

New Measures for Shaping Trajectories in Dynamic Optimization

Joshua L. Pulsipher[†], Benjamin R. Davidson[‡], and Victor M. Zavala[‡]

[†]*Department of Chemical Engineering,
Carnegie Mellon University, Pittsburgh, PA 15213 USA
(e-mail: jpulsiph@andrew.cmu.edu)*

[‡]*Department of Chemical and Biological Engineering,
University of Wisconsin-Madison, Madison, WI 53706 USA.*

Abstract: We propose a new class of measures for shaping time-dependent trajectories in dynamic optimization (DO). The proposed measures are analogous to risk measures used in stochastic optimization (SO) and are inspired by a recently-proposed unifying abstraction for infinite-dimensional optimization. Risk measures are summarizing statistics (e.g., average, variance, quantiles, worst-case values) that are used to shape the probability density of random objectives and constraints. We show that this extensive collection of measures can be applied in DO for computing and manipulating interesting features of time-dependent trajectories (e.g., excursion costs and quantiles). We also discuss how to implement these measures in the Julia modeling package `InfiniteOpt.jl`.

Keywords: Dynamic Optimization, Measures, Trajectories, Infinite-Dimensional Optimization

1. INTRODUCTION

We consider the continuous-time dynamic optimization (DO) problem:

$$\begin{aligned} \min_{y(\cdot) \in \mathcal{Y}} \quad & M_t f(\dot{y}(t), y(t), t) \\ \text{s.t.} \quad & g(\dot{y}(t), y(t), t) \leq 0, \quad t \in \mathcal{D}_t. \end{aligned} \quad (1)$$

This representation follows from a recently-proposed abstraction that unifies infinite-dimensional optimization (InfiniteOpt) problems (e.g., dynamic/stochastic/PDE optimization) (Pulsipher et al., 2022). In the context of DO, the infinite domain (which terminology arises from the infinite parameter modeling constructs presented in the unifying abstraction) \mathcal{D}_t is the time domain $[t_0, t_f]$, $y(t) \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$ are time-valued decision functions (e.g., state/control variables), $f(\cdot)$ is an infinite-dimensional cost function, and $g(\cdot)$ is a vector-valued infinite-dimensional constraint functions $g_j(\cdot)$, $j \in \mathcal{J} \subseteq \mathbb{R}^{n_g}$. The use of derivative variables $\dot{y}(t) \in \mathbb{R}^{n_y}$ (which denote $dy(t)/dt$) is done for convenience in notation such that these can be augmented via auxiliary state variables, as is common to DAE modeling approaches, to enhance the generality of Problem (1) (allowing for constraints on derivative values and not requiring explicit ODEs) (Biegler, 2010).

The measure operator $M_t : \mathcal{D}_t \mapsto \mathbb{R}$ is the focus of this work; this operator seeks to scalarize infinite-dimensional cost functions (summarizing them over the time domain \mathcal{D}_t) to form a well-posed objective function. Such measures can also be used to handle constraints, but here we focus on objectives to simplify the presentation. Problem (1) is general and captures a wide range of DO problems (e.g.,

model predictive control and state/parameter estimation) where $y(\cdot)$ consists of state/control/parameter variables and $g(\cdot) \leq 0$ can include DAE/path/point constraints.

Classical DO formulations minimize the integral of the cost trajectory (known as the Bolza objective). In other words, these formulations use the measure operator:

$$M_t f(t) = \int_{t \in \mathcal{D}_t} f(t) dt \quad (2)$$

where we write $f(t) := f(\dot{y}(t), y(t), t)$ for convenience. Minimizing (2) has the effect of *uniformly* shaping the cost trajectory $f(t)$ over the domain \mathcal{D}_t . This also amounts to minimizing the total cost (which is equivalent to the average cost with a normalization factor of $(t_f - t_0)^{-1}$). However, we can envision more advanced DO formulations that aim to shape the cost trajectory $f(\cdot)$ with other measures (e.g., peak or excursion costs). For example, Risbeck and Rawlings recently proposed an MPC objective that penalizes the total and peak costs (Risbeck and Rawlings, 2019). Minimizing the peak of a cost-trajectory is often a desirable feature in systems, as the peak can be associated with unsafe behavior and/or economic penalties.

Under the InfiniteOpt representation, Problem (1) is analogous to the SO problem:

$$\begin{aligned} \min_{z \in \mathcal{Z}, y(\cdot) \in \mathcal{Y}} \quad & M_\xi f(z, y(\xi), \xi) \\ \text{s.t.} \quad & g(z, y(\xi), \xi) \leq 0, \quad \xi \in \mathcal{D}_\xi. \end{aligned} \quad (3)$$

where $z \in \mathcal{Z} \subseteq \mathbb{R}^{n_z}$ are here-and-now decision variables, $\xi : \Omega \mapsto \mathcal{D}_\xi$ is a random parameter in accordance with the probability space (Ω, \mathcal{F}, F) , $\mathcal{D}_\xi \subseteq \mathbb{R}^{n_\xi}$ is the co-domain of its distribution, and $y(\xi) \in \mathcal{Y} \subseteq \mathbb{R}^{n_y}$ are recourse decision variables (Birge and Louveaux, 2011). We observe that (1) is a special case of (3) if we define $n_z = 0$, $n_\xi = 1$, $\mathcal{D}_\xi = \mathcal{D}_t$,

* This work was supported by the U.S. Department of Energy under grant DE-SC0014114.

