RECOVERING AND CERTIFYING INCENTIVES IN ZERO-SUM STOPPING GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

We study an offline inverse reinforcement learning problem in which either adversarial player can terminate a game by stopping a shared Markov process. A demonstration records which player stopped and where, while the terminal incentives rationalizing these decisions are unknown. Our goal is to recover both players' incentives and the game value and verify whether they implement the observed boundaries within a transfer and smooth-fit class. Local boundary matching is insufficient because both stopping decisions depend on the same continuation value and must remain mutually optimal from every initial state. Our method solves once for the common continuation value between the observed boundaries, recovers both incentives from its boundary values, and applies a global verification step to rule out profitable deviations. We establish necessity, sufficiency, and uniqueness within a smooth-fit class, and characterize stationary contact and dynamics uncertainty through admissible transfer and value sets. Across 24 synthetic games, direct recovery matches converged adaptive calibration without repeated game solves, while a stopping-action IRL baseline evaluates transfer recovery and unilateral deviations. Propagating volatility uncertainty yields certified transfer and value sets at all tested budgets. Each sixteenfold increase in data reduces mean transfer width by about 75%, whereas plug-in sets miss the true transfers. On held-out industry returns with 24 planted contracts, full-path monitoring reduces mean transfer- and value-set widths by 40.3% and 36.7%, respectively.

1 INTRODUCTION

Bilateral stopping contracts allow either participant to terminate an interaction, with payment determined by who stops first. In game options, for example, a buyer may exercise while a seller may cancel by making a contractual payment (Kifer, 2000). Because the parties have opposing objectives, their decisions form a zero-sum stopping game in which each player's incentive to wait depends on the other's stopping right. Given trajectories recording who stopped and at what state, we ask which terminal incentives make the observed stopping rules mutually optimal.

Recovering such termination payments connects observed stopping behavior to incentive design. Inverse reinforcement learning (IRL) provides the broader framework for recovering rewards from observed decisions (Ng & Russell, 2000), while inverse stopping methods provide optimality conditions and set-valued recovery of stopping costs (Pattanayak & Krishnamurthy, 2023). Our setting differs in that two adversarial players hold competing stopping rights, so the goal is to recover a pair of incentives that jointly supports their observed equilibrium behavior.

We formulate this task as *inverse stopping-contract design*. In the exact-data setting, the state dynamics, running payoff, terminal-horizon payoff, and base settlements are specified. The observations provide two stopping boundaries, each labeled by the player who owns the corresponding stopping right, while the unknowns are two additive owner-specific transfers. Inverse optimal stopping asks which time-dependent transfer implements a prescribed boundary for a single decision maker (Kruse & Strack, 2019). Here, both first-hit rules must instead form a saddle point from every initial state: neither player should benefit from deviating while the other follows the prescribed rule.

Because stopping sends the process to an absorbing terminal state, the problem also fits within zero-sum Markov-game IRL. Existing methods recover rewards from strategic behavior (Wang & Klabjan, 2018; Liao et al., 2025). By exploiting the additional structure of first-hit policies and

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

additive transfers, we obtain a direct inverse construction, a global implementation certificate, and uniqueness within the specified transfer and smooth-fit class.

Our method first solves for the common continuation value between the prescribed boundaries. Smooth fit fixes its boundary derivatives, yielding a linear Neumann problem whose boundary values determine the two transfers. Building on double-obstacle verification for Dynkin games (Bayraktar & Sîrbu, 2014), we apply a global gate to verify both stopping rules, returning transfers and the game value only upon success. In stationary games, this gives an explicit boundary-to-transfer map for Brownian dynamics and a linear boundary-value problem for general diffusions. The same sensitivity structure propagates uncertainty in contacts and dynamics into transfer and value sets, with abstention when the full uncertainty set cannot be certified. Our main contributions are:

- First, we show that both terminal transfers can be recovered from a single continuation-value solve. A global verification certificate then determines whether the recovered transfers implement the prescribed stopping behavior. Within the smooth-fit class considered here, the implementing transfers are unique.
- Second, we derive a cross-owner response law: each transfer responds to the other player's boundary through its stopping residual, implying monotonicity on certified contact rectangles. This converts uncertainty in the stopping boundaries and state dynamics into set-valued recovery with an abstain-or-cover guarantee, rather than treating them as exact.
- Finally, we evaluate the method against adaptive forward-game calibration and a stopping-action IRL baseline, certify transfer and value sets under volatility uncertainty, and test full-path boundary localization on historical industry-return paths with planted contracts.

2 Related Work

Forward stopping games. Dynkin games are zero-sum optimal stopping games characterized by double-obstacle variational inequalities, with established results on values, verification, and first-contact saddle points (Bayraktar & Sîrbu, 2014; Tang & Yang, 2011; Ooi & Uemura, 2023). Threshold and hidden-state models derive stopping boundaries from the underlying diffusion (De Angelis et al., 2018; Gapeev, 2026). Unordered contracts require different equilibrium analysis (Guo, 2020), while game options are a canonical bilateral example (Kifer, 2000). We reverse this direction: given labelled stopping boundaries, we recover the owner-specific transfers that implement them. The Jacobian then separates into contact residuals and diffusion influence terms.

Inverse stopping design. Inverse optimal stopping asks which incentives implement prescribed stopping behavior. Kruse & Strack (2019) recover a time-dependent transfer for one boundary, while Klimmek (2013) studies one-sided sensitivity. A related protocol is Wu (2018), where one agent chooses a stopping time and one of two actions under time–action transfers. Our setting is bilateral: opposing players have separate stopping rights but share one continuation value. Boundary matching is therefore insufficient; the recovered transfers must also satisfy the global saddle inequalities. Appendix E gives a side-by-side comparison.

Inverse reinforcement learning in games. Our problem is related to strategic IRL. Pattanayak & Krishnamurthy (2023) derive rationality tests and set-valued costs for Bayesian stopping, while competitive and zero-sum IRL methods infer rewards from observed actions, policies, or trajectories (Wang & Klabjan, 2018; Liao et al., 2025; Jabbari Asl et al., 2026; Zhang et al., 2026). Reward confidence sets are also constructed by Liao et al. (2025). Here, deterministic first-hit boundaries and the bilateral residual signs give a direct transfer map and diagonal-corner envelopes after global verification.

Unknown dynamics have been studied in IRL and optimal stopping (Herman et al., 2016; Dai et al., 2025); Kuchko (2025), for example, jointly estimate dynamics and value quantities from single-agent stopped trajectories. We instead retain a confidence set for the state generator, so boundary and dynamics uncertainty enters the implementation certificate. Parabolic existence theory (Lieberman, 1996), local-time formulas on curves (Peskir, 2005), and double-obstacle verification provide the forward tools. Our main inverse contributions are the bilateral inverse map, global verification gate, and cross-residual factorization.

3 PROBLEM SETUP

Forward zero-sum stopping game. We consider two players, Max and Min, who observe the same state process and decide whether to terminate the interaction or continue and face its future evolution. Fix a horizon $0 < T < \infty$, a compact state interval $I = [\underline{x}, \bar{x}]$ with $\underline{x} < \bar{x}$, and an initial time-state pair $(t, x) \in [0, T] \times I$. Starting from $X_t = x$, the scalar state evolves as the Itô diffusion:

$$dX_s = \mu(s, X_s)\, ds + \sigma(s, X_s)\, dW_s, \qquad \mathcal{A}_s = \mu(s, x)\partial_x + A(s, x)\partial_{xx}, \quad A = \tfrac{1}{2}\sigma^2, \tag{1}$$

where W is a standard one-dimensional Brownian motion, $\sigma > 0$, and $\mathcal{A}_s$ is the spatial generator. The process exits the state interval at $\eta = \inf\{s \geq t : X_s \notin (\underline{x}, \bar{x})\}$.

