}(x)$, and then transforms this policy into a (what they call) posterior decision policy, i.e. a probability distribution $P_{\mathcal{X}}(x)$, be defined as:

$$\begin{aligned} P_{\mathcal{X}}(x) &= P_{\mathcal{X}}(a,o) = P(o|a)P_{\mathcal{A}}(a), \\ P_{0,\mathcal{X}}(x) &= P(o|a)P_{0,\mathcal{A}}(a). \end{aligned} \tag{1}$$

During this transformation process, the decision–maker is not allowed to reason about the costs of transforming a prior decision policy into a posterior decision policy. Furthermore, $U : \mathcal{X} \to \mathbb{R}$ is a real–valued mapping of the outcomes, called the utility function. The decision maker's goal is to find the optimal posterior decision policy $P^*_{\mathcal{X}}(x)$ by optimizing this utility function over the probability distribution $P_{\mathcal{X}}(x)$, while facing limited information processing resources in the deliberation and planning process. Ortega et al. showed that this limitedness for an outsider will appear as if the decision–maker were explicitly optimizing the explicit objective function $-\Delta F_\beta[P]$, known as the functional for negative free energy difference due to its origin in thermodynamics:

$$-\Delta F_\beta[P] := \underbrace{\sum_{x \in \mathcal{X}} P(x)U(x)}_{\text{Expected Utility}} - \underbrace{\frac{1}{\beta} \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{P_0(x)}}_{\text{Information Cost}}. \tag{2}$$

The second term in the formula expresses information cost due to limited resources measured in units of utility (i.e., utiles), and the boundedness parameter $\beta \in \mathbb{R}$ acts as a conversion factor between units of information and utiles. The functional in Equation (2) expresses information–theoretic bounded rationality as a tradeoff between utilities and information cost, and is reflecting the decision–maker's net utility. In the literature, this cost term also goes under other names, such as KL–control cost, and has been motivated in numerous ways [7,10]. The boundedness parameter not only acts as a conversion factor, but also scales how far $P(x)$ can deviate from $P_0(x)$, measured in terms of the KL-divergence.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

The parameter therefore controls how much a decision maker is in control of the action of selecting a choice alternative (see the limit cases in Table 3 below).

Table 3. Limit cases for the boundedness parameter.

Limit case	Actions
$\beta \to \infty$	Perfectly rational decision maker with unlimited resources
$\beta \to 0$	Decision maker without resources simply selects an action according to the prior decision policy $P_0(x)$
$\beta \to -\infty$	Perfectly anti-rational decision maker, which always selects the action with the worst outcome.

To find the bounded rational optimal posterior decision policy, Ortega et al. have formulated a variational principle for maximizing the functional over probability distributions $P(x)$. The general solution of this variational principle is the optimal posterior decision policy:

$$P^*_\beta(x) = \frac{1}{Z_\beta} P_0(x) e^{\beta U(x)}, \text{ with } Z_\beta = \sum_{x \in \mathcal{X}} P_0(x) e^{\beta U(x)}. \tag{3}$$

For the probabilistic scenario in the decision–making of the patent application process (i.e. choice alternative selection decision A before observation of decision outcome of decision B), the particular optimal solution over a finite action space $\mathcal{A}$ becomes:

$$P^*_{\beta,\mathcal{A}}(a) = \frac{1}{Z_{\beta,\mathcal{A}}} P_{0,\mathcal{A}}(a) e^{\beta \mathbb{E}[U|a]}, \tag{4}$$

with $\mathbb{E}[U|a] = \sum_{o \in \mathcal{O}} [P(o|a)U(o,a)]$ and $Z_{\beta,\mathcal{A}} = \sum_{a \in \mathcal{A}} P_{0,\mathcal{A}}(a) e^{\beta \mathbb{E}[U|a]}$. The reader can find the derivation of this formula in Appendix A.

As of now, the case of decision–making in complete absence of adversarial influencing is referred to as the *default case*, denoted by the superscript d. In addition, R and C represent the choice alternatives “region office" and “country office", respectively. And 0 and 1 represent the decision outcomes “patent request not granted" and “patent request granted", respectively.

Example 1. *Let $\mathcal{A} = \{R, C\}$ and $\mathcal{O} = \{0, 1\}$. Let the utility function be defined as $U^d(0,R) = U^d(0,C) = 0$ if a patent has not been granted, and $U^d(1,R) = 45$ and $U^d(1,C) = 1$ if a patent has been granted. Let $P^d(0|R) = 0.92 = 1 - P^d(1|R)$ be the probability of a patent not being granted by a regional office in the default case, and $P^d(0|C) = 0.2 = 1 - P^d(1|C)$ the probability of a patent not being granted by a country office. Let the prior decision policy that the patent will be requested at a regional office be given by $P^d_{o,\mathcal{A}}(a = R) = 0.1 = 1 - P^d_{0,\mathcal{A}}(a = C)$. Choose a value for the boundedness parameter β^d (here $\beta^d = 1$), representing how much the bounded rational decision maker X is in control of selecting choice alternative R or C. Now, the optimal decision policy becomes: $P^{d,*}_{\beta^d,\mathcal{A}}(a) = P^{d,*}_{\mathcal{A}}(a; \beta^d = 1) = \frac{P^d_{0,\mathcal{A}}(a) e^{\mathbb{E}[U^d|a]}}{Z^d_{\beta^d=1,\mathcal{A}}}$, where $\mathbb{E}[U^d|a] = \sum_{o \in \mathcal{O}} P^d(o|a) U^d(o,a)$ is the expected utility in case choice action a is selected, and the partition function $Z^d_{\beta^d=1,\mathcal{A}} = \sum_{a \in \mathcal{A}} P^d_{0,\mathcal{A}}(a) e^{\mathbb{E}[U^d|a]}$. Then $\mathbb{E}[U^d|R] = 0.08 * 45 = 3.6$, $\mathbb{E}[U^d|C] = 0.8 * 1 = 0.8$, and the partition function $Z^d_{\beta^d=1,\mathcal{A}} = 0.1e^{3.6} + 0.9e^{0.8}$. The optimal decision policy components then become $P^{d,*}_{\mathcal{A}}(a = R; \beta^d = 1) = \frac{0.1e^{3.6}}{0.1e^{3.6} + 0.9e^{0.8}} = 0.65$ and $P^{d,*}_{\mathcal{A}}(a = C; \beta^d = 1) = 0.35$. The below decision function is used to determine the outcome of decision A:*

$$D_A(a; \beta^d) = \begin{cases} R & \text{if } P^{d,*}_{\mathcal{A}}(a = R; \beta^d) > 0.5, \\ C & \text{otherwise.} \end{cases} \tag{5}$$

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

As expected, with differences of utility values for both types of patent offices as above, decision maker X adapts his/her strategy according to the optimal decision policy and decision rule, and decides to send the patent request to a regional patent office.

