OPEN ACCESS
Vietnam Journal of Computer Science
Vol. 7, No. 4 (2020) 453–470
© The Author(s)
DOI: 10.1142/S2196888820500256

World Scientific
www.worldscientific.com

Refining Indeterministic Choice: Imprecise Probabilities and Strategic Thinking

Jorge Castro*, Joaquim Gabarro† and Maria Serna‡

Computer Science Department
Universitat Politècnica de Catalunya
Barcelona 08034, Spain
**castro@cs.upc.edu*
†gabarro@cs.upc.edu
‡mjserna@cs.upc.edu

Received 28 January 2020
Accepted 8 May 2020
Published 6 August 2020

Often, uncertainty is present in processes that are part of our routines. Having tools to understand the consequences of unpredictability is convenient. We introduce a general framework to deal with uncertainty in the realm of distribution sets that are descriptions of imprecise probabilities. We propose several non-biased refinement strategies to obtain sensible forecasts about results of uncertain processes. Initially, uncertainty on a system is modeled as the non-deterministic choice of its possible behaviors. Our refinement hypothesis translates non-determinism into imprecise probabilistic choices. Imprecise probabilities allow us to propose a notion of uncertainty refinement in terms of set inclusions. Later on, unpredictability is tackled through a strategic approach using uncertainty profiles and angel/daemon games ($\mathfrak{a}/\mathfrak{d}$-games). Here, imprecise probabilities form the set of mixed strategies and Nash equilibria corresponds to natural uncertainty refinements. We use this approach to study the performance of Web applications — in terms of response times — under stress conditions.

Keywords: Imprecise and uncertain reasoning; non-deterministic choice; refinement; $\mathfrak{a}/\mathfrak{d}$-games; Web orchestrations.

1. Introduction

Words like uncertainty, imprecision or non-determinism express several facets of a lack of knowledge. They are commonly used in many areas of human reasoning. In an economical context an approach to *uncertainty*, as a lack of probabilistic information, was studied by Knight.[1] In a mathematical setting, the term *imprecision*,

† Corresponding author.

reshaped as *imprecise probabilities* is considered by Cooman *et al.*[2] Non-deterministic choices were introduced by Tony Hoare[3] to describe the semantics of concurrent process when the agents have no full control on the decision making.

In this paper, we provide a general framework to deal with uncertainty. We model full uncertainty as a non-deterministic choice on the possible behaviors. Inspired by McIver and Morgan,[4] we state the *refinement hypothesis* that relates non-determinism to imprecise probabilistic choices. The refinement hypothesis claims that when non-determinism is replaced by a totally arbitrary probabilistic choice the system uncertainty does not increase. Once this first refinement step is done and depending on our system knowledge, additional steps to lower the unpredictability can be considered by restricting arbitrariness in probabilistic elections.

In the area of social sciences, strategic approaches are a commonly used tools in the fight against uncertainty, see for instance Porter work.[5] Gabarró *et al.*[6,7] extend strategic thinking to deal with Web apps and short time macroeconomics models. They introduce the *uncertainty profile* framework that sets up the components, the potential modifications, and the intensity of the actions of two agents, the angel and the daemon ($\mathfrak{a}/\mathfrak{d}$), perturbing a system. Both agents represent the *animal spirits*[8] modeling antagonistic interests rather than "a weighted average of benefits".[9] In the context of Web apps, a first link with imprecise probabilities was provided in Ref. 10. We extend here this approach to general systems. Uncertainty profiles allow us to consider strategic settings through angel/daemon-games.[6] In this strategic condition, it is assumed that the angel modifies the systems in such a way that the damage to the system is minimized, while the daemon tries to make that the system works at its worst. Nash equilibria of the $\mathfrak{a}/\mathfrak{d}$-games define imprecise probabilities that can drastically shrink the system uncertainty.

We apply the framework to Web applications defined as orchestrations of basic Web services or sites. An orchestration compounds the results of site calls to perform some computation. We assume that orchestrations are defined using the Orc language.[11,12] We note that the language itself has operators providing uncertain results. For instance, the *asymmetric parallelism* operation in Orc, that waits until a result arrives and then triggers the execution of another part of the orchestration is inherently uncertain. In an unreliable Web additional uncertainty sources may emerge and orchestrations have to manage both inherent uncertainty and uncertainty on site calls. In this scenario, semantic characterizations through imprecise probabilities were proposed in Ref. 10. We ground on these semantic representations to analyze the response delay time of an orchestration. We provide information about orchestration response time assuming different knowledge levels about perturbations.

The paper is structured as follows. In Sec. 2, we present our general approach to deal with uncertainty in a step by step presentation. In Sec. 3, we consider Web apps. Finally, in Sec. 4, we post our conclusions and raise some open question.

2. Refining Non-determinism by Imprecise Probabilities

We introduce a framework to deal with uncertainty consisting of different unsureness levels. In this framework, non-deterministic systems (or processes) are considered at the highest unpredictability level. Random processes described by imprecise probabilities are placed at lower uncertainty levels. Finally, probabilistic systems, defined by distributions with known weights, are at the lowest unpredictability levels. We also introduce several unbiased criteria to improve the predictability of a process. The refinement hypothesis and the strategic thinking tools explained below will play a role when reshaping a system into a more predictable one.

First, we introduce some notation.

- $\sqcap$ is the non-deterministic choice operator introduced by Tony Hoare.[3] If P and Q are processes, the expression $P \sqcap Q$ denotes a process which arbitrarily behaves either like P or Q.
- Δ_k, for $k > 1$, denotes the probability set

$$\Delta_k = \{(p_1, \ldots, p_k) | 0 \leq p_i \quad \text{for } 0 \leq i \leq k \quad \text{and} \quad p_1 + \cdots + p_k = 1\}.$$

- $\|$ is the *choice operator*. We use it to represent a *probabilistic choice*. The notation[4] $\mathbb{S}_1 @ p_1 \| \mathbb{S}_2 @ p_2$, for $(p_1, p_2) \in \Delta_2$ denotes a system $\mathbb{S}$ which behaves as system $\mathbb{S}_1$ with probability p_1 and as $\mathbb{S}_2$ with probability p_2.
- Imprecise probabilistic systems. We extend the previous notation to represent imprecise probabilistic choice systems. For instance,

$$\{\mathbb{S}_1 @ p_1 \| \mathbb{S}_2 @ p_2 | (p_1, p_2) \in \Delta_2\}$$

  denotes an imprecise probabilistic choice system, containing a probabilistic choice for each $(p_1, p_2) \in \Delta_2$.
- $\sqsubseteq$ denotes *refinement*.[4] In our setting $\mathbb{S} \sqsubseteq \mathbb{S}'$ means that all the outputs of system $\mathbb{S}'$ are valid results of $\mathbb{S}$, but $\mathbb{S}'$ is at least as predictable as $\mathbb{S}$. For instance, a system consisting of throwing n times an unbiased coin is a refinement of throwing n times a coin with an imprecise bias in a small range around zero.

2.1. *Refinement hypothesis*

Let us develop the interplay between non-deterministic choice and imprecise probabilities through an algebraic approach. Let us start with the simplest system $\mathbb{S}$ consisting of the non-deterministic choice among two subsystems $\mathbb{S}_1$ and $\mathbb{S}_2$. We express this as

$$\mathbb{S} = \mathbb{S}_1 \sqcap \mathbb{S}_2.$$

In words, system $\mathbb{S}$ behaves arbitrarily either as system $\mathbb{S}_1$ or $\mathbb{S}_2$.

Following McIver and Morgan,[4] we associate to such a pure indeterministic system a first refinement in terms of imprecise probabilities. This refinement is a set of probabilistic behaviors, one behavior for each possible probability distribution. We use the symbol $\oplus_{\mathsf{IP}}$ to denote the corresponding universe of imprecise probabilities.

Formally, we write

$$\mathbb{S}_1 \oplus_{\mathsf{IP}} \mathbb{S}_2 = \{\mathbb{S}_1 @ p_1 \| \mathbb{S}_2 @ p_2 | (p_1, p_2) \in \Delta_2\}$$

and we claim that $\mathbb{S}_1 \oplus_{\mathsf{IP}} \mathbb{S}_2$ provides a refinement on the actual process, so we write

$$\mathbb{S} \sqsubseteq \mathbb{S}_1 \oplus_{\mathsf{IP}} \mathbb{S}_2.$$

The intuitive explanation is the following: as $\mathbb{S}$ performs an arbitrary choice between $\mathbb{S}_1$ and $\mathbb{S}_2$, the uncertainty does not increase when the choice is done according to an unknown probability distribution.