Both players observe the filtration generated by X. Their strategic choices are stopping times. We write $\mathfrak{T}_{t,T}$ for the set of $[t, T]$-valued stopping times. Max chooses $\tau \in \mathfrak{T}_{t,T}$ and owns the lower settlement L, while Min chooses $\vartheta \in \mathfrak{T}_{t,T}$ and owns the upper settlement U. The same owner-specific contracts govern forced exit: reaching the lower endpoint settles the lower contract, while reaching the upper endpoint settles the upper contract. Accordingly, the endpoint payment is determined as:

$$e_{L,U}(s, x) = \begin{cases} L(s, \underline{x}), & x = \underline{x}, \\ U(s, \bar{x}), & x = \bar{x}. \end{cases} \tag{2}$$

This is not a separate reward channel. Once the transfers introduced below are specified, they determine both voluntary stopping payments and the corresponding forced-exit payments.

Let $h : I \to \mathbb{R}$ be the horizon payoff. The game ends at $\zeta = \tau \wedge \vartheta \wedge \eta \wedge T$, so the first among Max stopping, Min stopping, state exit, and the horizon determines the terminal settlement:

$$\Phi_{\tau,\vartheta} = \begin{cases} h(X_T), & \zeta = T, \\ e_{L,U}(\eta, X_\eta), & \zeta = \eta < T, \\ L(\tau, X_\tau), & \zeta = \tau < \eta \wedge T,\ \tau \leq \vartheta, \\ U(\vartheta, X_\vartheta), & \zeta = \vartheta < \eta \wedge T,\ \vartheta < \tau. \end{cases} \tag{3}$$

While the game remains active, Max accumulates a discounted running payoff $f : [0, T] \times I \to \mathbb{R}$ at discount rate $\varrho > 0$. Writing $\mathbb{E}_{t,x}$ for conditional expectation given $X_t = x$, Max's total payoff is

$$J_{t,x}(\tau, \vartheta) = \mathbb{E}_{t,x}\Big[\int_t^\zeta e^{-\varrho(s-t)} f(s, X_s)\, ds + e^{-\varrho(\zeta - t)} \Phi_{\tau,\vartheta} \Big]. \tag{4}$$

Min receives the negative of this payoff. The lower and upper game values are

$$\underline{V}(t, x) = \sup_\tau \inf_\vartheta J_{t,x}(\tau, \vartheta), \qquad \overline{V}(t, x) = \inf_\vartheta \sup_\tau J_{t,x}(\tau, \vartheta). \tag{5}$$

All extrema are taken over $\mathfrak{T}_{t,T}$. When the two values coincide, their common value is denoted by V.

Observed behavior and inverse objective. An observed episode records the state path, termination time and state, and which player stopped, but not the terminal incentive making the decision optimal. In the exact-data setting, behavior determines a lower stopping boundary $b_-(t)$ for Max and an upper stopping boundary $b_+(t)$ for Min, with $\underline{x} < b_-(t) < b_+(t) < \bar{x}$. In the stationary case, these boundaries reduce to constants a and b. The corresponding first-hit rules are

$$\tau_-^* = \inf\{s \in [t, T) : X_s \leq b_-(s)\} \wedge T, \qquad \vartheta_+^* = \inf\{s \in [t, T) : X_s \geq b_+(s)\} \wedge T. \tag{6}$$

The state dynamics, discount rate, running payoff, horizon payoff, and base settlements ℓ_0 and u_0 are assumed known. The unknown incentives are two owner-specific time-only transfers,

$$L(s, x) = \ell_0(s, x) + \alpha(s), \qquad U(s, x) = u_0(s, x) + \beta(s), \tag{7}$$

where the transfers, α and β, may represent bonuses, penalties, or side payments attached to the corresponding stopping events. In the stationary model, they reduce to constants.

The inverse problem is to recover α, β, and the game value V so that the observed first-hit rules form a saddle point from every initial state. That is, for every (t, x) and all $\tau, \vartheta \in \mathfrak{T}_{t,T}$,

$$J_{t,x}(\tau, \vartheta_+^*) \leq J_{t,x}(\tau_-^*, \vartheta_+^*) = V(t, x) \leq J_{t,x}(\tau_-^*, \vartheta). \tag{8}$$

Intuitively, if either player follows the observed stopping rule, the other cannot benefit by deviating. The recovered transfers must therefore make both boundaries jointly optimal.

4 FINITE-HORIZON MOVING-BOUNDARY IMPLEMENTATION

To solve the inverse problem, we impose standard smoothness, nondegeneracy, and separation conditions. In particular, the diffusion remains uniformly elliptic, the two stopping boundaries stay a fixed positive distance from each other and from the state-space endpoints, and the terminal payoff is slope-compatible with the two stopping boundaries. Appendix A states these assumptions more precisely and explains their role in the analysis. For a fixed dynamics model and boundary pair, our following 3-step method returns two terminal transfers and a game value, or ABSTAIN.

Step 1: solve for the common continuation value. Let $D = \{(s, x) : 0 \le s < T,\ b_-(s) < x < b_+(s)\}$ be the region in which neither player stops. Smooth fit requires the spatial derivative of the continuation value to match the derivative of the corresponding settlement at each boundary. Because the unknown transfers depend only on time, their spatial derivatives vanish. Therefore, the required boundary slopes are determined entirely by the known base settlements ℓ_0 and u_0. Under the above conditions, the continuation candidate v is obtained from the linear Neumann problem

$$\begin{cases} (\partial_s + \mathcal{A}_s - \varrho)v + f = 0, & (s, x) \in D, \\ v_x(s, b_-(s)) = \ell_{0,x}(s, b_-(s)), & 0 \le s < T, \\ v_x(s, b_+(s)) = u_{0,x}(s, b_+(s)), & 0 \le s < T, \\ v(T, x) = h(x), & b_-(T) \le x \le b_+(T). \end{cases} \tag{9}$$

The first equation balances continuation, discounting, and the running payoff, and determines how the continuation value evolves inside D. The remaining conditions specify the two boundary slopes and the horizon value. Importantly, no search over candidate transfers is needed to perform this solve.

Step 2: recover the transfers by value matching. Value matching requires each settlement to equal v at its owner's stopping boundary. Consequently, each transfer is obtained as the difference between the continuation value at that boundary and the corresponding known base settlement:

$$\begin{aligned} \alpha(s) &= v(s, b_-(s)) - \ell_0(s, b_-(s)), & \beta(s) &= v(s, b_+(s)) - u_0(s, b_+(s)), \\ L &= \ell_0 + \alpha, & U &= u_0 + \beta. \end{aligned} \tag{10}$$

These are candidate incentives, both obtained from the same continuation value v. We then form the full candidate value using L below the lower boundary, v between the boundaries, and U above the upper boundary. The same settlements also determine forced-exit payments through Equation (2).

Step 3: certify the prescribed equilibrium. Equation (10) identifies candidate transfers, but it does not by itself establish that the players actually prefer to stop in the prescribed stopping regions. Smooth fit and value matching impose local conditions at the stopping boundaries, whereas equilibrium requires global optimality. Step 3 provides these missing global checks.

Specifically, we verify that $L < U$ everywhere before the horizon, that $L < v < U$ throughout the continuation region, that the appropriate generator inequalities hold in the two stopping regions, and that the horizon payoff is consistent with the prescribed settlements (Theorem 1). If these conditions hold, then neither player can benefit from any unilateral stopping-time deviation.