3.2. Repeated patent application decision–making

Repeated patent application decision–making involves repeating the decision–making process that is described in Subsection 3.1. Let r denote an individual patent request. To capture variety in repeated patent requests in the simulation study, the following bounded rational decision–making model will be applied in the default case:

Model M_1 :

$$
\begin{aligned}
\mathcal{A} &= \{R, C\} \\
\mathcal{O} &= \{0, 1\} \\
\beta^{d,r} &= \beta^d + 0.05 * randint(1,5) - 0.15, \text{ with } 0.30 < \beta^d < 0.60, \\
U^{d,r}(0,R) &= U^{d,r}(0,C) = 0.0 \text{ for all } r, \\
U^{d,r}(1,C) &= U^{d}(1,C) = 1.0 \text{ for all } r, \\
U^{d,r}(1,R) &= U^{d}(1,R) + 0.05 * randint(1,5) - 0.15, \\
&\quad \text{with } 1.1 < U^{d}(1,R) < 5.0, \\
U^{d,r}(1,R) &> U^{d,r}(1,C) \text{ for all } r, \\
P^{d,r}(1|C) &= 0.40 \text{ for all } r, \\
P^{d,r}(1|R) &= 0.20 \text{ for all } r, \\
P^{d,r}_{0,\mathcal{A}}(a=R) &= 0.43 + 0.02 * randint(1,5), \\
\mathbb{E}[U^{d,r}|a; U^d(1,R)] &= \sum_{o \in \mathcal{O}} P^{d,r}(o|a) U^{d,r}(o,a), \text{ with } a \in \mathcal{A}, \\
Z^{d,r}_{\beta^{d,r},\mathcal{A}} &= \sum_{a \in \mathcal{A}} P^{d,r}_{0,\mathcal{A}}(a) e^{\mathbb{E}[U^{d,r}|a;U^d(1,R)]}, \\
P^{d,r,*}_{\mathcal{A}}(a=R;\beta^d, U^d(1,R)) &= \frac{P^{d,r}_{0,\mathcal{A}}(a=R) e^{\mathbb{E}[U^{d,r}|a=R;U^d(1,R)]}}{Z^{d,r}_{\beta^{d,r},\mathcal{A}}}, \\
D^r_A(a;\beta^d, U^d(1,R)) &= \begin{cases} \text{R} & \text{if } P^{d,r,*}_{\mathcal{A}}(a=R;\beta^d, U^d(1,R)) > 0.5 \text{ for all } r, \\ \text{C} & \text{otherwise.} \end{cases}
\end{aligned}
\tag{6}
$$

where the values of the parameters β^d and $U^d(1,R)$ remain to be specified (as stated in the introduction), and $randint(a,b)$ represents a uniform drawing from the interval $[a,b]$. The quotient $\frac{U^{d,r}(1,R)}{U^{d,r}(1,C)}$ expresses decision maker X's preference with regard to how much more important a grant at the regional office R is for patent a request r than a grant at a country office C. By lowering the boundedness parameter $\beta^{d,r}$ compared to β^d, decision maker X admits to have less control of selecting choice alternative R or C for patent request r. By raising the value of the prior decision policy parameter $P^{d,r}_{0,\mathcal{A}}(a=R)$ compared to 0.43, decision maker X is more confident that patent request r will be granted by a patent officer at regional office R. By raising the utility value $U^{d,r}(1,C)$ compared to the value $U^d(1,C)$, decision maker X lowers his/her preference for getting the patent request granted at office R.

3.3. Observed proportions of patent application decision–making outcomes in the default case

The repeated decision–making process described in Subsection 3.2 yields two equally sized time ordered sequences (i.e. samples) of $1 < N(W) \in \mathbb{N}$ decision outcomes from individual patent requests r in a time window W (expressed in years). A sample of R/C

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

outcomes (from decision A) and a sample of $0/1$ outcomes (from decision B). This study focuses on decision A outcomes R, and subsequent conditional decision B outcomes $1|R$, and defines two *observed* unaffected proportions of decision outcomes R and $1|R$ for the default case:

$$\begin{aligned} \widehat{p}^{\,d}_{\mathrm{R}}(W;\beta^d,U^d(1,R)) &= \frac{\#\text{ unaffected } R\text{–outcomes in sample decision } A}{N(W)} \\ \widehat{p}^{\,d}_{1\,|\,\mathrm{R}}(W;\beta^d,U^d(1,R)) &= \frac{\#\text{ unaffected } 1|R\text{–outcomes in sample decision } B}{N(W)} \end{aligned} \tag{7}$$

3.4. Observed proportions of adversarial influenced patent application decision–making outcomes

In this study, six scenarios are defined by which adversarial actors may negatively influence patent application process decision–making outcomes, compared to the default case. An adversarial influencing scenario is denoted by the superscript $\bullet$, followed by the number of the scenario.