$M_\xi = M_t$, and set $y(\xi)$ to contain both $\dot{y}(t)$ and $y(t)$. This equivalence might not be immediately obvious since $y(\xi)$ is not differentiable in general, but this special case corresponds to $|\Omega| = 1$ such that ξ becomes deterministic.

In SO, the measure M_ξ can be drawn from a wide collection of risk measures (summarizing statistics) that aim to shape the probability density function of $f(\xi)$; some examples include the expected value, mean-variance, absolute deviation, quantiles, and conditional-value-at-risk (CVaR) (Ruszczyński and Shapiro, 2006). Risk measures enable greater flexibility in shaping the cost function $f(\xi)$ and are commonly used to penalize extreme events (i.e., high costs). By leveraging the connection between (1) and (3), we have recently shown that the expected value and CVaR measures can be used as measures in DO to shape the time-dependent cost trajectory (Pulsipher et al., 2022).

In this work, we formalize the use of risk measures to DO problems that follow (1). The contributions that distinguish this work from that presented in (Pulsipher et al., 2022) are:

- the consideration of additional risk measures (e.g., mean-variance and disutility)
- the formal demonstration of time-valued measures as special cases of their SO counterparts
- the establishment of properties for risk measures in a DO context
- the more complete empirical demonstration of diverse time-valued measures
- the identification of avenues of future research for this new class of dynamic measures.

This new class of dynamic measures M_t enable us to shape trajectories and policies in new and interesting ways. Moreover, we highlight how the Julia-based modeling package `InfiniteOpt.jl` provides a useful interface to compactly express such measures.

2. MAIN RESULTS

2.1 Time-Valued Density Functions

We begin by establishing time analogues of probability density functions (pdf) and of cumulative density functions (cdf) typically used in SO (i.e., pdfs and cdfs defined over the time domain \mathcal{D}_t). These provide key constructs to facilitate the use of risk measures in DO.

In SO, ξ is described by the pdf $p_\xi : \mathcal{D}_\xi \mapsto \mathbb{R}_{\geq 0}$, which satisfies $\int_{\xi \in \mathcal{D}_\xi} p_\xi(\xi) d\xi = 1$. The pdf is used to compute statistics such as the expectation:

$$\mathbb{E}_\xi[f(\xi)] = \int_{\xi \in \mathcal{D}_\xi} f(\xi) p_\xi(\xi) d\xi. \quad (4)$$

This measure $M_\xi = \mathbb{E}_\xi$ is commonly used in SO; here, we can see that $p_\xi(\cdot)$ acts as a weighting function that places varied emphasis over the domain \mathcal{D}_ξ . The expectation summarizes the cost function in a single scalar value.

Following the analogy of SO and DO, we define a weighting function $p_t : \mathcal{D}_t \mapsto \mathbb{R}_{\geq 0}$ (i.e., a time-valued pdf) to yield the time average (expectation):

$$\mathbb{E}_t[f(t)] := \int_{t \in \mathcal{D}_t} f(t) p_t(t) dt. \quad (5)$$

This provides us flexibility in prioritizing different regimes in the domain \mathcal{D}_t . We obtain (2) as a special case by setting $p(t) = 1$; however, we can envision choosing from a wide range of candidate pdfs, such as:

$$p_t(t) = \frac{1}{t_f - t_0} \quad (6)$$

This is the pdf of a uniform random variable and places equal emphasis on different parts of the time domain. This inspires the consideration of other weighting functions, such as the exponential pdf:

$$p_t(t) = \gamma e^{-\gamma t} \quad (7)$$

where $\gamma \in \mathbb{R}_{>0}$ is the decay rate parameter. Using this weighting function to compute the time average places emphasis on the beginning of the domain and quickly decays over time. In other words, this can be seen as a discount factor (Shin and Zavala, 2020).

The cdf $P : \mathcal{D}_\xi \mapsto [0, 1]$ of a random variable ξ is:

$$P(\xi; \hat{\xi}) := \mathbb{P}_\xi(\xi \leq \hat{\xi}) = \int_{\xi \in \{\xi \in \mathcal{D}_\xi : \xi \leq \hat{\xi}\}} p_\xi(\xi) d\xi \quad (8)$$