In the most general case of a system defined as a non-deterministic choice among n possible systems, arguing similarly, we extend the former refinement to

$$\mathbb{S}_1 \sqcap \cdots \sqcap \mathbb{S}_n \sqsubseteq \mathbb{S}_1 \oplus_{\mathsf{IP}} \cdots \oplus_{\mathsf{IP}} \mathbb{S}_n,$$

where

$$\mathbb{S}_1 \oplus_{\mathsf{IP}} \cdots \oplus_{\mathsf{IP}} \mathbb{S}_n$$

denotes the set of all probabilistic choices between $\mathbb{S}_1, \ldots, \mathbb{S}_n$, that is

$$\{\mathbb{S}_1 @ p_1 \| \cdots \| \mathbb{S}_n @ p_n | (p_1, \ldots, p_n) \in \Delta_n\}.$$

Our *refinement hypothesis* states that *encoding non-determinism as imprecise probabilities does not increase the uncertainty.*

2.2. *Irreversibility of the refinement hypothesis*

We show in this section the effects of refinements on a system analysis. When studying a system behavior an its predictability, applying the refinement hypothesis to non-deterministic choices is a non-reversible step. For instance, assume that X_1, X_2, Y_1, Y_2 are four processes each one computing an integer and system $\mathbb{S}$ computes the maximum of two non-deterministic choices $X_1 \sqcap X_2$ and $Y_1 \sqcap Y_2$. We write $\mathbb{S} = \max(X_1 \sqcap X_2, Y_1 \sqcap Y_2)$. First, we analyze the system $\mathbb{S}$ allowing non-deterministic choices to evolve according to the CSP[3] distributive law that states the equality $f(X \sqcap Y) = f(X) \sqcap f(Y)$ for a function *f*. In our case

$$\mathbb{S} = \max(X_1, Y_1) \sqcap \max(X_1, Y_2) \sqcap \max(X_2, Y_1) \sqcap \max(X_2, Y_2).$$

Finally, refining non-deterministic choices, it holds

$$\mathbb{S} \sqsubseteq \max(X_1, Y_1) \oplus_{\mathsf{IP}} \max(X_1, Y_2) \oplus_{\mathsf{IP}} \max(X_2, Y_1) \oplus_{\mathsf{IP}} \max(X_2, Y_2). \quad (1)$$

In contrast, if the refinement hypothesis is applied on both non-deterministic choices, $X_1 \sqcap X_2$ and $Y_1 \sqcap Y_2$, before applying distributivity, we get

$$\mathbb{S} \sqsubseteq \max(X_1 \oplus_{\mathsf{IP}} X_2, Y_1 \oplus_{\mathsf{IP}} Y_2).$$

In such a case, writing $X_1 \oplus_{\mathsf{IP}} X_2 = \{X_1 @ p_1 \| X_2 @ p_2 | (p_1, p_2) \in \Delta_2\}$ and $Y_1 \oplus_{\mathsf{IP}} Y_2 = \{Y_1 @ q_1 \| Y_2 @ q_2 | (q_1, q_2) \in \Delta_2\}$. As the pair (X_1, Y_1) is chosen with probability $p_1 q_1$, the $\max(X_1, Y_1)$ is also chosen with this probability. Similarly in the other cases.

Therefore, $\max(X_1 \oplus_{\mathsf{IP}} X_2, Y_1 \oplus_{\mathsf{IP}} Y_2)$ can be rewritten as,

$$\max(X_1, Y_1)@p_1q_1 \,\|\, \max(X_1, Y_2)@p_1q_2 \,\|\, \max(X_2, Y_1)@p_2q_1 \,\|\, \max(X_2, Y_2)@p_2q_2,$$

where (p_1, p_2) and (q_1, q_2) are arbitrary distributions in Δ_2. The point is that this refinement of system $\mathbb{S}$ is different from the previous one. In fact, probability distributions in the former expression are a subset of the total probability space Δ_4 involved in refinement (1). It can be proved that distributions (w_1, w_2, w_3, w_4) that can be expressed as a pair in $\Delta_2 \times \Delta_2$ must hold $w_1w_4 = w_2w_3$, an equality that puts a strong constrain on Δ_4. For instance, distribution $(1/2, 0, 0, 1/2)$ cannot be written as $(p_1q_1, p_1q_2, p_2q_1, p_2q_2)$.

The former paragraph discussion shows that refinement steps are non-return points of the predictability analysis. It seems to reinforce the following claim: *dealing with uncertainty as soon as possible*, *is the best way to try to decrease the global unpredictability.*

Let us consider a detailed example of refinement steps motivated by Web apps. We use an Orc-inspired notation.[11] In Orc, given two Web sites $\mathbb{S}$ and $\mathbb{T}$, the symmetric parallelism $\mathbb{S}|\mathbb{T}$ calls both sites at the same time and returns an interleaving of the results delivered by both sites.

Example 2.1. Consider $\mathbb{C}razyChannel$, a TV channel offering two consecutive films, one of $\mathbb{A}ction$ and another one which is $\mathbb{M}onarchy$ related. This can be expressed in Orc as

$$\mathbb{C}razyChannel = (\mathbb{A}ction|\mathbb{M}onarchy).$$

The channel has the inconvenience of being non-deterministic, i.e. we do not known in advance which films will be broadcast. Assume that we know that the action film is drawn either from the Rambo or the Superman series and that the monarchy one is taken either from the Sissi or The Crown series. Formally,

$$\begin{aligned} &\mathbb{A}ction = \mathbb{R}ambo \sqcap \mathbb{S}uperman,\\ &\mathbb{M}onarchy = \mathbb{S}issi \sqcap \mathbb{T}heCrown. \end{aligned}$$

As the parallel composition produces an interleaving of the two subsystems the channel offers two films, one in each category. As a consequence of the $\sqcap$-distributive law, what couple of films will be broadcast is the result of a non-deterministic choice

$$\begin{aligned} \mathbb{C}razyChannel &= (\mathbb{A}ction|\mathbb{M}onarchy)\\ &= (\mathbb{R}ambo|\mathbb{S}issi) \sqcap (\mathbb{R}ambo|\mathbb{T}heCrown)\\ &\quad\ \sqcap\ (\mathbb{S}uperman|\mathbb{S}issi) \sqcap (\mathbb{S}uperman|\mathbb{T}heCrown). \end{aligned} \tag{2}$$

Finally, we can refine the channel in terms of imprecise probabilities

$$\begin{aligned} \mathbb{C}razyChannel \sqsubseteq \{&(\mathbb{R}ambo|\mathbb{S}issi)@r_1 \| (\mathbb{R}ambo|\mathbb{T}heCrown)@r_2\\ &\| (\mathbb{S}uperman|\mathbb{S}issi)@r_3 \| (\mathbb{S}uperman|\mathbb{T}heCrown)@r_4\\ &|(r_1, r_2, r_3, r_4) \in \Delta_4\}. \end{aligned}$$

A better refinement — or a less imprecise — can be inferred if we know that both non-deterministic choices have been refined into imprecise probabilities internally at, respectively, $\mathbb{A}ction$ and $\mathbb{M}onarchy$ processes. In this case,

$$\begin{aligned}\mathbb{C}razyChannel &= (\mathbb{A}ction|\mathbb{M}onarchy)\\ &\sqsubseteq ((\mathbb{R}ambo \oplus_{\mathsf{IP}} \mathbb{S}uperman)|(\mathbb{S}issi \oplus_{\mathsf{IP}} \mathbb{T}heCrown))\\ &= \{(\mathbb{R}ambo|\mathbb{S}issi)@p_1q_1 \| (\mathbb{R}ambo|\mathbb{T}heCrown)@p_1q_2\\ &\quad \| (\mathbb{S}uperman|\mathbb{S}issi)@p_2q_1 \| (\mathbb{S}uperman|\mathbb{T}heCrown)@p_2q_2\\ &\quad |(p_1, p_2) \in \Delta_2, (q_1, q_2) \in \Delta_2)\}.\end{aligned}$$

2.3. *Uncertainty profiles*

Another way to ascertain information about a system is through uncertainty profiles.[10] Given a system $\mathbb{S}(X_1, \ldots, X_n) = \mathbb{S}_1 \sqcap \cdots \sqcap \mathbb{S}_k$, let us study its behavior under angel–daemon stress. We use below symbol $[n]$ to denote the set of integers $\{1, \ldots, n\}$.

Definition 2.1. Given a system $\mathbb{S}(X_1, \ldots, X_n)$, an $\mathfrak{a}/\mathfrak{d}$ *stress* $\mathcal{S}$ is formed by a collection of processes $(X_{i,\mathfrak{a}}, X_{i,\mathfrak{d}}, X_{i,\mathfrak{a},\mathfrak{d}})_{i\in[n]}$. The behavior of component X_i under angelic stress is denoted by $X_{i,\mathfrak{a}}$. Respectively, $X_{i,\mathfrak{d}}$ and $X_{i,\mathfrak{a},\mathfrak{d}}$ denote the behavior under daemonic stress or both angelic and daemonic stress, respectively.

For a set $a \subseteq [n]$ of parameters that suffer *angelic stress* and a set $b \subseteq [n]$ of parameters under daemonic stress, we define the processes $X_i[a, d]$ as follows:

$$X_i[a,d] = \begin{cases} X_i & \text{if } i \notin a \cup d,\\ X_{i,\mathfrak{a}} & \text{if } i \in a \backslash d,\\ X_{i,\mathfrak{d}} & \text{if } i \in d \backslash a,\\ X_{i,\mathfrak{a},\mathfrak{d}} & \text{if } i \in a \cap d.\end{cases}$$

The system $\mathbb{S}$ under *stress action* (a, d) is $\mathbb{S}[a, d] = \mathbb{S}(X_1[a, d], \ldots, X_n[a, d])$.

Observe that a stress action describes a global change on the system behavior. Let us consider a first example of a system under a $\mathfrak{a}/\mathfrak{d}$-stress.