3.4.1. Adversarial influencing scenario $\bullet = 1$

An adversarial specialized patent officer in the regional patent office is able to create the opportunity to assess all the patent requests that are sent to the office by a company. The patent officer knows the observed granting chance $p^{d,r}(1|R)$ the decision maker X in the company is counting on for the sent patent requests in the default case, and tries to negatively influence the value of this chance, without raising suspicion. He/she first determines the lowest number of patent requests $N^{\text{lowest}} \in \mathbb{N}$ that will approximately result in the chance value $p^{d,r}(1|R)$ if just one of the N^{lowest} patent requests is granted: $N^{\text{lowest}} = \text{round to the nearest lowest integer}\left(\frac{1}{p^{d,r}(1|R)}\right)$. Suppose $p^{d,r}(1|R) = 0.40$, then $N^{\text{lowest}} = \text{round to the nearest lowest integer}(\frac{1}{0.40}) = 2$. This means that with just one granted patent on $N^{\text{lowest}} + 1$ patent requests, the resulting adversarial influenced granting chance $p^{\bullet=1,r}(1|R) = \frac{1}{N^{\text{lowest}}+1} = 0.33$. So, the officer's strategy is to not grant N^{lowest} patents on every $N^{\text{lowest}} + 1$ patent requests. This strategy takes into account that it may be hard to prove for company security analysts that there is a statistically significant difference between the observed unaffected proportion $\widehat{p}^{\,d}_{1\,|\,\mathrm{R}}(W;\beta^d,U^d(1,R))$, defined in Equation (7), and the affected proportion:

$$\widehat{p}^{\,\bullet=1}_{1\,|\,\mathrm{R}}(W;\beta^d,U^d(1,R)) = \frac{\#\text{ affected } 1|R\text{–outcomes in sample decision } B}{N(W)}. \tag{8}$$

Note that $\widehat{p}^{\,\bullet=1}_{\mathrm{R}}(W;\beta^d,U^d(1,R)) = \widehat{p}^{\,d}_{\mathrm{R}}(W;\beta^d,U^d(1,R))$.

3.4.2. Adversarial influencing scenario $\bullet = 2$

In this scenario, company decision maker X's decision A for individual patent requests r is influenced by either one of the below two scenario options:

- *Scenario option* $S^{\bullet=2}_{COW}$: An adversarial co–worker in the company's intellectual property department tries to persuade decision maker X to raise the value of the utility component $U^{d,r}(1,C)$ for a patent request with an integer value, with X being unaware of this.
- *Scenario option* $S^{\bullet=2}_{X}$: Decision maker X is the adversarial actor and raises the value of the utility component $U^{d,r}(1,C)$ for a patent request with an integer value himself/herself.

To capture both options, the following mathematical formulation is used:

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

$$
\begin{aligned}
&U^{\bullet=2,r}(1,C) = U^{d,r}(1,C) + v^r,\\
&\text{with } v^r \sim (P(v^r=0)=p_0, P(v^r=1)=p_1, P(v^r=2)=p_2, P(v^r=3)=p_3),\\
&\text{and } p_0+p_1+p_2+p_3 = 1.
\end{aligned} \tag{9}
$$

In scenario option $S^{\bullet=2}_{COW}$, the value of v^r is drawn from the distribution $(P(v^r=0) = 0.50, P(v^r=1) = 0.30, P(v^r=2) = 0.15, P(v^r=3) = 0.05)$, meaning that the adversarial co–worker has a 50% chance that decision maker X is willing to accept a proposed raise of $U^{d,r}(1,C)$. In scenario option $S^{\bullet=2}_{X}$, the value of v^r is drawn from the distribution $(P(v^r=0) = 0.20, P(v^r=1) = 0.45, P(v^r=2) = 0.27, P(v^r=3) = 0.08)$. Raising the value of $U^{d,r}(1,C)$ leads to a value of the chance component $P^{\bullet=2,r,*}_{\mathcal{A}}(a=R;\beta^d,U^d(1,R))$ that is lower than the value of the chance component $P^{d,r,*}_{\mathcal{A}}(a=R;\beta^d,U^d(1,R))$, and the more likely it is that the number of decision outcomes R will drop. Therefore, the value of the below defined affected observed proportion is expected to be lower than the value of the corresponding observed unaffected proportion:

$$
\widehat{p}^{\bullet=2}_{\mathrm{R}}(W;\beta^d,U^d(1,R)) = \frac{\#\text{ affected } R\text{–outcomes in sample decision } A}{N(W)}. \tag{10}
$$

Though this scenario does not affect repeated decision B outcomes, in simulation runs the below defined observed affected proportion may well differ from the value of the corresponding observed unaffected proportion:

$$
\widehat{p}^{\bullet=2}_{\mathrm{R}}(W;\beta^d,U^d(1,R)) = \frac{\#\text{ affected } 1|R\text{–outcomes in sample decision } B}{N(W)}. \tag{11}
$$

3.4.3. Adversarial influencing scenario $\bullet = 3$

In this scenario, company decision maker X's decision A for individual patent requests r is influenced by either one of the below two scenario options:

- *Scenario option* $S^{\bullet=3}_{COW}$: An adversarial co–worker in the company's intellectual property department tries to persuade decision maker X to decrease the value of the boundedness parameter $\beta^{d,r}$ for a patent request with some percentage, with X being unaware of this.
- *Scenario option* $S^{\bullet=3}_{X}$: Decision maker X is the adversarial actor and decreases the value of the boundedness parameter $\beta^{d,r}$ for a patent request with some percentage himself/herself.

To capture both scenario options, the following mathematical formulation is used:

$$
\begin{aligned}
\beta^{\bullet=3,r} \quad = \quad & \beta^{d,r}(1-\frac{p^r}{100}), \text{ with } p^r \text{ being a drawing from the distribution}\\
& (P(p^r=0)=p_0, P(p^r=30)=p_1, P(p^r=40)=p_2) \text{ and,}\\
& p_0+p_1+p_2 = 1.
\end{aligned} \tag{12}
$$

