

FIRMGrasp: A Friction-Informed Risk Margin for Robust Grasp Synthesis

Clinton Enwerem¹, John S. Baras¹, and Calin Belta¹

Abstract—Classical grasp quality metrics assume a single deterministic friction coefficient, so they cannot predict whether a grasp retains force closure across the range of friction values the contacting surfaces may exhibit. To predict these failures, we present FIRMGrasp, a family of friction-volatility-aware grasp quality metrics grounded in the Conditional Value-at-Risk (CVaR) risk measure. Unlike standard grasp quality assessors that assume a single friction realization, our metric evaluates the force-closure margin at the CVaR-discounted mean of the adverse friction tail, yielding a risk-adjusted margin $\varepsilon^{(\beta)}$, the inscribed-ball radius of the risk-adjusted wrench space. We establish its monotonicity in the confidence level β , its differentiability in the grasp parameters, and a probabilistic closure certificate that guarantees force closure with probability at least β whenever $\varepsilon^{(\beta)}$ is positive. Under a calibrated friction distribution, analytic evaluation shows our $\varepsilon^{(\beta)}$ metric identifies friction-sensitive grasps that the nominal Ferrari-Canny epsilon rates as high-quality, and we compare against the nominal epsilon and recent differentiable baselines. Across 1,599 LEAP Hand and Allegro Hand grasps, 53% of the grasps the nominal Ferrari-Canny margin certifies lose force closure in the adverse friction tail. On the same set, the nominal margin separates realized shake and pick success with probabilities of only 0.53 and 0.67, near chance on shake success, whereas $\varepsilon^{(\beta)}$ orders the pair correctly with probabilities of 0.63 and 0.78, respectively. In simulated lift trials with gravity enabled at an adverse friction coefficient of 0.2, grasps $\varepsilon^{(\beta)}$ certifies reach a 70% success rate under lateral pull, against 25% for grasps the nominal margin certifies but $\varepsilon^{(\beta)}$ rejects.

I. INTRODUCTION

Grasping a rigid object with a rigid-link robotic hand entails choosing contacts, a wrist frame, and hand configurations that stabilize the object against arbitrary exogenous wrenches including gravity, and one may further require that the grasp be dexterous, in equilibrium, or stable. Since many grasps meet these desiderata, *quality measures* rank them so that synthesis reduces to an optimization problem. The most widely used such measure is the Ferrari-Canny, or epsilon, metric, the radius of the largest ball inscribed in the grasp wrench space [1], alongside a range of alternatives that instead evaluate the grasp matrix, finger force limits, or hand-object geometry [2], [3], [4], [5], [6], [7].

However, these metrics score a grasp at the single instant of synthesis, under a friction coefficient, object pose, and contact configuration all assumed fixed and known, so none of them bounds what happens once execution departs from that instant. A grasp that is force-closed and positive-epsilon at synthesis can destabilize within seconds under lateral exogenous forces and an execution-time friction coefficient




Fig. 1. **Predicting Lift Success under Adverse Friction.** Two Allegro Hand grasps of the same object that the nominal Ferrari-Canny margin certifies almost identically ($\varepsilon_{\text{nom}} = +0.008$ and $+0.013$) part ways once friction uncertainty is introduced. Our risk-adjusted grasp quality margin $\varepsilon^{(\beta)}$ (Definition 7) rejects the first ($\varepsilon^{(\beta)} < 0$) and certifies the second ($\varepsilon^{(\beta)} > 0$) under the adverse friction prior $0.5\mathcal{U}(0.7, 1.0) + 0.5\mathcal{U}(0.1, 0.3)$ at confidence $\beta = 0.9$. A gravity-on lift at Coulomb friction $\mu=0.3$ realizes exactly this split, with the friction-sensitive grasp (a, b) losing the object, while the certified-robust grasp (c, d) lifts it and retains it under lateral pull.

distinct from the one assumed at synthesis [8]. Figure 1 previews this failure on a toy grasp-and-lift example, two Allegro Hand grasps that the nominal epsilon margin certifies as force-closed at synthesis time. Only the grasp our risk-sensitive certificate also certifies retains the object upon lift under adverse friction, while the other slips. The pattern is not specific to friction. The same epsilon metric is a poor predictor of robustness to pose error, with its force-closure verdict flipping under pose perturbations the metric does not model [9]. We therefore treat grasp quality as the probability that a grasp keeps force closure across the distribution of friction coefficients the physical contact could plausibly realize.

Given this probabilistic setting, and motivated by the demonstrated effectiveness of risk-sensitive formulations built on the Conditional Value-at-Risk (CVaR) and related coherent risk measures for safe decision-making under distributional uncertainty in motion planning [10], reinforcement learning [11], and belief-space control under latent uncertainty [12], we ground grasp quality in the same CVaR construction by formulating it as a risk-sensitive functional over the friction distribution. Thus, our metric enables a direct quantification of robustness, capturing failure modes

¹The authors are with the Department of Electrical & Computer Engineering and Institute for Systems Research, University of Maryland, College Park, MD, USA. Emails: {enwerem, baras, calin}@umd.edu.

that arise only after interaction with the environment, a capability standard quality metrics do not provide.

In practice, evaluating force closure through CVaR at confidence level β gives a margin, tunable by β , that quantifies the likelihood a grasp keeps force closure under the adverse tail of the friction prior instead of at a single nominal coefficient. This margin reduces to the classical epsilon margin under a degenerate, point-mass prior, so it generalizes the classical desiderata instead of replacing them. A grasp with a positive margin under the friction tail also has a positive margin at the nominal coefficient, a monotonicity property we prove in Theorem 2. The resulting metric complements recent grasp-synthesis methods [6], [13], [14], since any of them can supply candidate grasps for our metric to score and rank by friction sensitivity, a ranking our simulated execution studies show predicts lift success and post-execution stability (Table II, Sec. IX-G, and fig. 11).

Contributions:

- 1) *New Risk-Sensitive Quality Metrics.* We define two new grasp quality margins, the CVaR margin Q_β (Equation (12)) and its closed-form analytic counterpart $\varepsilon^{(\beta)}$ (Definition 7). Each margin scores a grasp by its friction-tail behavior rather than by a single nominal friction coefficient, and $\varepsilon^{(\beta)}$ reduces to the classical Ferrari-Canny margin under a point-mass friction prior, so both generalize the nominal metric rather than restating it. We prove their monotonicity, differentiability, and probabilistic-closure properties (Theorems 2 to 1, Proposition 2).
- 2) *Probabilistic Closure Certificate.* We provide a formal guarantee that a positive risk-sensitive margin implies force closure with probability at least β , for Q_β through the CVaR tail bound (Theorem 1) and for $\varepsilon^{(\beta)}$ through friction monotonicity (Corollary 2).
- 3) *Risk-Aware Grasp Synthesis and Evaluation.* We implement a sample-based grasp synthesis pipeline that uses the CVaR metric to select robust grasps from a candidate pool, and we demonstrate on seven objects that the CVaR metric surfaces a friction sensitivity that the classical epsilon metric and the min-weight metric of FRoGGeR [6] do not model, since neither treats friction as uncertain.

II. RELATED WORK

A. Classical Analytical Grasp Metrics

Analytical grasp quality assessment dates to Ferrari and Canny’s epsilon metric [1], alongside the Grasp Wrench Space volume [2], the grasp isotropy index [3], [4], and the largest minimum resisted wrench [5], each quantifying a different algebraic property of the same grasp map. Uncertainty-aware treatments of these metrics have, until recently, stayed largely empirical, scoring robustness with Monte Carlo statistics over the nominal metric [9]. Such statistics predict grasp success reasonably well, but the sampling is expensive, and both the friction coefficient and the pose distribution stay fixed at their nominal values throughout.

By contrast, a more recent line of work departs from the epsilon metric analytically. The min-weight metric of FRoGGeR [6] reports the minimum among a set of scaled friction generators at the linearized friction cones. Min-weight is almost-everywhere differentiable, and its linear-program formulation, differentiability included, drives grasp synthesis in under a second, fast enough for dynamic manipulation [15], [14]. It shares with the epsilon metric it approximates, however, a fixed Coulomb friction coefficient, so a grasp it scores highly is precision-optimal only under that nominal coefficient. Our metric extends the same line of work to a friction coefficient treated as uncertain rather than fixed.

B. Uncertainty Representations for Dexterous Grasping

Other recent work characterizes grasp uncertainty at a different moment in a grasp’s lifetime. PONG bounds the probability of force closure analytically under uncertain contact normals [16], Gaussian Process Implicit Surfaces represent the object surface itself as a distribution rather than a fixed mesh [17], [18], and a learned variational belief over contact parameters certifies an already-synthesized grasp under multimodal uncertainty [8]. Each of these certifies that an already-synthesized grasp withstands execution-time or post-execution perturbation under an assumed uncertainty model, applying robustness only *after* the grasp is fixed. The uncertainty we target instead acts *before* and *during* synthesis. We retain Coulomb’s friction condition but allow the friction coefficient to be variable. As a consequence, at each grasp contact point, we consider the true friction cone to be unknown and instead allow the contact force to lie within a *set* of friction cones, the size of which is determined by the underlying uncertainty in the friction coefficient. Our metric stands as a complement to existing force-closure margins, in particular FRoGGeR’s min-weight ℓ^* and the robustness measures of Li *et al.* [7], and each member of the resulting metric family answers a different robustness question about a grasp.

C. Robust Grasp Synthesis

A separate line of work targets wrench capability directly. [7] prove that a force-closed grasp synthesized under a nominal basis-wrench hull withstands any disturbance smaller than the residual between that hull and the hull the perturbed basis wrenches sweep out. Grasp stability tracks wrench resistance closely enough that this residual, which we term the *wrench capability residual*, is a useful quantity in its own right. The residual measures the distance a friction perturbation can move the primitive wrenches before closure fails.

In addition to analytic synthesis, a substantial body of learning-based work targets dexterous grasp generation directly. Differentiable formulations of force closure treat contact selection as a smooth optimization objective, propagating gradients through a grasp planner [19] or through a differentiable simulator that scores candidate grasps by how well they withstand contact-rich perturbation [20], [21], a




Fig. 2. **Grasping Frames.** Contact, object, and hand coordinate frames used in our uncertainty-aware grasp planning framework. These include the hand base frame, $\{\mathcal{F}_H\}$, the object frame, $\{\mathcal{F}_O\}$, and the i^{th} contact frame, $\{\mathcal{F}^{C_i}\}$ attached to contact point i at $p^{(i)} \in \mathbb{R}^3$. Admissible contact forces at a contact point are expressed in $\{\mathcal{F}^{C_i}\}$, their corresponding positions are expressed in $\{\mathcal{F}_O\}$, and we assume knowledge of the contact-to-object frame force rotation matrix (R_{OC}) and the object-to-hand homogeneous transformation (T_{OH}).

differentiable force-closure estimator built on the classical epsilon metric extends the same treatment to arbitrary hand structures [22], and a bilevel formulation pairs a conditional generative model with a differentiable force-closure constraint [23]. A separate line of work learns generative models that sample grasp poses conditioned on object geometry or point clouds, spanning variational [24], adversarial [25], [26], [27], real-time regression [28], and diffusion-based [29], [30], [31] formulations, alongside SE(3)-equivariant dexterous grasp flows [32]. Earlier discriminative approaches instead score or synthesize grasps directly from depth images or contact maps [33], [34], [35], while hand-agnostic frameworks learn grasping policies that transfer across gripper morphologies [36], [37], [38]. These generators are efficient at producing candidate grasps but do not certify or rank robustness under friction uncertainty, so our metric is agnostic to the generator, scoring grasps from any such source.