Example 2.2. Consider the system $\mathbb{S}(X_1, X_2) = X_1X_2$ where X_1 and X_2 are processes that read integer data from an associated device. Both components truly report the readings under normal working conditions. Assume a stress $\mathcal{S}$ that on component X_1 behaves as $X_{1,\mathfrak{a}} = 2X_1$, $X_{1,\mathfrak{d}} = \frac{1}{2}X_1$, $X_{1,\mathfrak{a},\mathfrak{d}} = X_1$ and on the second component the action is $X_{2,\mathfrak{a}} = \frac{2}{3}X_2$, $X_{2,\mathfrak{d}} = \frac{1}{4}X_2$ and $X_{2,\mathfrak{a},\mathfrak{d}} = \frac{3}{4}X_2$. Let (a, d) be the stress action $(\{1\}, \{1, 2\})$. As $1 \in a \cap d$, $X_1[a, d] = X_{1,\mathfrak{a},\mathfrak{d}} = X_1$. As X_2 is just only under $\mathfrak{d}$ control, we have $X_2[a, d] = X_{2,\mathfrak{d}} = \frac{1}{4}X_2$. Then $\mathbb{S}[a, d] = \mathbb{S}(X_1[a, d], X_2[a, d]) = X_{1,\mathfrak{a},\mathfrak{d}}X_{2,\mathfrak{d}} = X_1\frac{1}{4}X_2 = \frac{1}{4}X_1X_2$. When $(a, d) = (\{2\}, \{1\})$, we have $\mathbb{S}[a, d] = X_{1,\mathfrak{d}}X_{2,\mathfrak{a}} = \frac{1}{2}X_1\frac{2}{3}X_2 = \frac{1}{3}X_1X_2$.

Suppose that we know that $\mathbb{S}$ will suffer stress, but the extent and location of the perturbations are uncertain. Following Castro *et al.*,[10] we can describe this knowledge by an uncertainty profile.

Definition 2.2. An *uncertainty profile* is a tuple $\mathcal{U} = \langle \mathbb{S}(X_1, \ldots, X_n), \mathcal{S}, \mathcal{A}, \mathcal{D}, b_{\mathfrak{a}}, b_{\mathfrak{d}}, u\rangle$, where $\mathcal{S}$ is a stress model, $\mathcal{A} \subseteq [n]$ points out the parameters that can suffer angelic stress, $\mathcal{D} \subseteq [n]$ the ones that can suffer daemonic stress, $b_{\mathfrak{a}}, b_{\mathfrak{d}} \in \mathbb{N}$, and $u : 2^{\mathcal{A}} \times 2^{\mathcal{D}} \to \mathbb{N}$ is the utility — or value — function.

Agents $\mathfrak{a}$ and $\mathfrak{d}$ have the *capability to act*, respectively, on the given subsets of parameters. The integers $b_{\mathfrak{a}}$ and $b_{\mathfrak{d}}$ measure the *intensity* of the real $\mathfrak{a}$ and $\mathfrak{d}$ actions, giving the *number of parameters* that $\mathfrak{a}$ or $\mathfrak{d}$ can really stress. Finally, u is a performance measure on the system $\mathbb{S}([a, d])$. The possible actions to be undertaken by $\mathfrak{a}$ and $\mathfrak{d}$ are given by the sets

$$A_{\mathfrak{a}} = \{a \subseteq \mathcal{A} | \#a = b_{\mathfrak{a}}\}, \quad A_{\mathfrak{d}} = \{d \subseteq \mathcal{D} | \#d = b_{\mathfrak{d}}\}.$$

With no further information about agents $\mathfrak{a}$ and $\mathfrak{d}$, the system can behave as $\mathbb{S}[a, d]$, for any $(a, d) \in A_{\mathfrak{a}} \times A_{\mathfrak{d}}$, and it can be described as the non-deterministic choice among all the possible ways to perturb $\mathbb{S}$ according to the profile, i.e.

$$\mathbb{S}(\mathcal{U}) = \sqcap_{(a,d)\in A_{\mathfrak{a}} \times A_{\mathfrak{d}}} \mathbb{S}[a, d].$$

Let us reconsider the system $\mathbb{S}(X_1, X_2) = X_1 X_2$.

Example 2.3. Let $\mathbb{S}$ be as in Example 2.2. Assume a high lack of knowledge about which components will be perturbed. In such case $\mathcal{A} = \mathcal{D} = \{1, 2\}$. Moreover, suppose we know that when perturbation arrives it will be focused on a single component, so $b_{\mathfrak{a}} = b_{\mathfrak{d}} = 1$, meaning that both $\mathfrak{a}$ and $\mathfrak{d}$ can perturb just one component. Thus, $A_{\mathfrak{a}} \times A_{\mathfrak{d}} = \{(\{1\}, \{1\}), (\{1\}, \{2\}), (\{2\}, \{1\}), (\{2\}, \{2\})\}$, and the system under $\mathfrak{a}/\mathfrak{d}$ perturbation is

$$\begin{aligned}\mathbb{S}(\mathcal{U}) &= \mathbb{S}(\{1\}, \{1\}) \sqcap \mathbb{S}(\{1\}, \{2\}) \sqcap \mathbb{S}(\{2\}, \{1\}) \sqcap \mathbb{S}(\{2\}, \{2\}) \\ &= \mathbb{S}(X_{1,\mathfrak{a},\mathfrak{d}}, X_2) \sqcap \mathbb{S}(X_{1,\mathfrak{a}}, X_{2,\mathfrak{d}}) \sqcap \mathbb{S}(X_{1,\mathfrak{d}}, X_{2,\mathfrak{a}}) \sqcap \mathbb{S}(X_1, X_{2,\mathfrak{a},\mathfrak{d}}).\end{aligned}$$

Tuple $\mathcal{U} = \langle \mathbb{S}(X_1, X_2), \mathcal{S}, \{1, 2\}, \{1, 2\}, 1, 1, u\rangle$ where the utility function can be chosen to be any performance measure of interest.

Choices for $\mathfrak{a}$ and $\mathfrak{d}$ can be defined probabilistically. Mixed strategies for $\mathfrak{a}$ and $\mathfrak{d}$ are, respectively probability distributions $\alpha : A_{\mathfrak{a}} \to [0, 1]$ and $\beta : A_{\mathfrak{d}} \to [0, 1]$. The utility of the mixed strategy (α, β) is defined as $u(\alpha, \beta) = \sum_{(a,d)\in A_{\mathfrak{a}} \times A_{\mathfrak{d}}} \alpha(a) u(a, d) \beta(d)$. To simplify notation, let $\Delta_{\mathfrak{a}} = \Delta_{b_{\mathfrak{a}}}$ and $\Delta_{\mathfrak{d}} = \Delta_{b_{\mathfrak{d}}}$ denote the set of mixed strategies for $\mathfrak{a}$ and $\mathfrak{d}$, respectively. We associate to $\mathcal{U}$ the probabilistic multiset

$$\mathbb{S}(\mathcal{U})_{\mathsf{Mixed}} = \{\Vert_{(a,d)\in A_{\mathfrak{a}} \times A_{\mathfrak{d}}} \mathbb{S}[a, d] @ \alpha(a) \beta(d) | (\alpha, \beta) \in \Delta_{\mathfrak{a}} \times \Delta_{\mathfrak{d}}\}.$$

This set describes, in terms of imprecise probabilities, all the possible behaviors of the system under the stressed environment described by the joint actions of the $\mathfrak{a}/\mathfrak{d}$ agents. By the refinement hypothesis in Sec. 2.1, $\mathbb{S}(\mathcal{U}) \sqsubseteq \mathbb{S}(\mathcal{U})_{\mathsf{Mixed}}$. We can associate the set of expected values for mixed strategies

$$\mathsf{E}(\mathbb{S}(\mathcal{U})_{\mathsf{Mixed}}) = \left\{ \sum_{(a,d)\in A_{\mathfrak{a}}\times A_{\mathfrak{d}}} \alpha(a)u(a,d)\beta(d)) | (\alpha,\beta) \in \Delta_{\mathfrak{a}} \times \Delta_{\mathfrak{d}} \right\}$$
$$= \{u(\alpha,\beta) | (\alpha,\beta) \in \Delta_{\mathfrak{a}} \times \Delta_{\mathfrak{d}}\}.$$

Informally, $\mathsf{E}(\mathbb{S}(\mathcal{U})_{\mathsf{Mixed}})$ describes the set of possible average measurements corresponding to the $\mathfrak{a}/\mathfrak{d}$ choices.

2.4. *Strategic refinement using angel-daemon games*

In many fields and specially in business, strategic approaches are considered to model and restrain unsureness.[5] We incorporate this strategic view into our model through angel–daemon games based on uncertainty profiles. In order to model the strategic interaction between $\mathfrak{a}$ and $\mathfrak{d}$, we transform a *static description* provided by an uncertainty profile $\mathcal{U}$, into a *strategic situation* called $\mathfrak{a}/\mathfrak{d}$-game.

Definition 2.3. To each uncertainty profile $\mathcal{U}$, we associate an $\mathfrak{a}/\mathfrak{d}$-game $\Gamma(\mathcal{U}) = \langle A_{\mathfrak{a}}, A_{\mathfrak{d}}, u(a,d)\rangle$. The players are $\mathfrak{a}$ and $\mathfrak{d}$ and the set of strategy profiles is $A_{\mathfrak{a}} \times A_{\mathfrak{d}}$. With no more information, playing the game consists on choosing a tuple $(a,d) \in A_{\mathfrak{a}} \times A_{\mathfrak{d}}$. Players $\mathfrak{a}$ and $\mathfrak{d}$ are antagonistic; $\mathfrak{a}$ tries to maximize the utility and $\mathfrak{d}$ tries to minimize the value. Formally, $\Gamma(\mathcal{U})$ is a zero-sum game where the utility of $\mathfrak{a}$ is $u_{\mathfrak{a}}(a,d) = u(a,d)$ and the utility of $\mathfrak{d}$ is $u_{\mathfrak{d}}(a,d) = -u(a,d)$.