In scenario option $S^{\bullet=3}_{COW}$, the value of p^r is drawn from the distribution $(P(p^r=0) = 0.40, P(p^r=30) = 0.40, P(p^r=40) = 0.20)$, and in scenario option $S^{\bullet=3}_{X}$ from the distribution $(P(p^r=0) = 0.20, P(p^r=30) = 0.50, P(p^r=40) = 0.30)$. A decrease of $\beta^{d,r}$ may drop the value of the chance component $P^{\bullet=3,r,*}_{\mathcal{A}}(a=R;\beta^d,U^d(1,R))$, and may result in a lower number of decision A outcomes R. Therefore, the value of the below defined affected observed proportion is expected to be lower than the value of the corresponding observed unaffected proportion:

$$
\widehat{p}^{\bullet=3}_{\mathrm{R}}(W;\beta^d,U^d(1,R)) = \frac{\#\text{ affected } R\text{–outcomes in sample decision } A}{N(W)}. \tag{13}
$$

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

As is the case for influencing scenario $\bullet = 2$, this scenario does not affect repeated decision B outcomes. However, in simulation runs the below defined observed affected proportion may well differ from the value of the corresponding observed unaffected proportion:

$$\widehat{p}_{\mathrm{R}}^{\bullet=3}(W;\beta^d,U^d(1,R)) = \frac{\#\text{ affected } 1|R\text{–outcomes in sample decision } B}{N(W)}. \tag{14}$$

3.4.4. Combined influencing scenario $\bullet = 2+3$

In this combined influencing scenario, either the combination of scenario options $S_{COW}^{\bullet=2}$ and $S_{COW}^{\bullet=3}$ is active, or the combination of scenario options $S_{X}^{\bullet=2}$ and $S_{X}^{\bullet=3}$. Combining the individual influencing scenarios offers the adversarial actor the advantage that smaller value changes of ev^r and p^r may be more effective. However, the risk of exposure may be higher than in case a single influencing scenario is deployed. The definitions of the two observed affected proportions for this scenario are identical to the definitions for the individual scenarios $\bullet = 2$ and $\bullet = 3$.

3.4.5. Combined influencing scenarios $\bullet = 1+2$ or $\bullet = 1+3$

In these two combined scenarios, the patent office–side adversarial actor and the company–side adversarial actor do cooperate. Scenario $\bullet = 1+2$ is a combination of scenario $\bullet = 1$ with either scenario option $S_{COW}^{\bullet=2}$ or scenario option $S_{X}^{\bullet=2}$. And scenario $\bullet = 1+3$ is a combination of scenario $\bullet = 1$ with either scenario option $S_{COW}^{\bullet=3}$ or scenario option $S_{X}^{\bullet=3}$. If the company–side adversarial actor succeeds in dropping the number of patent requests that is send to the regional patent office, then it may be statistically harder for company security annalists to test for the presence of scenario $\bullet = 1$, being deployed by the patent office–side adversarial actor.

3.5. Testing for the presence of an adversarial influencing scenario

To find out whether an adversarial influencing scenario $\bullet$ has been active on decision A–outcomes in a time window W, or not, a paired proportions test will be conducted. To be precise, the *asymptotic McNemar-test* without continuity correction [8]. The distinguishable case outcomes for decision A are captured by the 2x2 contingency table shown in Table 4 below, where $(n_{11}, n_{12}, n_{21}, n_{22})$ denotes a combination of outcome pairs on a total of $N(W)$ pairs.

Table 4. 2 x 2 contingency table with the distinguishable case outcomes for decision A.

		Case influence scenario $\bullet$ be active		
		Outcome R	Outcome C	Totals
Default case d	Outcome R	n_{11}	n_{12}	$n_{11}+n_{12}$
	Outcome C	n_{21}	n_{22}	$n_{21}+n_{22}$
	Totals	$n_{11}+n_{21}$	$n_{12}+n_{22}$	$N(W)$

The McNemar test procedure that is followed in the simulation study is explained below by means of Example 2.

Example 2. *Suppose, the simulation framework has generated $N(W) = 11$ drawn pairs of binary decision outcomes R/C for the default case and the case adversarial influencing $\bullet$ has been deployed, with $(n_{11}, n_{12}, n_{21}, n_{22}) = (6, 3, 0, 2)$ being the generated combination of drawn outcome pairs. This corresponds with the observed proportions $\widehat{p}_R^d(W;\beta^d,U^d(1,R)) = \frac{n_{11}+n_{12}}{N(W)} = \frac{9}{11} = 0.818$ and $\widehat{p}_R^{\bullet}(W;\beta^d,U^d(1,R)) = \frac{n_{11}+n_{21}}{N(W)} = \frac{6}{11} = 0.545$. Do these proportions significantly differ from each other or not, under the null hypothesis H_0 that the proportions of R– outcomes in the population are equal for both cases?*

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

Based on the criterion $n_{12} + n_{21} = 3 + 0 = 3 < 20$, *the exact binomial version of the McNemar test will be conducted, otherwise the normal* (χ^2_1) *approximation. Choose a significance level (here* $\alpha = 0.05$*), and let a software package compute the McNemar score statistic* $M = \frac{(n_{12}-n_{21})}{\sqrt{n_{12}+n_{21}}}$ *and the two-tailed P-value, according to* H_0 *and by using the exact binomial distribution. The computed (two-sided) P-value is equal to 0.25. Because this value is greater than* $\alpha = 0.05$, *we fail to reject* H_0 *and assume there is no significant difference between the proportions.*

To assure that the conducted McNemar test does not lack sufficient power to demonstrate and prove adversarial influencing for small and moderate sample sizes, as well as for larger sample sizes, a power analysis will be performed by means of a software package. The power value β *will be computed given the sample size* $N(W) = 11$, *the significant level* $\alpha = 0.05$ *and the observed effect size* $\Delta(W;\beta^d,U^d(1,R)) = |\widehat{p}^{\bullet}_R(W;\beta^d,U^d(1,R)) - \widehat{p}^{\bullet}_R(W;\beta^d,U^d(1,R))| = 0.273$. *The computed power value* $\beta = 0.701$. *In the performed simulation experiments, the calculated power value should not be far away from the value 0.8, which is normally imposed on a statistic hypothesis test. Instead of using the real difference between the two proportions as effect size, the* Cohen difference *will be used, be defined as the absolute difference between the arcsine-root-transformed values of the proportions [9], i.e.* $\Delta_{Cohen} = |2arcsine(\sqrt{\widehat{p}^{d}_R(W;\beta^d,U^d(1,R))}) - 2arcsine(\sqrt{\widehat{p}^{\bullet}_R(W;\beta^d,U^d(1,R))})|$. *Cohen suggested that* $\Delta_{Cohen} = 0.2$ *can be considered a small effect size,* $\Delta_{Cohen} = 0.5$ *represents a medium effect size and* $\Delta_{Cohen} = 0.8$ *a large effect size. This means that if two proportions do not differ by 0.2 (threshold) standard deviations or more, the difference is trivial, even if it is statistically relevant. Here, the Cohen distance is equal to* 0.599, *representing a more than medium effect.*

