

Sufficient Decision Proxies for Decision-Focused Learning

Noah Schutte¹, Grigori Vevurko¹, Krzysztof Postek² and Neil Yorke-Smith¹

¹Delft University of Technology

²Independent Researcher

{n.j.schutte, n.yorke-smith}@tudelft.nl, {grivurko, krzysztof.postek}@gmail.com

Abstract

When solving optimization problems under uncertainty with contextual data, utilizing machine learning to predict the uncertain parameters' values is a popular and effective approach. Decision-focused learning (DFL) aims at learning a predictive model such that decision quality, instead of prediction accuracy, is maximized. Common practice is to predict a single scenario representing the uncertain parameters, implicitly assuming that there exists a deterministic problem approximation (proxy) that allows for optimal decision-making. The opposite has also been considered, where the underlying distribution is estimated with a parameterized distribution. However, little is known about when either choice is valid. This paper investigates for the first time problem properties that justify using a certain decision proxy. Using this, we present alternative decision proxies for DFL, with little or no compromise on the complexity of the learning task. We show the effectiveness of presented approaches in experiments on continuous and discrete problems, as well as problems with uncertainty in the objective function and in the constraints.

1 Introduction

Both in our daily life and in large-scale problems, we often make decisions under uncertainty. Imagine packing your suitcase or finding the best travel route when going to a conference. To make the right decision, we use available data that correlate with what is uncertain: e.g. we use seasonal weather patterns to predict the weather and decide if an umbrella is needed. We are trying to solve a contextual optimization or 'predict-then-optimize' problem: making predictions based on contextual information, which we then use for optimization (i.e., decision making). Since our goal is to make the best decision, we do not care about the predictions by themselves, but only about the decision they lead to. This is the main premise of *decision-focused learning* (DFL): learning a predictive model that optimizes for decision quality instead of prediction accuracy.

Defining our setting more formally: Observing contextual information in the form of feature values $z \in Z$, the goal of

the decision maker is to find optimal decision x from constrained set X by solving the following problem:

$$x^*(z) = \operatorname{argmin}_{x \in X} \mathbb{E}_{c \sim \mathcal{C}_z} [f(c, x) | z], \quad (1)$$

where $\mathcal{C}_z$ is an unknown probability distribution. Since both c and $\mathcal{C}_z$ are unknown, learning a predictive model $\phi_\theta(\cdot)$ for c is a natural approach to the problem, assuming we have data with z, c pairs. After training, once new contextual information z is observed, we can predict an estimated *scenario* $\hat{c} = \phi_\theta(z)$. We do this because we would like to predict the realized value of c , i.e., the specific value that c will take in this empirical observation of z . If we do this successfully, we will get the empirical optimal decision by solving the *deterministic approximation* (proxy) of the true optimization problem in Equation (1): $\arg \min_{x \in X} f(\hat{c}, x)$. However, since the true c is stochastic and its distribution $\mathcal{C}_z$ is unknown, the empirical optimal decision is not necessarily the same as the *true* optimal decision based on Equation (1).

In certain settings it has been proven that a single estimated scenario and deterministic proxy allow for optimal decision-making, such as when uncertainty is linear in the objective [Elmachtoub and Grigas, 2022], or for certain fixed recourse, fixed costs two-stage stochastic problems [Homem-de Mello *et al.*, 2026]. However, it has not been studied when a deterministic proxy leads to suboptimal decision-making, and how we can effectively apply DFL in this case. We address this gap in literature by contributing the following:

1. We provide sufficient and necessary conditions for when deterministic decision proxies are suboptimal and present problem classes where these conditions are met.
2. We present alternative decision proxies that adhere to sufficient conditions allowing for optimal decision-making, with little or no compromise on complexity.
3. We analyze continuous and discrete problems where deterministic proxies are suboptimal and show the effectiveness of the proposed proxies.

2 Problem Formulation

The aim in a contextual optimization setting is to go from context z to optimal decision $x^*(z)$, by finding a mapping or policy $\pi : Z \rightarrow X$. Due to the structure of the (potentially combinatorial) optimization problem $x^*(\cdot)$, it is challenging




Figure 1: Schematic representation of contextual optimization and the difference between the true optimal decision based on the unknown contextual distribution (orange/top) and the empirical optimal decision based on the deterministic proxy (blue/bottom).

to learn this complete end-to-end mapping directly [Bengio *et al.*, 2021]. Considering it as separate prediction and optimization problems (predict-then-optimize) is practical, as both problems have been studied extensively [Sadana *et al.*, 2025]. From the prediction problem perspective, assuming data with feature-label pairs z, c available, a typical learning would be to train predictor $\phi_\theta : Z \rightarrow C$ in a supervised manner by minimizing its predictive error (e.g., mean squared error: $\|c - \hat{c}\|_2^2$). We will refer to this approach as *prediction-focused learning* (PFL). By contrast, DFL aims at learning a predictive model that minimizes decision error, taking the downstream optimization problem into account. Decision error is a measure of the difference between empirical optimal decision $x^D(c)$ and the obtained decision based on the prediction $x^D(\hat{c})$ (e.g., regret: $f(c, x^D(\hat{c})) - f(c, x^D(c))$). Here $x^D : C \rightarrow X$ is the (single scenario) *deterministic proxy* of stochastic optimization problem $x^*(\cdot)$ from (1):

$$x^D(c) = \operatorname{argmin}_{x \in X} f(c, x).$$

Both learning methods lead to a policy $\pi^D = x^D \circ \phi_\theta$. It is natural to consider this deterministic proxy, since in practice we only have data available and therefore do not have C_z to compute $x^*(z)$ directly. We observe realized pairs z, c , that have an accompanied empirical optimal $x^D(c)$. However, at inference time, $x^*(z)$ is the optimal decision and there is no guarantee that $x^D(\cdot)$ is able to equal $x^*(z)$. Proxy $x^D(\cdot)$ can be limited to a subset of the feasible space $X^D \subset X$ due to its structure, which means that $x^*(z)$ is unattainable if it lies outside X^D . Figure 1 shows a schematic representation of this potential problem. We give an illustrative example.