A pure strategy profile (a,d) is a special case of mixed strategy (α,β) in which $\alpha(a) = 1$ and $\beta(d) = 1$. A mixed strategy profile (α,β) is a *Nash equilibrium* if, for any $\alpha' \in \Delta_{\mathfrak{a}}$, $u(\alpha,\beta) \geq u(\alpha',\beta)$ and, for any $\beta' \in \Delta_{\mathfrak{d}}$, $u(\alpha,\beta) \leq u(\alpha,\beta')$. Let $\mathbb{N}ash$ denote the set of Nash equilibria. We associate to $\Gamma(\mathcal{U})$ the following multiset:

$$\mathbb{S}(\mathcal{U})_{\mathsf{Nash}} = \{\Vert_{(a,d)\in A_{\mathfrak{a}}\times A_{\mathfrak{d}}} \mathbb{S}[a,d]@\alpha(a)\beta(d) | (\alpha,\beta) \in \mathbb{N}ash\}.$$

As $\mathbb{N}ash \subseteq \Delta_{\mathfrak{a}} \times \Delta_{\mathfrak{d}}$, we have the refinement $\mathbb{S}(\mathcal{U}) \sqsubseteq \mathbb{S}(\mathcal{U})_{\mathsf{Mixed}} \sqsubseteq \mathbb{S}(\mathcal{U})_{\mathsf{Nash}}$. Therefore, the strategic approach trough $\mathfrak{a}/\mathfrak{d}$-games can decrease uncertainty in $\mathbb{S}$. A *pure Nash equilibrium*, PNE, is a Nash equilibrium (a,d) with pure strategies. It holds that all Nash equilibrium of a zero-sum game Γ have the same value $u(\mathcal{U})$ corresponding to the utility of the row player.[13] For an $\mathfrak{a}/\mathfrak{d}$ game $\Gamma(\mathcal{U})$ we write the *value of the game* as

$$u(\mathcal{U}) = \max_{\alpha\in\Delta_{\mathfrak{a}}} \min_{\beta\in\Delta_{\mathfrak{d}}} u(\alpha,\beta) = \min_{\beta\in\Delta_{\mathfrak{d}}} \max_{\alpha\in\Delta_{\mathfrak{a}}} u(\alpha,\beta).$$

Nash equilibria, and therefore, the value of a zero-sum game, can be characterized as follows (from Osborne's book[14]).

Lemma 2.1. *Given an uncertainty profile $\mathcal{U}$, let (α,β) be a Nash equilibrium of $\Gamma(\mathcal{U})$. It holds that, for any $a \in A_{\mathfrak{a}}$ such that $\alpha(a) > 0$, we have $u(a,\beta) = u(\mathcal{U})$ and for any $d \in A_{\mathfrak{d}}$ such that $\beta(d) > 0$, we have $u(\alpha,d) = u(\mathcal{U})$.*

Note that from above $\mathsf{E}(\mathbb{S}(\mathcal{U})_{\mathsf{Nash}}) = \{u(\mathcal{U})\}$. This fact is important because even if $\mathbb{S}(\mathcal{U})_{\mathsf{Nash}}$ contains many imprecise systems, the expected behavior of all of them is the same. This fact opens a door to dramatically remove uncertainty on the expected behavior.

Let us consider an example of application of these ideas. Let $\mathcal{U} = \langle \mathbb{S}(X_1, X_2), \mathcal{S}, \{1,2\}, \{1,2\}, 1, 1, u\rangle$ be the uncertainty profile introduced in Example 2.3. The $\mathfrak{a}$ utilities of the corresponding $\mathfrak{a}/\mathfrak{d}$ game $\Gamma(\mathcal{U})$ are the following:

		$\mathfrak{d}$: $\{1\}$	$\mathfrak{d}$: $\{2\}$
$\mathfrak{a}$	$\{1\}$	$u(\mathbb{S}(X_{1,\mathfrak{a},\mathfrak{d}}, X_2))$	$u(\mathbb{S}(X_{1,\mathfrak{a}}, X_{2,\mathfrak{d}}))$
	$\{2\}$	$u(\mathbb{S}(X_{1,\mathfrak{d}}, X_{2,\mathfrak{a}}))$	$u(\mathbb{S}(X_1, X_{2,\mathfrak{a},\mathfrak{d}}))$

About Nash equilibria of this game, we can show the following.

Lemma 2.2. *There is a single Nash equilibria defined by the mixed strategies* $\alpha = (\alpha_1, 1-\alpha_1)$ *for the angel and* $\beta = (\beta_1, 1-\beta_1)$ *for the daemon. It holds*

$$\alpha_1 = \frac{1}{z}(u(\{2\},\{2\}) - u(\{2\},\{1\})), \quad \beta_1 = \frac{1}{z}(u(\{2\},\{2\}) - u(\{1\},\{2\}))$$

when the normalizing factor z

$$z = u(\{1\},\{1\}) + u(\{2\},\{2\})) - (u(\{1\},\{2\}) + u(\{2\},\{1\}) \neq 0.$$

Moreover, the value of the game is

$$u(\mathcal{U}) = \frac{1}{z}(u(\{1\},\{1\})u(\{2\},\{2\}) - u(\{2\},\{2\})u(\{2\},\{1\})).$$

Proof. According to Lemma 2.1 in order to compute $\alpha = (\alpha_1, 1-\alpha_1)$, we have to solve the equation $u(\alpha, \{1\}) = u(\alpha, \{2\})$. Similarly, $\beta = (\beta_1, 1-\beta_1)$ is the solution of $u(\{1\}, \beta) = u(\{2\}, \beta)$. Moreover, $u(\mathcal{U}) = u(\{1\}, \beta)$. All calculus are straightforward. □

3. Web Applications

The content of this section is based on the preliminary work,[15] where we introduce uncertainty refinements in the field of Web services. In this setting, calls to sites that trigger service executions are the basic resources. Additional operators are offered to describe complex computations — so called *orchestrations* — that involve interaction of processes in different ways. In any real scenario, uncertainty is inherent in this type of Web computations. We show how to deal with orchestration unpredictability introducing semantic characterizations described in terms of imprecise probabilistic choices.

3.1. *Web under stress: Imprecision*

We introduce now the basic concepts concerning orchestrations and Web uncertainty. An *orchestration* is a user-defined program that uses services on the Web.

A basic service is called a *site*. A site is *silent* if it does not publish any result. A site call publish *at most one response*. An orchestration which composes a number of site calls into a complex computation can be represented by an Orc[11,12] expression.[a]

An orchestration publishes a data stream. We only deal here with orchestrations generating a finite number of results. Two Orc expressions (our systems) $\mathbb{E}$ and $\mathbb{F}$ can be combined using the following operators.[11,12] The *symmetric parallelism* $\mathbb{E}|\mathbb{F}$: $\mathbb{E}$ and $\mathbb{F}$ are evaluated in parallel. $\mathbb{E}|\mathbb{F}$ publishes some interleaving of the streams published by $\mathbb{E}$ and $\mathbb{F}$. The *asymmetric parallelism* $\mathbb{E}(x) < x < \mathbb{F}$: $\mathbb{E}$ and $\mathbb{F}$ are evaluated in parallel. Some sub-expressions in $\mathbb{E}$ may become blocked by a dependency on x. The first result published by $\mathbb{F}$ is bound to x, the remainder of $\mathbb{F}$'s evaluation is terminated and evaluation of the blocked residue of $\mathbb{E}$ is resumed. Finally, the *sequence* $\mathbb{E} > x > \mathbb{F}(x)$: $\mathbb{E}$ is evaluated and, for each value v published by $\mathbb{E}$, an instance $\mathbb{F}(v)$ is executed. Given an orchestration $\mathbb{E}$, we denote by $\mathsf{sites}(\mathbb{E})$ the set of sites in the definition of $\mathbb{E}$ (the parameters). Information on delays is given by an *evaluation function* δ providing the return time of each orchestration site.

Example 3.1. In $\mathbb{M}aryNews$ orchestration sites $\mathbb{C}NN$ and $\mathbb{B}BC$[b] are called in parallel and the result of the first one to answer is emailed to Mary. This procedure can be described in Orc as $\mathbb{M}aryNews = \mathbb{M}ary(x) < x < \mathbb{T}woNews$ where $\mathbb{T}woNews = \mathbb{C}NN|\mathbb{B}BC$. Orchestration sites are $\mathbb{C}NN$ and $\mathbb{B}BC$ in $\mathbb{T}woNews$ and site $\mathbb{M}ary(x)$, providing an email service to Mary. Roughly, the information received by Mary depends on the response times (the delays) of $\mathbb{C}NN$ and $\mathbb{B}BC$, denoted as $\delta(\mathbb{C}NN)$ and $\delta(\mathbb{B}BC)$. Depending on the delays several cases arise as we will see in subsequent examples.