In a similar way as for decision A outcomes, the McNemar test can be conducted for decision B outcomes, that is for the paired observed proportions $\widehat{p}^{d}_{1|R}(W;\beta^d,U^d(1,R))$ and $\widehat{p}^{\bullet}_{1|R}(W;\beta^d,U^d(1,R))$, under the null hypothesis H_0 that the proportions of $1|R$– outcomes in the population are equal for both cases. The two McNemar tests for decision A and B will be referred to as MNT(X,W) and MNT(Y,W), respectively (see the test setup shown in Figure 9 below).

Figure 9. Setup to test for the six considered adversarial influencing scenarios in time window W.

figure_test_setup_eps-converted-to.pdf

Depending on the particular influencing scenario $\bullet$, either McNemar test MNT(X,W) or McNemar test MNT(Y,W) has to be conducted, or both McNemar tests have to be conducted to test for the presence of the scenario in time window W. If the null hypothesis in a particular McNemar test is rejected, this is called a positive test result (denoted by +), otherwise the test result is negative (-). A positive test result is translated into the binary score 1, and a negative test result in the score 0. Table 5 below shows for each influencing scenario (option) the applicable McNemar test(s), together with the above score mechanism.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

Table 5. Applicable McNemar tests for the six adversarial influencing scenarios, and the test score mechanism.

Scenario	Option	McNemar test (H_0 rejected)	Test result (+/-)	Score (1/0)
$\bullet = 2$	$S_{COW}^{\bullet=2}$ or $S_{X}^{\bullet=2}$	MNT(X,W)	+/-	1/0
$\bullet = 3$	$S_{COW}^{\bullet=3}$ or $S_{X}^{\bullet=3}$	MNT(X,W)	+/-	1/0
$\bullet = 2+3$	$S_{COW}^{\bullet=2}$ + $S_{COW}^{\bullet=3}$	MNT(X,W)	+/-	1/0
$\bullet = 2+3$	$S_{X}^{\bullet=2}$ + $S_{X}^{\bullet=3}$	MNT(X,W)	+/-	1/0
$\bullet = 1$		MNT(Y,W)	+/-	1/0
$\bullet = 1+2$	$S_{COW}^{\bullet=2}$	MNT(X,W)	+/-	1/0
		MNT(Y,W)	+/-	1/0
$\bullet = 1+2$	$S_{X}^{\bullet=2}$	MNT(X,W)	+/-	1/0
		MNT(Y,W)	+/-	1/0
$\bullet = 1+3$	$S_{COW}^{\bullet=3}$	MNT(X,W)	+/-	1/0
		MNT(Y,W)	+/-	1/0
$\bullet = 1+3$	$S_{X}^{\bullet=3}$	MNT(X,W)	+/-	1/0
		MNT(Y,W)	+/-	1/0

3.6. Building test statistics in the simulation study

In order to statistically examine the presence of a specific adversarial influencing scenario $\bullet$ in the three considered time windows in more detail, test statistics need to be build in the simulation study. Therefore, $N_{sim} = 50$ simulation runs with $N_s = 50$ sub–runs s per simulation run will be performed for each time window. For each simulation run, the number of patent requests in an individual sub–run for the time window, $N^{(s)}(W)$, will be determined by the drawing $N^{(s)}(W) \sim W(2 + randint(1,4))$. For both McNemar tests MNT(X,W) and MNT(Y,W), the percentage of positive test results of a simulation run (denoted by ptr) can be computed over the 50 sub–runs, as well as the mean power and mean Cohen distance. Over the 50 simulation runs, the sample means of these quantities and their associated standard deviations can be computed for the two McNemar tests. Let the latter quantities for the test MNT(X,W) be represented by: $\mu_{ptr^{\bullet}_{(X,W)}}$, $\sigma_{ptr^{\bullet}_{(X,W)}}$, $\mu_{\Delta^{Cohen,\bullet}_{(X,W)}}$, $\sigma_{\Delta^{Cohen,\bullet}_{(X,W)}}$, $\mu_{\beta^{\bullet}_{(X,W)}}$ and $\sigma_{\beta^{\bullet}_{(X,W)}}$, and for test MNT$(Y,W)$ by: $\mu_{ptr^{\bullet}_{(Y,W)}}$, $\sigma_{ptr^{\bullet}_{(Y,W)}}$, $\mu_{\Delta^{Cohen,\bullet}_{(Y,W)}}$, $\sigma_{\Delta^{Cohen,\bullet}_{(Y,W)}}$, $\mu_{\beta^{\bullet}_{(Y,W)}}$ and $\sigma_{\beta^{\bullet}_{(Y,W)}}$. Though not shown in the notation, all these statistical quantities are parameterized by the parameter pair $(\beta^d, U^d(1,R))$, because they are based on an observed unaffected and affected proportion pair and each proportion in this pair is parameterized by β^d and $U^d(1,R)$.