Definition 1 (Strict piecewise-classical implementation)**.** *We say that a time-only transfer pair* $(\alpha, \beta) \in C([0, T])^2$ implements *the prescribed curves* (b_-, b_+) *if the resulting game admits a bounded continuous value function* V *such that: (i)* V *is* $C^{1,2}$ *on each open stopping and continuation region, where* $C^{1,2}$ *means once continuously differentiable in time and twice continuously differentiable in the state variable; (ii) the spatial derivative* V_x *extends continuously across both stopping curves; and (iii) for every* $s < T$, V *has the strict stop–continue–stop partition*

$$V = L < U\ (x \le b_-(s)), \qquad L < V < U\ (b_-(s) < x < b_+(s)), \qquad L < V = U\ (x \ge b_+(s)).$$

The declared exit/terminal settlements must hold, and the first-hit rules (6) *must form a saddle point.*

The resulting implementation and uniqueness guarantees are understood relative to the smooth-fit class introduced in Definition 1. A proved violation rules out implementation in that class. An unresolved numerical check also causes ABSTAIN, but does not establish impossibility. The following theorem characterizes implementation in this class.

Theorem 1 (Bilateral moving-bound implementation)**.** *Let v denote the Neumann solution in Equation* (9)*. Using the transfers in Equation* (10)*, define the piecewise candidate value as:*

$$V(s,x) = \begin{cases} L(s,x), & x \le b_-(s), \\ v(s,x), & b_-(s) < x < b_+(s), \\ U(s,x), & x \ge b_+(s). \end{cases} \tag{11}$$

Under the regularity and separation assumptions, the curves (b_-, b_+) are implemented in the sense of Definition 1 if and only if all of the following global conditions hold:

(i) $L < U$ *on* $[0,T) \times I$*; and* $L < v < U$ *throughout* D*;*

(ii) $(\partial_s + \mathcal{A}_s - \varrho)L + f \le 0$ *when* $s < T$ *and* $x < b_-(s)$*;*

(iii) $(\partial_s + \mathcal{A}_s - \varrho)U + f \ge 0$ *when* $s < T$ *and* $x > b_+(s)$*; and*

(iv) the terminal settlement is obstacle-consistent (all states below lie in I): (a) $h(x) = L(T,x) \le U(T,x)$ *when* $(x \le b_-(T))$*; (b)* $L(T,x) \le U(T,x) = h(x)$ *when* $(x \ge b_+(T))$*; and (c)* $L(T,x) \le h(x) \le U(T,x)$ *when* $(b_-(T) < x < b_+(T))$*.*

When these conditions hold, the transfers in (10) *are the unique time-only transfers in this class, V is the game value, and $(\tau_-^*, \vartheta_+^*)$ is a saddle point from every initial state (t,x).*

5 Infinite-Horizon Stationary Brownian Specialization

Section 4 gave the inverse construction for time-dependent dynamics and moving contacts. We now apply the same construction to its infinite-horizon stationary Brownian counterpart. The moving curves become constant contacts, the time-dependent transfers become constants, and the continuation PDE becomes an ODE solvable in closed form. This specialization makes the bilateral coupling explicit and provides the analytic chart used in the sensitivity and uncertainty analysis of Section 6.

Inverse Problem Setup. Throughout this section, we set $f \equiv 0$ and consider $X_t = x + \sigma W_t,\ \sigma > 0$, on $I = [\underline{x}, \bar{x}]$, with infinite horizon and discount $\rho = \varrho > 0$. In the notation of Section 4, this means $\mu = 0,\ A = \frac{\sigma^2}{2},\ b_-(s) \equiv a,\ b_+(s) \equiv b$, for prescribed stationary contacts $\underline{x} < a < b < \bar{x}$, and $\alpha(s) \equiv \alpha,\ \beta(s) \equiv \beta$. Thus the full stopping payoffs remain exactly the same objects as before, $L(x) = \ell_0(x) + \alpha,\ U(x) = u_0(x) + \beta$, and the continuation strip becomes the fixed interval (a,b). The owner-matched exit convention is also unchanged: forced exit at $\underline{x}$ pays $L(\underline{x})$ and forced exit at $\bar{x}$ pays $U(\bar{x})$. We write $J_x(\tau, \vartheta)$ for the corresponding infinite-horizon stationary specialization of (4), with $\zeta = \tau \wedge \vartheta \wedge \eta$ and no terminal-horizon settlement. The forward game can be now written as:

$$\max\{L - V,\ \min\{U - V,\ \mathcal{H}V\}\} = 0. \tag{12}$$

Here $\mathcal{H}w := \frac{\sigma^2}{2} w'' - \rho w$ for any stationary smooth function w.

Step 1: solve the common continuation value. On the fixed continuation interval (a,b), the stationary Neumann problem is formulated as follows: $\mathcal{H}v = 0,\ v'(a) = \ell_0'(a),$ and $v'(b) = u_0'(b)$. The two slope conditions are known because the transfers are constant in x and therefore do not change the spatial derivatives of L and U. Let $\kappa = \frac{\sqrt{2\rho}}{\sigma},\ p_a = \ell_0'(a),\ q_b = u_0'(b),\ z = \kappa(b-a)$. The continuation equation $\mathcal{H}v = 0$ becomes $v'' - \kappa^2 v = 0$. Centering its general solution at the lower contact gives the solution: $v(x) = C_1 \cosh(\kappa(x-a)) + C_2 \sinh(\kappa(x-a))$.

The lower smooth-fit condition $v'(a) = p_a$ fixes $C_2 = p_a/\kappa$ while the upper smooth-fit condition $v'(b) = q_b$ fixes $C_1 = (q_b - p_a \cosh z)/(\kappa \sinh z)$, so the final solution becomes:

$$v_{a,b}(x) = \frac{q_b - p_a \cosh z}{\kappa \sinh z} \cosh(\kappa(x-a)) + \frac{p_a}{\kappa} \sinh(\kappa(x-a)), \tag{13}$$

Appendix B derives Equation (13) step by step.

Step 2: recover the constant transfers by value matching. Equation (10) now reduces to: $\alpha = v_{a,b}(a) - \ell_0(a),\ \beta = v_{a,b}(b) - u_0(b)$. Substituting Equation (13) gives

$$\alpha(a,b) = \frac{q_b - p_a \cosh z}{\kappa \sinh z} - \ell_0(a), \qquad \beta(a,b) = \frac{q_b \cosh z - p_a}{\kappa \sinh z} - u_0(b). \tag{14}$$

Therefore the moving-boundary inverse construction $(b_-, b_+) \mapsto v \mapsto (\alpha, \beta)$ becomes the explicit stationary chart $(a, b) \mapsto v_{a,b} \mapsto (\alpha(a, b), \beta(a, b))$. As before, these formulas identify the unique *candidate* transfers; global verification is still required.

Step 3: certify the prescribed stationary equilibrium. Paste the candidate value exactly as in Equation (11): $V_{a,b}(x) = \begin{cases} L(x), & x \le a, \\ v_{a,b}(x), & a < x < b, \\ U(x), & x \ge b. \end{cases}$ The infinite-horizon model has no terminal-horizon settlement. Its verification uses the four stationary obstacle and residual conditions below.

Theorem 2 (Bilateral stationary prescribed-contact implementation)**.** *The prescribed contacts (a, b) are implemented in the strict disjoint-contact class of Definition 3 if and only if*

(i) $L < U$ *on* $[\underline{x}, \bar{x}]$; *(ii)* $L < v_{a,b} < U$ *on* (a, b);
(iii) $\mathcal{H}L \le 0$ *on* $[\underline{x}, a]$; *(iv)* $\mathcal{H}U \ge 0$ *on* $[b, \bar{x}]$.

When these conditions hold, Equation (14) *gives the unique constant transfers implementing the contacts, $V_{a,b}$ is the game value, and $\tau_a = \inf\{t \ge 0 : X_t \le a\}$, $\vartheta_b = \inf\{t \ge 0 : X_t \ge b\}$ form a saddle point. The owner-matched exit settlement agrees with $V_{a,b}$ at $\underline{x}$ and $\bar{x}$ by construction.*

Figure 1 shows the three steps on the certified witness of Corollary 2. Appendix C provides details.