Example 1. Suboptimal investing based on the stock-exchange paradox by Kibzun and Kan [1996]. Assume we have some capital that we want to grow in the coming t years. Every year, we can deposit capital in a bank with fixed rate of return $\beta = 0.05$ and we can purchase shares with an uncertain rate of return $c \sim \mathcal{U}[-1, 1.7]$ (continuous uniformly distributed). If our aim is to maximize the next period's return, we will invest our whole capital in shares, since

$\mathbb{E}[c] = 0.35 > 0.05$. However, this leads to a negative *expected growth rate*:

$$\mathbb{E}[\ln(1 + c)] = \ln(2.7) - 1 < 0,$$

Because of this, investing all capital in shares every period will make the probability of ruin, i.e., losing all capital at some point, converge to 1.

This is considered the stock-exchange paradox, where the Kelly strategy [Kelly, 1956] is being considered as solving:

$$x^* = \operatorname{argmax}_{0 \leq x \leq 1} \mathbb{E}[\ln(1 + \beta x + c(1 - x))].$$

Given $c \sim \mathcal{U}[-1, 1.7]$ and $\beta = 0.05$, we get $x^* \approx 0.465$.

Now, if we look at the deterministic proxy of the problem, it will become clear that we cannot find the optimal decision:

$$x^D(\hat{c}) = \operatorname{argmax}_{0 \leq x \leq 1} \ln(1 + \beta x + \hat{c}(1 - x)).$$

Since $\ln(\cdot)$ is a monotonically increasing function, we get: $x^D(\hat{c}) = 1$ when $\hat{c} < \beta$ & $x^D(\hat{c}) = 0$ when $\hat{c} > \beta$. Hence true optimal decision $x^* \approx 0.465$ cannot be obtained.

2.1 Contextual Value of the Stochastic Solution

To more clearly present theory on the suboptimality of deterministic proxies, we extend the concept of *value of the stochastic solution* [Birge, 1982] to a contextual optimization setting. This concept is introduced as a measure for the value of solving a stochastic optimization problem as compared to solving a deterministic proxy with the uncertain variables replaced by their expectation. This is useful because in practice it is both harder to estimate the uncertain variables distribution (compared to its expectation) and more complex to solve a stochastic optimization problem. Here it gives us some clear structure to prove strict differences between solving a deterministic or stochastic contextual optimization problem. We first define the value of the true optimal decision based on Equation (1):

$$V^* = \min_{x \in X} \mathbb{E}_{c \sim C_z} [f(c, x) | z] = \mathbb{E}_{c \sim C_z} [f(c, x^*(z)) | z]$$

When we apply PFL, the predictive error is minimized if we are able to exactly estimate $\bar{c}_z = \mathbb{E}_{c \sim C_z} [c | z]$. Assuming that we can find this estimate, the value of a PFL decision is:

$$V_{\text{PFL}} = \mathbb{E}_{c \sim C_z} [f(c, x^D(\bar{c}_z)) | z].$$

In the case of DFL the decision error is minimized. Assuming that we can find $\hat{c}_z$ that minimizes the decision error, the value of the DFL decision is:

$$V_{\text{DFL}} = \min_{\hat{c}_z \in C} \mathbb{E}_{c \sim C_z} [f(c, x^D(\hat{c}_z)) | z]$$

We have $V^* \leq V_{\text{DFL}} \leq V_{\text{PFL}}$. The first inequality holds because in V^* we directly minimize over the input of $f(c, \cdot)$, $x \in X$, compared to minimizing over the argument of a function with codomain X , resulting in the same input of $f(c, \cdot)$. The second inequality holds because we take minimizer $\hat{c}_z \in C$ instead of $\bar{c}_z \in C$. Both inequalities hold for any z . This means that if there exists a z with non-zero probability mass for which one of the inequalities is strict, this inequality remains strict in expectation.

The value of using DFL over PFL $V_{\text{PFL}} - V_{\text{DFL}}$ is highly relevant for predict-then-optimize problems. A general consensus is that this value increases with predictive error [Mandi *et al.*, 2024], as the higher the predictive error the more important it is to learn such that decision quality is least affected. Additionally, Cameron *et al.* [2022] show that $V_{\text{PFL}} - V_{\text{DFL}}$ is also influenced by dependencies between multiple uncertain variables: Predicting marginal expectations using PFL can cause theoretically unbounded worse performance compared to DFL.

In this work we are interested in the difference between the value of using DFL with a deterministic proxy V_{DFL} and true optimal value V^* : $V_{\text{DFL}} - V^*$. Due to DFL’s end-to-end learning approach, proxy $x^D(\cdot)$ has not been considered as limiting in its performance, even in more complex stochastic settings [Silvestri *et al.*, 2026; van den Houten *et al.*, 2024]. In the following we will identify cases when $V_{\text{DFL}} - V^* > 0$.

3 Suboptimal Deterministic Decision Proxies

This section presents cases where DFL with a deterministic proxy is provably suboptimal ($V_{\text{DFL}} > V^*$). We first present theorems for the cases where $V^* = V_{\text{DFL}} = V_{\text{PFL}}$ and $V^* = V_{\text{DFL}} < V_{\text{PFL}}$. Proofs are presented in the Appendix.