3.7. The attractiveness of an adversarial influencing scenario from the perspective of an adversarial actor

Combining the sample mean of positive test results $\mu_{ptr^{\bullet}_{(\cdot,W)}}$ (with $\cdot = X$ or Y) and the associated sample mean Cohen distance $\mu_{\Delta^{Cohen,\bullet}_{(\cdot,W)}}$ is a way to express the *attractiveness* of influencing scenario (option) $\bullet$ from the perspective of an adversarial actor who considers to deploy $\bullet$ in a time window. The lower the $\mu_{ptr^{\bullet}_{(\cdot,W)}}$ value, the more attractive the scenario (option)/window combination is for an adversarial actor. And, the higher the $\mu_{\Delta^{Cohen,\bullet}_{(\cdot,W)}}$ value, the higher is the gain for an adversarial actor, and the more attractive is the combination. The most attractive combination for an adversarial actor would be a low value of $\mu_{ptr^{\bullet}_{(\cdot,W)}}$ against a high value of $\mu_{\Delta^{Cohen,\bullet}_{(\cdot,W)}}$. Based on this directive and on Table 5, the attractiveness $A_{\bullet,W}$ of a scenario (option) in a time window is defined as:

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

$$A_{\bullet,W} := \begin{cases} 0.0 & \text{if } \bullet = 2, 3, \text{ or } 2{+}3, \\ & \mu_{\Delta^{Cohen,\bullet}_{(X,W)}} = 0 \text{ and} \\ & \mu_{ptr^{\bullet}_{(X,W)}} = 0, \\ \frac{\mu_{\Delta^{Cohen,\bullet}_{(X,W)}}}{\mu_{ptr^{\bullet}_{(X,W)}} + \epsilon} & \text{if } \bullet = 2, 3, \text{ or } 2{+}3 \\ & \text{otherwise,} \\ 0.0 & \text{if } \bullet = 1, \\ & \mu_{\Delta^{Cohen,\bullet}_{(Y,W)}} = 0 \text{ and} \\ & \mu_{ptr^{\bullet}_{(Y,W)}} = 0, \\ \frac{\mu_{\Delta^{Cohen,\bullet}_{(Y,W)}}}{\mu_{ptr^{\bullet}_{(Y,W)}} + \epsilon} & \text{if } \bullet = 1 \text{ otherwise,} \\ 0.0 & \text{if } \bullet = 1{+}2, \text{ or } 1{+}3, \\ & \mu_{\Delta^{Cohen,\bullet}_{(X,W)}} = 0, \\ & \mu_{ptr^{\bullet}_{(X,W)}} = 0, \\ & \mu_{\Delta^{Cohen,\bullet}_{(Y,W)}} = 0 \text{ and} \\ & \mu_{ptr^{\bullet}_{(Y,W)}} = 0, \\ \frac{\mu_{\Delta^{Cohen,\bullet}_{(X,W)}} + \mu_{\Delta^{Cohen,\bullet}_{(Y,W)}}}{\mu_{ptr^{\bullet}_{(X,W)}} + \mu_{ptr^{\bullet}_{(Y,W)}} + \epsilon} & \text{if } \bullet = 1{+}2, \text{ or } 1{+}3 \\ & \text{otherwise,} \end{cases} \tag{15}$$

where $\mu_{ptr^{\bullet}_{(X,W)}} = 0$ if $\mu_{\Delta^{Cohen,\bullet}_{(X,W)}} = 0$, $\mu_{ptr^{\bullet}_{(Y,W)}} = 0$ if $\mu_{\Delta^{Cohen,\bullet}_{(Y,W)}} = 0$, and $\epsilon \in \mathbb{R}$ is a small factor to prevent dividing by zero. The higher the value $A_{\bullet,W}$, the more attractive the scenario (option)/time window combination is for an adversarial actor, and the higher will be the risk for a patent requesting company. Though not shown in the notation, the attractiveness value $A_{\bullet,W}$ is parameterized by the parameter pair $(\beta^d, U^d(1,R))$, because the statistical quantities on the right side of Equation (15) are parameterized by this parameter pair (see Subsection 3.6).

3.8. Procedure for making patent application decision–making outcomes on average less vulnerable to negative adversarial influencing

As stated before, the values of the parameters β^d and $U^d(1,R)$ of model M_1, defined in Equation (6), remain to be specified. Instead of choosing some pair of parameter values within the specified bounds, a procedure is provided to determine the most favorable parameter values in a time window with regard to the six adversarial influencing scenarios. By considering the two parameters as variables and by formulating the attractiveness objective function defined in Equation (15) for each of the six considered influencing scenarios, the multi–objective optimization problem stated below is used to determine a set of pairs of optimal parameter values. On this set, a refinement procedure will be applied to determine the most favorable parameter pair $(\beta^{d,*}_W, U^{d,*}_W(1,R))$ for a time window W.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

Model M_2 :

Multi–objective optimization problem:

$$\begin{aligned}\textbf{min}\Big[&A_{\bullet=1,W}(\beta^d_W, U^d_W(1,R)),\ A_{\bullet=2,W}(\beta^d_W, U^d_W(1,R)),\ A_{\bullet=3,W}(\beta^d_W, U^d_W(1,R)),\\ &A_{\bullet=2+3,W}(\beta^d_W, U^d_W(1,R)),\ A_{\bullet=1+2,W}(\beta^d_W, U^d_W(1,R)),\ A_{\bullet=1+3,W}(\beta^d_W, U^d_W(1,R))\Big]\end{aligned}$$

$$\textbf{s.t.}\quad 0.30 < \beta^d_W < 0.60 \text{ and } 1.1 < U^d_W(1,R) < 5.0, \tag{16}$$

over

N_{sim} simulation runs, with N_s sub–runs per simulation run, according to the test setup shown in Figure 9 above,

where

$W \in \{1,2,3\}$, the objective functions $A_{\bullet=\cdot,W}(\beta^d_W, U^d_W(1,R))$ are computed according to Equation (15), $\epsilon = 0.0001$ and where the

output

is a finite set of favorable parameter pairs $\{(\beta^{(i),d}_W, U^{(i),d}_W(1,R))\}_{i=1}^{N_{\text{pairs}}}$, $1 \leq N_{\text{pairs}} \in \mathbb{N}$.

Selection procedure :

A selection procedure will be applied on the output set, in order to arrive at a single most favorable pair of optimal parameters $(\beta^{d,*}_W, U^{d,*}_W(1,R))$ for the time window.