6 From Exact Stationary Contacts to Certified Transfer Sets

We retain the infinite horizon, zero running payoff, and constant contacts of Section 5, while allowing state-dependent drift and diffusion. Unlike the exact inverse problem in previous sections, finite data do not reveal exact boundaries or generators, so we keep the same objects $(L, U, v, V, \alpha, \beta)$ and study how uncertainty in stationary contacts and dynamics propagates through the construction. In the stationary setting, a contact pair $c = (a, b)$ specializes the boundaries to $b_-(s) \equiv a$ and $b_+(s) \equiv b$.

Contact sensitivity under a fixed stationary generator We first keep the dynamics fixed and extend the Brownian analysis in Section 5 to a general stationary diffusion. Assume $A, B \in C^1(I)$ and $\ell_0, u_0 \in C^3(I)$, and define the stationary operator

$$\mathcal{H}_{A,B} w := A(x) w'' + B(x) w' - \rho w, \qquad A(x) > 0, \tag{15}$$

where w is an arbitrary stationary smooth function, $A(x) = \sigma^2(x)/2$, and $B(x) = \mu(x)$. The Brownian case in Section 5 is recovered by setting $A \equiv \sigma^2/2$ and $B \equiv 0$.

For each stationary contact pair $a < b$, we repeat Steps 1–2 from Section 4. First, we solve for the common continuation value $\mathcal{H}_{A,B} v_{a,b} = 0,\ v'_{a,b}(a) = \ell'_0(a),\ v'_{a,b}(b) = u'_0(b)$, and then recover the constant transfers by value matching: $\alpha(a, b) = v_{a,b}(a) - \ell_0(a),$ and $\beta(a, b) = v_{a,b}(b) - u_0(b)$. The corresponding stopping payoffs are $L_{a,b}(x) = \ell_0(x) + \alpha(a, b),$ and $U_{a,b}(x) = u_0(x) + \beta(a, b)$, which we abbreviate by L and U when the contact pair is clear. Thus, as in Section 5, the contacts determine an inverse map $F(a, b) := (\alpha(a, b), \beta(a, b))$, under a general fixed stationary diffusion.

We next ask how the recovered transfers change when a contact moves. Since both transfers are boundary values of the *same* continuation solution, moving one contact generally affects both. To describe this dependence, we define the normalized influence functions $\mathcal{H}_{A,B} y_a = 0, y_a(a) = 1, y'_a(a) = 0, ; \mathcal{H}_{A,B} y_b = 0, y_b(b) = 1, y'_b(b) = 0,$ and $K_+ := \frac{1}{y'_a(b)}, K_- := -\frac{1}{y'_b(a)}$. These describe how a perturbation at one contact propagates through the continuation region to the other.

Proposition 1 (Cross-residual and Jacobian factorization)**.** *On a verified strict stationary implementation branch, we have $K_+, K_- > 0$ and*

$$\partial_b \alpha = K_+ \frac{\mathcal{H}_{A,B} U(b)}{A(b)}, \qquad \partial_a \beta = -K_- \frac{\mathcal{H}_{A,B} L(a)}{A(a)}. \tag{16}$$

Moreover, $\partial_a \alpha = (\partial_a \beta) y_b(a),\ \partial_b \beta = (\partial_b \alpha) y_a(b)$, with $y_a(b) > 1$ and $y_b(a) > 1$. Hence,

$$\mathsf{J}_{a,b} := DF(a, b) = \begin{pmatrix} (\partial_a \beta) y_b(a) & \partial_b \alpha \\ \partial_a \beta & (\partial_b \alpha) y_a(b) \end{pmatrix}, \ \det \mathsf{J}_{a,b} = (\partial_a \beta)(\partial_b \alpha)\big(y_a(b) y_b(a) - 1\big). \tag{17}$$

A continuously differentiable local transfer-to-contact inverse exists when both residuals are nonzero.

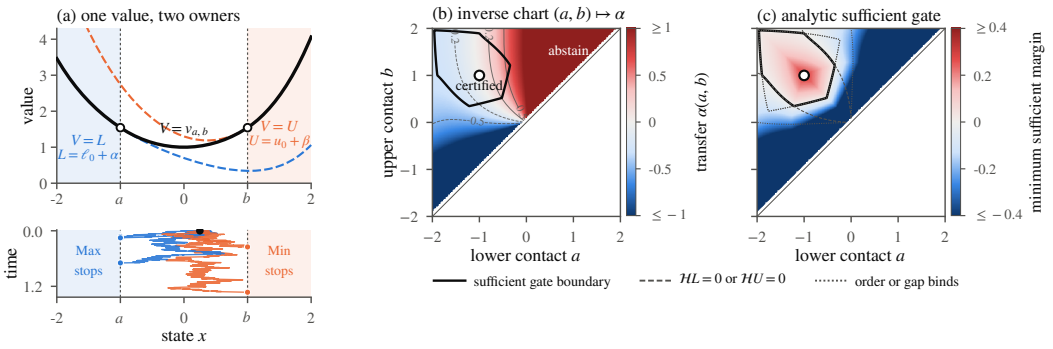


Figure 1: Inverse construction for the certified witness of Corollary 2 ($\sigma = \sqrt{2}$, $\varrho = 1$, $c = 0.3$). (a) The continuation value $v_{a,b}$ pasted between owner settlements and tangent at both contacts; below, owner-labelled Brownian paths stop at the first contact hit. (b) The chart $\alpha(a, b)$ from Equation (14) for $b - a \geq 0.08$; here $\beta(a, b) = -\alpha(-b, -a)$ (Figure 3). (c) The minimum sufficient margin from Appendix C; positive values certify the full state interval. Thick contours mark zero; thin contours show residual and order/gap bounds. The white marker is the witness.

Proposition 1 separates transfer sensitivity into two components. The residuals $\mathcal{H}_{A,B}L(a)$ and $\mathcal{H}_{A,B}U(b)$ measure how strongly the stopping conditions hold at the contacts, while K_- and K_+ measure how a contact perturbation propagates through the continuation solution. On a verified branch, $\mathcal{H}_{A,B}L(a) \leq 0 \leq \mathcal{H}_{A,B}U(b)$, so the cross derivatives $\partial_b\alpha$ and $\partial_a\beta$ are nonnegative, as are the own-contact derivatives. Thus, moving either contact can affect the other player's transfer, because both incentives are linked through the same continuation value.

From uncertain contacts to transfer and value envelopes With finite owner-labelled data, the stationary analogue of the prescribed pair (b_-, b_+) is not known exactly. Suppose the data yield a contact rectangle $\mathcal{C} = [a_\ell, a_u] \times [b_\ell, b_u]$, $a_u < b_\ell$. We require $\mathcal{C} \subset \mathcal{D} = \{(a, b) : \underline{x} < a < b < \bar{x}\}$. The condition $a_u < b_\ell$ guarantees a nonempty *common continuation core* $[a_u, b_\ell]$. Every $x \in (a_u, b_\ell)$ lies strictly in continuation for every contact pair in $\mathcal{C}$, while the endpoints are included by continuity and may coincide with an extreme contact. We call $\mathcal{C}$ *gate-verified* if every $(a, b) \in \mathcal{C}$ satisfies the conditions of Theorem 2, with $\mathcal{H}$ replaced by $\mathcal{H}_{A,B}$. We call it *residual-regular* if the contact residuals remain uniformly strict throughout the rectangle.