III. MATHEMATICAL PRELIMINARIES

We consider a multi-fingered hand with configuration $q_h \in \mathbb{R}^{n_h}$ grasping a rigid object whose pose in the hand frame $\{\mathcal{F}_H\}$ is $X^o \in SE(3)$. The hand forward kinematics $FK(\cdot)$ generates a contact set $\mathcal{C} := \{(p^{(i)}, \{\mathcal{F}^{C_i}\}) \mid p^{(i)} := FK(q_h, X^o) \in \mathbb{R}^3, i = 1, \dots, n_c, n_c \leq N_{\max}\}$, where n_c is the contact count and $N_{\max}$ the fingertip count of the hand. Each contact comprises a point $p^{(i)}$ and a frame $\{\mathcal{F}^{C_i}\} := \{t_x^{(i)}, t_y^{(i)}, n^{(i)}\}$ of orthonormal tangent directions and the unit surface normal $n^{(i)} \in \mathbb{R}^3$.

Concretely, our risk-sensitive margins evaluate the adverse tail of a friction-dependent quantity through the lower-tail Conditional Value-at-Risk (CVaR) [39]. For a scalar random variable Z with quantile function F_Z^{-1} and confidence $\beta \in (0, 1)$,

$$\text{CVaR}_\beta(Z) := \frac{1}{1-\beta} \int_0^{1-\beta} F_Z^{-1}(u) du, \quad (1)$$

the mean of the worst $(1-\beta)$ -fraction of outcomes, so $\beta \rightarrow 1$ is the risk-averse limit and $\text{VaR}_\beta(Z) := F_Z^{-1}(1-\beta)$ is the matching lower quantile, with $\text{CVaR}_\beta(Z) \leq \text{VaR}_\beta(Z)$.

A. The Geometry of Parametric Grasps

We assume hard-finger contacts with friction, with contact forces satisfying Coulomb's condition [40] and lying in a linearized polyhedral friction cone of a predetermined number of edge rays or *generators*, where each generator is a primitive contact force in $\mathbb{R}^3$.

Definition 1 (Grasp Parameterization): We denote a grasp by the tuple $g := (\mathcal{C}, q_h, X^o)$ comprising the contact set, hand configuration, and object pose. Let $\mathcal{G}$ denote the set of all such grasps.

Given this parameterization, we measure the quality of a grasp g (Definition 1) by a signed scalar whose sign reports force closure. Both the classical Ferrari-Canny margin and the risk-adjusted margin we develop in Sec. VI instantiate the underlying quality map, and every closure certificate below turns on its sign.

Definition 2 (Grasp Quality Function): A grasp quality function is a scalar map $Q : \mathcal{G} \rightarrow \mathbb{R}$, positive when the grasp attains the encoded property (here, force closure) and negative otherwise.

Definition 3 (Primitive Contact Forces & Induced Wrench Cone): The polyhedral cone approximating admissible contact forces at the i^{th} contact point is

$$\mathfrak{C}^{(i)} := \left\{ F = \sum_{j=1}^{n_s} \alpha_j^{(i)} f_j^{(i)} \mid \alpha_j^{(i)} \geq 0 \right\}, \quad (2)$$

where $j = 1, \dots, n_s$ indexes the generator set and $f_j^{(i)}$ is the primitive force corresponding to contact i and generator j , expressed in the object frame:

$$f_j^{(i)} = f_{n,j}^{(i)} n^{(i)} + f_{t,j}^{(i)}, \quad f_{t,j}^{(i)} \in \text{span}\{t_x^{(i)}, t_y^{(i)}\}. \quad (3)$$

The generators satisfy the friction limit specified by Coulomb's condition, i.e., $\|f_t\| = \mu f_n$. Thus, from (3), each generator is an affine function of the friction coefficient μ :

$$f_j^{(i)} := \phi_{ij}(\mu) = f_{n,j}^{(i)} n^{(i)} + \mu f_{n,j}^{(i)} \hat{t}_{ij}. \quad (4)$$

The primitive wrench about the object's center-of-mass o is then

$$w_{ij}(\mu) = \begin{bmatrix} \phi_{ij}(\mu) \\ (p^{(i)} - o) \times \phi_{ij}(\mu) \end{bmatrix}. \quad (5)$$

We assemble the GWS as the convex hull of the primitive wrenches of Definition 3, $\mathcal{W}(\mu) := \text{conv}(\{w_{ij}(\mu)\}_{i,j})$, the total-force-normalized construction of [1], [2] with a unit normal force per generator ($f_{n,j}^{(i)} = 1$) and unscaled moment coordinates. Because the tangential components of the generators come in opposing pairs, $\mathcal{W}(\mu') \subseteq \mathcal{W}(\mu)$ for $0 \leq \mu' \leq \mu$. Every generator at μ' is a convex combination of the corresponding pair at μ , so the GWS nests in the friction coefficient. Splitting each generator at $\mu = 0$ also fixes two auxiliary bodies used repeatedly below, the *normal wrench body* $\mathcal{W}_n := \mathcal{W}(0)$ and the *tangential wrench body* $\mathcal{W}_t := \text{conv}(\{(f_{n,j}^{(i)} \hat{t}_{ij}, (p^{(i)} - o) \times f_{n,j}^{(i)} \hat{t}_{ij})\}_{i,j})$.

IV. RISK-ADJUSTED GRASP WRENCH SPACE

The friction coefficient assumed at planning time may deviate from its realized value at execution, so the grasp wrench space is a random polytope rather than a deterministic set. We scale the underlying polytope with friction, discount its adverse tail, and evaluate a force-closure margin on the discounted set. Under friction uncertainty, the polytope scales linearly with the coefficient (Lemma 1)¹ Let the friction coefficient be a random scalar $\mu \sim p_\mu$. The resulting polyhedral friction cone (2) thus becomes a convex set of uncertain generators whose magnitudes depend on the realization of μ .

Lemma 1 (GWS Linear Scaling Under Friction Uncertainty): The primitive wrench $w_{ij}(\mu)$ is affine in μ . As a result, the GWS $\mathcal{W}(\mu)$ satisfies $\mathcal{W}(\mathbb{E}[\mu]) \subseteq \mathbb{E}[\mathcal{W}(\mu)]$ in the Minkowski-expectation sense, and its diameter scales linearly in μ .

Proof: From (4), $\phi_{ij}(\mu) = f_{n,j}^{(i)} n^{(i)} + \mu f_{n,j}^{(i)} \hat{t}_{ij}$, where $\hat{t}_{ij}$ is the unit tangent direction. This is affine in μ , so $w_{ij}(\mu)$ is affine in μ by (5). The convex hull of affinely-parameterized vertices is Minkowski-affine in μ , yielding the stated scaling. The diameter $\text{diam}(\mathcal{W}(\mu)) = 2\mu \max_{ij} \|f_{n,j}^{(i)}\| \cdot \|\hat{t}_{ij}^\top (p^{(i)} - o) \times \hat{t}_{ij}\|$ is linear in μ . ■

In particular, for simplicity, we focus on variable friction. Extension to variable object pose follows analogously. Throughout, $\theta \in \mathbb{R}^{n_\theta}$ collects the grasp parameters, i.e., the wrist pose and hand configuration that determine the grasp g of Definition 1, and we write $\mathcal{W}(\mu) \equiv \mathcal{W}(\mu; \theta)$, suppressing θ where it is fixed. Because each primitive wrench is affine in μ (Lemma 1), the friction CVaR v_β of (1) collapses the underlying friction prior into a single adverse coefficient. The next definition names this friction value and the wrench body assembled at it.

Definition 4 (Risk-Adjusted Friction and Wrench Body): Let $\mu \sim p_\mu$. The *risk-adjusted friction* at confidence β is $v_\beta := \text{CVaR}_\beta(\mu)$, the mean of the adverse friction tail, and the *risk-adjusted wrench body* is the GWS assembled at the risk-adjusted friction v_β ,

$$\mathcal{W}^{(\beta)}(\theta) := \mathcal{W}(v_\beta; \theta). \quad (6)$$

Since $v_\beta \leq \mathbb{E}[\mu]$ and the GWS nests in μ , $\mathcal{W}^{(\beta)}$ shrinks as β rises. $\mathcal{W}^{(\beta)}$ is the wrench capability the grasp retains once the friction coefficient is discounted to its adverse tail.

Definition 5 (Risk-Sensitive Directional Support): For each direction $d \in \mathbb{S}^5$, the unit sphere in $\mathbb{R}^6$, the *risk-sensitive directional support* of the uncertain GWS is

$$h_\beta(d; \theta) := \text{CVaR}_\beta(h_{\mathcal{W}(\mu)}(d)), \quad (7)$$

where $h_{\mathcal{W}(\mu)}(d) := \sup_{w \in \mathcal{W}(\mu)} d^\top w$ is the support function of the uncertain GWS.

Proposition 1 (Support Structure of the Risk Adjustment): The map $d \mapsto h_\beta(d; \theta)$ is a valid support function, convex

and positively homogeneous of degree one in d , and it is the support function of the convex body

$$\mathcal{W}^{(\beta)}(\theta) := \{w \in \mathbb{R}^6 : d^\top w \leq h_\beta(d; \theta), \forall d \in \mathbb{S}^5\}, \quad (8)$$

which contains the risk-adjusted wrench body: $\mathcal{W}^{(\beta)} \subseteq \mathcal{W}^{(\beta)}$, with $h_{\mathcal{W}^{(\beta)}}(d) = h_\beta(d; \theta)$ in every direction d whose maximizing primitive wrench does not change over the friction tail.

Proof: Every $h_{\mathcal{W}(\mu)}(d)$ is non-decreasing in μ , since the GWS nests in the friction coefficient. The tail event realizing the lower-tail CVaR is therefore $\{\mu \leq \text{VaR}_\beta(\mu)\}$ for every direction simultaneously, so $h_\beta(d; \theta) = \mathbb{E}[h_{\mathcal{W}(\mu)}(d) \mid \mu \leq \text{VaR}_\beta(\mu)]$ is a tail-conditional expectation of support functions and hence convex and positively homogeneous in d . For the inclusion, each vertex support $d^\top w_{ij}(\mu)$ is affine in μ by Lemma 1, so $h_{\mathcal{W}(\mu)}(d) = \max_{ij} \mathbb{E}[d^\top w_{ij}(\mu) \mid \mu \leq \text{VaR}_\beta(\mu)] \leq \mathbb{E}[\max_{ij} d^\top w_{ij}(\mu) \mid \mu \leq \text{VaR}_\beta(\mu)] = h_\beta(d; \theta)$, with equality when one primitive wrench attains the maximum throughout the tail. ■

Beyond this directional support, assessing force closure in the distributional setting demands a probabilistic definition over the whole friction prior, so we formalize probabilistic (force) closure next as a scalar condition.

Definition 6 (Probabilistic Closure): A grasp $g \in \mathcal{G}$ satisfies *probabilistic closure at level β* if $\text{CVaR}_\beta(\varepsilon(g, \mu)) > 0$, which implies force closure in all but the worst $(1 - \beta)$ -fraction of uncertainty realizations under the assumed model.

Remark 1 (Necessity of a Distribution-Dependent Quality): A quality that ignores the friction distribution cannot certify probabilistic closure. Take a grasp that is force-closed at the nominal friction μ_{nom} but relies on friction, so $\varepsilon(g, \mu_{\text{nom}}) > 0$ while its friction threshold $\mu^* := \inf\{\mu : \varepsilon(g, \mu) > 0\}$ is strictly positive and the grasp loses closure, $\varepsilon(g, \mu) \leq 0$, for every $\mu < \mu^*$. Fix $\delta \in (0, 1)$ and $\eta \in (0, 1 - \delta)$ and pick any friction value $\mu_- < \mu^*$. The two-point distribution p_δ that places mass $\delta + \eta$ at μ_- and mass $1 - \delta - \eta$ at μ_{nom} has closure-loss probability $\Pr_{p_\delta}[\varepsilon(g, \mu) \leq 0] = \delta + \eta > \delta$. A quality $Q(g)$ that does not depend on p_δ returns the same value here as under a distribution concentrated near μ_{nom} with negligible closure loss, so it cannot separate the two and cannot guarantee closure with probability at least β . The risk-sensitive condition of Definition 6 depends on the adverse tail of the distribution and avoids the gap described above.

Given this, the probabilistic closure definition of Definition 6 yields a force-closure certificate. The following theorem states, formally, the probability with which a grasp is force-closed once its quality margin turns positive.