This denotes the cumulative probability of finding ξ below the threshold $\hat{\xi} \in \mathcal{D}_\xi$. In the context of the DO, we can write the cdf of the time trajectory $f(t)$ by using the *excursion sets*:

$$\begin{aligned} \mathcal{D}_t^+(f(t); \hat{f}) &:= \{t \in \mathcal{D}_t : f(t) \geq \hat{f}\} \\ \mathcal{D}_t^-(f(t); \hat{f}) &:= \{t \in \mathcal{D}_t : f(t) \leq \hat{f}\} \end{aligned} \quad (9)$$

where $\mathcal{D}_t^+(f(t); \hat{f}) \subseteq \mathcal{D}_t$ and $\mathcal{D}_t^-(f(t); \hat{f}) \subseteq \mathcal{D}_t$ are the positive and negative function excursion sets, respectively. We can use the negative excursion set to establish:

$$P(f(t); \hat{f}) = \int_{t \in \mathcal{D}_t^-(f(t); \hat{f})} p_t(t) dt. \quad (10)$$

In a DO context, the cdf measures the fraction of time that the trajectory $f(t)$ is below a threshold \hat{f} . We will see that expressing the cdf in terms of excursion set allows us to interpret measures as mechanisms to bound trajectories.

2.2 Dynamic Measures

Here we illustrate the interpretation of risk measures in a DO context. Numerous risk measures have been proposed in the SO community and analyzing them comprehensively is beyond the scope of this work (we refer the reader to (Krokhmal et al., 2013) for a review). This section establishes key constructs and exemplifies the steps needed to interpret risk measures in a time setting. We discuss properties of these proposed measures in Section 2.3.

Expectation The time expectation measure \mathbb{E}_t shown in Equation (5) allows us to assess a weighted average of our cost trajectory $f(t)$ in accordance with the weighting function $p_t(t)$. In the context of DO, minimizing the expected cost provides an easily interpretable objective and the key modeling choice lies in the selection of $p_t(t)$. Reasonable candidates for many DO applications are (6) and (7), as demonstrated in Section 3; however, other weighting functions are possible (e.g., Gaussian and Gamma). The selection of the weighting function dictates how much emphasis is placed on different parts of the time domain. The expectation serves as a core construct in defining the more sophisticated measures.

Proposition 1. The time expectation \mathbb{E}_t in (5) is a special case of \mathbb{E}_ξ in (4) if $n_z = 0$, $n_\xi = 1$, $\mathcal{D}_\xi = \mathcal{D}_t$.

Mean-Variance The mean-variance $\mathbb{E}\text{-}\mathbb{V}_\xi$ is a classical measure used in portfolio optimization (Leland, 1999). For SO problems (3), this minimizes the variance (i.e., spread) of the cost outcomes in an attempt to mitigate high-cost events:

$$\mathbb{E}\text{-}\mathbb{V}_\xi[f(\xi)] := \mathbb{E}_\xi[f(\xi)] + \lambda \mathbb{V}_\xi[f(\xi)] \quad (11)$$

where $\mathbb{V}_\xi = \mathbb{E}_\xi[(f(\xi) - \mathbb{E}_\xi[f(\xi)])^2]$ is the variance and $\lambda \in \mathbb{R}_{\geq 0}$ is a tradeoff parameter. Transferring this to a DO setting we obtain the time-valued measure:

$$\mathbb{E}\text{-}\mathbb{V}_t[f(t)] := \mathbb{E}_t[f(t)] + \lambda \mathbb{V}_t[f(t)]. \quad (12)$$

Minimizing $\mathbb{E}\text{-}\mathbb{V}_t[f(t)]$ in a DO problem provides a tradeoff problem that seeks to minimize the magnitude of the cost trajectory (the expectation) and the variability/fluctuations (variance) of the cost trajectory. A disadvantage of this measure is that it penalizes cost variability equally for low and high costs (symmetric penalty). However, this property can be advantageous for cost functions that seek to enforce smooth control trajectories. This bares relation to the classical DO practice of using a quadratic penalty terms to dampen oscillation; however, we note that (12) is more general since it can embed an arbitrary pdf.

Proposition 2. The measure operator $\mathbb{E}\text{-}\mathbb{V}_t$ from (12) is a special case of $\mathbb{E}\text{-}\mathbb{V}_\xi$ from (11) under the same conditions of Proposition 1.

Quantile The quantile $Q_\xi(f(\xi); \alpha)$ (also referred to as the value-at-risk) denotes the threshold value \hat{f} for $f(\xi)$ such that the cumulative probability of incurring costs below the threshold is at least $\alpha \in [0, 1]$:

$$Q_\xi(f(\xi); \alpha) := \inf_{\hat{f} \in \mathbb{R}} \left\{ P(f(\xi); \hat{f}) \geq \alpha \right\}. \quad (13)$$

Constraining (13) is equivalent to enforcing a probabilistic constraint (Sarykalin et al., 2008):

$$Q_\xi(f(\xi); \alpha) \leq 0 \iff \mathbb{P}_\xi(f(\xi) \leq 0) \geq \alpha. \quad (14)$$

We can use the cdf (10) in combination with (13) to define the time-valued quantile:

$$Q_t(f(t); \alpha) := \inf_{\hat{f} \in \mathbb{R}} \left\{ \int_{t \in \mathcal{D}_t^-(f(t); \hat{f})} p_t(t) dt \geq \alpha \right\}. \quad (15)$$

Using this measure in (1) minimizes the excursion threshold of the cost trajectory such that the fraction of time that exceed it is no more than $1 - \alpha$. Unlike the $\mathbb{E}\text{-}\mathbb{V}_t$ measure, the quantile measure only penalizes high cost values, making it an attractive alternative in certain cases. However, the potential disadvantages of this measure are that it does not strongly discourage high cost peaks in the positive function excursion $\mathcal{D}_t^+(f(t); Q_t(f(t); \alpha))$, and it is nonconvex and difficult to compute in general.