In order to characterize ex-ante the execution of an orchestration $\mathbb{E}$, we introduced in the work[16] the *meaning* or *semantics* of $\mathbb{E}$, denoted by $[\![\mathbb{E}]\!]$. When there is no information about return times, but we know that orchestration results are $m_1, \ldots, m_k$, the semantics $[\![\mathbb{E}]\!] = \lfloor\!\lfloor m_1, \ldots, m_k \rfloor\!\rfloor$ where, abstracting away any time order, $\lfloor\!\lfloor m_1, \ldots, m_k \rfloor\!\rfloor$ is the multiset of results. As we will see later on, a non-deterministic choice of multisets M_i may be necessary to express the semantics of an orchestration. We write, in this case, $[\![\mathbb{E}]\!] = \sqcap_i M_i$. First of all, we consider an example with no information about delays. When delay time δ is used to analyze semantics of $\mathbb{E}$ we write $[\![\mathbb{E}]\!]_\delta$.

Example 3.2. Let consider Example 3.1 under lack of information about delays. Suppose $[\![\mathbb{C}NN]\!] = \lfloor\!\lfloor \mathtt{cnn} \rfloor\!\rfloor$ and $[\![\mathbb{B}BC]\!] = \lfloor\!\lfloor \mathtt{bbc} \rfloor\!\rfloor$. Then $[\![\mathbb{T}woNews]\!] = \lfloor\!\lfloor \mathtt{cnn}, \mathtt{bbc} \rfloor\!\rfloor$. Assume that both delays are unknown, encoded as $\delta(\mathbb{C}NN) =\perp$ and $\delta(\mathbb{B}BC) =\perp$. Parameter x in $\mathbb{M}aryNews = \mathbb{M}ary(x) < x < \mathbb{T}woNews$ will get either `cnn` or `bbc`, in fact the first one to arrive. As we do not have any prior knowledge of the first

[a] Although it seems that Orc is unnecessary to present the problem under investigation, Orc is an useful tool in order to properly deal with Web uncertainty. In particular, Orc allows us to develop the interplay between non-determinism and imprecise probabilities in a clean mathematical way.

[b] A call to $\mathbb{C}NN$ or $\mathbb{B}BC$ can be interpreted as a call to https://edition.cnn.com/ or https://www.bbc.com/news.

arrival, we assume $[\![x]\!]_\delta = \lfloor\lfloor \texttt{cnn} \rfloor\rfloor \sqcap \lfloor\lfloor \texttt{bbc} \rfloor\rfloor$, i.e. a non-deterministic choice of two small multisets. Then $[\![\mathbb{M}aryNews]\!]_\delta = \lfloor\lfloor \texttt{mary\_cnn} \rfloor\rfloor \sqcap \lfloor\lfloor \texttt{mary\_bbc} \rfloor\rfloor$, pointing out that Mary gets an email with news provided by either $\mathbb{C}NN$ or $\mathbb{B}BC$.

Having full knowledge means that the delay function δ is known and is defined on each orchestration site. We borrow from Hoare[17] the notation $P \triangleleft Q \triangleright R$. It should be read: P if Q else R.

Example 3.3 We assume that the true values $\delta(\mathbb{C}NN)$ and $\delta(\mathbb{B}BC)$ are known. Under the hypothesis $\delta(\mathbb{C}NN) \neq \delta(\mathbb{B}BC)$, we have

$$[\![\mathbb{M}aryNews]\!]_\delta = \lfloor\lfloor \texttt{mary\_cnn} \rfloor\rfloor \triangleleft \delta(\mathbb{C}NN) < \delta(\mathbb{B}BC) \triangleright \lfloor\lfloor \texttt{mary\_bbc} \rfloor\rfloor.$$

For function $\delta(\mathbb{C}NN) = 5$ and $\delta(\mathbb{B}BC) = 6$, it holds $[\![x]\!]_\delta = \lfloor\lfloor \texttt{cnn} \rfloor\rfloor$ and $[\![MaryNews]\!]_\delta = \lfloor\lfloor \texttt{mary\_cnn} \rfloor\rfloor$.

When (consistent) information increases, the imprecision reduces. However, it is impossible to avoid completely the non-determinism.

Example 3.4. Assume that $\delta(\mathbb{C}NN)$ and $\delta(\mathbb{B}BC)$ could have the same value, a case that assuming discrete time is feasible. Here, a race condition give us

$$\begin{aligned}[\![\mathbb{M}aryNews]\!]_\delta = {} & \lfloor\lfloor \texttt{mary\_cnn} \rfloor\rfloor \sqcap \lfloor\lfloor \texttt{mary\_bbc} \rfloor\rfloor \triangleleft \delta(\mathbb{C}NN) = \delta(\mathbb{B}BC) \triangleright \\ & (\lfloor\lfloor \texttt{mary\_cnn} \rfloor\rfloor \triangleleft \delta(\mathbb{C}NN) < \delta(\mathbb{B}BC) \triangleright \lfloor\lfloor \texttt{mary\_bbc} \rfloor\rfloor).\end{aligned}$$

When a site $\mathbb{S}$ is *under stress*, the information about delay is uncertain. Assume that $\delta_\mathfrak{a}(\mathbb{S}) \geq 0$, $\delta_\mathfrak{d}(\mathbb{S}) \geq 0$ are the corresponding levels of stress. We take

$$\delta(\mathbb{S}_\mathfrak{a}) = \delta(\mathbb{S}) + \delta_\mathfrak{a}(\mathbb{S}), \delta(\mathbb{S}_\mathfrak{d}) = \delta(\mathbb{S}) + \delta_\mathfrak{d}(\mathbb{S}), \delta(\mathbb{S}_{\mathfrak{a},\mathfrak{d}}) = \delta(\mathbb{S}) + \delta_\mathfrak{a}(\mathbb{S}) + \delta_\mathfrak{d}(\mathbb{S}).$$

In the case of Web apps, the proposed *knowledge framework* is the equivalent to our stress model.

Definition 3.1. A *knowledge framework* is a tuple $\mathcal{K} = \langle \mathbb{E}, \delta, \delta_\mathfrak{a}, \delta_\mathfrak{d} \rangle$ where δ is the delay function and $\delta_\mathfrak{a}$ and $\delta_\mathfrak{d}$ provide the delay bounds under stress Let $(a, d) \subseteq \mathsf{sites}(\mathbb{E}) \times \mathsf{sites}(\mathbb{E})$ be a pair of site subsets under, respectively, $\mathfrak{a}$ stress (subset a) and $\mathfrak{d}$ stress (subset d). We evaluate $\mathbb{S}[a, d]$ as follows:

$$\mathbb{S}[a,d] = \begin{cases} \mathbb{S} & \text{if } \mathbb{S} \notin a \cup d, \\ \mathbb{S}_\mathfrak{a} & \text{if } \mathbb{S} \in a \backslash d, \\ \mathbb{S}_\mathfrak{d} & \text{if } \mathbb{S} \in d \backslash a, \\ \mathbb{S}_{\mathfrak{a},\mathfrak{d}} & \text{if } \mathbb{S} \in a \cap d. \end{cases}$$

We denote $\mathbb{E}[a, d]$ the orchestration under stress where each $\mathbb{S} \in \mathsf{sites}(\mathbb{E})$ has been replaced by $\mathbb{S}[a, d]$.

When we want to emphasize $\mathcal{K}$, we write $\mathbb{S}_\mathcal{K}[a, d]$ and $\mathbb{E}_\mathcal{K}[a, d]$. The delay cost function $t(\mathbb{E}_\mathcal{K}[a, d])$ is the delay of the first return based on $\delta(\mathbb{S}_\mathcal{K}[a, d])$.

Our next example borrows many ideas from a typical fuzzy approach.

Example 3.5. Consider the following knowledge framework $\mathcal{K}$ for $\mathbb{M}aryNews$:

	δ	$\delta_{\mathfrak{a}}$	$\delta_{\mathfrak{d}}$		δ	$\delta_{\mathfrak{a}}$	$\delta_{\mathfrak{d}}$
$\mathbb{C}NN$	5	-2	3	$\mathbb{B}BC$	6	-4	2

Take $(a,d) = (\{\mathbb{C}NN, \mathbb{B}BC\}, \{\})$. Then $\mathbb{T}woNews_{\mathcal{K}}[a,d] = (\mathbb{C}NN_{\mathfrak{a}} | \mathbb{B}BC_{\mathfrak{a}})$ and $\mathbb{M}aryNews_{\mathcal{K}}[a,d] = \mathbb{M}ary(x) < x < \mathbb{T}woNews_{\mathcal{K}}[a,d]$. Then $t(\mathbb{T}woNews_{\mathcal{K}}[a,d]) = \min\{5-2, 6-4\} = 2$. As $\mathbb{B}BC_{\mathfrak{a}}$ returns before than $\mathbb{C}NN_{\mathfrak{a}}$ the result of the orchestration is $[\![\mathbb{M}aryNews_{\mathcal{K}}[a,d]]\!] = \lfloor\!\lfloor \texttt{mary\_bbc} \rfloor\!\rfloor$. Other results are also possible depending on (a,d), for instance $[\![\mathbb{M}aryNews_{\mathcal{K}}[\{\mathbb{C}NN\},\{\}]]\!] = \lfloor\!\lfloor \texttt{mary\_cnn} \rfloor\!\rfloor$.

3.2. *Imprecise probability and refinement*

A semantics for orchestrations where non-determinism is refined to imprecise probabilities was proposed in Castro *et al.*,[10] following our imprecise refinement hypothesis. We adapt here the semantics to deal with delays.