For each time window value W, the evolutionary optimization method NSGA–II [11,12] will be applied to Equation (16), with a population size of $N_{pop} = 50$ and $N_{gen} = 30$ generations as a termination criterium. This results in a set of $1 \leq N_{\text{pairs}} \leq N_{pop}$ favorable parameter pairs $(\beta^{(i),d}_W, U^{(i),d}_W(1,R))$, from which a single most favorable parameter value pair $(\beta^{d,*}_W, U^{d,*}_W(1,R))$ is selected. It is expected that implementing the latter parameter pair in model M_1 for a time window, will discourage adversaries from deploying one of the six considered scenarios. Especially, because it is hard for them to not only guess the underlying mathematical decision support model a patent applying company is using, but also the most favorable parameter pair the company has selected.

4. Discussion

In the literature, a lot of attention is drawn to adversaries trying to explore vulnerabilities of IT systems that are supporting crucial business processes or infrastructure, and how to detect attempts to manipulate such systems. Considerably less attention is drawn to adversaries trying to manipulate the decision outcomes of repeated decision–making processes with underlying parameterized decision support models. And no serious attention at all is drawn to incorporating simulated statistics of repeated decision outcomes affected by a set of well–defined possible influencing scenarios into the parametrization of mathematical decision support models. The purpose of this study is to draw attention to this deficiency, and set a stage for the topic by means of the proposed general simulation framework. The decision support model underlying the patent application decision–making process serves as an example, because of its interesting structure: a non–Bayesian bounded rational *action–reward* model with two successive binary decisions. Most mathematical decision support models have some parameters that remain to be specified, and usually

an optimization problem is formulated to find the optimal parameter values with regard to some general objective function or loss function. A crucial attribution of the proposed simulation framework is that it provides a general definition of a measure that is feeded by simulated statistical test outcomes and that expresses the attractiveness of a defined influencing scenario (from the perspective of an adversary), in terms of the decision support model parameters that remain to be specified. The present study has demonstrated that by considering this measure as an objective function, a multi–objective optimization problem can be formulated for a set of well–defined adversarial influencing scenarios. And that solving the optimization problem for a chosen time window, and applying some selection procedure on its solution set, will provide the *most favorable* (for adversaries hard to guess) support model parameter values for the time window. Parameterizing the decision support model according to these parameter values, will *on average* make the considered set of influencing scenarios less attractive for adversaries to deploy in the chosen time window.

Of course, company security analysts cannot be accounted for preventing adversaries from crafting and deploying adversarial influencing scenarios to manipulate decision outcomes of repeated decision–making processes they are supposed to protect. However, they can be accounted for taking countermeasures, such as implementing the proposed approach, that will make such scenarios on average less effective and that on average will raise the chance that adversarial influencing of decision outcomes will be detected. Once adversaries suspect that company security analysts themselves craft and simulate influencing scenarios to make them less effective and that this may raise the chance of being exposed, this may discourage them from crafting and deploying such scenarios in the future.

The statistical theory underlying the presented mathematical model M_2 needs to be further developed. Company security analysts should be given stricter guarantees than that the effect of considered scenarios on average will be less and the detection chance of a deployed scenario will on average be higher. But, this is left to future research. The approach presented in this study is general and can be applied to a variety of repeated decision–making processes and underlying mathematical decision support models. For instance, to decision support of repeated decision–making by means of machine learning models, in which case the presented approach needs to be included into the hyperparameter tuning of the used machine learning model. In forthcoming work, a case study of this will be presented. As well as a case study of detection of adversarial influencing of repeated decision outcomes of a repeated decision–making process supported by a bounded rational Bayesian decision support model.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Acknowledgments: I gratefully thanks my employer, the department DMO of the Dutch Ministry of Defense, for granting and supporting my position of part time external doctoral candidate. Furthermore, my warm thanks go to Marie–Colette van Lieshout for the time she spent to ensure the mathematical rigor and quality of this paper. Also, I acknowledge Rob van der Mei and Sandjai Bhulai for their valuable comments and suggestions to improve the readability of this manuscript.

Conflicts of Interest: The author declares no conflict of interest. The department DMO of the Dutch Ministry of Defense had no role in the design and execution of the study, in the writing of the manuscript, or in the decision to publish the results.

Appendix A Derivation of the optimal posterior decision policy for the action before observation probabilistic scenario

This appendix describes how to derive the expression for the optimal posterior decision policy (i.e., the distribution $p^*_{\mathcal{A}}(a)$) in the theory of Ortega et al. for the probabilistic scenario in which $P_{\mathcal{X}}(x) = P(o,a) = P(o|a)P_{\mathcal{A}}(a)$. In this scenario, there is interaction

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

with the environment in that the choice action the decision maker takes according to the optimized $P_{\mathcal{A}}(a)$ has a stochastic effect on the environment according to the distribution $P(o|a)$.

$$\begin{aligned}\mathcal{X} &= \mathcal{A} \times \mathcal{O}, \\ &= \underbrace{\{a_1, \cdots, a_M\}}_{\text{Actions}} \times \underbrace{\{o_1, \cdots, o_N\}}_{\text{Observations}},\end{aligned}$$

$$\begin{aligned}P_{\mathcal{X}}(x) &= P_{\mathcal{X}}(a,o) = P(o|a)P_{\mathcal{A}}(a), \\ Q_{\mathcal{X}}(x) &= P(o|a)P_{0,\mathcal{A}}(a),\end{aligned}$$

$$\begin{aligned}-\Delta F_{\beta}[P] &:= \underbrace{\sum_{x\in\mathcal{X}} P_{\mathcal{X}}(x)U(x)}_{\text{Expected utility}} - \underbrace{\frac{1}{\beta}\sum_{x\in\mathcal{X}} P_{\mathcal{X}}(x)\log\frac{P(x)}{P_{0,\mathcal{X}}(x)}}_{\text{Information cost}} \\ &= \sum_{x\in\mathcal{X}}\{P_{\mathcal{X}}(x)U(x) - \frac{1}{\beta}P_{\mathcal{X}}(x)\log\frac{P_{\mathcal{X}}(x)}{P_{0,\mathcal{X}}(x)}\},\end{aligned}$$