Proposition 3. The quantile measure Q_t (15) is a special case of its analogue Q_ξ in (13) under the same conditions of Proposition 1.

Conditional-Value-at-Risk The conditional-value-at-risk (CVaR) measure seeks to address the limitations of the quantile measure Q_t by penalizing the expected value of the $1 - \alpha$ largest cost values:

$$\text{CVaR}_\xi(f(\xi); \alpha) := \min_{\hat{f} \in \mathbb{R}} \left\{ \hat{f} + \frac{1}{1 - \alpha} \mathbb{E}_\xi[f(\xi) - \hat{f}]_+ \right\} \quad (16)$$

where $\mathbb{E}_\xi[f(\xi) - \hat{f}]_+ := \mathbb{E}_\xi[\max(0, f(\xi) - \hat{f})]$ and $\alpha \in [0, 1]$. CVaR is also known as the superquantile; under mild assumptions, CVaR can be represented as:

$$\text{CVaR}_\xi(f(\xi); \alpha) = \mathbb{E}_\xi[f(\xi) : f(\xi) \geq Q_t(f(t); \alpha)] \quad (17)$$

since the minimizer \hat{f}^* is $Q_t(f(t); \alpha)$ (Rockafellar et al., 2000). With this observation, the time-valued CVaR measure can be expressed as:

$$\text{CVaR}_t(f(t); \alpha) := \min_{\hat{f} \in \mathbb{R}} \left\{ \hat{f} + \frac{1}{1 - \alpha} \mathbb{E}_t[f(t) - \hat{f}]_+ \right\}. \quad (18)$$

This provides a convex measure that penalizes the high (peak) costs incurred in the positive function excursion set $\mathcal{D}_t^+(f(t); Q_t(f(t); \alpha))$. Note that this penalizes *multiple* peak costs (and not just the peak cost, as done in typical DO formulation). Moreover, one can show that:

$$\begin{aligned} \lim_{\alpha \rightarrow 0} \text{CVaR}_t(f(t); \alpha) &= \mathbb{E}_t[f(t)] \\ \lim_{\alpha \rightarrow 1} \text{CVaR}_t(f(t); \alpha) &= \max_{t \in \mathcal{D}_t} f(t) \end{aligned} \quad (19)$$

which both follow from (17). As such, CVaR is highly versatile measure for use in DO that can capture both average and extreme features of a time trajectory.

Proposition 4. The CVaR measure in (18) is a special case of CVaR_ξ in (16) under the same conditions of Proposition 1.

Disutility Disutility risk measures are another prevalent measure class used in SO; this employs an expectation over a disutility function $s : \mathbb{R} \mapsto \mathbb{R}$ (typically a convex increasing function) that penalizes unfavorable values of $f(\xi)$:

$$D_\xi(f(\xi)) := \mathbb{E}_\xi[s(f(\xi))]. \quad (20)$$

In an effort to define a translation invariant measure, the measure is often expressed as:

$$\tilde{D}_\xi(f(\xi)) := \inf_{\hat{f} \in \mathbb{R}} \mathbb{E}_\xi[f(\xi) + s(f(\xi) - \hat{f})]. \quad (21)$$

One can show that $\text{CVaR}_\xi(f(\xi); \alpha)$ is a special case of (21). This follows by letting $s(x; \alpha) = (1 - \alpha)^{-1} \max(0, x) - x$ where $x \in \mathbb{R}$. Then by substituting $s(x; \alpha)$ in (21) we obtain (16). We transfer (21) to DO by using the time-valued expectation \mathbb{E}_t :

$$\tilde{D}_t(f(t)) := \inf_{\hat{f} \in \mathbb{R}} \mathbb{E}_t[f(t) + s(f(t) - \hat{f})]. \quad (22)$$

This measure class provides great flexibility in shaping dynamic trajectories as there are diverse choices of $s(\cdot)$ (in addition to the flexibility provided via selecting the weighting function $p_t(\cdot)$). A useful survey on the properties of disutility functions in the context of SO is provided in (Fulga, 2016).

Proposition 5. The time-valued disutility measure $\tilde{D}_t(f(t))$ is a special case of $\tilde{D}_\xi(f(\xi))$ under the same conditions of Proposition 1.

2.3 Measure Properties

Here we formalize some key mathematical properties of the dynamic measures presented in Section 2.2. This provides some interesting and useful insights on the behavior that these measures induce. These properties have been studied in the SO community, and we will show that this rich theory can be readily applied to DO.