Theorem 3 (Verified propagation of contact uncertainty)**.** *Let $\mathcal{C}$ be gate-verified, with diagonal corners $c_\ell = (a_\ell, b_\ell)$ and $c_u = (a_u, b_u)$. Then both coordinates of F are nondecreasing in both contacts. For every x in $[a_u, b_\ell]$, $v_{a,b}(x)$ is also nondecreasing in both contacts. Consequently,*

$$F(c_\ell) \preceq F(a, b) \preceq F(c_u), \qquad v_{c_\ell}(x) \leq v_{a,b}(x) \leq v_{c_u}(x). \tag{18}$$

If $\mathcal{C}$ is residual-regular, all four contact derivatives are positive and $\mathsf{J}_{a,b}$ has principal minors all positive) throughout $\mathcal{C}$. Hence, F is one-to-one there, with a finite uniform inverse condition bound.

Theorem 3 converts uncertainty in the prescribed stopping geometry into uncertainty in the quantities recovered in Section 4: the two transfers and common continuation value. Once the rectangle passes the global implementation gate, no dense search over contact pairs is needed. The componentwise extrema of the transfers and the common-core value occur at the two diagonal corners c_ℓ and c_u.

Joint contact and dynamics uncertainty We next consider the uncertain stationary generator. Let ξ index $\mathcal{H}_\xi w = A_\xi(x)w'' + B_\xi(x)w' - \rho w$, and let $\mathcal{Q}$ be a nonempty compact outer cover of the generator confidence set. For each (ξ, c), define $v_{\xi,c}$, $L_{\xi,c}$, $U_{\xi,c}$, and $F_\xi(c)$ using the same continuation-solve and value-matching construction. Thus, dynamics uncertainty changes the Step-1 operator, but not the transfer definition or Step-3 equilibrium conditions. Joint verification requires obstacle ordering, continuation gaps, residual signs, ellipticity, contact separation, and Neumann denominators bounded away from zero (ensuring a unique continuation solve) uniformly over $\mathcal{Q} \times \mathcal{C}$.

Theorem 4 (Joint contact–dynamics certificate)**.** *Assume the model, measurability, and confidence-set conditions of Appendix D.1. Let $(\alpha^\star, \beta^\star) = F_{\xi^\star}(c^\star)$ be the true implementing transfers and $V^\star$*

their game value. Suppose $\xi^\star \in \mathcal{Q}$ with probability at least $1-\delta$, $c^\star \in \mathcal{C}$ with probability at least $1-\delta_c$, and $\delta + \delta_c \le 1$. If $\mathcal{Q} \times \mathcal{C}$ is jointly verified, every retained pair has unique constant transfers in the stationary smooth-fit class and a first-hit saddle. Define

$$\mathcal{T}_i = \big[\inf_{\xi\in\mathcal{Q}} F_{\xi,i}(c_\ell), \sup_{\xi\in\mathcal{Q}} F_{\xi,i}(c_u)\big], \quad i = 1, 2, \tag{19}$$

$$\mathcal{V}(x) = \big[\inf_{\xi\in\mathcal{Q}} v_{\xi,c_\ell}(x), \sup_{\xi\in\mathcal{Q}} v_{\xi,c_u}(x)\big], \text{ for } x \in [a_u, b_\ell]. \tag{20}$$

A procedure that reports these sets after joint verification and otherwise returns ABSTAIN *satisfies*

$$\Pr\{\text{ABSTAIN or } \alpha^\star \in \mathcal{T}_1,\ \beta^\star \in \mathcal{T}_2,\ V^\star(x) \in \mathcal{V}(x)\ \forall x \in [a_u, b_\ell]\} \ge 1 - \delta - \delta_c. \tag{21}$$

No independence between the dynamics and contact data is required for this union-bound guarantee.

Eq. (21) gives an *abstain-or-cover* guarantee, not coverage conditional on return. With probability at least $1-\delta-\delta_c$, the procedure either abstains or all transfer and value sets contain their truths. Its usefulness therefore depends on both coverage and return rate; Section 7 evaluates both. Theorem 3 still reduces contact optimization to the two diagonal corners; only generator uncertainty requires optimization over $\mathcal{Q}$ (Figure 4 in Appendix C). Appendix D.4 derives the complementary (A, B) response and local contact correction needed to keep both transfers fixed. In the Brownian model, the chart has bounded derivatives on any fixed compact verified contact–volatility product. Thus, for a contact rectangle $\mathcal{C}$ and volatility interval Σ_n in that product, the mean value theorem gives a constant C such that $\operatorname{diam}(\mathcal{T}_i) \le C\big(\operatorname{diam}(\mathcal{C}) + \operatorname{diam}(\Sigma_n)\big),\ \ i = 1, 2$. A fixed-level chi-square volatility interval from n Gaussian increments has length $O_p(n^{-1/2})$. If the contact diameter and numerical-enclosure excess width are also $O_p(n^{-1/2})$, the reported transfer widths have the same rate.

7 EXPERIMENTS

We evaluate numerical correctness and rejection (Q1), matched inversion (Q2), the value of full trajectories (Q3), robustness to historical increments (Q4), and volatility-calibration uncertainty (Q5). Appendix E gives full protocols and results. The *owner-only* method infers boundaries from each player's final continuation–stop bracket, whereas the *full-path* method also uses earlier continuation states. Point recovery is scored by mean transfer RMSE across games. Set recovery either *returns* a certified transfer/value set or *abstains* when the full uncertainty set cannot be verified. Returned outputs are validated outer enclosures of the transfer/value images; since the truth is planted, we report inclusion and width. All value errors and widths in Q2–Q5 are evaluated at the common continuation state $x_q = 0$. Behavior records use planted threshold rules at monitoring times, while the saddle certificate concerns the associated continuous first-hit game. Figure 2 shows the per-game evidence summarized below and tabulated in Appendix E.

Q1: Does the inverse construction reproduce the intended game? We first compare the proposed construction with independent numerical solvers. Finite-difference and Chebyshev methods reproduce the moving-boundary solution, with best value and transfer errors of 1.905×10^{-7} and 1.799×10^{-7}. Independent variational-inequality (VI) solvers converge on all 32 stationary test cases. The predicted response to finite contact perturbations agrees to within 6.164×10^{-6} relative error, and the general-diffusion cross-residual identity has discrepancy 1.453×10^{-13}. Detailed results are in Table 2 in the appendix. The verification step also rejects all seven targeted violations; among trajectory checks, all 24 valid controls return while none of 96 invalid transformations does.

Q2: How does direct inversion compare with solving the game repeatedly? We compare the stationary inverse chart with adaptive VI calibration and a stopping-action IRL baseline adapted from Wang & Klabjan (2018) on 24 games. As the VI grid is refined, its recovered transfers converge to the chart; at 2,049 states, the mean RMSE distance is 2.59×10^{-7}, using 274.2 forward solves per game. Action IRL achieves similar transfer accuracy but requires substantially more forward solves (Table 5 in the appendix). Under fitted dynamics, independent best responses give mean unilateral gains of 2.36×10^{-7} for the chart and 3.21×10^{-6} for action IRL. Under true dynamics, both are about 2.3×10^{-5} and the paired difference interval contains zero. Figure 2(c) shows near-exact gamewise agreement for adaptive VI and comparable error for action IRL. The chart recovers candidates from one continuation solve, and the global gate verifies the observed first-hit rules.