Theorem 1. *There exists at least one optimal single-scenario w.r.t. Equation (1) if $\forall z \in Z, \mathbb{E}_{c \sim \mathcal{C}_z}[f(c, x)|z] = f(\mathbb{E}_{c \sim \mathcal{C}_z}[c|z], x)$, i.e., Jensen’s inequality renders an equality. In this case, an optimal single-scenario is $\bar{c}_z = \mathbb{E}_{c \sim \mathcal{C}_z}[c|z]$. We have $V^* = V_{\text{DFL}} = V_{\text{PFL}}$.*

The most common application of Theorem 1 is when $f(c, x)$ is linear in c , $\forall x \in X$. If Jensen’s inequality is strict, and hence both functions have different minimizers, PFL becomes provably suboptimal:

Theorem 2. *If $\arg\min_{x \in X} \mathbb{E}_{c \sim \mathcal{C}_z}[f(c, x)|z] \neq \arg\min_{x \in X} f(\mathbb{E}_{c \sim \mathcal{C}_z}[c|z], x)$, we have $V^* < V_{\text{PFL}}$ ¹.*

Now we related this to V_{DFL} , by presenting a sufficient condition for when a decision proxy allows for optimal decision-making, i.e., an optimal single-scenario exists.

Theorem 3. *Deterministic proxy x^D allows for optimal decision-making if $x^D(\cdot) : C \rightarrow X$ is surjective, i.e., $\forall x \in X, \exists \hat{c} \in C : x^D(\hat{c}) = x$. We have $V^* = V_{\text{DFL}}$.*

In other words: When x^D is surjective, it is expressive enough to attain any feasible decision $x \in X$, which includes optimal $x^*(z)$. Theorem 2 and Theorem 3 demonstrate a theoretical case where PFL is inferior to DFL. The result that $V^* = V_{\text{DFL}}$, based on $x^D(\cdot)$ being surjective, shows the flexibility of DFL: it allows for learning predictions that work well w.r.t. the true optimization problem. On the other hand, if the deterministic proxy is not able to reach stochastically optimal decisions, using this proxy is guaranteed to be sub-optimal. This is the main intuition behind suboptimality of the deterministic proxy, and occurs when we have decision dominance:

Definition 1. *Decision $x \in X$ is (scenario-wise) dominated if $\forall c \in C, \exists$ (different) $\hat{x} \in X$ s.t. $f(c, \hat{x}) < f(c, x)$.*

¹Notation is slightly abused here, as the minima’s arguments are not necessarily unique. Technically these must be disjoint sets.

Note that this definition is based on the objective function $f(\cdot)$, which means that this property does not relate decisions w.r.t. the true stochastic problem in Equation (1).

Lemma 1. *If there exists a scenario-wise dominated $x \in X$, proxy $x^D(c) = \arg\min_{x \in X} f(c, x)$ is not surjective.*

Lemma 1 holds because dominated decisions can not be reached, violating the property of surjectivity. However, non-surjectivity is not a sufficient condition for a proxy to be sub-optimal. What is a sufficient and necessary condition is true optimal decision $x^*(z)$ being scenario-wise dominated, as the next theorem expresses.

Theorem 4. *Proxy $x^D(\cdot)$ is suboptimal if and only if there exists a feature value $z \in Z$ for which true optimal $x^*(z)$ is scenario-wise dominated. We have $V^* < V_{\text{DFL}}$.*

Despite not having access to $x^*(z)$, based on objective function f ’s characteristics we can prove that dominance occurs for certain continuous decision problems, i.e., when the feasibility set is compact (closed and bounded) and objective function f is continuous. After this we present a class of discrete problems where it is more natural to occur.

3.1 Continuous Optimization

The main result for *continuous optimization problems* is the following: Depending on the shape of objective function $f(c, x)$ and expected objective function $g(x) := \mathbb{E}_{c \sim \mathcal{C}_z}[f(c, x)|z]$, it can occur that $\forall c \in C$ the minimizer of $f(c, x)$ lies on the boundary of X , while the minimizer of $g(x)$ lies in the interior of X . In this case, all interior decisions are dominated by boundary decisions. In Example 1, we observe that the true objective function is convex non-monotonic in x , while the deterministic proxy objective function is strictly monotonic in x . Because of this the optimal decision according to the proxy lies on the boundary of the feasibility space, while the true optimal decision lies in the interior. This results in both decisions never coinciding. Example 1 has single-variable functions, but since we are working with multivariate functions in general, we first give a monotonicity definition for multivariate functions.

Definition 2. *Let $f : X \rightarrow \mathbb{R}, X \subset \mathbb{R}^m$ be a multivariate function. f is (strictly) coordinate-wise (non-)monotonic if $\forall i \in \{1, \dots, m\}$ we have that for any fixed variables $\{\bar{x}_1, \dots, \bar{x}_{i-1}, \bar{x}_{i+1}, \dots, \bar{x}_m\} \in \mathbb{R}^{m-1}$, $f(\bar{x}_1, \dots, x_i, \dots, \bar{x}_m)$ is (strictly) (non-)monotonic in x_i .*

Coordinate-wise monotonicity allows the function to be decreasing in some variables while increasing in others.

Theorem 5. *Given two continuous (objective) functions $f : C, X \rightarrow \mathbb{R}, g : X \rightarrow \mathbb{R}$, with $g(x) := \mathbb{E}_{c \sim \mathcal{C}_z}[f(c, x)|z]$, and (feasibility set) X is compact, i.e., closed and bounded, and has a nonempty interior. If $f(c, x)$ is strictly coordinate-wise monotonic in x for all $c \in C$, and $g(x)$ is convex and coordinate-wise non-monotonic, then $V^* < V_{\text{DFL}}$.*

In practice the constraints that define X can cause a lower intrinsic dimension and therefore an empty interior. However, in the Appendix we prove that the same result holds.