Example 3.6. Let us revisit Example 3.2 where $\delta(\mathbb{C}NN) = \delta(\mathbb{B}BC) = \perp$. By the imprecise refinement hypothesis,

$$[\![x]\!]_{\delta} = \lfloor\!\lfloor \texttt{cnn} \rfloor\!\rfloor \sqcap \lfloor\!\lfloor \texttt{bbc} \rfloor\!\rfloor \sqsubseteq \{\lfloor\!\lfloor \texttt{cnn} \rfloor\!\rfloor @\mu_1 \| \lfloor\!\lfloor \texttt{bbc} \rfloor\!\rfloor @\mu_2 | (\mu_1,\mu_2) \in \Delta_2\} = [\![x]\!]_{ip}.$$

Therefore, $[\![MaryNews]\!]_{ip}$ is represented by

$$\{\lfloor\!\lfloor \texttt{mary\_cnn} \rfloor\!\rfloor @\mu_1 \| \lfloor\!\lfloor \texttt{mary\_bbc} \rfloor\!\rfloor @\mu_2 | \mu_1, \mu_2 \geq 0 \quad \text{and} \quad \mu_1 + \mu_2 = 1\}.$$

Given ℓ multisets $M_1, \ldots, M_\ell$ and the Cartesian product of probability spaces $\Delta = \Delta_{m_1} \times \cdots \times \Delta_{m_k}$, we introduce the imprecise probabilistic choice of multisets $M = \|_{1 \leq i \leq \ell} M_i @ P_i(\mu)$ where multiset M_i is chosen with probability $P_i(\mu)$. Here, μ is any element of Δ and P_i's are arithmetic expressions on μ adding up one. As before, we isolate the probabilities, $\texttt{prbs}(M) = \{(P_1(\mu), \ldots, P_\ell(\mu)) | \mu \in \Delta\}$. This allow us to compare (as before) multiset probabilistic choice. Given ℓ multisets $M_1, \ldots, M_\ell$ and two multiset choices with imprecise weights on them, M is more imprecise than M' (or M' is more precise than M), $M \sqsubseteq M'$, if and only if $\texttt{prbs}(M') \subseteq \texttt{prbs}(M)$. The next result is an adaptation of Theorem 3 in Ref. 10 to our setting.

Theorem 3.1. *Let $\mathbb{F}$ be a sub-orchestration of $\mathbb{E} = \mathbb{E}(\mathbb{F})$. Suppose that our knowledge $[\![\mathbb{F}]\!]$ improves into $[\![\mathbb{F}']\!]$, i.e. $[\![\mathbb{F}]\!] \sqsubseteq [\![\mathbb{F}']\!]$. Then, the increase of knowledge goes smoothly through the whole orchestration; $[\![\mathbb{E}(\mathbb{F})]\!] \sqsubseteq [\![\mathbb{E}(\mathbb{F}')]\!]$.*

Proof. We proceed by structural induction showing that refinement is monotone through Orc operators. As an illustration, let us consider the parallel composition case. Take $\mathbb{E} = \mathbb{F}|\mathbb{G}$ where

$$[\![\mathbb{F}]\!] = \|_{1 \leq i \leq \ell} M_i @ P_i(\mu), \quad [\![\mathbb{G}]\!] = \|_{1 \leq j \leq m} N_j @ Q_j(\gamma)$$

for imprecise μ and γ in Δ. According to Ref. 10, $[\![\mathbb{F}|\mathbb{G}]\!] = \|_{i,j}(M_i + N_j)@(P_i(\mu)Q_j(\gamma))$. We write $\texttt{prbs}(\mathbb{F}|\mathbb{G})$ when $\ell = m = 2$,

$$\{(P_1(\mu)Q_1(\gamma), P_1(\mu)Q_2(\gamma), P_2(\mu)Q_1(\gamma), P_2(\mu)Q_2(\gamma)) | \mu, \gamma \in \Delta\}.$$

Suppose $[\![\mathbb{G}]\!] \sqsubseteq [\![\mathbb{H}]\!]$ with $[\![\mathbb{H}]\!] = \|_{1\le j\le m} N_j @ R_j(\tau)$ for imprecise τ in Δ, so $\texttt{prbs}([\![\mathbb{H}]\!]) \subseteq \texttt{prbs}([\![\mathbb{G}]\!])$. For $m = 2$ this inclusion can be written as,

$$\{(R_1(\tau), R_2(\tau) | \tau \in \Delta\} \subseteq \{(Q_1(\gamma), Q_2(\gamma) | \gamma \in \Delta\}.$$

For each value $P_i(\mu)$, it holds that $\{(P_i(\mu)R_1(\tau), P_i(\mu)R_2(\tau) | \tau \in \Delta\}$ is contained into $\{(P_i(\mu)Q_1(\gamma), P_i(\mu)G_2(\gamma) | \gamma \in \Delta\}$. Therefore, $\texttt{prbs}([\![F|H]\!]) \subseteq \texttt{prbs}([\![F|G]\!])$ and $[\![F|G]\!] \sqsubseteq [\![F|H]\!]$ holds. From definitions given in Ref. 10, we prove similar results for the asymmetric parallelism and sequence operators. □

3.3. *Uncertainty profiles and refinement*

Stressed orchestrations can deliver different results depending on the location of the stress, see for instance the Example 3.5. If we bound the spread of the stress but it is not possible to locate it, what can be said about the delay? To answer this question, we use uncertainty profiles. Under knowledge framework $\mathcal{K} = \langle \mathbb{E}, \delta, \delta_{\mathfrak{a}}, \delta_{\mathfrak{d}} \rangle$ the effects of the joint interaction of $\mathfrak{a}$ and $\mathfrak{d}$ are measured by the *cost* function $t(a, d) = t(\mathbb{E}_{\mathcal{K}}[a, d])$. The uncertainty profile is then a tuple $\mathcal{U} = \langle \mathcal{K}, \mathcal{A}, \mathcal{D}, b_{\mathfrak{a}}, b_{\mathfrak{d}}, t \rangle$. As we are dealing with a cost function, for $\Gamma(\mathcal{U})$, we extend t in the usual way to deal with mixed strategies. Thus we have

$$t(\mathcal{U}) = \min_{\alpha \in \Delta_{\mathfrak{a}}} \max_{\beta \in \Delta_{\mathfrak{d}}} t(\alpha, \beta) = \max_{\beta \in \Delta_{\mathfrak{d}}} \min_{\alpha \in \Delta_{\mathfrak{a}}} t(\alpha, \beta).$$

As we have seen, in this way we associate a delay value to an uncertain situation looking at it through the associated zero-sum game. When (α, β) is a Nash equilibrium it holds $t(\mathcal{U}) = t(\alpha, \beta)$.

Example 3.7. We consider for orchestration $\mathbb{T}woNews = \mathbb{C}NN | \mathbb{B}BC$ the knowledge profile $\mathcal{K}$ in Example 3.5. We examine the uncertainty profile $\mathcal{U}$ defined as

$$\langle \mathcal{K}, \textsf{sites}(\mathbb{T}woNews), \textsf{sites}(\mathbb{T}woNews), 1, 1, t \rangle,$$

where both sites can be stressed but, the angel (respectively, the daemon) affects only one site. Actions of $\mathfrak{a}$ in $\Gamma(\mathcal{U})$ are given by $A_{\mathfrak{a}} = \{a \subseteq \textsf{sites}(\mathbb{T}woNews) | \#a = 1\} = \{\{\mathbb{C}NN\}, \{\mathbb{B}BC\}\}$. As $A_{\mathfrak{d}} = A_{\mathfrak{a}}$, the set of strategy profiles $A_{\mathfrak{a}} \times A_{\mathfrak{d}}$ is

$$\{(\{\mathbb{C}NN\}, \{\mathbb{C}NN\}), (\{\mathbb{C}NN\}, \{\mathbb{B}BC\}), (\{\mathbb{B}BC\}, \{\mathbb{C}NN\}), (\{\mathbb{B}BC\}, \{\mathbb{B}BC\})\}.$$

Then, $t(\{\mathbb{C}NN\}, \{\mathbb{C}NN\}) = \min\{\delta(\mathbb{C}NN) + \delta_{\mathfrak{a}}(\mathbb{C}NN) + \delta_{\mathfrak{d}}(\mathbb{C}NN), \delta(\mathbb{B}BC)\} = 6$. Other cases can be computed similarly. The $\mathfrak{a}/\mathfrak{d}$-game is shown on the left table in Fig. 1. A strategy giving the Nash equilibrium for $\mathfrak{a}$, $\alpha = (\alpha(\{\mathbb{C}NN\}), \alpha(\{\mathbb{B}BC\}))$, is a solution to the equation $t(\alpha, \{\mathbb{C}NN\}) = t(\alpha, \{\mathbb{B}BC\})$. Similarly, to get β for $\mathfrak{d}$, we have to solve $t(\{\mathbb{C}NN\}, \beta) = t(\{\mathbb{B}BC\}, \beta)$. Then $\alpha = (2/5, 3/5)$ and $\beta = (1/5, 4/5)$. Thus, the expected return time on the first output of *TwoNews* in equilibrium is

$$t(\mathcal{U}) = \frac{2}{5} \times 6 \times \frac{1}{5} + \frac{2}{5} \times 3 \times \frac{4}{5} + \frac{3}{5} \times 2 \times \frac{1}{5} + \frac{3}{5} \times 4 \times \frac{4}{5} = \frac{18}{5}.$$

		$\mathfrak{d}$: $\{\mathbb{C}NN\}$	$\mathfrak{d}$: $\{\mathbb{B}BC\}$
$\mathfrak{a}$	$\{\mathbb{C}NN\}$	6	3
	$\{\mathbb{B}BC\}$	2	4

		$\mathfrak{d}$: $\{\mathbb{C}NN\}$	$\mathfrak{d}$: $\{\mathbb{B}BC\}$
$\mathfrak{a}$	$\{\mathbb{C}NN\}$	$\mathbb{C}NN_{\mathfrak{a},\mathfrak{d}} \sqcap \mathbb{B}BC$	$\mathbb{C}NN_{\mathfrak{a}}$
	$\{\mathbb{B}BC\}$	$\mathbb{B}BC_{\mathfrak{a}}$	$\mathbb{B}BC_{\mathfrak{a},\mathfrak{d}}$

Fig. 1. Take $\mathcal{U} = \langle \mathcal{K}, \mathsf{sites}(\mathbb{T}woNews), \mathsf{sites}(\mathbb{T}woNews), 1, 1, t\rangle$. The upper table corresponds to $\Gamma(\mathcal{U})$ and the lower one to the values of the indicator function.