Q3: Does observing the full trajectory improve inverse recovery? On 24 synthetic games with 32,768 trajectories per game, full-path monitoring reduces mean contact, known-volatility transfer,

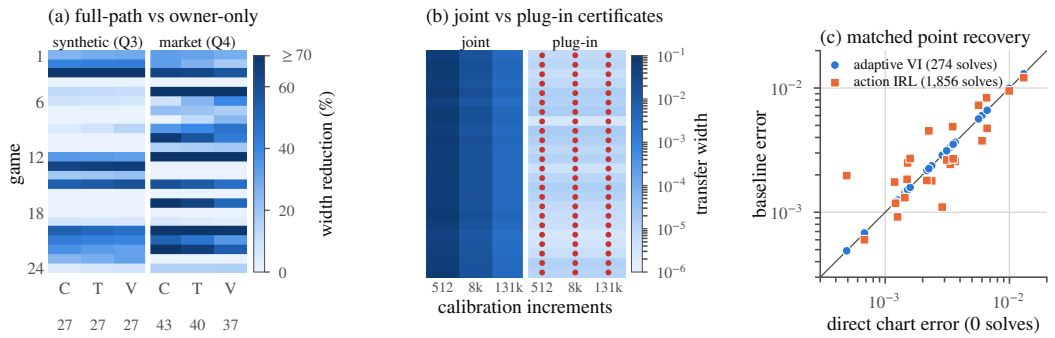


Figure 2: Per-game evidence. (a) Q3/Q4: contact (C), transfer (T), and value (V) width reductions from owner-only to full-path recovery; reductions in mean widths (%) appear below. Q3 uses RMS contact/transfer widths, Q4 their maximum; full-path is never wider. (b) Q5: certified transfer-set width averaged over four calibration repeats. Joint sets shrink with calibration data, while plug-in returns narrow sets excluding the truth. A dot marks cells where at least one repeat misses the truth. (c) Q2: adaptive VI (2,049 states) and action IRL (1,025 states) versus the chart, with all transfer RMSEs measured against planted transfers; labels give mean forward solves per game.

and value widths by 26.6%, 26.8%, and 26.9%, respectively (Figure 2(a); Table 3). Both methods contain all planted transfers and query values in all 24 games. When the same estimated-volatility set is propagated through the inverse problem, the difference between owner-only and full-path widths becomes negligible ($< 0.003\%$), indicating that dynamics uncertainty, rather than contact uncertainty, is then the dominant source of error.

Q4: Does the full-path advantage persist beyond Gaussian trajectories? We next replace Gaussian increments with held-out returns from 49 industry portfolios (Fama & French, 2026), with 24 planted games and contracts fixed. Both methods return on all 24 games and contain every planted contact, transfer, and query value. Full-path monitoring reduces mean contact, transfer, and value widths by 42.7%, 40.3%, and 36.7%, respectively, and is never wider than the owner-only method (Figure 2(a); Table 4). None of 96 negative controls passes the verification gate. Thus the benefit of observing continuation decisions persists under different state-increment distributions.

Q5: Why retain uncertainty in the dynamics? Finally, we estimate volatility from independent calibration data and propagate it jointly with contact uncertainty. For each of 24 games, we use four calibration samples at three budgets. The certificate returns on all 96 records at 512, 8,192, and 131,072 increments, with mean transfer widths 0.0584, 0.0146, and 0.0037 and value widths 0.0258, 0.0061, and 0.0015 (Figure 2(b)). Each sixteenfold increase in data reduces mean transfer width by about 75% over these budgets, consistent with $n^{-1/2}$ scaling. The sets contain both true transfers and the value at $x_q = 0$ in 96/96, 94/96, and 95/96 records. Table 7 reports widths and verification times at both resolutions. A delta-method comparator using the same volatility estimate and contact rectangle achieves joint inclusion in 91/96, 87/96, and 92/96 records. The certified maximum-coordinate transfer width is 1.15–1.20 times the delta-method width. Unlike our method, it uses a normal approximation without checking global equilibrium conditions. By contrast, a plug-in method treating volatility as exact always returns but contains both true transfers in 0/96 records at every budget.

8 Conclusion

One continuation solve between the observed boundaries determines both transfers. We establish conditions under which these transfers make the observed stopping rules a saddle equilibrium, characterize how they respond to boundary changes, and show how uncertainty in the boundaries and state dynamics propagates into certified transfer and value sets, with abstention when certification fails. Recovered games agree with independent forward solvers, full trajectories yield tighter sets, and propagating volatility uncertainty produces verified sets that shrink as calibration data increase. On held-out returns, full-path monitoring cuts mean transfer- and value-set widths by 40.3% and 36.7%, and direct recovery matches converged adaptive calibration without repeated game solves.

STATEMENT OF USE OF AI

The research questions, the main ideas, the theoretical claims, the methodology, and the evaluation protocols of this paper were chosen by the authors, who divided the work among themselves, supervised every stage of the project, and made all decisions on claims and conclusions. Under the authors' direction, we used generative AI tools as assistants for the following tasks with required disclosure. In the theoretical component of the work, generative AI was used in part as a "theory attacker": it assisted in implementing computational searches and large-scale simulations over extensive collections of generated examples and synthetic games in order to look for potential counterexamples, boundary cases, or weaknesses in candidate theorem statements. These computational stress tests helped the authors examine, refine, and gain additional confidence in the formulation of the theoretical claims. However, the absence of a counterexample in such searches does not constitute a proof, and no theorem in this paper is considered established on the basis of AI-assisted simulation, search, or testing. All theorem statements and formal proofs were ultimately checked and accepted by the authors based on their mathematical validity.

Generative AI tools were also used to discuss the experimental design, assist with implementing the methods and the generators of the synthetic games, and support the interpretation of results. The authors inspected the primary sources and executed and independently verified all analyses. No experimental data or results were produced by generative AI tools themselves: every reported number comes from experiments that were actually executed or from previously published public data, and translation and qualitative data analysis are not applicable to this work. Additionally, we used generative AI tools for literature-search triage, drafting and editing the text and LaTeX for readability, checking notation and the consistency of reported numbers, and figure layout. The authors take full responsibility for every claim, proof, citation, and experimental result in this paper; no AI system is an author.

ETHICS AND REPRODUCIBILITY

This mathematical work uses synthetic games and publicly downloadable, copyrighted aggregate industry returns; it neither models investor decisions nor recovers investor preferences. Human applications require consent and uncertainty reporting. The appendices prove the formal claims and specify the evaluation protocols. The supplementary material provides the source code and configurations for all experiments, together with summaries of the recorded outputs.

REFERENCES

Erhan Bayraktar and Mihai Sîrbu. Stochastic perron's method and verification without smoothness using viscosity comparison: Obstacle problems and dynkin games. *Proceedings of the American Mathematical Society*, 142(4):1399–1412, 2014. doi: 10.1090/S0002-9939-2014-11860-0.

Min Dai, Yu Sun, Zuo Quan Xu, and Xun Yu Zhou. Learning to optimally stop diffusion processes, with financial applications. *arXiv preprint arXiv:2408.09242*, 2025. doi: 10.48550/arXiv.2408.09242.

Tiziano De Angelis, Giorgio Ferrari, and John Moriarty. Nash equilibria of threshold type for two-player nonzero-sum games of stopping. *The Annals of Applied Probability*, 28(1):112–147, 2018. doi: 10.1214/17-AAP1301. URL https://arxiv.org/abs/1508.03989.

Eugene F. Fama and Kenneth R. French. 49 Industry Portfolios [Daily]. Kenneth R. French Data Library, Tuck School of Business at Dartmouth, 2026. URL https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_49_ind_port.html. Daily value- and equal-weighted returns; data through 2026-06-30; accessed 2026-08-31.

David Gale and Hukukane Nikaidô. The jacobian matrix and global univalence of mappings. *Mathematische Annalen*, 159:81–93, 1965. doi: 10.1007/BF01360282.

Pavel V. Gapeev. Optimal stopping zero-sum games in continuous hidden markov models. *Advances in Applied Probability*, 58(1):170–213, 2026. doi: 10.1017/apr.2025.10018.

Ivan Guo. On dynkin games with unordered payoff processes. *arXiv preprint arXiv:2008.06882*, 2020. URL https://arxiv.org/abs/2008.06882.