Theorem 1 (Probabilistic Closure Certificate): Let $\mu \sim p_\mu$. If $\text{CVaR}_\beta(\varepsilon(g, \mu)) > 0$, i.e., if g satisfies probabilistic closure at level β per Definition 6, then

$$\Pr[\varepsilon(g, \mu) > 0] \geq \beta, \quad (9)$$

i.e., the grasp is force-closed with probability at least β .

Proof: By the lower-tail definition of the CVaR risk measure [39], $\text{CVaR}_\beta(Z) > 0$ implies $\text{VaR}_\beta(Z) > 0$, i.e., the $(1 - \beta)$ -quantile of Z is positive, so $\Pr[Z \leq 0] < 1 - \beta$ and $\Pr[Z > 0] > \beta$. Applying this to $Z = \varepsilon(g, \mu)$ and

¹Under object-pose uncertainty, the GWS polytope rotates instead, and the active uncertainty modality determines which transformation applies.

recalling that $\varepsilon(g, \mu) > 0$ certifies force closure gives the result. ■

V. PROBLEM FORMULATION

Problem 1 states the synthesis target our metric family serves.

Problem 1 (Risk-Sensitive Robust Grasp Synthesis): Given an object model, a robotic hand model with configuration space $\mathbb{R}^{n_h}$, a friction prior p_μ , and a confidence level $\beta \in (0, 1)$, find

$$g^* \in \arg \max_{g \in \mathcal{G}_{\text{feas}}} \text{CVaR}_\beta(\varepsilon(g, \mu)), \quad \mu \sim p_\mu, \quad (10)$$

where $\mathcal{G}_{\text{feas}} \subseteq \mathcal{G}$ collects the grasps satisfying joint limits, self-collision avoidance, and contact-formation feasibility for the arm-hand-object system of Figure 2, and $\varepsilon(g, \mu)$ is the Ferrari-Canny margin of grasp g under friction realization μ . By Theorem 1, any g^* attaining a positive risk-sensitive margin admits probabilistic closure at level β in the sense of Definition 6 and therefore secures force closure with probability at least β over the friction realizations of p_μ .

Specifically, our objective in (10) couples grasp parameters to the tail of the friction-induced wrench-margin distribution, so the optimizer prefers grasps whose wrench capacity withstands the adverse $(1 - \beta)$ -fraction of friction realizations rather than the nominal one. Sec. VI develops our metric family that scores candidate grasps, and Sec. VII establishes the monotonicity, differentiability, and closure-certificate properties that ground (10). Sec. VIII then describes the sample, certify, and score procedure that approximates g^* . The studies instantiate two computable quantities, our analytic margin $\varepsilon^{(\beta)}$ for scoring and selection, and, as the synthesis surrogate of (10), a risk-adjusted extension of FROGGER’s min-weight metric [6], $\ell^{*(\beta)}$, the min-weight program of (15) evaluated on the risk-adjusted body $\mathcal{W}^{(\beta)}$ (Sec. X-B).

VI. RISK-SENSITIVE GRASP QUALITY

Classical grasp quality metrics such as force closure and ε -metrics evaluate feasibility at a single instant under nominal assumptions, and provide no guarantee once the realized contact friction departs from the value assumed at synthesis. We therefore treat grasp quality as a property of a distribution of friction realizations rather than a single value. Published friction tables [41] report a range of values for a given material pair, motivating a treatment via intervals or distributions rather than a point estimate. This motivates the following definition.

Definition 7 (Risk-Adjusted Margin): The risk-adjusted margin at confidence $\beta \in (0, 1)$, denoted $\varepsilon^{(\beta)}$, is the inscribed-ball radius of the risk-adjusted wrench body of Definition 4,

$$\varepsilon^{(\beta)}(g) := \min_{\|d\|=1} h_{\mathcal{W}^{(\beta)}}(d) = \varepsilon(g, v_\beta), \quad (11)$$

the Ferrari-Canny margin of the GWS assembled at the risk-adjusted friction v_β (Figure 3). We define $\varepsilon^{(\beta)}$ as a signed margin, taking the signed origin-to-facet distance when the




Fig. 3. **Constructing the Risk-Adjusted Margin.** From a contact set and friction prior $p(\mu)$, the risk-adjusted grasp wrench space $\mathcal{W}^{(\beta)} = \mathcal{W}_n + v_\beta \mathcal{W}_t$ scales the tangential wrench body by the friction CVaR v_β ; a single linear program inscribes the largest ball of radius $\varepsilon^{(\beta)}$ in $\mathcal{W}^{(\beta)}$ (Equation (11)), whose positivity certifies force closure with probability at least β (Theorem 1).

TABLE I

COMMON GRASP QUALITY FUNCTIONS AND THEIR RISK-SENSITIVE ADAPTATIONS. ADAPTED FROM TABLE 2 IN [42]. PROB. = PROBABILISTIC UNCERTAINTY MODEL (SEE SEC. IX-D). THE SHADED ROW IS THE MEMBER THE PRESENT PAPER INSTANTIATES ANALYTICALLY. THE $\ell^{*(\beta)}$ MEMBER DRIVES THE SYNTHESIS STUDY OF SEC. X-B.

Quality Q	Risk-sens. Q_β	Unc. variable	Unc. model
ε (Ferrari-Canny)	$\varepsilon^{(\beta)}$	μ	Prob.
Min-weight [6]	$\ell^{*(\beta)}$	μ	Prob.

origin leaves the body, so a negative $\varepsilon^{(\beta)}$ measures the margin by which the adverse friction tail violates closure. Our metric recovers the classical ε under a point-mass prior, approaches the mean-friction margin $\varepsilon(g, \mathbb{E}[\mu])$ as $\beta \rightarrow 0$, and the worst-case margin $\varepsilon(g, \text{ess inf } \mu)$ as $\beta \rightarrow 1$.

Concretely, we treat grasp robustness as a quasi-static property of the synthesized grasp under uncertain contact friction. The realized grasp-quality margin is the Ferrari-Canny metric $\varepsilon(g, \mu)$, positive when the grasp is force-closed and negative otherwise, and we write $\varepsilon_{\text{nom}} := \varepsilon(g, \mu_{\text{nom}})$ for its value at the nominal friction μ_{nom} . Rather than evaluating the realized margin in expectation, we adopt a risk-sensitive formulation,

$$Q_\beta(\theta) = \text{CVaR}_\beta(\varepsilon(g, \mu)), \quad \mu \sim p_\mu, \quad (12)$$

the mean of the worst $(1 - \beta)$ -fraction of margin outcomes under the lower-tail convention of (1). Two risk summaries of the same margin distribution therefore coexist: $\varepsilon^{(\beta)}$ discounts the friction coefficient first and evaluates the margin once, while Q_β evaluates the margin at every friction realization and averages its adverse tail. Proposition 2 relates the two, and each admits the closure certificate of Sec. VII. For N Monte Carlo friction samples $\{\mu_i\}_{i=1}^N \sim p_\mu$, $Q_\beta(\theta)$ is computed by sorting $\{\varepsilon(g, \mu_i)\}$ in ascending order and averaging the worst $\lceil (1 - \beta)N \rceil$ values, which penalizes low-probability, high-impact failures. The analytic protocol of Sec. IX-E avoids the Monte Carlo sampling entirely by evaluating $\varepsilon^{(\beta)}$. A dynamic extension that evaluates the margin over a finite-horizon execution rollout, rather than at the grasp instant, is the domain of the execution-time controller and lies outside the scope of the present study.

VII. THEORETICAL PROPERTIES

Three properties make Q_β usable in practice. The closure certificate of Theorem 1, established in Sec. IV, already turns a positive margin into a force-closure guarantee that holds with probability at least β . Monotonicity in β (Theorem 2) lets a user dial conservatism, and differentiability (Theorem 3) exposes a gradient for optimization.

Theorem 2 (Monotonicity of Q_β in β): For fixed grasp parameters θ , $Q_\beta(\theta)$ is non-increasing in β .

Proof: Under the lower-tail convention, $Q_\beta(\theta) = \text{CVaR}_\beta(Z)$ with $Z = \varepsilon(g, \mu)$ admits the quantile-integral representation

$$\text{CVaR}_\beta(Z) = \frac{1}{1-\beta} \int_0^{1-\beta} F_Z^{-1}(u) du,$$

where F_Z^{-1} is the quantile function of Z . The right-hand side is the running average of F_Z^{-1} over $[0, 1-\beta]$, and F_Z^{-1} is non-decreasing on $[0, 1]$. For $\beta' > \beta$, the upper endpoint shrinks from $1-\beta$ to $1-\beta'$, so the average moves toward the left end of the quantile function, which is smaller. Hence $\text{CVaR}_{\beta'}(Z) \leq \text{CVaR}_\beta(Z)$. ■

Corollary 1: Higher risk aversion ($\beta \nearrow 1$) yields more conservative synthesis objectives, at the cost of potentially overpenalizing grasps that are robust under typical disturbances.

Theorem 3 (Differentiability of Q_β): If $\varepsilon(g, \mu)$ is differentiable in θ for almost every μ , then $Q_\beta(\theta) = \text{CVaR}_\beta(\varepsilon(g, \mu))$ is differentiable in θ almost everywhere, with

$$\nabla_\theta Q_\beta(\theta) = \mathbb{E}[\nabla_\theta \varepsilon(g, \mu) \mid \varepsilon(g, \mu) \leq \text{VaR}_\beta(\varepsilon(g, \mu))]. \quad (13)$$

Proof: The Ferrari-Canny margin $\varepsilon(g, \mu)$ is the value of a linear program and is piecewise smooth in θ , differentiable away from a measure-zero set of active-vertex changes. For a continuous friction distribution, CVaR is differentiable by the envelope (Danskin) theorem applied to the Rockafellar-Uryasev infimum representation [39], with the infimizing v^* equal to VaR_β , which yields (13). In the sample-average approximation, CVaR is piecewise linear in the sorted tail values, with subgradients available via the sorting permutation, which is a.e. differentiable. ■

Taken together, Theorem 1 certifies grasps through Q_β , whereas the analytic protocol of Sec. IX-E evaluates $\varepsilon^{(\beta)}$. The next two results connect these margins. The first ties them to the risk-sensitive directional support of Definition 5, and the second gives $\varepsilon^{(\beta)}$ its own closure certificate.

Proposition 2 (Relations Among the Risk-Sensitive Margins): For any θ and $\beta \in (0, 1)$,

$$Q_\beta(\theta) \leq \min_{\|d\|=1} h_\beta(d; \theta) \quad \text{and} \quad \varepsilon^{(\beta)}(g) \leq \min_{\|d\|=1} h_\beta(d; \theta), \quad (14)$$

while Q_β and $\varepsilon^{(\beta)}$ admit no general ordering between themselves. The minimum over directions pulls Q_β below $\varepsilon^{(\beta)}$, and Jensen's inequality along any single direction pulls it above.

Proof: For every d , $\varepsilon(g, \mu) \leq h_{\mathcal{W}(\mu)}(d)$ pointwise in μ ; applying the monotone functional CVaR_β and then

Algorithm 1: Risk-Sensitive Grasp Synthesis

Input: Object model, hand model, friction prior p_μ , confidence level β , seed budget S

Output: Scored grasp pool $\mathcal{G}_{\text{pool}}$, best grasp g^*

$\mathcal{G}_{\text{pool}} \leftarrow \emptyset$; $v_\beta \leftarrow \text{CVaR}_\beta(\mu)$, $\mu \sim p_\mu$

for $s = 1, \dots, S$ **do**

 Sample a seeded palm pose and finger preshape

 Refine the sample with the force-closure

 synthesis NLP

if the refined grasp certifies force closure via (15), $\ell^* > 0$ **then**

 Add g_s , with contact set $\mathcal{C}_s$ and primitive wrenches, to $\mathcal{G}_{\text{pool}}$

foreach $g_s \in \mathcal{G}_{\text{pool}}$ **do**

 Assemble $\mathcal{W}^{(\beta)} = \mathcal{W}(v_\beta)$ from the stored primitive wrenches

$\varepsilon^{(\beta)}(g_s) \leftarrow$ signed inscribed-ball radius of $\mathcal{W}^{(\beta)}$

$g^* \leftarrow \arg \max_{g_s \in \mathcal{G}_{\text{pool}}} \varepsilon^{(\beta)}(g_s)$

return ($\mathcal{G}_{\text{pool}}$, g^*)

minimizing over d yields the first inequality. The second inequality restates Proposition 1. Per direction, $h_{\mathcal{W}^{(\beta)}}(d) \leq h_\beta(d; \theta)$, and minimizing both sides over d preserves the bound. ■

Corollary 2 (Closure Certificate for $\varepsilon^{(\beta)}$): If $\varepsilon^{(\beta)}(g) > 0$, then $\Pr[\varepsilon(g, \mu) > 0] \geq \beta$. The guarantee mirrors Theorem 1 but follows from friction monotonicity rather than from the CVaR tail bound. The GWS nests in μ , so $\varepsilon(g, \cdot)$ is non-decreasing, $\varepsilon^{(\beta)} > 0$ implies closure for every $\mu \geq v_\beta$, and $v_\beta \leq \text{VaR}_\beta(\mu)$ gives $\Pr[\mu \geq v_\beta] \geq \beta$.