We are interested in modeling how the $\mathfrak{a}/\mathfrak{d}$-games are able to refine the imprecise knowledge on asymmetric parallelism. Consider $\mathbb{E}(x) < x < \mathbb{F}$ where $\mathbb{F}$ is a parallel composition of sites $\mathbb{S}_1|\cdots|\mathbb{S}_k$. Assume $[\![\mathbb{S}_i]\!] = \lfloor\!\lfloor \mathsf{s}_i \rfloor\!\rfloor$, then $[\![\mathbb{F}]\!] = \lfloor\!\lfloor \mathsf{s}_1, \ldots, \mathsf{s}_k \rfloor\!\rfloor$, and with no delay time information ($\delta = \perp$), we can only infer that parameter x will hold any of the values in $[\![\mathbb{F}]\!]$. So, in this case $[\![x]\!] = \lfloor\!\lfloor \mathsf{s}_1 \rfloor\!\rfloor \sqcap \cdots \sqcap \lfloor\!\lfloor \mathsf{s}_k \rfloor\!\rfloor$. By the imprecise refinement hypothesis

$$[\![x]\!] \sqsubseteq \{\lfloor\!\lfloor \mathsf{s}_1 \rfloor\!\rfloor @\mu_1 \| \cdots \| \lfloor\!\lfloor \mathsf{s}_k \rfloor\!\rfloor @\mu_k | (\mu_1, \ldots, m_k) \in \Delta_k\} = [\![x]\!]_{ip}.$$

Therefore,

$$[\![\mathbb{E}(x) < x < \mathbb{F}]\!]_{ip} = \|_{1 \le i \le k} [\![\mathbb{E}(\mathsf{s}_i)]\!] @\mu_i.$$

In order to provide an expression for the refinement of $[\![x]\!]_{ip}$ through $\mathfrak{a}/\mathfrak{d}$-games, we introduce some additional concepts.

Definition 3.2. Let $\mathcal{U} = \langle \mathcal{K}, \mathcal{A}, \mathcal{D}, b_{\mathfrak{a}}, b_{\mathfrak{d}}, t\rangle$, where $\mathcal{K} = \langle \mathbb{F}, \delta, \delta_{\mathfrak{a}}, \delta_{\mathfrak{d}}\rangle$ and $\mathbb{F}$ is a parallel composition of sites $\mathbb{S}_1|\cdots|\mathbb{S}_k$. For each strategy profile (a, d) we consider $\mathbb{F}[a, d] = \mathbb{T}_1|\cdots|\mathbb{T}_k$ where $\mathbb{T}_\ell = \mathbb{S}_\ell[a, d]$, for $1 \le \ell \le k$. The *indicator function* of strategy profile (a, d) in $\Gamma(\mathcal{U})$ is the set consisting of all $\mathbb{S}_\ell$ sites in which $\delta(\mathbb{T}_\ell)$ is minimum among $\{\delta(\mathbb{T}_1), \ldots, \delta(\mathbb{T}_k)\}$. Formally,

$$I_{\mathcal{U}}(a, d) = \{\mathbb{S}_\ell | \delta(\mathbb{T}_\ell) = t(a, d))\}.$$

Proposition 3.1. *Let (α, β) be a Nash equilibrium in $\Gamma(\mathcal{U})$. For any $\mathbb{S} \in I_{\mathcal{U}}(a, d)$,*

$$t(\mathcal{U}) = \sum_{a,d} \alpha(a)\delta(\mathbb{S}[a, d])\beta(d).$$

Proof. By definition, $t(\mathcal{U}) = \sum_{a,d} \alpha(a) t(\mathbb{F}[a, d])\beta(d)$. As $\mathbb{S} \in I_{\mathcal{U}}(a, d)$, it holds $\delta(\mathbb{S}[a, d]) = t(\mathbb{F}[a, d])$ and we get the result. □

Let us introduce refinements $[\![x]\!]_{\mathcal{U}}$, provided by the game $\Gamma(\mathcal{U})$, on the $[\![x]\!]_{ip}$ semantics.

Example 3.8. We continue with Example 3.7 emphasizing the stress suffered by the sites, see the bottom table in Fig. 1. This table points out the sites giving minimum return time in $(\mathbb{C}NN|\mathbb{B}BC)[a,d]$. For instance, when $(a,d) = (\{\mathbb{C}NN\},\{\mathbb{C}NN\})$, the table value is a non-deterministic choice between $\mathbb{C}NN_{\mathfrak{a},\mathfrak{d}}$ and $\mathbb{B}BC$ representing that $I_{\mathcal{U}}(\mathbb{C}NN,\mathbb{C}NN)$ is $\{\mathbb{C}NN,\mathbb{B}BC\}$. The following refinement of $[\![x]\!]_{ip}$ can be introduced from a Nash equilibrium and the indicators functions of $\Gamma(\mathcal{U})$

$$\left(([\![\mathbb{C}NN]\!] \sqcap [\![\mathbb{B}BC]\!])@\frac{2}{5}\times\frac{1}{5}\,\|\,[\![\mathbb{C}NN]\!]@\frac{2}{5}\times\frac{4}{5}\,\|\,[\![\mathbb{B}BC]\!]@\frac{3}{5}\times\frac{1}{5}\,\|\,[\![\mathbb{B}BC]\!]@\frac{3}{5}\times\frac{4}{5}\right)$$

Note that in this case $t(\mathcal{U})$ is $18/5$.

The $\mathcal{U}$ imprecise set that fixes the stress exerted by $\mathfrak{a}$ and $\mathfrak{d}$ over the sites, according to weighted strategy profiles in $\Gamma(\mathcal{U})$ in a Nash equilibrium, is defined as follows.

Definition 3.3. The $\mathcal{U}$ *imprecise* set contains the expressions

$$[\![x]\!]_{\mathcal{U}} = \|_{a,d}(\sqcap_{\mathbb{S}\in I_{\mathcal{U}}(a,d)}[\![S]\!])@\alpha(a)\times\beta(d),$$

for each Nash equilibrium (α,β) of $\Gamma(\mathcal{U})$.

$[\![x]\!]_{\mathcal{U}}$ is a refinement of $[\![x]\!]_{ip}$ obtained by applying the imprecise refinement hypothesis in the daemonic choice appearing in $\sqcap_{S\in I_{\mathcal{U}}(a,d)}[\![\mathbb{S}]\!]$.

Example 3.9. For the Nash equilibrium in Example 3.7, the $\mathcal{U}$ refinement for $[\![x]\!]_{ip}$ is

$$\begin{aligned}[\![x]\!]_{\mathcal{U}} &= ([\![\mathbb{C}NN]\!] \sqcap [\![\mathbb{B}BC]\!])@\frac{2}{5}\times\frac{1}{5}\,\|\,[\![\mathbb{C}NN]\!]@\frac{2}{5}\\ &\quad\times\frac{4}{5}\,\|\,[\![\mathbb{B}BC]\!]@\frac{3}{5}\times\frac{1}{5}\,\|\,[\![\mathbb{B}BC]\!]@\frac{3}{5}\times\frac{4}{5}\\ &= (\lfloor\lfloor\texttt{cnn}\rfloor\rfloor \sqcap \lfloor\lfloor\texttt{bbc}\rfloor\rfloor)@\frac{2}{5}\times\frac{1}{5}\,\|\,\lfloor\lfloor\texttt{cnn}\rfloor\rfloor@\frac{2}{5}\times\frac{4}{5}\,\|\,\lfloor\lfloor\texttt{bbc}\rfloor\rfloor@\frac{3}{5}.\end{aligned}$$

Based on $[\![x]\!]_{\mathcal{U}}$, we can define $[\![\mathbb{M}aryNews]\!]_{\mathcal{U}}$. By, the imprecise refinement hypothesis, refining the daemonic choice $\sqcap$ by imprecise probabilities and regrouping, we have

$$\lfloor\lfloor\texttt{cnn}\rfloor\rfloor@\left(\frac{2}{5}\times\frac{4}{5}+\mu_1\times\frac{2}{5}\times\frac{1}{5}\right)\|\,\lfloor\lfloor\texttt{bbc}\rfloor\rfloor@\left(\frac{3}{5}+\mu_2\times\frac{2}{5}\times\frac{1}{5}\right)$$

therefore, we get the refinement

$$[\![x]\!]_{\mathcal{U}} \sqsubseteq \left\{\lfloor\lfloor\texttt{cnn}\rfloor\rfloor@\left(\frac{8}{25}+\mu_1\frac{2}{25}\right)\|\,\lfloor\lfloor\texttt{bbc}\rfloor\rfloor@\left(\frac{3}{5}+\mu_2\frac{2}{25}\right)\middle|(\mu_1,\mu_2)\in\Delta_2\right\} = [\![x]\!]_{\mathcal{U},ip}.$$