In the context of SO, four main properties are typically considered for risk measures: convexity, monotonicity, translation invariance, and positive homogeneity. Moreover, a measure operator is said to be coherent if it satisfies all these properties (Artzner et al., 1999; Ruszczyński and Shapiro, 2006).

Convexity asserts that a measure operator M_ξ satisfy:

$$M_\xi(\beta f + (1 - \beta)h) \leq \beta M_\xi(f) + (1 - \beta)M_\xi(h) \quad (23)$$

for all measurable functions $f(\xi), h(\xi) : \mathcal{D}_\xi \mapsto \mathbb{R}$ in the linear function space \mathcal{F} and all $\beta \in [0, 1]$. This property is key for creating optimization objectives that are well-posed and guarantees that the measure of a convex cost is also convex.

Under monotonicity, we have that if $f_1(\xi) \succeq f_2(\xi)$ ($f_1(\xi)$ dominates $f_2(\xi)$), then the measure M_ξ satisfies:

$$M_\xi(f_1(\xi)) \geq M_\xi(f_2(\xi)). \quad (24)$$

This ensures that, if a cost function dominates another cost function, then the measure former will also be greater than that of the latter. The concept of dominance (comparing whether a random variable is better than another random variable) is an interesting and important concept that has not been explored in DO. In a DO context, dominance of first-order ($f_1(t) \succeq f_2(t)$) requires that $P(f_1(t) > \hat{f}) \geq P(f_2(t) > \hat{f})$ for any threshold value \hat{f} . In other words, the fraction of time that the trajectory $f_1(t)$ remains above the threshold is greater or equal than the fraction of time that the trajectory $f_2(t)$ remains above the same threshold. A monotonic measure is such that, if $f_1(t) \succeq f_2(t)$, then $M_t(f_1(t)) \geq M_t(f_2(t))$ holds. Note that dominance holds trivially if $f_1(t) \geq f_2(t)$ for all $t \in \mathcal{D}_t$. These concepts are important because comparisons (benchmarks) of time trajectories are not as straightforward (Renteria et al., 2018), as the trajectories are functions (not scalar values) and thus a trajectory might be better in some parts of the time domain but not in others

A translation invariant measure satisfies:

$$M_\xi(f(\xi) + a) = M_\xi(f(\xi)) + a \quad (25)$$

if $a \in \mathbb{R}$ and $f(\xi) \in \mathcal{F}$. In a DO context, this property ensures that offsetting the cost function will not change the shape of the optimal cost trajectory.

The positive homogeneity property is given by:

$$M_\xi(\tau f(\xi)) = \tau M_\xi(f(\xi)) \quad (26)$$

if $\tau > 0$ and $f(\xi) \in \mathcal{F}$. In the context of DO, a positive homogeneous measure provides the property that uniformly scaling the cost by τ will not affect the shape of the optimal cost trajectory.

The analysis of the stochastic risk measures featured in Section 2.2 is well-established in the SO literature and Table 1 provides a summary of these (Artzner et al., 1999; Ruszczyński and Shapiro, 2006).

Because time-valued measures are special cases of the SO counterparts, they inherit the properties of Table 1. This illustrates how we can transfer rich theory from SO (with respect to these measure operators) to a DO context. The time-valued expectation measure is a coherent risk measure and this might explain why this has been the classical measure used in DO. It is particularly important to observe that convexity ensures that the use of this

M_ξ	(23)	(24)	(25)	(26)
\mathbb{E}_ξ	Yes	Yes	Yes	Yes
$\mathbb{E}\text{-}\mathbb{V}_\xi$	Yes	No	Yes	No
Q_ξ	No	Yes	Yes	Yes
CVaR_ξ	Yes	Yes	Yes	Yes
\tilde{D}_ξ	Yes	Yes	Yes	Yes

Table 1. A summary of the properties satisfied by certain measures M_ξ . Note that \tilde{D}_ξ only satisfies (26) if $s(\cdot)$ is positive homogeneous.

measure yields a convex objective if the cost function is convex. Interestingly, the convexity of the objective has key implications for establishing stability conditions (e.g., closed-loop stability of MPC) (Rawlings et al., 2017). From this, we observe that other non-convex measures such as $\mathbb{E}\text{-}\mathbb{V}_t$ and Q_t may not yield stability. On the other hand, CVaR_t and \tilde{D}_t are convex and thus might inherit stability properties (this is an interesting topic of future work).

2.4 Modeling in *InfiniteOpt.jl*

The unifying abstraction for infinite-optimization (which facilitated the connection between Problems (1) and (3)) is implemented in a Julia package called *InfiniteOpt.jl* (Pulsipher et al., 2022). This enables us to intuitively model continuous DO problems following a simple symbolic syntax (see Code Snippet 1). It is measure-centric and readily enables to quickly implement new candidate measure operators (such as those proposed in Section 2.2). Moreover, the unifying abstraction behind *InfiniteOpt.jl* facilitates the incorporation of random constructs and/or PDE constraints. By default, these models are solved via direct transcription, but other methodologies can be implemented. For more information on how to install and use *InfiniteOpt.jl*, visit <https://github.com/pulsipher/InfiniteOpt.jl> where we provide tutorials, examples, API guides, and more.