VIII. SYNTHESIS ALGORITHM

Theorem 3 provides a basis for gradient-based grasp optimization. In the present work, however, we approximate the maximizer of Problem 1 with a sample, certify, and score strategy layered on an established force-closure synthesizer. Algorithm 1 summarizes the procedure, and we describe its stages next.

The procedure begins with a sampling stage that draws seeded palm poses and finger preshapes around the object and refines each sample through the FRoGGeR min-weight nonlinear program [6], which certifies force closure ($\ell^* > 0$) at convergence and serializes the contact set and primitive wrenches of every accepted grasp. Next, the scoring stage assembles the risk-adjusted body $\mathcal{W}^{(\beta)}$ once per grasp at the risk-adjusted friction v_β and computes $\varepsilon^{(\beta)}$ as its signed inscribed radius (Definition 7), so the risk adjustment scores an entire pool with one convex-hull computation per grasp and no simulation. The selection stage keeps the scored pool for the aggregate studies and returns the risk-ranked best grasp where a single grasp is executed, as in the per-object analysis of Table VI. Because Algorithm 1 scores a sampled pool rather than computing gradients (e.g., $\nabla_\theta Q_\beta$), it does not exploit the differentiability established in Theorem 3. In

the synthesis study of Sec. X-B, we take one step in the differentiability direction by optimizing the differentiable risk-adjusted min-weight objective inside the synthesizer itself. Nevertheless, we leave a gradient-based synthesis extension of our pipeline to future work.

IX. EXPERIMENTS

The protocol below evaluates our risk-sensitive metrics on stored grasp datasets across three robotic hands.

A. Robotic Hand Platforms and Grasp Datasets

We evaluate our metrics on grasp datasets for several anthropomorphic hands, the LEAP [43] and Allegro hands for the aggregate and synthesis studies, the RealHand L6 for the per-object analysis, and the Shadow Hand for the cross-dataset study (Sec. X-D). Grasps come from a Drake-backed pipeline derived from FRoGGeR [6], which generates candidate contacts, filters them by opposition geometry, and certifies force closure analytically through the min-weight linear program of (15) at synthesis time. Each accepted grasp serializes its grasp map and contact set, and our metrics run analytically from the stored logs without re-running the simulator. The dynamic six-axis shake [44], [45] in Sec. X-E is a separate predictive-validity study on these stored grasps, not a synthesis-time verification. We run headless for synthesis and use Drake’s offscreen rendering for visualization. The same pipeline also retargets to other arm-hand systems. The appendix documents a ZArm-mounted LEAP variant with arm-reachability and fixture-collision gates, together with a qualitative gallery of representative certified-robust grasps (Figure 14).

B. Object Set

Each study draws on a study-specific object slate. The aggregate study (Sec. X-A) evaluates a combined set of 1,599 force-closed grasps, 800 on LEAP and 799 on Allegro, over eight objects, four geometric primitives (sphere, cube, box, cylinder) and four YCB objects [46] (tennis ball, master chef can, potted meat can, mustard bottle). The broader dexterous grasping literature draws on a wider range of standardized object sets and benchmarks beyond YCB, including the Columbia Grasp Database [47], the GraspIt! simulator’s object library [48], the Yale-CMU-Berkeley manipulation benchmark [49], and hand-pose-annotated datasets such as DexYCB [50] and GRAB [51], together with the GRASP taxonomy of human grasp types [52] and the MANO hand model [53] used to parameterize several of these resources. The predictive-validity study (Sec. X-E) drops the sphere, since the rotational symmetry of a sphere collapses the six shake directions to a single representative outcome and stops the test from discriminating on orientation, and reports 641 LEAP grasps over the remaining seven objects. The per-object analysis (Table VI) reports one representative RealHand L6 grasp for each of seven objects (cube, box, cylinder, potted meat can, mustard bottle, tomato soup can, tennis ball), a set that swaps the tomato soup can in for the master chef can to keep every object representable in the L6

aperture. The primitives provide controlled baselines with known symmetry, while the YCB objects introduce realistic surface geometry and scale variation.

C. Grasp Synthesis

For each object we draw 100 seeded synthesis attempts. Each attempt samples a palm pose and finger preshape around the object, refines it through the FRoGGeR nonlinear program, and keeps the result only when the solver converges to a force-closed grasp ($\ell^* > 0$ via (15)). The resulting protocol yields the 800 LEAP and 799 Allegro grasps of the aggregate study and the RealHand L6 pool behind Table VI. Scoring follows the analytic protocol of Sec. IX-E, one risk-adjusted hull per grasp, with no Monte Carlo sampling at synthesis time.

D. Friction Priors

We model friction uncertainty with two priors. The nominal prior is a tight Gaussian, $\mu \sim \mathcal{N}(0.7, 0.1^2)$, centered at a typical rubber-on-plastic coefficient with mild spread. The adverse prior is a bimodal mixture, $0.5\mathcal{U}(0.7, 1.0) + 0.5\mathcal{U}(0.1, 0.3)$, that places half its mass in a slippery low-friction tail to model surfaces with mixed-friction populations. The nominal prior probes the regime where the risk adjustment is mild, and the adverse prior stresses friction sensitivity.

E. Analytic Evaluation Protocol

We evaluate every grasp in the stored dataset analytically, without Monte Carlo sampling. By Lemma 1 the primitive wrenches are affine in the friction coefficient, so the risk-adjusted quality $\varepsilon^{(\beta)}$ of Definition 7 reduces to one inscribed-radius computation per grasp. We assemble the polytope $\mathcal{W}^{(\beta)} = \mathcal{W}(v_\beta)$ at the risk-adjusted friction $v_\beta = \text{CVaR}_\beta(\mu)$ of the prior, estimated once per prior from 2×10^5 samples, build its convex hull, and take the minimum origin-to-facet distance. We report $\varepsilon^{(\beta)}$ at confidence $\beta = 0.9$ under each prior, taking $\varepsilon^{(\beta)}$ as a continuous signed margin, negative when the risk tail violates closure, so it ranks grasps without a clamp. The sign of $\varepsilon^{(\beta)}$ is the closure certificate of Corollary 2.

F. Baselines

We compare against two baselines. The first is the nominal Ferrari-Canny epsilon ε_{nom} , which evaluates the GWS inscribed-ball radius at the fixed nominal friction $\mu = 0.7$. The aggregate and per-object analyses evaluate ε_{nom} with each cone edge normalized to a unit contact force, the convention of [1], while the predictive and cross-dataset studies evaluate it on the same unit-normal-force cone that defines $\varepsilon^{(\beta)}$. At a fixed friction the two conventions differ by the constant factor $(1 + \mu_{\text{nom}}^2)^{-1/2} \approx 0.82$, a global rescaling that leaves every sign, ranking, correlation, and certified fraction unchanged. The second is the min-weight metric ℓ^* from FRoGGeR [6], computed by solving the linear program

$$\max_{\alpha, \ell} \ell \quad \text{s.t. } W\alpha = 0, \mathbf{1}^\top \alpha = 1, \alpha_i \geq \ell \forall i, \quad (15)$$

where W is the $6 \times m$ matrix of primitive wrenches, $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_m]^\top \in \mathbb{R}^m$ is a vector of non-negative weights, one per GWS friction-cone generator, and $m = n_c n_s$ for n_c contacts and n_s cone edges. We report the raw ℓ^* , which Li et al. [6] normalize to $m \cdot \ell^*$ as a practical quality indicator for precision grasps.

G. Simulated Execution Studies

Our execution testbed replays every certified grasp through the Drake stack of Sec. A, which transfers contact through a small pad welded at each synthesized fingertip, closes the hand to its certified configuration, and subjects the held object to the gravity-off six-axis shake and, for a per-object subset, to the gravity-on lift with a graded lateral pull of Sec. X-E.3. We run two studies on the execution testbed. The first establishes that the certified configurations execute, in the sense that the closure acquires the object and the acquired object stays held at the nominal friction. The second tests whether our risk metric $\varepsilon^{(\beta)}$ predicts object retention under execution-time friction uncertainty, in particular under slips arising from low-friction surfaces.

First, we evaluate grasp executability. We replay the seven-object certified pool of Sec. X-E through the execution stack, where each grasp transfers contact through its certified pads, the hand closes to the synthesized configuration, and the simulator applies the six-axis shake at the nominal friction $\mu = 0.7$. Note that the present study measures whether the certified configurations acquire and hold the objects but does not isolate the friction robustness our metric targets, since every grasp executes at the friction it was certified against. Table II reports the per-object outcomes. Contact transfer is the dominant attrition stage, 302 of the 641 replayed grasps establish contact, while execution after transfer is reliable, 279 of the 302 established grasps hold through the nominal shake and all 51 established grasps in the lift subset hold through the nominal lift.

Second, we test our metric’s claim through a controlled friction contrast on the same executed pool. Rather than synthesizing one grasp per objective, we retain the established grasps whose nominal margin ε_{nom} sits at or above its per-object median, so object identity does not drive the split, and within this high-nominal pool the per-object median of $\varepsilon^{(\beta)}$ separates a high- $\varepsilon^{(\beta)}$ group, high nominal margin and high risk margin, from a low- $\varepsilon^{(\beta)}$ group, high nominal margin and low risk margin. The two groups are the executed analog of the nominal-versus-risk contrast, both look equally certifiable to ε_{nom} , and only $\varepsilon^{(\beta)}$ distinguishes them before execution. Table III reports the retention of each group across the executed friction sweep $\mu \in \{0.2, 0.3, 0.4, 0.7\}$. At the nominal friction the arms hold at matching rates, 0.93 against 0.96, which is the regime where the nominal margin certifies both. As the friction drops the arms separate, 0.69 against 0.50 at $\mu = 0.4$, 0.57 against 0.36 at $\mu = 0.3$, and 0.41 against 0.32 at $\mu = 0.2$, so among grasps the nominal margin rates equally, $\varepsilon^{(\beta)}$ identifies before execution which ones fail once the friction assumption breaks.

TABLE II

SIMULATED GRASP EXECUTABILITY PER OBJECT. ATTEMPTED: CERTIFIED GRASPS REPLAYED THROUGH THE EXECUTION STACK; GRIP: THE PAD TRANSFER ESTABLISHES CONTACT; SHAKE HOLD: THE ESTABLISHED GRASP HOLDS THE GRAVITY-OFF SIX-AXIS SHAKE AT $\mu = 0.7$; LIFT GRIP AND LIFT HELD: CONTACT ESTABLISHED AND OBJECT HELD AT $\mu = 0.7$ IN THE 15-GRASP PER-OBJECT GRAVITY-ON LIFT SUBSET. MCC: MASTER CHEF CAN. SALMON MARKS THE LOWEST GRIP RATE AND MINT THE HIGHEST; BOLD MARKS LIFT HELD, WHICH EQUALS LIFT GRIP IN EVERY ROW.