Then, we have the following refinement for $\mathbb{M}aryNews$:

$$[\![\mathbb{M}aryNews]\!]_{\mathcal{U},ip} = \lfloor\lfloor\texttt{mary\_cnn}\rfloor\rfloor@\left(\frac{8}{25}+\mu_1\frac{2}{25}\right)\|\,\lfloor\lfloor\texttt{mary\_bbc}\rfloor\rfloor@\left(\frac{3}{5}+\mu_2\frac{2}{25}\right).$$

As $\texttt{prbs}(\llbracket x \rrbracket_{\mathcal{U},ip}) \subseteq \texttt{prbs}(\llbracket x \rrbracket_{ip})$, we have

$$\llbracket \mathbb{M}aryNews \rrbracket \sqsubseteq \llbracket \mathbb{M}aryNews \rrbracket_{ip} \sqsubseteq \llbracket \mathbb{M}aryNews \rrbracket_{\mathcal{U}} \sqsubseteq \llbracket MaryNews \rrbracket_{\mathcal{U},ip}.$$

Observe that we get a chain of improving refinements.

Moreover, if we assume that the environment behaves as predicted by $\mathcal{U}$, a more precise behavior can be announced.

Theorem 3.2. *Given $\mathbb{E}(x) < x < \mathbb{F}$ where $\mathbb{F} = (\mathbb{S}_1 | \cdots | \mathbb{S}_k)$ and an uncertainty profile $\mathcal{U}$ on $\mathbb{F}$ we have $\llbracket x \rrbracket_{ip} \sqsubseteq \llbracket x \rrbracket_{\mathcal{U}} \sqsubseteq \llbracket x \rrbracket_{\mathcal{U},ip}$ and*

$$\llbracket \mathbb{E}(x) < x < \mathbb{F} \rrbracket_{ip} \sqsubseteq \llbracket \mathbb{E}(x) < x < \mathbb{F} \rrbracket_{\mathcal{U}} \sqsubseteq \llbracket \mathbb{E}(x) < x < \mathbb{F} \rrbracket_{\mathcal{U},ip}.$$

Assuming that $\mathbb{F}$ behaves as predicted by $\Gamma(\mathcal{U})$, the arrival times to x of the different possible values follow the $\llbracket x \rrbracket_{\mathcal{U},ip}$ distribution. Moreover, assuming that the execution of $\mathbb{E}$ is triggered by x, $\mathsf{E}(t(\mathbb{E}(x) < x < \mathbb{F})) = t(\mathbb{E}) + t(\mathcal{U})$.

In this way, using the imprecise refinement hypothesis, induction and additional rules for more complex composition, we can associate a meaning to any E under uncertainty profile $\mathcal{U}$ such that $\llbracket \mathbb{E} \rrbracket \sqsubseteq \llbracket \mathbb{E} \rrbracket_{ip} \sqsubseteq \llbracket \mathbb{E} \rrbracket_{\mathcal{U}} \sqsubseteq \llbracket \mathbb{E} \rrbracket_{\mathcal{U},ip}$.

4. Conclusions

We have extended previous works[10,15] studying relationships between uncertainty, non-determinism and imprecision. In order to obtain sensible forecasts of uncertain systems, we have proposed to replace non-determinism by an arbitrary probabilistic choice, through the refinement hypothesis. Depending on our system knowledge, additional refinements shrinking arbitrariness can be considered. Our first conclusion is that addressing uncertainty when it emerges is no worse than, in terms of predictability, postponing the challenge.

Uncertainty profiles are formal resources to set up the knowledge about working conditions of a system. To each profile, we associated a zero sum game that allows a strategic analysis. We have proposed the distributions in the Nash equilibrium as an unbiased set of probabilities where to limit arbitrariness. This approach allows us to consider more realistic scenarios and provides an additional analysis tool to support the decisions of the system managers. As a proof of the viability of the proposal, we have analyzed some families of Web orchestrations.

Nash equilibria and the value of the $\mathfrak{a}/\mathfrak{d}$-game is a natural way for refinements, however, finding Nash equilibria in $\mathfrak{a}/\mathfrak{d}$-games can be computationally difficult, in fact it is an EXP-complete problem.[18] The approach through mixed strategies, appearing in the Nash equilibria, seems also to suggest the possibility to develop algorithms-based Monte Carlo techniques, that could be more efficient. As $\mathfrak{a}/\mathfrak{d}$-games are zero-sum games, it is also possible to find the game values through iterative methods.[19] These computational aspects need also to be explored.

It will be of interest to analyze whether the uncertainty analysis performed here, in particular the $\mathfrak{a}/\mathfrak{d}$-refinement can be applies to other settings. One area of interest is short-term economic systems, like the IS-LM or IS-MP models,[7] or other decision support models. We can also apply the approach to study perturbations in systems coming from real politics[20] or climate change.[21]

Acknowledgments

We wish to express our thanks to the reviewer for the careful reading of the paper and the useful suggestions. J. Castro was partially supported by MINECO and EU FEDER funds grant TIN2017-89244-R and AGAUR grant 2017 SGR-856. J. Gabarro and M. Serna were partially supported by MINECO and EU FEDER funds grant TIN2017-86727-C2-1-R and AGAUR 2017 SGR 786.

References

1. F. Knight, *Risk, Uncertainty and Profit* (Houghton Mifflin, 1921).
2. G. D. C. Thomas Augustin, F. Coolen and M. Troffaes, *Introduction to Imprecise Probabilities* (Wiley, 2014).
3. C. Hoare, *Communicating Sequential Processes* (Prentice-Hall, London, 1985).
4. A. McIver and C. C. Morgan, *Abstraction, Refinement and Proof for Probabilistic Systems* (Springer, New York, 2005).
5. M. E. Porter, *Competitive Strategy: Techniques for Analyzing Industries and Competitors* (Free Press, New York, 1980).
6. J. Gabarro, M. Serna and A. Stewart, Analysing web-orchestrations under stress using uncertainty profiles, *Comput. J.* **57**(11) (2014) 1591–1615.
7. J. Gabarro and M. Serna, Uncertainty in basic short-term macroeconomic models with angel-daemon games, *Int. J. Data Anal. Tech. Strat.* **9**(4) (2017) 314–330.
8. G. Akerlof and R. Schiller, *Animal Spirits* (Princeton University Press, Princeton, Oxford, 2009).
9. J. Keynes, *The General Theory of Employment, Interest and Money* (Macmillan and Co., 1936. Re-edition, Palgrave-Macmillan for the Royal Economic Society, 2007).
10. J. Castro, J. Gabarro and M. Serna, Web apps and imprecise probabilitites, in *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Theory and Foundations — 17th IPMU, Proceedings, Part II*, eds. J. Medina, M. Ojeda-Aciego, J. L. V. Galdeano, D. A. Pelta, I. P. Cabrera, B. Bouchon-Meunier and R. R. Yager, Communications in Computer and Information Science, Vol. 854 (Springer, 2018), pp. 226–238.
11. J. Misra and W. Cook, Computation orchestration: A basis for wide-area computing, *Softw. Syst. Model.* **6**(1) (2007) 83–110.
12. D. Kitchin, A. Quark, W. Cook and J. Misra, The Orc programming language, in *Proceedings of FMOODS/FORTE 2009*, eds. D. Lee, A. Lopes and A. Poetzsch-Heffter, Lecture Notes in Computer Science, Vol. 5522 (Springer, 2009), pp. 1–25.
13. J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, 60th Anniversary Commemorative edn. (Princeton University Press, Princeton, Oxford, 1953).
14. M. Osborne, *An Introductions to Game Theory* (Oxford University Press, New York, Oxford, 2004).

15. J. Castro, J. Gabarro and M. Serna, Refining the imprecise meaning of non-determinism in the web by strategic games, in *11th Computational Collective Intelligence, Part I*, eds., N. T. Nguyen, R. Chbeir, E. Exposito, P. Aniorté and B. Trawinski, Lecture Notes in Computer Science, Vol. 11683 (Springer, 2019), pp. 566–578.
16. J. Gabarro, S. Leon-Gaixas and M. Serna, The computational complexity of QoS measures for orchestrations, *J. Combin. Optim.* **34**(4) (2017) 1265–1301.
17. C. A. R. Hoare, A couple of novelties in the propositional calculus, *Math. Log. Q.* **31**(9–12) (1985) 173–178.
18. J. Gabarro, A. Garcia and M. Serna, Computational aspects of uncertainty profiles and angel-daemon games, *Theory Comput. Syst.* **54**(1) (2014) 83–110.
19. J. Robinson, An iterative method of solving a game, *Ann. Math. Sec. Ser.* **54**(2) (1951) 296–301.
20. M. D. García-Sanz, I. Llamazares and M. A. Manrique, Ideal and real party positions in the 2015–2016 spanish general elections, in *IPMU 2018, Part III*, Communications in Computer and Information Science, Vol. 855 (Springer, 2018), pp. 52–62.
21. W. Nordhaus, *Climate Casino: Risk, Uncertainty, and Economics for a Warming World* (Yale, 2013).