3. CASE STUDY

We compare the time-valued measures proposed in Section 2.2 in the context of optimal control. We adapt the pandemic control problem that seeks to choose an isolation policy to control the spread of a contagion that minimally imposes economic impact (induced by mandated isolation). We model the spread of the disease via the SEIR model which defines the populations of individuals susceptible to infection $y_s : \mathcal{D}_t \rightarrow [0, 1]$, exposed individuals that are not yet infectious $y_e : \mathcal{D}_t \rightarrow [0, 1]$, infectious individuals $y_i : \mathcal{D}_t \rightarrow [0, 1]$, and recovered individuals $y_r : \mathcal{D}_t \rightarrow [0, 1]$ (considered immune to future infection). Moreover, these satisfy $y_s(t) + y_e(t) + y_i(t) + y_r(t) = 1$. Thus, our state variables are comprised of $y_s(t)$, $y_e(t)$, $y_i(t)$, and $y_r(t)$. Moreover, we exhibit control by imposing an isolation policy $y_u(t) \in [0, \bar{y}_u] \subseteq [0, 1]$ that entails the separation of susceptible and exposed individuals ($y_u(t) = 0$ denotes no separation and $y_u(t) = 1$ denotes complete separation). The formulation seeks to minimize the isolation policy function $y_u(t)$ (i.e., $f(t) = y_u(t)$ for simplicity) while enforcing that the amount of infectious individuals $y_i(t)$ remains below $\bar{y}_i \in (0, 1]$:

$$\begin{aligned}
& \min M_t y_u(t) \\
& \text{s.t. } \dot{y}_s(t) = (y_u(t) - 1) \beta y_s(t) y_i(t), \quad t \in \mathcal{D}_t \\
& \quad \dot{y}_e(t) = (1 - y_u(t)) \beta y_s(t) y_i(t) - \xi y_e(t), \quad t \in \mathcal{D}_t \\
& \quad \dot{y}_i(t) = \xi y_e(t) - \gamma y_i(t), \quad t \in \mathcal{D}_t \\
& \quad \dot{y}_r(t) = \gamma y_i(t), \quad t \in \mathcal{D}_t \\
& \quad y_s(0) = s_0, y_e(0) = e_0, y_i(0) = i_0, y_r(0) = r_0 \\
& \quad y_i(t) \leq \bar{y}_i, \quad t \in \mathcal{D}_t \\
& \quad y_u(t) \in [0, \bar{y}_u], \quad t \in \mathcal{D}_t
\end{aligned} \tag{27}$$

where $s_0, e_0, i_0, r_0 \in [0, 1]$ are initial conditions and $\beta, \gamma, \xi \in \mathbb{R}$ are the rates of infection, recovery, and incubation, respectively, which are specific to the disease in question. Here we specify initial conditions at $s_0 = .9999$, $e_0 = 10^{-5}$, and $i_0 = r_0 = 0$. The disease parameters are taken to be $\beta = 0.727$, $\gamma = 0.303$, and $\xi = 0.3$. We choose limits $\bar{y}_u = 0.8$ and $\bar{y}_i = 0.02$. Finally, we set $\mathcal{D}_t = [0, 200]$. We model Formulation (27) in `InfiniteOpt.jl` and use backward finite-difference to evaluate the derivatives using 101 discretization points. Code Snippet 1 highlights the compact syntax required to model Formulation (27) in `InfiniteOpt.jl` under these conditions. For the objective $M_t y_u(t)$, we consider the following time-valued measures: $\int_{t \in \mathcal{D}_t} y_u(t) dt$, $\mathbb{E}_t[y_u(t)]$, $\mathbb{E}\text{-}\mathbb{V}_t(y_u(t); \lambda)$, and $\text{CVaR}_t(y_u(t); \alpha)$. We also investigate the implications of using the uniform time-valued pdf in (6) against the exponential pdf of (7).

```

1 using InfiniteOpt, Ipopt
2
3 # Set the parameters
4  $\gamma, \beta, \xi = 0.303, 0.727, 0.3$ 
5  $s_0, e_0, i_0, r_0 = 1 - 1e-5, 1e-5, 0, 0$ 
6
7 # Define the model
8  $m = \text{InfiniteModel}(\text{Ipopt.Optimizer})$ 
9
10 # Define the time parameter
11  $\text{@infinite\_parameter}(m, t \in [0, 200], \text{num\_supports} = 101)$ 
12
13 # Add the variables
14  $\text{@variable}(m, y_s, \text{Infinite}(t))$ 
15  $\text{@variable}(m, y_e, \text{Infinite}(t))$ 
16  $\text{@variable}(m, y_i \leq 0.02, \text{Infinite}(t))$ 
17  $\text{@variable}(m, y_r, \text{Infinite}(t))$ 
18  $\text{@variable}(m, 0 \leq y_u \leq 0.8, \text{Infinite}(t))$ 
19
20 # Set the time expectation objective
21  $\text{@objective}(m, \text{Min}, \mathbb{E}(y_u, t))$ 
22
23 # Define the SEIR equations
24  $\text{@constraint}(m, \partial(y_s, t) == -(1 - y_u) * \beta * y_s * y_i)$ 
25  $\text{@constraint}(m, \partial(y_e, t) == (1 - y_u) * \beta * y_s * y_i - \xi * y_e)$ 
26  $\text{@constraint}(m, \partial(y_i, t) == \xi * y_e - \gamma * y_i)$ 
27  $\text{@constraint}(m, \partial(y_r, t) == \gamma * y_i)$ 
28  $\text{@constraint}(m, y_s(0) == y_s0)$ 
29  $\text{@constraint}(m, y_e(0) == y_e0)$ 
30  $\text{@constraint}(m, y_i(0) == y_i0)$ 
31  $\text{@constraint}(m, y_r(0) == y_r0)$ 
32
33 # Solve the model and retrieve results
34  $\text{optimize!}(m)$ 
35  $u\_opt = \text{value}(y_u)$ 
36  $ts = \text{value}(t)$ 