Example 2. *Continuous objective functions with(out) sub-optimal deterministic proxy. (Assuming X is compact.)*

Objective $f(c, x) = c^T x$ does not satisfy the conditions of Theorem 5, since $g(x) = \mathbb{E}_{c \sim \mathcal{C}_z}[c^T x | z] = \mathbb{E}_{c \sim \mathcal{C}_z}[c | z]^T x$ is linear in x for any distribution $\mathcal{C}_z$ and therefore not coordinate-wise non-monotonic. Similarly $\frac{1}{c^T x}$ and $(c^T x)^n$ with even n 's do not satisfy the conditions of Theorem 5 as they are also not coordinate-wise non-monotonic.

Objective $f(c, x) = -\log(c^T x)$ does satisfy the conditions of Theorem 5 for at least some distributions $\mathcal{C}_z$. This is because $-\log(c^T x)$ is strictly coordinate-wise monotonic $\forall c \in C$, because for any $c_i > 0$, increasing x_i increases $\log(c^T x)$, and for any $c_i < 0$, decreasing x_i decreases $\log(c^T x)$. On the other hand, $g(x)$ can be strictly convex with a minimum in the interior, depending on the distribution $\mathcal{C}_z$ (see Example 1). Similarly $\sqrt{c^T x}$, $e^{c^T x}$ and $(cx)^n$ with odd n 's satisfy the conditions of Theorem 5.

So far we have derived properties for the structure of the optimization problem to show suboptimality. Based on the function examples, one could argue that the types of functions for which this occurs are not common. However, there is an important class of practical optimization problem that typically have complex, non-smooth objective functions.

3.2 Two-Stage Stochastic Optimization Problems

A practical class of optimization problems where a deterministic proxy can be suboptimal is the class of *two-stage stochastic optimization problems*. This two-stage formulation is of particular interest in optimization literature [Birge and Louveaux, 2011], as it is effective in modeling real-world problems under uncertainty. These problems consist of making first-stage decisions x that are taken before uncertain parameters c are realized and second-stage decisions y that denote some recourse action after uncertain parameters c are realized. The general formulation is:

$$\begin{aligned} & \min_{x \in X} f(x) + \mathbb{E}_c[Q(c, x)] \\ & \text{with } Q(c, x) = \min_{y \in Y(c, x)} g(c, x, y) \end{aligned} \quad (2)$$

The relevance of this formulation comes from the fact that in practice the realization of uncertainty leads to some action: if we did not bring an umbrella, we can buy a new one. These *recourse actions* include additional costs compared to what was initially decided. One can model the problem in a way it has *relatively complete recourse*, i.e., the second stage is feasible for any feasible first-stage decision. This is one of the main benefits of two-stage stochastic optimization, as other uncertain problem formulations like (distributionally) robust optimization or chance constraints [Charnes and Cooper, 1962] can only ensure feasibility for a specific set of uncertain value realizations or with a certain probability. Specifically, models with a linear second stage are commonly used: Equation (2) with $g(c, x, y) = q(c)^T y$ and $Y(c, x) = \{y : T(c)x + U(c)y = h(c)\}$.

Homem-de Mello *et al.* [2026] showed that for these problems with fixed costs and fixed recourse, i.e., $q(c) = q$ and $U(c) = U$ being independent of c , an optimal single-scenario exists. The proof is based on the assumption that each entry of vector $h(c)$ is dependent on c , but when it has independent fixed entries, an optimal single-scenario does not necessarily exist. We demonstrate this with an example based on

the weighted-set multi-cover problem in the Appendix. This shows that even for this class of problems it is relevant to consider the used decision proxy.

4 Sufficient Decision Proxies

In the previous section we saw that a sufficient condition for a proxy to be able to lead to an optimal decision is that it is expressive enough. In this section we present proxies that can adhere to this condition, when a deterministic proxy does not. We also discuss why an intuitive choice, using a parameterized distribution proxy, is most often not effective. Recall that in a contextual optimization problem, the input to the chosen proxy $x(\cdot)$ is the output of the predictive model $\phi_\theta(\cdot)$, so each proxy asks for a different output from the predictive model. Applying DFL we train this predictive model to get a decision-quality maximizing policy $\pi : Z \rightarrow X$, $\pi = x \circ \phi_\theta$ w.r.t. Equation (1), instead of aiming for the most accurate estimation of uncertain parameters c or distribution $\mathcal{C}_z$.

4.1 Scenario-Based Proxy

Based on Theorem 4 we know that in certain cases proxy $x^D(\cdot)$ is suboptimal. Since this deterministic variant is introduced to approximate decision problem of interest in Equation (1), we consider an alternative approximation based on the principle of *sample average approximation* (SAA) [Kleywegt *et al.*, 2002]:

$$x_n^S(\hat{c}_1, \dots, \hat{c}_n) = \argmin_{x \in X} \sum_{i=1}^n f(\hat{c}_i, x).$$

Note that $x^D(\cdot) = x_1^S(\cdot)$. SAA is used to approximately solve stochastic optimization problems, when the uncertain parameter distribution is known but does not lead to a solvable optimization problem (non closed-form objective). This approach works by sampling and subsequently solving the SAA. The SAA converges to the true stochastic optimization problem as the number of samples goes to infinity. In the contextual optimization setting this results in the following theorem:

Theorem 6 (Kim *et al.* [2015]). *Suppose $A \subset X$ is nonempty and compact, such that (i) $x^*(z) \in A$, (ii) for sufficiently large n , $x_n^S \in A$ and (iii) suppose that $f(c, \cdot)$ is continuous at x a.s. for any $x \in X$, and that there exists $\delta > 0$ such that the family of random variables $\{f(c, y) : \|y - x\| < \delta\}$ is uniformly integrable. Then $x_n^S(\cdot) \rightarrow x^*(z)$.*

Based on this theorem, as long as we pick n large enough we can learn to closely approximate the true distribution and get the optimal decision. In a DFL pipeline however, it would be too costly to pick a high value of n . Fortunately, the main benefit of DFL is that it can learn to produce high-quality decisions as long as the decision proxy is expressive enough. To achieve this, we predict the scenarios directly instead of sampling them as done in SAA. This results in the following predictor $\phi_\theta^{S_n} : Z \rightarrow C^n$, and scenario-based policy $\pi^{S_n} = x_n^S \circ \phi_\theta^{S_n}$. In the experimental results we see that with as few as $n = 2$ scenarios the scenario-based proxy leads to significantly better decisions than the deterministic proxy. This shows that DFL eliminates the need of using many samples. Figure 2 shows a schematic representation of the scenario-based proxy.




Figure 2: Scenario-based proxy: Making decisions based on multi-predicted scenarios, given features.

4.2 Quadratic Proxy

Instead of trying to approximate Equation (1) better, we can also look at the problem from the other perspective: Based on the earlier-developed theory we can design a proxy that has sufficient conditions that allow for optimal decision-making. We can do this by replacing the objective function $f(\cdot)$. First, we generalize the first part of Theorem 3:

Theorem 7. Consider a function $x(\cdot) : \Xi \rightarrow X$. An optimal prediction w.r.t. Equation (1) exists if $x(\cdot)$ is surjective, i.e., $\forall x \in X, \exists \hat{\xi} \in \Xi : x(\hat{\xi}) = x$.

Proof. Take an arbitrary z , with optimal decision $x^*(z)$. Given that $x(\cdot)$ is surjective, $\exists \hat{\xi} \in \Xi$ s.t. $x(\hat{\xi}) = x^*(z)$. $\square$

We now introduce a surjective function for this setting, which coincides with the *quadratic approximation* proposed by Vevurko *et al.* [2024] to deal with the zero-gradient problem and is the projection predictions on the decision space.

$$x^Q(\hat{\xi}) = \operatorname{argmin}_{x \in X} \|\hat{\xi} - x\|_2^2.$$

This proxy is surjective by design, as for arbitrary $x \in X$, taking $\hat{\xi} = x$ leads to $x^Q(\hat{\xi}) = x$. This leads us to using predictive model $\phi_\theta^Q : Z \rightarrow \Xi$, where $\Xi = \mathbb{R}^{\dim(X)}$, and quadratic policy $\pi^Q = x^Q \circ \phi_\theta^Q$. Figure 3 shows a schematic representation of the quadratic proxy.

We elaborate more on the rationale of this proxy $x^Q(\cdot)$: it has a strictly convex objective function, while adhering to the constraints of the problem (feasibility set X). This means we can train with DFL and always reach feasible decisions. Based on the assumption of the predictive model being a universal approximator, the objective function does not have to mimic the true objective function in Equation (1). This is because an optimal decision always lies in some minimum, and this minimum can be approximated by a quadratic function. The DFL pipeline should ensure the predictor learns to predict in the right area. In this case, this area is in the same space as the decisions, where the proxy projects the prediction in the feasibility space. *So effectively we are predicting directly from features to approximate decisions, without considering the uncertain parameters in between.* Most approaches in this category have some correction step to make the decision feasible [Kotary *et al.*, 2024]. In this case this can be considered to be the decision proxy itself.




Figure 3: Quadratic proxy: Making decisions based on predictions in the unconstrained decision space, given features.

4.3 Parameterized Distribution Proxy

Since we argue that deterministic proxies and therefore single-scenario predictions can be suboptimal, an intuitive approach would be to predict a distribution and use this to approximate $\mathcal{C}_z$ in Equation (1). This leads to the following decision proxy, with $\beta \in B$ the parameters of the chosen distribution $\mathbb{P}$ and $p_\beta(\cdot)$ its density function:

$$x^\mathbb{P}(\beta) = \operatorname{argmin}_{x \in X} \int_{\mathcal{C}} f(c, x) p_\beta(c) dc.$$

This leads to using predictive model $\phi_\theta^\mathbb{P} : Z \rightarrow B$ and policy $\pi^\mathbb{P} = x^\mathbb{P} \circ \phi_\theta^\mathbb{P}$. This approach is discussed as it is effective when the decision proxy has a closed-form, as in one of the experiments by Donti *et al.* [2017] where the underlying distribution is Gaussian. In practice however, the chosen distribution is unlikely to result in a closed-form integral, especially when problems are more complex, such as the class of problems we aim to tackle. It is still possible to approximate this proxy by discretizing the integral and using SAA. However, this indirect approach of first predicting a distribution and then extracting scenarios is a more restricted version of the scenario-based approach, always resulting in poorer performance. Therefore, we did not include the results.

4.4 Complexity

All presented policies have a different combination of predictors and decision proxies. Table 1 summarizes the policies by denoting the prediction space and the number of decision variables in a two-stage stochastic optimization setting. This number correlates highly with the number of problem constraints, as any constraint involving a second-stage decision is considered for each scenario. In general more decision variables and constraints increase the complexity of the decision proxy. Despite the focus here being on achieving high-quality policies, the complexity of the proxy is highly relevant as it is solved for each data point and epoch in a DFL pipeline.