Object	Attempted	Grip	Shake Hold	Lift Grip	Lift Held
Cube	100	36	31	4	4
Box	98	45	43	7	7
Cylinder	100	57	52	10	10
MCC	96	52	49	7	7
Mustard	82	33	27	5	5
Meat Can	98	28	26	4	4
Tennis Ball	67	51	51	14	14

TABLE III

FRICTION-CONTRAST RETENTION ON EXECUTED GRASPS. RETENTION RATE OF THE TWO HIGH- ε_{nom} GROUPS, SPLIT AT THE PER-OBJECT $\varepsilon^{(\beta)}$ MEDIAN INTO A HIGH- $\varepsilon^{(\beta)}$ AND A LOW- $\varepsilon^{(\beta)}$ GROUP, AT EACH EXECUTED FRICTION OF THE SHAKE SWEEP. THE TWO GROUPS TIE AT THE NOMINAL FRICTION AND SEPARATE AS THE FRICTION DROPS. MINT MARKS THE RISK-FAVORED GROUP, SALMON THE GROUP THE RISK MARGIN EXCLUDES.

Group	n	$\mu = 0.2$	$\mu = 0.3$	$\mu = 0.4$	$\mu = 0.7$
High ε_{nom} , high $\varepsilon^{(\beta)}$	109	0.41	0.57	0.69	0.93
High ε_{nom} , low $\varepsilon^{(\beta)}$	44	0.32	0.36	0.50	0.96

X. RESULTS AND DISCUSSION

The aggregate analysis spans a combined dataset of 1,599 force-closed grasps, 800 on the LEAP hand and 799 on the Allegro hand over the eight core test objects. Across the grasp data, we compare our risk metric $\varepsilon^{(\beta)}$ to the nominal Ferrari-Canny margin ε_{nom} and to the FRoGGeR min-weight baseline ℓ^* under two friction priors, a tight Gaussian $\mathcal{N}(0.7, 0.1^2)$ and an adverse mixture with equal weights $0.5\mathcal{U}(0.7, 1.0) + 0.5\mathcal{U}(0.1, 0.3)$, at confidence $\beta = 0.9$. Every number traces to a stored evaluation log that records the convention, the priors with their risk-adjusted frictions, and the method.

A. Aggregate Metric Divergence

Under a tight Gaussian friction prior our risk metric $\varepsilon^{(\beta)}$ tracks the nominal Ferrari-Canny margin, with a combined Pearson correlation of 0.95 (Table IV). The risk adjustment at the nominal prior is mild, so the two metrics rank grasps almost identically. Under the adverse prior the picture changes. The correlation of $\varepsilon^{(\beta)}$ with ε_{nom} falls to 0.32 and with ℓ^* to 0.52, so our risk metric separates from both friction-unaware baselines (Figure 4). Figure 5 traces the resulting degradation as the confidence level β rises, the nominal-prior correlation

TABLE IV
AGGREGATE METRIC RELATIONSHIPS (PEARSON r , 1,599 GRASPS).
THE SHADED ROWS SHOW THE ADVERSE PRIOR, WHERE OUR RISK METRIC SEPARATES FROM BOTH FRICTION-VOLATILITY-UNAWARE BASELINES.

Metric Pair	LEAP	Allegro	Combined
ε_{nom} vs ℓ^*	0.62	0.62	0.60
$\varepsilon^{(\beta)}$ (nominal prior) vs ε_{nom}	0.95	0.95	0.95
$\varepsilon^{(\beta)}$ (nominal prior) vs ℓ^*	0.72	0.71	0.69
$\varepsilon^{(\beta)}$ (adverse prior) vs ε_{nom}	0.38	0.30	0.32
$\varepsilon^{(\beta)}$ (adverse prior) vs ℓ^*	0.57	0.52	0.52




Fig. 4. **Aggregate Divergence Under the Adverse Prior.** Risk metric $\varepsilon^{(\beta)}$ under the adverse friction prior against the nominal Ferrari-Canny ε_{nom} , one point per grasp over the combined 1,599-grasp dataset. Every grasp is force-closed at nominal friction, so all points sit at positive ε_{nom} , yet the shaded population falls below the closure boundary in the adverse tail. The nominal margin overrates these grasps.

holding near 0.95 while the adverse-prior correlation settles near 0.30.

Specifically, the divergence is concentrated in a friction-sensitive majority. Of the 1,599 grasps that are force-closed at nominal friction, 849 (53%) lose closure in the adverse tail, where $\varepsilon^{(\beta)}$ goes non-positive (57% on LEAP, 49% on Allegro). The set is geometric rather than a synthesizer defect. Rounded objects stay robust, the sphere at 5% and the tennis ball at 13%, while irregular objects are the most friction-sensitive, the mustard bottle at 86% and the potted meat can at 80%. Our metric, therefore, identifies friction-sensitive grasps in that ε_{nom} overrates such grasps while $\varepsilon^{(\beta)}$ correctly marks them, adding a unique robustness axis friction-volatility-unaware metrics do not encode.

B. Synthesis Under the Risk Objective

Table V and Figure 6 demonstrate the risk adjustment’s use in a synthesis objective, beyond evaluation. We run FRoGGeR synthesis under the differentiable risk objective, the min-weight program of (15) solved on the risk-adjusted body $\mathcal{W}^{(\beta)}$ (the $\ell^{*(\beta)}$ member of Table I), and under the nominal min-weight objective at matched budget over the eight core objects on both hands, then score every synthesized grasp with our metric $\varepsilon^{(\beta)}$ under the adverse prior. At matched synthesis time the risk objective lowers the friction-sensitive




Fig. 5. **Correlation Degradation and Object Stratification.** Left: the Pearson correlation of $\varepsilon^{(\beta)}$ with ε_{nom} as the confidence β rises, holding near 0.95 under the nominal prior and falling toward 0.30 under the adverse prior. Right: the per-object fraction of nominally force-closed grasps that lose closure in the adverse tail, from the rounded sphere to the irregular mustard bottle.

TABLE V
SYNTHESIS UNDER THE RISK OBJECTIVE VERSUS THE MIN-WEIGHT OBJECTIVE (COMBINED LEAP AND ALLEGRO, ADVERSE PRIOR, $\beta = 0.9$). THE SHADED ROW IS THE RISK OBJECTIVE.

Synth. Obj.	Friction-Sens. ↓	Median $\varepsilon^{(\beta)}$ ↑	Solve (s) ↓
Risk $\ell^{*(\beta)}$	72/125 (57.6%)	−0.00119	19.3
Min-weight ℓ^*	86/127 (67.7%)	−0.00322	14.9

fraction from 67.7 to 57.6% combined, and on LEAP alone from 87.3 to 75.4%, while raising the median adverse-tail margin from −0.00322 to −0.00119. The differentiable objective the FRoGGeR pipeline exposes is parameterized by a Gaussian friction model, so it optimizes a mild Gaussian-prior risk while the evaluation applies the harsher adverse prior. The improvement is therefore best understood as a transfer result. Grasps tuned for a mild risk withstand the harsher adverse tail more often than the min-weight grasps the FRoGGeR baseline produces.

Next, Table VI reports a representative RealHand L6 grasp for each of seven objects under the analytic protocol of Sec. IX-E, the force-closed grasp that maximizes $\varepsilon^{(\beta)}$ under the nominal prior at $\beta = 0.9$. For each grasp it lists the nominal Ferrari-Canny margin ε_{nom} , the FRoGGeR min-weight ℓ^* , our risk metric $\varepsilon^{(\beta)}$ under the nominal prior across confidence levels and under the adverse prior at $\beta = 0.9$, and the adverse-prior closure certificate β_c .

C. Per-Object Analysis

The per-object data reproduce the aggregate divergence of Sec. X-A at the level of single grasps and confirm the analytic properties of Sec. VII. Under the nominal prior our risk metric $\varepsilon^{(\beta)}$ tracks the nominal Ferrari-Canny margin, so both rate the seven grasps as comparable (Figure 7). Under the adverse prior, $\varepsilon^{(\beta)}$ separates them, falling to a small positive margin for the rounded tennis ball and cylinder and below the closure boundary for the cube, box, and mustard bottle, exposing the per-object friction sensitivity that the friction-unaware margin cannot capture. Figure 11 shows twelve such RealHand L6 grasps spanning household items, primitive

TABLE VI

PER-OBJECT REALHAND L6 ANALYSIS (ANALYTIC, REPRESENTATIVE GRASP PER OBJECT). n_c : CONTACT COUNT; ε_{nom} : FERRARI-CANNY AT $\mu=0.7$; ℓ^* : FROGGER MIN-WEIGHT; $\varepsilon_{\mathcal{N}}^{(\beta)}$: RISK METRIC UNDER THE NOMINAL PRIOR $\mathcal{N}(0.7, 0.1^2)$ AT CONFIDENCE β ; $\varepsilon_{\text{adv}}^{(0.9)}$: RISK METRIC UNDER THE ADVERSE PRIOR AT $\beta=0.9$; β_c : LARGEST ADVERSE-PRIOR CONFIDENCE WITH $\varepsilon^{(\beta)} > 0$. SALMON CELLS LOSE CLOSURE IN THE ADVERSE TAIL ($\varepsilon_{\text{adv}}^{(0.9)} \leq 0$), MINT CELLS RETAIN IT, AND THE MUSTARD BOTTLE PAIRS THE WEAKEST CERTIFICATE WITH THE DEEPEST LOSS IN THE ADVERSE TAIL.

Object	n_c	ε_{nom}	ℓ^*	$\varepsilon_{\mathcal{N}}^{(0.5)}$	$\varepsilon_{\mathcal{N}}^{(0.8)}$	$\varepsilon_{\mathcal{N}}^{(0.9)}$	$\varepsilon_{\mathcal{N}}^{(0.95)}$	$\varepsilon_{\text{adv}}^{(0.9)}$	β_c
Cube	5	0.0075	0.0211	0.0083	0.0076	0.0072	0.0067	-0.0003	0.85
Box	5	0.0099	0.0268	0.0103	0.0090	0.0082	0.0075	-0.0001	0.87
Cylinder	5	0.0078	0.0266	0.0085	0.0078	0.0073	0.0069	0.0010	1.00
Potted Meat Can	5	0.0089	0.0307	0.0100	0.0092	0.0085	0.0079	0.0003	0.99
Mustard Bottle	5	0.0081	0.0259	0.0086	0.0077	0.0071	0.0065	-0.0036	0.46
Tomato Soup Can	5	0.0091	0.0328	0.0099	0.0087	0.0080	0.0075	0.0003	0.97
Tennis Ball	5	0.0070	0.0294	0.0075	0.0066	0.0061	0.0057	0.0010	1.00




Fig. 6. **Synthesis Under the Risk Objective.** Left: the fraction of synthesized grasps that are friction-sensitive under the adverse prior, lower for the risk objective than for the min-weight objective at matched budget. Right: the distribution of $\varepsilon^{(\beta)}$ under the adverse prior, with the risk objective shifted toward the closure boundary.




Fig. 7. **Per-Object Quality Comparison.** Nominal Ferrari-Canny ε_{nom} , our risk metric $\varepsilon^{(\beta)}$ under the nominal prior, and $\varepsilon^{(\beta)}$ under the adverse prior at $\beta = 0.9$, for each RealHand L6 grasp. The nominal prior tracks the friction-unaware margin, while the adverse prior collapses toward and below the closure boundary, sharply for the mustard bottle.

shapes, and rounded objects, each force-closed ($\ell^* > 0$) and satisfying the friction-tail requirement $\varepsilon^{(0.99)} > 0$ ($\varepsilon^{(\beta)}$ at $\beta = 0.99$ under the nominal prior, Definition 7).