```

Code Snippet 1. Formulation (27) implemented in `InfiniteOpt.jl`.

Figure 1 summarizes the results; the expectation $\mathbb{E}_t[y_u(t)]$ with the uniform pdf shapes $y_u(t)$ identically to the classical integral measure and establishes a baseline for comparison. we are able to place increased emphasis on the early time regime when we use the exponential pdf defined in (7) in combination with $\mathbb{E}_t[y_u(t)]$. In comparison to the other expectation measure, the exponentially weighted counterpart exhibits a policy trajectory that is significantly

reduced at early times while later times lead to increased isolation requirements. This highlights how the choice of pdf $p_t(t)$ enhances the flexibility of our proposed measures in accordance with the requirements of the problem. The optimal policy we obtain with the mean-variance measure $\mathbb{E}\text{-}\mathbb{V}_t(y_u(t); 8)$ demonstrates how placing increased priority on minimizing the variance of the cost function induces the trajectory to be increasingly smoothed (i.e., cost fluctuations are damped). This comes at the trade-off (controlled via specification of λ) of increasing the mean isolation policy, but helps to derive a more consistent policy. For this application, a smoother policy would likely be preferred since rapid policy changes can be highly disruptive and can lead to public dissatisfaction. Finally, in contrast to the mean-variance (which equally penalizes positive and negative cost deviations from the mean), the CVaR measure only penalizes the high cost deviations that surpass the threshold determined by the α -quantile. In Figure 1 we see that $\text{CVaR}_t(y_u(t); 0.9)$ flattens the peak isolation policy values observed with the standard integral/expectation policy. This hedging against high costs also induces a more substantial response at later times which results in a larger cumulative cost. Thus, we observe a trade-off (controlled via α) between penalizing cost peaks and minimizing the total cumulative cost.

The variety of optimal trajectories observed in Figure 1 clearly demonstrates how the choice of M_t can greatly impact the solution of (1). This diversity in solutions is enabled by our proposed toolbox of dynamic measures. The choice of an appropriate measure will depend on the particular application considerations. In this case, the use of mean-variance or CVaR would be reasonable choices depending on whether dampening policy fluctuation or peak policy values is of greater concern.

4. CONCLUSIONS

In this work, we have shown that risk measures used in SO can be transferred (along with their mathematical properties) to DO. This enables a new class of DO formulations that shape time trajectories in interesting and useful ways. The transfer of insights across the SO and DO disciplines is facilitated by a unifying infinite-dimensional abstraction. In future work, it will be interesting to investigate the analogy between SO and DO further in transferring more amenable measure operators and establishing their properties in a DO context (e.g., dominance and stability). Moreover, the establishment of time-valued pdfs provides a foundation from which the utility of transferring distributionally robust measure functions to DO can be investigated; which would potentially allow us to consider multiple weighting functions over the time horizon.

REFERENCES

- Artzner, P., Delbaen, F., Eber, J.M., and Heath, D. (1999). Coherent measures of risk. *Mathematical finance*, 9(3), 203–228.
- Biegler, L.T. (2010). *Nonlinear programming: concepts, algorithms, and applications to chemical processes*. SIAM.
- Birge, J.R. and Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.