The prediction space has less impact. The size of this space impacts the number of parameters of the predictive model. Since in general we (can) use neural networks with vastly more parameters than the dimension of the prediction space, this does not change much. What is more of an open question is whether the prediction space influences the complexity of learning a high-quality policy. In some problems $\dim(X)$ is

	policy	prediction space	decision var. #
deterministic	π^D	C	$d_y + d_x$
scenario-based	π^{S_n}	C^n	$nd_y + d_x$
quadratic	π^Q	$\mathbb{R}^{\dim(X)}$	$d_y + d_x$
distribution	π^P	B	$d_y + d_x$

Table 1: Summary of different policies, the output space of the predictor and the number of decision variables in the decision proxy. d_x and d_y denote the number of first- and second-stage variables.

significantly different from $\dim(C)$, which leads to the question if π^Q is harder to learn than π^{S_n} . We see small differences in the experiments that are in line with this hypothesis.

5 Experimental Evaluation

To experimentally evaluate the proposed approaches as alternatives to a deterministic proxy, we consider one continuous and two discrete problems. We use three baselines: deterministic proxy learned using DFL π^D , deterministic proxy learned using PFL π_{PFL}^D and a PFL learned model using SAA at inference $\pi_{PFL}^{S_n}$. In the PFL approaches we train on mean squared error, where $\pi_{PFL}^{S_n}$ uses the same learned predictive model as PFL ϕ_θ , but evaluates based on an SAA approximation x_n^S . This is done by using a residual-based distribution [Sadana *et al.*, 2025], using empirical additive errors [Deng and Sen, 2022]: taking the predictions as mean, and adding samples from the observed residuals from the validation set. We use $n = 16$, so $\pi_{PFL}^{S_{16}}$ in all problems. We compare the baselines with our proposed approaches: π^Q , π^{S_n} , with $n \in \{2, 8, 16\}$. For the continuous problem we apply DFL using differentiable convex layers [Agrawal *et al.*, 2019] as the problem is differentiable. For the other problems we use score function gradient estimation [Silvestri *et al.*, 2026], which approximates gradients by using stochastic sampling and is therefore applicable to any non-differentiable problem. Regarding the predictors we use neural networks with 2 hidden layers of size 256 and LeakyReLU activation functions, a learning rate of 5×10^{-4} and batch size of 32. Each approach has the same number of epochs, which we consider enough to converge. The best performing model on validation is used for test. Adam optimizer is used for gradient descent and Gurobi [Gurobi Optimization, 2025] as the solver. Code is available², based on PyDFLT [Schutte *et al.*, 2026]. For further experimental details please refer to the Appendix.

We consider problems where the deterministic proxy is not sufficient. Since literature so far has omitted these problems, we modified problems from literature to fit this requirement.

Portfolio optimization. As considered in earlier work on DFL [Wang *et al.*, 2020; Vevurko *et al.*, 2024], we consider a portfolio problem based on data from Quandl [2023]. Instead of balancing return and risk, we maximize the long-term expected value of the portfolio using the Kelly strategy [Kelly, 1956]. This makes the problem a multivariate version of Example 1. We determine the fixed rate of return β to be the

²<https://github.com/PyDFLT/PyDFLT/tree/sufficient-decision-proxies-ijcai2026>

median of all the fourth highest returns per training instance. Since data is limited (2898 data points) we do not use different data splits, instead we randomly pick 10 different securities for each seed, running 10 different seeds. We use a train, validation, test split of 70%, 15%, 15%.

Weighted-set multi-cover. We test on the WSMC problem as presented by Silvestri *et al.* [2026] with one adjustment: We allow removing single-item sets in the case of excess coverage for a partial return of costs, which makes the recourse similar to the common two-stage knapsack problem with uncertain weights [Kosuch and Lissner, 2011]. The deterministic proxy can be suboptimal in this case (see Appendix for example). We run 10 random seeds and for train, validation and test set sizes, we use 1000, 250 and 1250.

Probabilistic traveling salesperson. Finally we test on a two-stage probabilistic traveling salesperson (PTSP). Probabilistic vehicle routing problems have been studied since Tillman [1969], with PTSP specifically introduced by Jaillet, 1985], and are still highly relevant [Oyola *et al.*, 2018]. The goal of the PTSP is to plan a route along a known set of customers while minimizing transportation costs, but it is uncertain which customers require service. Missing a customer who requires service incurs a penalty (based on going back-and-forth to the depot, which is a typical recourse choice [Oyola *et al.*, 2018]). In this two-stage version, one is allowed to add direct trips to customers, similar to allowing crowdsourcing Santini *et al.* [2022]. After it becomes known which customers require service, these direct trips can be canceled in the second stage if they are not required. This makes the problem such that in a given scenario, direct trips are always dominated, but given uncertainty they have the benefit of incurring low costs if canceled. We run 10 seeds and for train, validation and test set sizes, we use 1000, 250 and 1250.

5.1 Results

Figure 4 shows the final results, while Figure 5 visualizes the learning. We use a paired t -test (5% significance level) to compare our proposed methods with the deterministic proxy baseline, and they performed significantly better in all cases.




Figure 4: Test set mean absolute regret normalized by DFL baseline π^D . Comparing with PFL baselines and proposed π^{S_n} , π^Q . Dashes denote individual runs. π_{PFL}^D bars for WSMC (2.08) and PTSP (2.63) were cut off for visibility. Due to effectiveness of $n = 2$ for WSMC, $n = 8$ for PTSP, and runtime, higher values of n were not included.




Figure 5: Validation curves per epoch (x -axis). The mean of all approaches’ minimum absolute regret was used to normalize. Error bars denote one standard deviation. PFL-approaches do not record absolute regret, but π_{PFL}^D was included for comparison.