Empirically, the data confirm the two analytic guarantees of Sec. VII. The margin $\varepsilon^{(\beta)}$ is non-increasing in the confidence β , since it evaluates the non-decreasing piecewise-linear quality curve of Remark 2 at a risk-adjusted friction v_β that falls as β rises (the $\varepsilon^{(\beta)}$ analog of Theorem 2), and this holds for every grasp across both priors in Table VI and Figure 8. By Corollary 2, the sign of $\varepsilon^{(\beta)}$ is a probabilistic closure certificate. Under the nominal prior every selected grasp is certified at all confidence levels, while under the




Fig. 8. **Risk Sensitivity to the Confidence Level.** Our risk metric $\varepsilon^{(\beta)}$ against confidence β , summarized over the seven representative RealHand L6 grasps of Table VI as a median line with a min-to-max band. Under the nominal prior $\varepsilon^{(\beta)}$ stays above the closure boundary. Under the adverse prior the band crosses into the loss-of-closure region as β rises. $\varepsilon^{(\beta)}$ is non-increasing in β (the $\varepsilon^{(\beta)}$ analog of Theorem 2, via Remark 2).

adverse prior our certificate β_c tightens with friction sensitivity, from $\beta_c = 1.0$ for the cylinder and tennis ball down to $\beta_c = 0.46$ for the mustard bottle. The FRoGGeR min-weight ℓ^* varies little across the seven grasps, consistent with its role as a friction-unaware force-closure margin that our risk metric complements rather than replaces.

Remark 2 (Piecewise-Linear Quality Curve): Because Lemma 1 establishes that each primitive wrench is affine in μ , the support function $h_{\mathcal{W}(\mu)}(d)$ is piecewise linear and convex in μ for every fixed direction d , a maximum of affine vertex supports. The Ferrari-Canny margin $\varepsilon(g, \mu) = \min_{\|d\|=1} h_{\mathcal{W}(\mu)}(d)$ is therefore piecewise linear in μ . For any grasp there exists the critical friction coefficient $\mu^* = \inf\{\mu : \varepsilon(g, \mu) > 0\}$ of the Necessity remark below which force closure is lost. Plotting $\varepsilon(g, \mu)$ against μ directly reveals μ^* and shows what fraction of the prior p_μ falls below it (Figure 9). The resulting curve is the most direct expression of the analytic reduction, a single sweep of μ recovers the full failure boundary without Monte Carlo sampling, and $\varepsilon^{(\beta)}$ evaluates the curve at the risk-adjusted friction v_β , the mean of the adverse tail that the confidence level selects.




Fig. 9. **Per-Grasp Quality Curve.** The Ferrari-Canny margin $\varepsilon(g, \mu)$ of the mustard-bottle and tennis-ball grasps of Table VI, swept over the friction coefficient. The curve is piecewise linear (Remark 2) and comes from the stored grasp maps in the unit-normal-force convention of $\varepsilon^{(\beta)}$, with no sampling. The zero crossing marks the critical friction μ^* below which closure is lost. The mustard bottle’s μ^* sits above the adverse-prior risk-adjusted friction $v_{0.9}$, so its adverse margin $\varepsilon_{\text{adv}}^{(0.9)}$ is negative, while the tennis ball’s μ^* sits below $v_{0.9}$ and its certificate holds across the sweep.

D. Cross-Dataset Validation on FRoGGeR and DexGraspNet

To test whether the divergence of Sec. X-A generalizes beyond the LEAP and Allegro datasets, we evaluate the same analytic signed $\varepsilon^{(\beta)}$ across two public grasp datasets and a third hand. We compute $\varepsilon^{(\beta)}$ on 800 LEAP and 799 Allegro grasps synthesized with a FRoGGeR-derived pipeline [6], and on 400 grasps from the DexGraspNet release [13] on the Shadow Hand, sampled over 50 core-bottle objects with eight grasps each. The FRoGGeR datasets expose contact points and inward normals directly. For DexGraspNet we obtain fingertip positions by forward kinematics on the Shadow Hand model and project them onto the scaled object mesh.

Specifically, every fraction uses the analytic signed $\varepsilon^{(\beta)}$ at confidence $\beta = 0.9$ under the two priors of Sec. IX-D, with a grasp certified robust when $\varepsilon^{(\beta)} > 0$, the same certificate the rest of the paper uses. On both FRoGGeR datasets our ε_{nom} correlates strongly with FRoGGeR’s own Ferrari-Canny field, $r = 0.96$ on LEAP and $r = 0.94$ on Allegro, an independent cross-check of the two pipelines. Table VII reports the certified-robust fractions. Under the nominal prior both FRoGGeR datasets certify at least 99% of grasps and the DexGraspNet subset 33.2%. Under the adverse prior the LEAP fraction falls to 42.9%, Allegro to 50.9%, and DexGraspNet to 3.8%, with the per-hand distributions in Figures 10a to 10c. The pattern matches the aggregate study, classical force closure is overconfident under realistic friction, and the certified fraction shrinks as the prior shifts mass into the slippery tail.

In particular, the cross-dataset pattern reflects a property of our metric family rather than a defect of any synthesizer. Both datasets were produced for purposes other than risk-sensitive evaluation, and the authors that released them did so on the basis of metrics (Ferrari-Canny ε , ℓ^* , simulator success rate) that do not encode the friction prior used here.

TABLE VII
FRACTION OF GRASPS CERTIFIED ROBUST ($\varepsilon^{(\beta)} > 0$ AT $\beta=0.9$)
ACROSS THREE HAND-DATASET COMBINATIONS AND TWO FRICTION
PRIORS. SALMON-COLORED ROWS SHOW THE ADVERSE-PRIOR
COLLAPSE OF THE CERTIFIED FRACTION.

Dataset	Hand	N	Gaussian Prior	Adverse Prior
FRoGGeR	LEAP	800	99.4%	42.9%
FRoGGeR	Allegro	799	100.0%	50.9%
DexGraspNet	Shadow	400	33.2%	3.8%

The presence of the gap across two independent synthesizers indicates that an explicit-belief margin like $\varepsilon^{(\beta)}$ adds a real new axis to the family of force-closure margins rather than restating an existing one.

E. Predictive Validity Under Dynamic Perturbation

The analytic divergence of Sec. X-A predicts that a grasp the nominal metric rates highly can fail once the contact friction departs from its assumed value. We test the aforementioned prediction in dynamic simulation with the six-axis shake protocol of GenDexGrasp [44], [45], and the risk margin orders realized robustness above both friction-volatility-unaware baselines.

1) *Shake-Based Simulation Stress Test:* We stress each certified grasp in the Drake physics-validation system. The FRoGGeR collision proxy recesses behind the visual fingertip, so we weld a small contact pad at every certified contact point and confirm that the grasp engages the object at the synthesized configuration. We then apply a gravity-off six-axis shake at a fixed contact friction and record how many of the six world-axis directions retain the object. We repeat the shake over the adverse friction sweep $\mu \in \{0.2, 0.3, 0.4\}$ at the nominal friction $\mu_{\text{nom}} = 0.7$. The object body takes on the mesh-and-density mass and moment of inertia, so the disturbance force scales per object and the rotation criterion uses the true inertia rather than a unit surrogate. Of the 641 certified grasps across the seven objects, 302 retain the object in at least one direction at the nominal friction, and we restrict the friction analysis to these, since a grasp that never engages cannot test friction sensitivity.

2) *Realized Robustness Versus the Risk Margin:* Figure 12 ranks realized robustness against three predictors, none of which has access to the realized friction. The risk margin $\varepsilon^{(\beta)}$ (Definition 7) is computed from the adverse prior of Sec. IX-D through its CVaR friction $v_\beta \approx 0.12$, so it scores each grasp once without knowledge of the operating friction. The nominal Ferrari-Canny margin $\varepsilon_{\text{nom}} = \varepsilon(g, \mu_{\text{nom}})$ evaluates the same functional at the assumed friction and does not model the friction uncertainty, and the FRoGGeR min-weight ℓ^* is the third predictor. By the area under the receiver operating characteristic curve (AUC) for the adverse-friction outcome, $\varepsilon^{(\beta)}$ reaches 0.629 against 0.578 for ℓ^* and 0.534 for ε_{nom} , and it ranks above the min-weight in 97.9% and above the nominal margin in 100% of grasp-level paired bootstrap resamples, each gap with a 95% interval clear of




Fig. 10. **Cross-Dataset Distribution of $\varepsilon^{(\beta)}$.** Per-grasp histograms of our risk metric $\varepsilon^{(\beta)}$ at $\beta=0.9$ under the nominal prior $\mathcal{N}(0.7, 0.1^2)$ (blue) and the adverse prior $0.5\mathcal{U}(0.7, 1.0) + 0.5\mathcal{U}(0.1, 0.3)$ (green), for each dataset. The nominal distribution sits above the closure boundary while the adverse distribution shifts across it. The title reports the certified-robust fraction ($\varepsilon^{(\beta)} > 0$) under the nominal and adverse priors, and the adverse mass at or below zero is the certification gap of Table VII.




Fig. 11. **Friction-Error Robust Grasps on the RealHand L6.** Twelve penetration-free, friction-uncertainty-aware grasps span household items (meat and soup cans, mustard and bleach bottles, foam brick), primitive shapes (box, cube, cylinder, sphere), and rounded objects (tennis ball, apple, strawberry), and exercise enveloping, tripod, and pinch closures on the RealHand L6 (rendered in Drake). Every rendered grasp is force-closed ($\ell^* > 0$) and satisfies the friction-tail requirement ($\varepsilon^{(0.99)} > 0$).

zero. The per-grasp Spearman correlation against the graded directions-retained rate orders the predictors the same way, 0.25 for $\varepsilon^{(\beta)}$ against 0.14 for ℓ^* and 0.05 for ε_{nom} . Realized robustness climbs across the $\varepsilon^{(\beta)}$ quartiles, from 0.34 to 0.64, whereas under ε_{nom} it stays near 0.50 across the four bins, so the nominal margin separates neither the least nor the most robust grasps well. A matched-friction oracle with access to the realized friction reaches AUC 0.639, and $\varepsilon^{(\beta)}$ recovers most of the corresponding separation without the friction knowledge.

3) Weight-Bearing Retention Under a Gravity-On Lift:

The shake measures wrench-space quality once contact is established. To probe weight-bearing retention under the same friction sweep, we complement the shake with a gravity-on prismatic lift on the same seven objects. The hand raises the object along a vertical mount, and the pick then applies a lateral pull at graded multiples of object weight while the grip holds under Coulomb friction. The pick grades realized retention by the highest weight multiple the grasp holds before slip, and we rank the resulting outcome against the same three predictors used for the shake. Of the 105

lifts attempted, 51 establish contact under gravity across the seven objects, and we restrict the retention analysis to these. Figure 13 ranks the three predictors against the adverse pick. The risk margin $\varepsilon^{(\beta)}$ reaches AUC 0.78 against 0.75 for the min-weight ℓ^* and 0.67 for the nominal margin ε_{nom} . The success rate climbs from 0.25 to 0.88 across the $\varepsilon^{(\beta)}$ quartiles, while under ε_{nom} the low quartile already sits at 0.40 and ε_{nom} fails to separate the least robust picks. The lead of $\varepsilon^{(\beta)}$ over ε_{nom} is decisive, and the ranking against ℓ^* is directional only, not significant, at the seven-object cluster count. The two dynamic axes rank the grasps consistently against the nominal margin, and the pick extends the evidence into a weight-bearing regime the free-space shake does not model. Table VIII places the two axes side by side.

In addition, a larger re-run reproduces the ordering on both dynamic axes. The gravity-off shake over 329 established-contact grasps scores $\varepsilon^{(\beta)}$ at AUC 0.635 against 0.590 for ℓ^* and 0.561 for ε_{nom} , and the gravity-on pick over 321 established-contact grasps scores 0.744 against 0.714 and 0.641, so $\varepsilon^{(\beta)}$ leads at the larger sample as well.




Fig. 12. **Predictive Validity Under Dynamic Perturbation.** Realized robustness of 302 established-contact grasps under a gravity-off six-axis shake over a friction sweep $\mu \in \{0.2, 0.3, 0.4\}$, pooled across the seven objects, at each object’s mesh-and-density mass and moment of inertia. *Left*: the mean fraction of shake directions that retain the object, binned by predictor quartile, rises with the risk margin $\varepsilon^{(\beta)}$ but saturates under the nominal Ferrari-Canny margin ε_{nom} . *Right*: the area under the receiver operating characteristic curve (AUC) for the adverse-friction outcome, where $\varepsilon^{(\beta)}$, computed only from the friction prior and never from the realized friction, ranks realized robustness above ε_{nom} and the FRoGGeR min-weight ℓ^* , neither of which models friction volatility. The hatched bar marks a matched-friction oracle with access to the realized friction, an upper bound, and the dotted line marks the chance level, an AUC of 0.5, where a predictor ranks robustness no better than a coin flip.