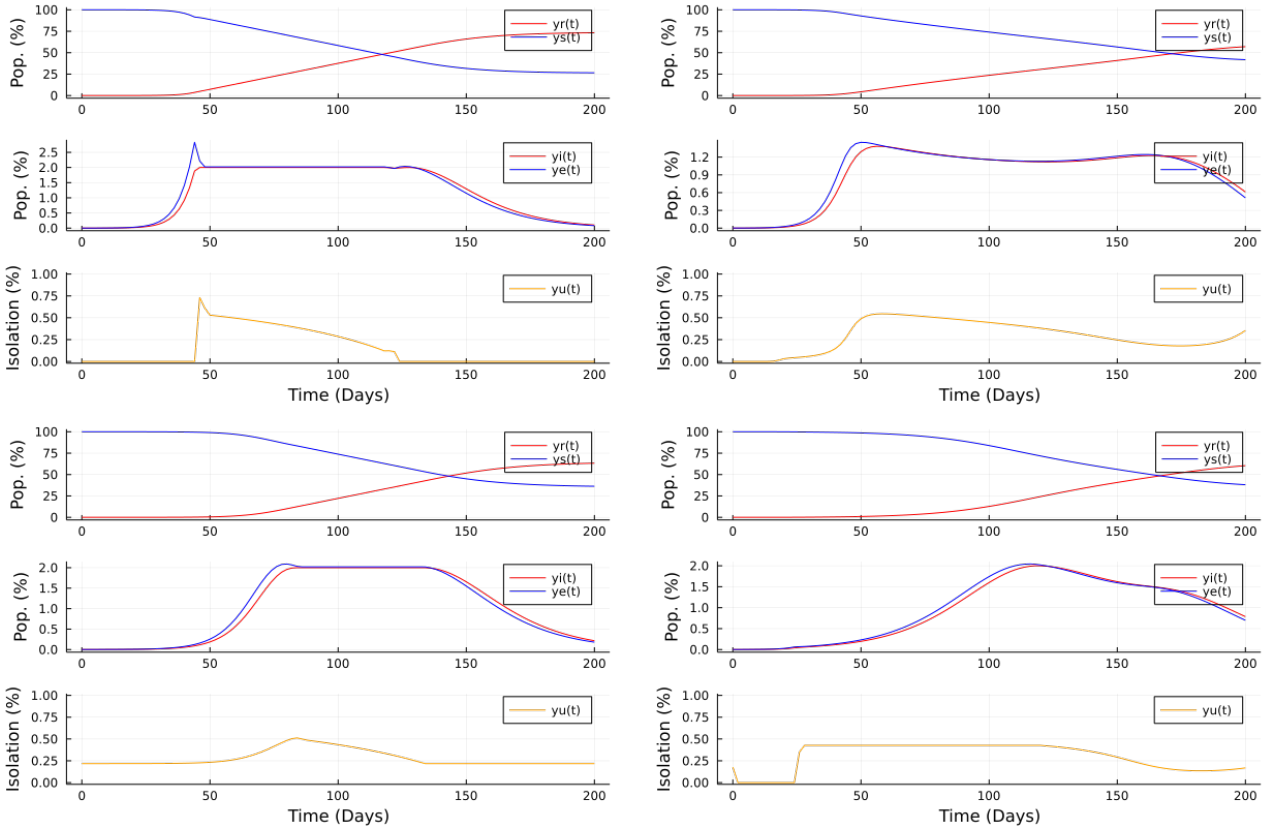


Fig. 1. The optimal policy and state trajectories. Top Left: \mathbb{E}_t with uniform pdf. Top Right: \mathbb{E}_t with exponential pdf ($\gamma = 1$). Bottom Left: $\mathbb{E}\text{-}\mathbb{V}_t$ with $\lambda = 8$ and uniform pdf. Bottom Right: CVaR_t with $\alpha = 0.9$ and uniform pdf.

Fulga, C. (2016). Portfolio optimization with disutility-based risk measure. *European Journal of Operational Research*, 251(2), 541–553.

Krokhmal, P., Zabarankin, M., and Uryasev, S. (2013). Modeling and optimization of risk. *Handbook of the fundamentals of financial decision making: Part II*, 555–600.

Leland, H.E. (1999). Beyond mean–variance: Performance measurement in a nonsymmetrical world (corrected). *Financial analysts journal*, 55(1), 27–36.

Pulsipher, J.L., Zhang, W., Hongisto, T.J., and Zavala, V.M. (2022). A unifying modeling abstraction for infinite-dimensional optimization. *Computers & Chemical Engineering*, 156, 107567.

Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI.

Renteria, J.A., Cao, Y., Dowling, A.W., and Zavala, V.M. (2018). Optimal pid controller tuning using stochastic programming techniques. *AIChE Journal*, 64(8), 2997–3010.

Risbeck, M.J. and Rawlings, J.B. (2019). Economic model predictive control for time-varying cost and peak demand charge optimization. *IEEE Transactions on Automatic Control*, 65(7), 2957–2968.

Rockafellar, R.T., Uryasev, S., et al. (2000). Optimization of conditional value-at-risk. *Journal of risk*, 2, 21–42.

Ruszczynski, A. and Shapiro, A. (2006). Optimization of risk measures. In *Probabilistic and randomized methods for design under uncertainty*, 119–157. Springer.

Sarykalin, S., Serraino, G., and Uryasev, S. (2008). Value-at-risk vs. conditional value-at-risk in risk management and optimization. In *State-of-the-art decision-making tools in the information-intensive age*, 270–294. Informis.

Shin, S. and Zavala, V.M. (2020). Diffusing-horizon model predictive control. *arXiv preprint arXiv:2002.08556*.