We see that π^Q performs best or second best in all cases. In both WSMC and PTSP the dimension of decision variables is greater than the dimension of uncertain variables. Since this makes π^Q learn the same relationship through a mapping with higher dimension, it could be the reason the scenario-based approach performs slightly better. In general we see that more scenarios for the scenario-based approach results in better performance. For portfolio, Figure 5a shows that π^D is not able to learn at all. This is primarily caused by having zero-gradients. A higher number of scenarios here results in more effective gradients. The most notable part about the WSMC is that the learning curves have high variance. This is inherent to the score function gradient estimation approach. For PTSP this is not as pronounced, due to the small problem size and therefore less complex decision landscape. However, for PTSP we do observe higher variance in the test results for π^Q and π_2^S , suggesting problem-dependent lower robustness.

6 Discussion and Related Work

We formally showed cases when deterministic proxies are suboptimal, even in a DFL setting, and introduce effective alternative proxies. This is another step in understanding the applicability of DFL, after Cameron *et al.* [2022] consider when PFL is worse compared to DFL and Homem-de Mello *et al.* [2026] show for a certain class of problems that an optimal single-scenario exists. Similarly, Schutte *et al.* [2024] discuss the importance of the decision proxy by highlighting the difference between the empirical (deterministic) problem and the true stochastic problem, and introduce loss functions that generalize better. Furthermore, [Wilder *et al.*, 2019] introduced a continuous relaxation proxy to speed up DFL, using a quadratic regularization term on the decision vector.

There are several works that consider the true problem as stochastic and use distribution-based decision proxies. These works differ from our work in that we provide theoretical backing on when a deterministic proxy is suboptimal, as well as propose alternatives based on this theory that keep complexity low. Donti *et al.* [2017] tackle a problem that has a differentiable formulation when assuming the uncertainty is Gaussian. Alternatively they assume the uncertain parameter

has finite discrete support, having the predictive model output probabilities for these discrete scenarios. This is a similar assumption as done by Qi *et al.* [2025], which is quite limiting as uncertain parameters are often continuous and multi-dimensional, requiring to decide on some discretization that is representative without having too many scenarios. Our scenario-based approach learns representative scenarios and only requires to decide the number of scenarios up front. Elmachtoub *et al.* [2025] also include a distribution-based approach, assuming a parameterized family of distributions. They show that when the distribution is misspecified, distributional DFL (contextual integrated-estimation-optimization) outperforms distributional PFL (estimate-then-optimize), while in a well-specified distribution case the result is opposite. Jia *et al.* [2025] use scenarios in a predict-then-optimize routing problem, but do not train end-to-end.

Stochasticity is also used to obtain gradients when they do not exist or are zero, but in these cases the deterministic proxy is applied at inference [Berthet *et al.*, 2020; Silvestri *et al.*, 2026]. The same is done at inference by Hu *et al.* [2023], who consider uncertainty in the constraints and model infeasibility using penalties. We refer the reader to [Qi and Shen, 2022] for an operations management perspective on DFL, and [Mandi *et al.*, 2024] for a general DFL survey.

7 Conclusions

A prevailing assumption in decision-focused learning is that there exists a single-scenario problem approximation that allows for optimal decision-making. While this is a valid assumption in a wide class of problems, this paper investigates for the first time general theoretical properties of problems for which this assumption is violated. Based on these properties we derive requirements for problem approximations that are sufficient to reach optimal decisions. On three problems we empirically showed the value of the proposed approaches. Future work includes studying problems where parameterized distributions can easily be assumed and problems from practice, to evaluate the applicability of DFL and different proxies. Given the strong performance of the quadratic proxy, it raises the question how relevant the true objective is in DFL.

Acknowledgments

We thank the reviewers for their suggestions. This work was partially supported by Epistemic AI and by TAILOR, both funded by EU Horizon 2020 under grants 964505 and 952215 respectively.