Fig. 13. **Pick Companion Under a Gravity-On Lift.** Realized weight-bearing retention of 51 established-contact grasps under a gravity-on prismatic lift with a graded lateral pull at 1, 3, 6 times object weight, over the friction sweep $\mu \in \{0.2, 0.3, 0.4\}$. *Left*: the success rate binned by predictor quartile climbs with $\varepsilon^{(\beta)}$ and saturates under the nominal Ferrari-Canny margin ε_{nom} . *Right*: the pooled area under the receiver operating characteristic curve (AUC) for the binary adverse pick, where $\varepsilon^{(\beta)}$ ranks above the min-weight ℓ^* and the nominal margin ε_{nom} without access to the realized friction, and the dotted line marks the chance level, an AUC of 0.5.

F. Discussion

The aggregate results expose a mismatch between nominal and distributional grasp assessment. Across the 1,599 LEAP and Allegro grasps the nominal Ferrari-Canny margin ε_{nom} correlates with the adverse-prior $\varepsilon^{(\beta)}$ at only 0.32, down from 0.95 under the tight nominal prior (Sec. X-A), so the two metrics agree when friction is well characterized and part ways when it is uncertain. The divergence is geometric rather than a synthesizer artifact, since it tracks object shape. Rounded objects such as the cylinder and tennis ball stay certified across the adverse tail ($\beta_c = 1.0$), while irregular objects concentrate the friction sensitivity, the mustard bottle losing closure at $\beta_c = 0.46$ (Table VI). A grasp that ε_{nom} rates as adequate can therefore hold most of its wrench capacity in the normal directions and collapse along the friction-dependent ones, a failure mode the scalar nominal margin cannot express but the risk-adjusted margin $\varepsilon^{(\beta)}$

TABLE VIII

PREDICTIVE VALIDITY ACROSS TWO DYNAMIC AXES. POOLED ROC AUC AND PER-GRASP SPEARMAN ρ FOR EACH PREDICTOR AGAINST THE ADVERSE-FRICTION OUTCOME, ON THE GRAVITY-OFF SHAKE ($n = 302$ ESTABLISHED-CONTACT GRASPS) AND THE GRAVITY-ON PICK ($n = 51$ ESTABLISHED-CONTACT GRASPS), OVER THE FRICTION SWEEP $\mu \in \{0.2, 0.3, 0.4\}$. THE SHADED ROW IS OUR RISK MARGIN, WHICH LEADS BOTH DYNAMIC AXES ON BOTH STATISTICS ALL WITHOUT ACCESS TO THE REALIZED FRICTION.

Predictor	Shake, $n = 302$		Pick, $n = 51$	
	AUC $\uparrow$	$\rho \uparrow$	AUC $\uparrow$	$\rho \uparrow$
$\varepsilon^{(\beta)}$	0.63	0.25	0.78	0.55
ℓ^*	0.58	0.14	0.75	0.50
ε_{nom}	0.53	0.05	0.67	0.33
Matched oracle	0.64	n/a	n/a	n/a

captures.

In addition, our risk objective extends the same distinction into synthesis. Table V shows that FRoGGeR run under the differentiable risk objective lowers the friction-sensitive fraction from 67.7% to 57.6% at matched budget, even though our objective targets the mild Gaussian prior while the evaluation applies the harsher adverse prior (Sec. X-B). The selected candidates favor contact configurations whose margin degrades gracefully as friction drops, so $\varepsilon^{(\beta)}$ serves as an evaluation metric and a synthesis signal from one linear program.

G. Limitations

The risk margin $\varepsilon^{(\beta)}$ orders realized dynamic robustness above both friction-unaware baselines and reproduces the friction-sensitivity pattern across three hands and a public dataset. Several aspects of the present evaluation admit principled refinements that we view as natural extensions rather than fundamental obstacles. Specifically, the predictive-validity study of Sec. X-E spans seven objects, which limits an object-level grasp failure prediction analysis. Furthermore, the pooled AUC lead of $\varepsilon^{(\beta)}$ stays stable at 0.629 against 0.578 and 0.534, and a grasp-level paired bootstrap places its 95% interval clear of zero, yet a resampling that groups grasps by object has less power at seven object clusters, so we take the per-object generalization to be directional rather than absolute.

In addition, since a static wrench-space margin only partly predicts a dynamic shake outcome, the absolute discrimination is moderate at AUC 0.629, so we lean towards a comparative interpretation of this result in lieu of a calibrated closure probability. The shake test welds a contact pad at each certified contact and runs gravity-off, so it measures wrench-space quality once contact is established, not the robot achieving the grasp or bearing the object weight against gravity. The pick companion of Sec. X-E.3 adds the weight-bearing axis on the same seven objects, and its object-cluster count is the same, so the two-axis ordering is directional rather than a per-object generalization.

XI. CONCLUSION

We derived a family of risk-sensitive grasp quality metrics grounded in CVaR and established three theoretical properties, monotonicity in the confidence level (Theorem 2), differentiability with respect to grasp parameters (Theorem 3), and a probabilistic closure certificate (Theorem 1). Across the aggregate, cross-dataset, dynamic shake, and synthesis studies, $\varepsilon^{(\beta)}$ consistently identifies friction-sensitive grasps that the nominal Ferrari-Canny margin overrates, reproduces the pattern across a third hand and a public dataset, orders realized adverse robustness above both baselines, and lowers the friction-sensitive fraction of synthesized grasps at matched budget (Secs. X-A, X-B, X-D and X-E), so $\varepsilon^{(\beta)}$ pays off both as an evaluation metric and as a synthesis objective.

Beyond this, several directions for future work remain open. The differentiability result of Theorem 3 provides a foundation for gradient-based synthesis via backpropagation through a differentiable contact simulator. The present sample-based approach does not exploit the resulting capability. Integration with learned contact models could extend the uncertainty distribution beyond parametric friction variation. Finally, hardware validation with real tactile feedback under natural contact uncertainty would confirm whether the risk-sensitive certificates transfer from simulation to physical grasping, a step we leave to future work.

REFERENCES

- [1] C. Ferrari and J. Canny, "Planning Optimal Grasps," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, pp. 2290–2295 vol.3, 1992.
- [2] Z. Li and S. S. Sastry, "Task-Oriented Optimal Grasping by Multi-fingered Robot Hands," *IEEE Journal on Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988.
- [3] J. K. Salisbury and J. J. Craig, "Articulated Hands: Force Control and Kinematic Issues," *The International Journal of Robotics Research*, vol. 1, no. 1, pp. 4–17, 1982.
- [4] B.-H. Kim, S.-R. Oh, B.-J. Yi, and I. H. Suh, "Optimal Grasping Based on Non-Dimensionalized Performance Indices," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, pp. 949–956, IEEE, 2001.
- [5] E. Rimon and J. Burdick, *The Mechanics of Robot Grasping*. Cambridge University Press, 2019.
- [6] A. H. Li, P. Culbertson, J. W. Burdick, and A. D. Ames, "FRoGGeR: Fast Robust Grasp Generation via the Min-Weight Metric," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6809–6816, IEEE, 2023.
- [7] A. H. Li, P. Culbertson, and A. D. Ames, "Toward an Analytic Theory of Intrinsic Robustness for Dexterous Grasping," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2992–2999, 2024. ISSN: 2153-0866.
- [8] C. Enwerem, S. Kalyanaraman, J. S. Baras, and C. Belta, "Variational Neural Belief Parameterizations for Robust Dexterous Grasping under Multimodal Uncertainty," 2026. arXiv:2604.25897.
- [9] J. Weisz and P. K. Allen, "Pose Error Robust Grasping from Contact Wrench Space Metrics," in *2012 IEEE International Conference on Robotics and Automation*, pp. 557–562, 2012.
- [10] C. Enwerem, E. Noorani, J. S. Baras, and B. M. Sadler, "Robust Stochastic Shortest-Path Planning via Risk-Sensitive Incremental Sampling," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pp. 1087–1094, IEEE, 2024.
- [11] C. Enwerem, A. G. Puranic, J. S. Baras, and C. Belta, "Safety-Aware Reinforcement Learning for Control via Risk-Sensitive Action-Value Iteration and Quantile Regression," in *2025 IEEE 64th Conference on Decision and Control (CDC)*, pp. 4890–4895, IEEE, 2025.
- [12] C. Enwerem, J. S. Baras, and C. Belta, "Risk-Constrained Belief-Space Optimization for Safe Control under Latent Uncertainty," 2026. arXiv:2604.03868.
- [13] R. Wang, J. Zhang, J. Chen, Y. Xu, P. Li, T. Liu, and H. Wang, "DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11359–11366, IEEE, 2023. arXiv:2210.02697.
- [14] H. J. T. Suh, T. Pang, T. Zhao, and R. Tedrake, "Dexterous Contact-Rich Manipulation via the Contact Trust Region," *The International Journal of Robotics Research*, 2026. Advance online publication.
- [15] A. H. Li, P. Culbertson, V. Kurtz, and A. D. Ames, "DROP: Dexterous Reorientation via Online Planning," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 14299–14306, 2025.
- [16] A. H. Li, P. Culbertson, and A. D. Ames, "PONG: Probabilistic Object Normals for Grasping via Analytic Bounds on Force Closure Probability," 2023. arXiv:2309.16930.
- [17] J. Mahler, S. Patil, B. Kehoe, J. van den Berg, M. T. Ciocarlie, P. Abbeel, and K. Goldberg, "GP-GPIS-OPT: Grasp Planning with Shape Uncertainty Using Gaussian Process Implicit Surfaces and Sequential Convex Programming," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4919–4926, IEEE, 2015.
- [18] A. Kumar, P. Mitrano, and D. Berenson, "Constraining Gaussian Process Implicit Surfaces for Robot Manipulation via Dataset Refinement," *IEEE Robotics and Automation Letters*, 2024. arXiv:2410.00157.
- [19] M. Liu, Z. Pan, K. Xu, K. Ganguly, and D. Manocha, "Deep Differentiable Grasp Planner for High-DOF Grippers," in *Robotics: Science and Systems (RSS)*, 2020.
- [20] D. Turpin, L. Wang, E. Heiden, Y.-C. Chen, M. Macklin, S. Tsogkas, S. Dickinson, and A. Garg, "Grasp'D: Differentiable Contact-rich Grasp Synthesis for Multi-fingered Hands," in *Computer Vision – ECCV 2022*, pp. 201–221, 2022.
- [21] D. Turpin, T. Zhong, S. Zhang, G. Zhu, E. Heiden, M. Macklin, S. Tsogkas, S. Dickinson, and A. Garg, "Fast-Grasp'D: Dexterous Multi-finger Grasp Generation Through Differentiable Simulation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8082–8089, IEEE, 2023.
- [22] T. Liu, Z. Liu, Z. Jiao, Y. Zhu, and S.-C. Zhu, "Synthesizing Diverse and Physically Stable Grasps with Arbitrary Hand Structures Using Differentiable Force Closure Estimator," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 470–477, 2022.
- [23] A. Wu, M. Guo, and K. Liu, "Learning Diverse and Physically Feasible Dexterous Grasps with Generative Model and Bilevel Optimization," in *Proceedings of The 6th Conference on Robot Learning*, vol. 205 of *PMLR*, pp. 1938–1948, PMLR, 2023.
- [24] A. Mousavian, C. Eppner, and D. Fox, "6-DOF GraspNet: Variational Grasp Generation for Object Manipulation," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2901–2910, IEEE, 2019.
- [25] E. Corona, A. Pumarola, G. Alenyà, F. Moreno-Noguer, and G. Rogez, "GanHand: Predicting Human Grasp Affordances in Multi-Object Scenes," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5031–5041, IEEE, 2020.
- [26] J. Lundell, E. Corona, T. N. Le, F. Verdoja, P. Weinzaepfel, G. Rogez, F. Moreno-Noguer, and V. Kyrki, "Multi-FinGAN: Generative Coarse-to-Fine Sampling of Multi-Finger Grasps," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4495–4501, IEEE, 2021.
- [27] J. Lundell, F. Verdoja, and V. Kyrki, "DDGC: Generative Deep Dexterous Grasping in Clutter," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6899–6906, 2021.
- [28] V. Mayer, Q. Feng, J. Deng, Y. Shi, Z. Chen, and A. Knoll, "FFHNet: Generating Multi-Fingered Robotic Grasps for Unknown Objects in Real-time," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 762–769, IEEE, 2022.
- [29] J. Lu, H. Kang, H. Li, B. Liu, Y. Yang, Q. Huang, and G. Hua, "UGG: Unified Generative Grasping," in *Computer Vision – ECCV 2024*, pp. 414–433, 2024.
- [30] Z. Weng, H. Lu, D. Kragic, and J. Lundell, "DexDiffuser: Generating Dexterous Grasps with Diffusion Models," *IEEE Robotics and Automation Letters*, vol. 9, no. 12, pp. 11834–11840, 2024.