Michael Herman, Tobias Gindele, Jörg Wagner, Felix Schmitt, and Wolfram Burgard. Inverse reinforcement learning with simultaneous estimation of rewards and dynamics. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 102–110. PMLR, 2016. URL https://proceedings.mlr.press/v51/herman16.html.

Hamed Jabbari Asl, Anh Vu Le, Bui Vu Minh, and Mohan Rajesh Elara. Model-free inverse reinforcement learning algorithms for continuous-time and discrete-time zero-sum games. *Neurocomputing*, 665:132101, 2026. doi: 10.1016/j.neucom.2025.132101.

Yuri Kifer. Game options. *Finance and Stochastics*, 4(4):443–463, 2000. doi: 10.1007/PL00013527.

Martin Klimmek. Parameter dependent optimal thresholds, indifference levels and inverse optimal stopping problems. *arXiv preprint arXiv:1302.2769*, 2013. doi: 10.48550/arXiv.1302.2769. URL https://arxiv.org/abs/1302.2769.

Thomas Kruse and Philipp Strack. An inverse optimal stopping problem for diffusion processes. *Mathematics of Operations Research*, 44(2):423–439, 2019. doi: 10.1287/moor.2018.0930.

Anna Kuchko. DO-IQS: Dynamics-aware offline inverse Q-learning for optimal stopping with unknown gain functions, 2025. URL https://arxiv.org/abs/2503.03515.

Junyi Liao, Zihan Zhu, Ethan X Fang, Zhuoran Yang, and Vahid Tarokh. Decoding rewards in competitive games: Inverse game theory with entropy regularization. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 37610–37622. PMLR, 2025. URL https://proceedings.mlr.press/v267/liao25i.html.

Gary M. Lieberman. *Second Order Parabolic Differential Equations*. World Scientific, Singapore, 1996. doi: 10.1142/3302.

Andrew Y. Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 663–670. Morgan Kaufmann, 2000. URL https://ai.stanford.edu/~ang/papers/icml00-irl.pdf.

Takumu Ooi and Toshihiro Uemura. Dynkin games for markov processes associated with semi-Dirichlet forms. *arXiv preprint arXiv:2304.11897*, 2023. URL https://arxiv.org/abs/2304.11897.

Kunal Pattanayak and Vikram Krishnamurthy. Necessary and sufficient conditions for inverse reinforcement learning of bayesian stopping time problems. *Journal of Machine Learning Research*, 24(52):1–64, 2023. URL https://www.jmlr.org/papers/v24/20-1202.html.

Goran Peskir. A change-of-variable formula with local time on curves. *Journal of Theoretical Probability*, 18(3):499–535, 2005. doi: 10.1007/s10959-005-3517-6.

Shanjian Tang and Zhou Yang. Dynkin game of stochastic differential equations with random coefficients, and associated backward stochastic partial differential variational inequality. *arXiv preprint arXiv:1109.5348*, 2011. URL https://arxiv.org/abs/1109.5348.

Xingyu Wang and Diego Klabjan. Competitive multi-agent inverse reinforcement learning with sub-optimal demonstrations. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5143–5151, 2018. URL https://proceedings.mlr.press/v80/wang18d.html.

Larry Wasserman, Aaditya Ramdas, and Sivaraman Balakrishnan. Universal inference. *Proceedings of the National Academy of Sciences*, 117(29):16880–16890, 2020. doi: 10.1073/pnas.1922664117.

Young Wu. *Dynamic Mechanism Design without Transfers*. PhD thesis, University of Toronto, 2018. URL https://library-archives.canada.ca/eng/services/services-libraries/theses/Pages/item.aspx?idNumber=1333975704. Chapter 3: Mechanism Design for Stopping Problems with Two Actions.

Yun Zhang, Zhengyu Guo, Lixiu Yao, and Yunze Cai. Inverse reinforcement learning for recovering objectives in two-player asymmetric zero-sum games via non-equilibrium opponent behaviors. *Neurocomputing*, 673:132827, 2026. doi: 10.1016/j.neucom.2026.132827.

Appendix outline. The appendix first proves the moving-strip, stationary implementation, contact-response, and joint contact-by-dynamics results. A final section gives the synthetic and historical-market protocols, independent validation criteria, and per-game evidence supporting the reported numbers.

A Proof of the moving-strip implementation theorem

Proof roadmap. The proof separates construction from game verification. We first straighten the moving strip and solve one linear Neumann problem. Dynamic programming and local delayed-stop deviations then show that any implementation must use the same traces and satisfy every global gate. For the converse, smooth fit removes the two local-time terms in the change-of-variable formula, leaving the drift signs needed for the saddle sandwich. Exit, ties, and the terminal edge are handled only after this interior argument is established.

Regularity and compatibility. Fix a Hölder exponent $\gamma \in (0,1)$. Here $C^{\gamma/2,\gamma}$ denotes parabolic Hölder regularity with exponents $\gamma/2$ in time and γ in space; $C^{1+\gamma/2,2+\gamma}$ is the corresponding parabolic Schauder class. The coefficients μ, σ are bounded, continuous, uniformly Lipschitz in x, and $\mu, A \in C^{\gamma/2,\gamma}([0,T]\times I)$, with $A \geq \lambda > 0$. The discount satisfies $\varrho > 0$, the running payoff belongs to $C^{\gamma/2,\gamma}$, and $\ell_0, u_0 \in C^{1+\gamma/2,2+\gamma}([0,T]\times I)$. The prescribed curves $b_-, b_+ \in C^{1+\gamma/2}([0,T])$ obey, for some $d_0 > 0$,

$$\underline{x} + d_0 \leq b_-(s), \qquad b_-(s) + d_0 \leq b_+(s), \qquad b_+(s) \leq \bar{x} - d_0. \tag{22}$$

The terminal payoff $h \in C^{2+\gamma}(I)$ satisfies

$$h_x(b_-(T)) = \ell_{0,x}(T, b_-(T)), \qquad h_x(b_+(T)) = u_{0,x}(T, b_+(T)). \tag{23}$$

After straightening the strip, the lateral Neumann data belong to $C^{(1+\gamma)/2}([0,T])$, the terminal trace belongs to $C^{2+\gamma}([0,1])$, and Equation (23) gives order-zero compatibility at both terminal corners. Time reversal and the global Schauder theorem for a linear uniformly parabolic oblique-derivative problem (Lieberman, 1996, Chs. 5–6) then give the regularity stated in Section 4 up to the lateral boundaries and terminal edge.

Definition 2 (Strict piecewise-classical implementation)**.** *This restates Definition 1 with the regularity and endpoint conventions explicit. The curves (b_-, b_+) are implemented in this class if the game has a bounded value $V \in C([0,T]\times I)$ that is $C^{1,2}$ on each open stopping or continuation region, whose spatial derivative extends continuously across both curves, and, for $s < T$, satisfies*

$$V = L < U \ (x \leq b_-(s)), \qquad L < V < U \ (b_-(s) < x < b_+(s)), \qquad L < V = U \ (x \geq b_+(s)).$$

It agrees with the endpoint contract in Equation (2) *and with the terminal settlement, and the first entrance rules into the two stopping regions form a saddle from every initial $(t,x) \in [0,T]\times I$. All moving if-and-only-if claims are relative to this class. Admissible time-only transfer pairs belong to $C([0,T])^2$, and residual conditions use the displayed classical derivatives, which the constructed traces possess.*

Lemma 1 (Moving-strip verification sandwich)**.** *Let the pasted candidate in Equation* (11) *satisfy conditions (i)–(iv) of Theorem 1. Then, for every pair of stopping times τ, ϑ and every initial (t,x),*

$$J_{t,x}(\tau, \vartheta_+^*) \leq V(t,x) \leq J_{t,x}(\tau_-^*, \vartheta). \tag{24}$$

The proof below tracks the drift sign and every tie, exit, and terminal settlement after localization at the parabolic corners.