where $P_{0,\mathcal{X}}(x)$ and $P_{\mathcal{X}}(x)$ represent the prior decision policy and the posterior decision policy with respect to the space $\mathcal{X}$, respectively. Optimizing the above objective function is equivalent to optimizing the objective function

$$\sum_{a\in\mathcal{A}}\sum_{o\in\mathcal{O}}\left\{P(o|a)P_{\mathcal{A}}(a)U(o,a) - \frac{1}{\beta}P(o|a)P_{\mathcal{A}}(a)\log\frac{P_{\mathcal{A}}(a)}{P_{0,\mathcal{A}}(a)}\right\} + \lambda\left\{\sum_{a\in\mathcal{A}}\sum_{o\in\mathcal{O}}P_{\mathcal{X}}(o,a) - 1\right\},$$

in terms of actions

$$P_{\mathcal{A}}(a) = \sum_{o\in\mathcal{O}} P_{\mathcal{X}}(o,a).$$

Take the derivative with respect to $P_{\mathcal{A}}(a)$ for fixed $a \in \mathcal{A}$:

$$\begin{aligned}&\frac{d}{dP_{\mathcal{A}}(a)}\left[\sum_{a\in\mathcal{A}}\sum_{o\in\mathcal{O}}\left\{P(o|a)P_{\mathcal{A}}(a)U(o,a) - \frac{1}{\beta}P(o|a)P_{\mathcal{A}}(a)\log\frac{P_{\mathcal{A}}(a)}{P_{0,\mathcal{A}}(a)}\right\}\right] \\ &+\frac{d}{dP_{\mathcal{A}}(a)}\left[\lambda\sum_{a\in\mathcal{A}}\sum_{o\in\mathcal{O}}P_{\mathcal{X}}(o,a) - \lambda\right] = 0 \\ &\sum_{o\in\mathcal{O}}\left[P(o|a)U(o,a) - \frac{1}{\beta}P(o|a)\log\frac{P_{\mathcal{A}}(a)}{P_{0,\mathcal{A}}(a)} - \frac{1}{\beta}P(o|a)\tilde{P}_{\mathcal{A}}(a)\frac{1}{P_{\mathcal{A}}(a)}\right] + \lambda = 0 \\ &\underbrace{\sum_{o\in\mathcal{O}}[P(o|a)U(o,a)]}_{\mathbb{E}[U|a]} - \frac{1}{\beta}\log\frac{P_{\mathcal{A}}(a)}{P_{0,\mathcal{A}}(a)} - \frac{1}{\beta} + \lambda = 0.\end{aligned}$$

Take the derivative with respect to λ:

$$\sum_{a\in\mathcal{A}} P_{\mathcal{A}}(a) = 1. \tag{A1}$$

This yields

$$\beta\mathbb{E}[U|a] - \log\frac{P_{\mathcal{A}}(a)}{P_{0,\mathcal{A}}(a)} = 1 - \beta\lambda$$

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 June 2021 doi:10.20944/preprints202106.0676.v1

$$P_{\mathcal{A}}(a) = P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}e^{\beta\lambda-1}.$$

Using equation (A1):

$$\sum_{a\in\mathcal{A}} P_{\mathcal{A}}(a) = \sum_{a\in\mathcal{A}} P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}e^{\beta\lambda-1} = 1$$
$$\sum_{a\in\mathcal{A}} P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]} = e^{1-\beta\lambda}.$$

This finally yields the optimal decision policy over the finite action space $\mathcal{A}$:

$$\begin{aligned} P^*_{\beta,\mathcal{A}}(a) &= \frac{P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}}{\sum_{a\in\mathcal{A}} P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}} \\ &= \frac{P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}}{Z_{\beta,\mathcal{A}}}, \end{aligned} \tag{A2}$$

with $\mathbb{E}[U|a] = \sum_{o\in\mathcal{O}}[P(o|a)U(o,a)]$ and $Z_{\beta,\mathcal{A}} = \sum_{a\in\mathcal{A}} P_{0,\mathcal{A}}(a)e^{\beta\mathbb{E}[U|a]}$.

References

1. Dalvi, N.; et al. Adversarial Classification. Report KDD'04, Seattle (USA), Department of Computer Science and Engineering, University of Washington, August 22–25, **2004**.
2. Banks, D.L.; Rios Aliaga, J.M. Adversarial Risk Analysis. Taylor and Francis Inc: Cambridge (MA), 1st ed., December **2015**.
3. World International Patent Organization, General information on patents. Available online: https://www.wipo.int (accessed on 1st April **2005**).
4. Aristodemou, L.; Tietze, F. The state–of–the–art on intellectual property analytics (IPA): A literature review on artifical intelligence, machine learning and deep learning methods for analysing intellectual property (IP) data. *World Patent Information* **2018**, *55*, 37-51.
5. Simon, H. The sciences of the artifical. MIT Press: Cambridge (MA), **1969**.
6. Ortega, P.A. A unified framework for resource-bounded autonomous agents interacting with unknown environments. PhD thesis, University of Cambridge, Cambridge, **2011**.
7. Ortega, P.A.; Braun, D.A. Thermodynamics as a theory of decision making with information-processing costs. Proceedings of the Royal Society A: Mathematical and Physical Engineering Sciences, **2013**; 469 (2153), 20210683.
8. McNemar, Q. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* **1947**, *12*, (*55*), 153–157.
9. Cohen, J. Statistical power analysis for the behavorial science. Routledge, ISBN: 1-134-74270-3, **1988**.
10. Ortega, P.A. An adversarial interpretation of information-theoretic bounded rationality. Proceedings of the Twenty-Eighth (AAAI) Conference on Artifical Intelligence, **2014**; 2483–2489.
11. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA–II. *IEEE Trans Evol Comput* **2002**, *6*, (*12*), 182–197.
12. Deb, K.; Jain, H. An evolutionary many–objective optimization algorithm using reference–point based non–dominated sorting approach, part (I): solving problems with box contraints. *IEEE Trans Evol Comput* **2014**, *18*, (*4*), 577–601.