[31] J. Urain, N. Funk, J. Peters, and G. Chaltatzaki, “SE(3)-DiffusionFields: Learning Smooth Cost Functions for Joint Grasp and Motion Optimization through Diffusion,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5923–5930, IEEE, 2023.

[32] C. Enwerem, J. S. Baras, and C. Belta, “EquiDexFlow: Contact-Grounded SE(3)-Equivariant Dexterous Grasp Generative Flows,” 2026. arXiv:2606.12728.

[33] J. Varley, J. Weisz, J. Weiss, and P. Allen, “Generating Multi-Fingered Robotic Grasps via Deep Learning,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4415–4420, IEEE, 2015.

[34] P. Schmidt, N. Vahrenkamp, M. Wächter, and T. Asfour, “Grasping of Unknown Objects Using Deep Convolutional Neural Networks Based on Depth Images,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6831–6838, IEEE, 2018.

[35] S. Brahmabhatt, A. Handa, J. Hays, and D. Fox, “ContactGrasp: Functional Multi-Finger Grasp Synthesis from Contact,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2386–2393, IEEE, 2019.

[36] L. Shao, F. Ferreira, M. Jorda, V. Nambiar, J. Luo, E. Solowjow, J. A. Ojea, O. Khatib, and J. Bohg, “UniGrasp: Learning a Unified Model to Grasp with Multifingered Robotic Hands,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2286–2293, 2020.

[37] K. Li, N. Baron, X. Zhang, and N. Rojas, “EfficientGrasp: A Unified Data-Efficient Learning to Grasp Method for Multi-Fingered Robot Hands,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8619–8626, 2022.

[38] Z. Xu, B. Qi, S. Agrawal, and S. Song, “AdaGrasp: Learning an Adaptive Gripper-Aware Grasping Policy,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021.

[39] R. T. Rockafellar and S. Uryasev, “Optimization of Conditional Value-at-Risk,” *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[40] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1 ed., 2017.

[41] Schneider & Company, “Coefficient of Friction Reference Chart,” 2026. Accessed: Feb. 25, 2026.

[42] M. A. Roa and R. Suárez, “Grasp Quality Measures: Review and Performance,” *Autonomous Robots*, vol. 38, no. 1, pp. 65–88, 2015.

[43] K. Shaw, A. Agarwal, and D. Pathak, “LEAP Hand: Low-Cost, Efficient, and Anthropomorphic Hand for Robot Learning,” in *Robotics: Science and Systems XIX*, 2023.

[44] P. Li, T. Liu, Y. Li, Y. Geng, Y. Zhu, Y. Yang, and S. Huang, “GenDexGrasp: Generalizable Dexterous Grasping,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8068–8074, 2023.

[45] T. Zhong and C. Allen-Blanchette, “GAGrasp: Geometric Algebra Diffusion for Dexterous Grasping,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6771–6778, 2025.

[46] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, “The YCB Object and Model Set: Towards Common Benchmarks for Manipulation Research,” in *2015 International Conference on Advanced Robotics (ICAR)*, pp. 510–517, 2015.

[47] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, “The Columbia Grasp Database,” in *2009 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1710–1716, IEEE, 2009.

[48] A. T. Miller and P. K. Allen, “Graspl! A Versatile Simulator for Robotic Grasping,” *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.

[49] B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, and A. M. Dollar, “Yale-CMU-Berkeley Dataset for Robotic Manipulation Research,” *The International Journal of Robotics Research*, vol. 36, no. 3, pp. 261–268, 2017.

[50] Y.-W. Chao, W. Yang, Y. Xiang, P. Molchanov, A. Handa, J. Tremblay, Y. S. Narang, K. Van Wyk, U. Iqbal, S. Birchfield, J. Kautz, and D. Fox, “DexYCB: A Benchmark for Capturing Hand Grasping of Objects,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2021.

[51] O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas, “GRAB: A Dataset of Whole-Body Human Grasping of Objects,” in *Computer Vision – ECCV 2020*, pp. 581–600, 2020.

[52] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic, “The GRASP Taxonomy of Human Grasp Types,” *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 66–77, 2016.

[53] J. Romero, D. Tzionas, and M. J. Black, “Embodied Hands: Modeling and Capturing Hands and Bodies Together,” *ACM Transactions on Graphics*, vol. 36, no. 6, pp. 1–17, 2017.

[54] D. Morrison, P. Corke, and J. Leitner, “EGAD! An Evolved Grasping Analysis Dataset for Diversity and Reproducibility in Robotic Manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.

APPENDIX

A. One-Shot Grasp-and-Lift Success in Dynamic Simulation

We confirm grasp-and-lift success in the Drake physics-validation system with gravity on, the protocol of Sec. X-E.3. Each certified grasp mounts on a vertical prismatic joint at the synthesized palm pose, welds a contact pad at every certified contact, and closes with a force-controlled feedforward grip along the inward contact normals. Once the grip holds contact, the mount raises the object 12 cm, after which a lateral pull grows through 1, 3, and 6 times the object weight while Coulomb friction alone retains the grasp, over the adverse friction sweep $\mu \in \{0.2, 0.3, 0.4\}$. A run counts as a success when the object stays within a 3 cm slip and a 30° rotation relative to the hand with at least two contacts through the raise. Of the 105 lifts attempted across the seven objects, 51 establish contact and enter the success grading, and the risk margin $\varepsilon^{(\beta)}$ orders their realized success rate above the nominal Ferrari-Canny margin (Figure 13). Table IX isolates the same comparison. Among the 51 established lifts, all of which the nominal margin certifies ($\varepsilon_{\text{nom}} > 0$), the 27 that the risk margin also certifies ($\varepsilon^{(\beta)} > 0$) achieve a 0.70 success rate under the most adverse friction $\mu = 0.2$, against 0.25 for the 24 grasps the nominal margin certifies but the risk margin rejects ($\varepsilon^{(\beta)} \leq 0$), though the two sets hold identically at $\mu = 0.7$. Table X breaks the same lifts down by object.

TABLE IX

LIFT SUCCESS RATE BY METRIC CERTIFICATION IN DRAKE. SUCCESS RATE OF THE 51 ESTABLISHED-CONTACT LIFTS AT EACH FRICTION, SPLIT BY THE MARGIN THAT CERTIFIES THE GRASP. ALL 51 ARE NOMINAL-CERTIFIED ($\varepsilon_{\text{nom}} > 0$). THE CERTIFIED-ROBUST SUBSET ($\varepsilon^{(\beta)} > 0$) HOLDS UNDER ADVERSE FRICTION WHERE THE NOMINAL-ONLY GRASPS ($\varepsilon^{(\beta)} \leq 0$) FAIL, THOUGH THE NOMINAL MARGIN RATES THE TWO SETS IDENTICALLY AT $\mu = 0.7$.

Grasp Set	n	Success Rate at Friction μ			
		0.7	0.4	0.3	0.2
Certified-robust ($\varepsilon^{(\beta)} > 0$)	27	1.00	0.89	0.81	0.70
Nominal only ($\varepsilon^{(\beta)} \leq 0$)	24	1.00	0.50	0.42	0.25

B. Extending the Risk-Adjusted Synthesis Pipeline to Other Hand Morphologies

Our risk-adjusted metrics use only the stored grasp map and contact set, so retargeting the synthesis pipeline of Sec. VIII to a new hand amounts to swapping the hand model, its fingertip geometry, and the mounting arm, while the metric layer stays unchanged. We demonstrate the resulting portability on a LEAP right hand mounted on a 6-DoF




Fig. 14. **Friction-Uncertainty-Aware Grasp Synthesis on the LEAP Hand.** Twenty-four friction-uncertainty-aware grasps synthesized by the retargeted pipeline for the LEAP hand attached to a 6-DoF FAIR Innovation FR3 cobot, spanning EGAD! shapes (top row and left half of the middle row), household objects, and fruit-scale items (bottom row), rendered in Drake at the stored arm configuration. Every rendered grasp is force-closed ($\ell^* > 0$), satisfies the friction-tail requirement ($\varepsilon^{(0.99)} > 0$), and reaches its grasp pose without table or pedestal collision. At synthesis time, we build a Drake `MultibodyPlant` with the object set on a 150 mm wide by 40 mm high cylindrical pedestal to provide additional support-surface clearance and improve arm reachability. We further confirm grasp-and-lift success for the LEAP hand in Drake dynamic simulation, reporting the one-shot lift outcomes in Sec. A and the per-object executability study in Table II.

TABLE X

PER-OBJECT GRASP-AND-LIFT SUCCESS IN DRAKE. FOR EACH OBJECT, THE NUMBER OF CERTIFIED GRASPS EVALUATED, THE NUMBER THAT ESTABLISH CONTACT AND COMPLETE THE GRAVITY-ON LIFT, AND THE SUCCESS RATE OF THOSE LIFTS UNDER THE GRADED LATERAL PULL AT EACH FRICTION. GREEN MARKS SUCCESS AT OR ABOVE 0.80 AND SALMON AT OR BELOW 0.40, AND THE BOLD ROW TOTALS THE SWEEP.

Object	Grasps	Lifted	Success Rate at Friction μ			
			0.7	0.4	0.3	0.2
Cube	15	4	1.00	1.00	1.00	1.00
Box	15	7	1.00	0.71	0.57	0.57
Cylinder	15	10	1.00	0.60	0.40	0.20
Master Chef Can	15	7	1.00	0.43	0.29	0.29
Mustard Bottle	15	5	1.00	0.40	0.40	0.00
Potted Meat Can	15	4	1.00	1.00	1.00	0.50
Tennis Ball	15	14	1.00	0.86	0.86	0.79
Total	105	51	1.00	0.71	0.63	0.49

FAIR Innovation FR3 robot arm. Our retargeted generator seeds fingertip-opposition candidates on the object surface, gates every candidate through an arm-reachability check (an inverse-kinematics solve on the arm-hand assembly subject to an arm manipulability threshold and table and pedestal collision constraints), and certifies the accepted grasps with the same force-closure and friction-tail requirements as the RealHand L6 set of Figure 11.

In particular, Figure 14 shows twenty-four of the resulting grasps drawn from an 81-object slate, thirteen EGAD! [54] shapes and eleven household and fruit-scale objects, with meshes larger than the hand’s working grasp aperture scaled to a 55 mm median extent. Beyond the aforementioned arm-in-the-loop slate, the hand-retargeted synthesis layer alone

sustains a 234-object, 160-seed set of 37,440 LEAP grasps at full synthesis yield, which feeds the extended evaluation noted in Sec. X-G.

C. Acknowledgments

The cross-dataset benchmark in Section X-D uses 1,599 grasps for the LEAP and Allegro hands that we synthesized with an extension of the FRoGGeR library [6]. We thank the FRoGGeR authors for releasing the underlying nonlinear-programming synthesis tooling as open-source software, on which our extension builds. The 400 Shadow-hand grasps come from the DexGraspNet release [13], and we thank the DexGraspNet authors for making their grasp dataset publicly available.