References

- [Agrawal *et al.*, 2019] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J. Zico Kolter. Differentiable convex optimization layers. *Advances in Neural Information Processing Systems*, 32, 2019.
- [Bengio *et al.*, 2021] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.
- [Berthet *et al.*, 2020] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. *Advances in Neural Information Processing Systems*, 33:9508–9519, 2020.
- [Birge and Louveaux, 2011] John R. Birge and François Louveaux. *Introduction to Stochastic Programming*. Springer, 2011.
- [Birge, 1982] John R. Birge. The value of the stochastic solution in stochastic linear programs with fixed recourse. *Mathematical Programming*, 24(1):314–325, 1982.
- [Cameron *et al.*, 2022] Chris Cameron, Jason Hartford, Taylor Lundy, and Kevin Leyton-Brown. The perils of learning before optimizing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 3708–3715, 2022.
- [Charnes and Cooper, 1962] A. Charnes and W. W. Cooper. Chance constraints and normal deviates. *Journal of the American Statistical Association*, 57(297):134–148, 1962.
- [Deng and Sen, 2022] Yunxiao Deng and Suvrajeet Sen. Predictive stochastic programming. *Computational Management Science*, 19(1):65–98, 2022.
- [Donti *et al.*, 2017] Priya Donti, Brandon Amos, and J. Zico Kolter. Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, 30:5484–5494, 2017.
- [Elmachtoub and Grigas, 2022] Adam N. Elmachtoub and Paul Grigas. Smart “predict, then optimize”. *Management Science*, 68(1):9–26, 2022.
- [Elmachtoub *et al.*, 2025] Adam N. Elmachtoub, Henry Lam, Haofeng Zhang, and Yunfan Zhao. Estimate-then-optimize versus integrated-estimation-optimization versus sample average approximation: A stochastic dominance perspective. *arXiv preprint arXiv:2304.06833*, 2025.
- [Gurobi Optimization, 2025] Gurobi Optimization. Gurobi Optimizer reference manual (12.0.0). <https://docs.gurobi.com>, 2025.
- [Homem-de Mello *et al.*, 2026] Tito Homem-de Mello, Juan Valencia, Felipe Lagos, and Guido Lagos. Forecasting outside the box: application-driven optimal pointwise forecasts for a class of two-stage stochastic programs: T. homem-de-mello et al. *Mathematical Programming*, pages 1–55, 2026.
- [Hu *et al.*, 2023] Xinyi Hu, Jasper Lee, and Jimmy Lee. Two-stage predict+optimize for mixed integer linear programs with unknown parameters in constraints. *Advances in Neural Information Processing Systems*, 36:14247–14272, 2023.
- [Jaillet, 1985] Patrick Jaillet. *Probabilistic traveling salesman problems*. PhD thesis, Massachusetts Institute of Technology, 1985.
- [Jia *et al.*, 2025] Menglei Jia, Albert H. Schrottenboer, and Feng Chen. Scenario predict-then-optimize for data-driven online inventory routing. *Transportation Science*, 59(5):1032–1056, 2025.
- [Kelly, 1956] John L. Kelly. A new interpretation of information rate. *The Bell System Technical Journal*, 35(4):917–926, 1956.
- [Kibzun and Kan, 1996] Andrey I. Kibzun and Yuri S. Kan. *Stochastic programming problems with probability and quantile functions*. John Wiley and Sons, 1996.
- [Kim *et al.*, 2015] Sujin Kim, Raghu Pasupathy, and Shane G. Henderson. A guide to sample average approximation. In *Handbook of simulation optimization*, pages 207–243. Springer, 2015.
- [Kleywegt *et al.*, 2002] Anton J. Kleywegt, Alexander Shapiro, and Tito Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002.
- [Kosuch and Lisser, 2011] S. Kosuch and A. Lisser. On two-stage stochastic knapsack problems. *Discrete Applied Mathematics*, 159(16):1827–1841, 2011.
- [Kotary *et al.*, 2024] James Kotary, Vincenzo Di Vito, Jacob Christopher, Pascal Van Hentenryck, and Ferdinando Fioretto. Learning joint models of prediction and optimization. In *ECAI 2024*, pages 2476–2483. IOS Press, 2024.
- [Mandi *et al.*, 2024] Jayanta Mandi, James Kotary, Senne Berden, Maxime Mulamba, Victor Bucarey, Tias Guns, and Ferdinando Fioretto. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research*, 80:1623–1701, 2024.
- [Oyola *et al.*, 2018] Jorge Oyola, Halvard Arntzen, and David L. Woodruff. The stochastic vehicle routing problem, a literature review, part I: models. *EURO Journal on Transportation and Logistics*, 7(3):193–221, 2018.
- [Qi and Shen, 2022] Meng Qi and Zuo Jun Max Shen. Integrating prediction/estimation and optimization with applications in operations management. In *Tutorials in oper-*

- ations research: emerging and impactful topics in operations*, pages 36–58. INFORMS, 2022.
- [Qi *et al.*, 2025] Meng Qi, Paul Grigas, and Zuo-Jun (Max) Shen. Integrated conditional estimation-optimization. *Operations Research*, 2025.
- [Quandl, 2023] Quandl. WIKI various end-of-day data. <https://www.quandl.com/data/WIKI>, 2023.
- [Sadana *et al.*, 2025] Utsav Sadana, Abhilash Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 320(2):271–289, 2025.
- [Santini *et al.*, 2022] Alberto Santini, Ana Viana, Xenia Klimentova, and João Pedro Pedroso. The probabilistic travelling salesman problem with crowdsourcing. *Computers & Operations Research*, 142:105722, 2022.
- [Schutte *et al.*, 2024] Noah Schutte, Krzysztof Postek, and Neil Yorke-Smith. Robust losses for decision-focused learning. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI’24*, pages 4868–4875, 2024.
- [Schutte *et al.*, 2026] Noah Schutte, Kim van den Houten, and Grigori Vevurko. PyDFLT: A Python-based decision-focused learning toolbox (0.1.1). <https://doi.org/10.5281/zenodo.20177131>, 2026.
- [Silvestri *et al.*, 2026] Mattia Silvestri, Senne Berden, Jayanta Mandi, Ali İrfan Mahmutoğulları, Maxime Mulamba, Allegra De Filippo, Tias Guns, and Michele Lombardi. Score function gradient estimation to widen the applicability of decision-focused learning. *Journal of Artificial Intelligence Research*, 85:289–330, 2026.
- [Tillman, 1969] Frank A. Tillman. The multiple terminal delivery problem with probabilistic demands. *Transportation Science*, 3(3):192–204, 1969.
- [van den Houten *et al.*, 2024] Kim van den Houten, David M. J. Tax, Esteban Freydehl, and Mathijs de Weerdt. Learning from scenarios for repairable stochastic scheduling. In *Proceedings of the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 234–242, 2024.
- [Vevurko *et al.*, 2024] Grigori Vevurko, Wendelin Böhmer, and Mathijs de Weerdt. You shall pass: Dealing with the zero-gradient problem in predict and optimize for convex optimization. *arXiv preprint arXiv:2307.16304*, 2024.
- [Wang *et al.*, 2020] Kai Wang, Bryan Wilder, Andrew Perrault, and Milind Tambe. Automatically learning compact quality-aware surrogates for optimization problems. *Advances in Neural Information Processing Systems*, 33:9586–9596, 2020.
- [Wilder *et al.*, 2019] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1658–1665, 2019.