Step 1 (continuation on the prescribed strip). Write $d(s) = b_+(s) - b_-(s)$ and straighten the strip with $x = b_-(s) + d(s)y$, $y \in [0,1]$, defining $\widetilde{v}(s,y) = v(s, b_-(s) + d(s)y)$. The transformed leading coefficient is $A(s, b_-(s) + d(s)y)/d(s)^2$ and is uniformly positive by uniform ellipticity and

Equation (22); all transformed coefficients have the declared parabolic Hölder regularity. The two boundary data become

$$\widetilde{v}_y(s,0) = d(s)\ell_{0,x}(s,b_-(s)), \qquad \widetilde{v}_y(s,1) = d(s)u_{0,x}(s,b_+(s)). \tag{25}$$

At $s = T$, their compatibility with $\widetilde{v}(T,y) = h(b_-(T) + d(T)y)$ is exactly Equation (23). The Schauder theory for linear uniformly parabolic equations with oblique (here Neumann) data applies after time reversal: the two traces in (25) are in $C^{(1+\gamma)/2}([0,T])$, the terminal datum is in $C^{2+\gamma}([0,1])$, and the displayed compatibility is the required order-zero corner condition. It gives $\widetilde{v} \in C^{1+\gamma/2,2+\gamma}([0,T]\times[0,1])$, including the lateral boundaries and terminal edge; pulling back gives Equation (9). Uniqueness follows from the corresponding homogeneous Schauder estimate, using the standard oblique-boundary result cited in Section 4.

Step 2 (continuation uniqueness and transfer traces). Suppose time-only transfers $\widetilde{\alpha}, \widetilde{\beta}$ implement the prescribed curves in Definition 2. The dynamic programming principle gives $(\partial_s + \mathcal{A}_s - \varrho)V + f = 0$ on D, while smooth fit gives the two derivative conditions in Equation (9): a time-only transfer does not alter an obstacle's x derivative. Terminal settlement gives $V(T,\cdot) = h$ on the terminal strip. Uniqueness of the Neumann problem therefore implies $V = v$ on $\overline{D}$. Value matching at the two curves forces $\widetilde{\alpha} = \alpha$ and $\widetilde{\beta} = \beta$ for $s < T$; continuity of the admissible transfer pairs extends both identities to T.

Step 3 (necessity of the global gates). It remains to justify the two residual signs. Fix an interior point of the lower stopping region and a compact space–time neighborhood that stays strictly below b_-. Against Min's prescribed upper first-hit rule, Max may replace immediate stopping by $\theta_\varepsilon = (s+\varepsilon)\wedge\nu$, where ν is first exit from this neighborhood. For $r \geq s$, write $D_{s,r} = e^{-\varrho(r-s)}$. Saddle optimality and Dynkin's formula give

$$L(s,x) \geq \mathbb{E}_{s,x}\left[\int_s^{\theta_\varepsilon} D_{s,r} f(r,X_r)dr + D_{s,\theta_\varepsilon} L(\theta_\varepsilon, X_{\theta_\varepsilon})\right]. \tag{26}$$

Divide by $\mathbb{E}_{s,x}[\theta_\varepsilon - s]$ and let $\varepsilon \downarrow 0$ to obtain $(\partial_s + \mathcal{A}_s - \varrho)L + f \leq 0$. At an interior point of the upper region, the symmetric delayed-stop deviation by Min reverses the value comparison and gives $(\partial_s + \mathcal{A}_s - \varrho)U + f \geq 0$. Continuity extends both signs to the closures of their declared open regions. The preterminal strict value partition gives condition (i). Sending $s \uparrow T$ and using the declared terminal settlement gives $h = L \leq U$ on the lower terminal region, $L \leq U = h$ on the upper terminal region, and $L \leq h \leq U$ on the terminal interior, which is condition (iv).

Step 4 (verification and sufficiency). Assume (i)–(iv), fix arbitrary stopping times τ, ϑ, and abbreviate ϑ_+^* and τ_-^* as in Equation (6). Localize first at compact stopping times strictly before the finitely many parabolic corners. Concretely, remove a parabolic-distance $\varepsilon_m \downarrow 0$ neighborhood of every corner, cap time at $T - \varepsilon_m$, and stop on first exit from the remaining compact pieces. These stopping times increase to the original termination time, and each stopped path lies in finitely many corner-free strips where the change-of-variable formula on curves applies. Let $\ell_s^{b_\pm}(X)$ denote the symmetric local time at zero of $X_s - b_\pm(s)$, and let $V_x(s, b_\pm(s)\pm)$ denote the corresponding one-sided spatial trace. Peskir's formula applied to $D_{t,s}V(s,X_s)$ there has the usual drift and martingale terms plus

$$\frac{1}{2}\int D_{t,s}\big(V_x(s,b_\pm(s)+) - V_x(s,b_\pm(s)-)\big)\, d\ell_s^{b_\pm}(X). \tag{27}$$

Both local-time coefficients vanish: value matching follows from Equation (10), and the Neumann data give smooth fit.

Until ϑ_+^*, the process is either in continuation, where the drift plus f is zero, or in the lower region, where it is nonpositive. Optional sampling and the payoff comparisons $L \leq V$, $U = V$ at a strict upper hit, $e_{L,U} = V$ at either endpoint, and $h = V(T,\cdot)$ yield

$$J_{t,x}(\tau, \vartheta_+^*) \leq V(t,x). \tag{28}$$

The tie convention causes no gap: at a simultaneous preterminal stop the payoff is $L \leq V$. Symmetrically, until τ_-^* the drift plus f is zero in continuation or nonnegative in the upper region. Now $U \geq V$, $L = V$ at the lower hit, and the same exit/terminal identities give

$$V(t,x) \leq J_{t,x}(\tau_-^*, \vartheta). \tag{29}$$

All functions and spatial derivatives are bounded on the compact pieces, so the stopped stochastic integrals are uniformly square integrable; dominated convergence removes localization. Taking $\tau = \tau^*_-$ and $\vartheta = \vartheta^*_+$ makes both comparisons equalities because the path remains in continuation until its first contact (or horizon). This proves Lemma 1, the saddle, and that the pasted candidate is the value. The preceding Neumann uniqueness and trace matching prove transfer uniqueness, completing Theorem 1.

B Proof of the stationary implementation theorem

B.1 A concrete stationary model

The stationary construction takes the labelled pair (a, b) as input and recovers two scalar transfers (α, β). This section gives the explicit Brownian chart used in the experiments and its complete verification.

Here the horizon is infinite, the running payoff is zero, and $\varrho > 0$ remains fixed. Both strategies are $[0, \infty]$-valued stopping times of the completed natural filtration. Let $X_t = x + \sigma W_t$, $\sigma > 0$, be a Brownian diffusion killed when it leaves $(\underline{x}, \bar{x})$, at time $\eta = \inf\{t \geq 0 : X_t \notin (\underline{x}, \bar{x})\}$. Both players choose stopping times. The first Max stop pays L, a strictly earlier Min stop pays U, and exit pays the matching endpoint settlement. Pre-exit ties go to Max; Max maximizes and Min minimizes the expected discounted payoff. This also defines immediate settlement when the process starts at an endpoint. Explicitly, set $\zeta_\infty = \tau \wedge \vartheta \wedge \eta$ and define

$$\Phi^\infty_{\tau,\vartheta} = \begin{cases} e_{L,U}(X_\eta), & \eta \leq \tau \wedge \vartheta, \\ L(X_\tau), & \tau \leq \vartheta,\ \tau < \eta, \\ U(X_\vartheta), & \vartheta < \tau \wedge \eta, \end{cases} \qquad J_x(\tau, \vartheta) = \mathbb{E}_x[e^{-\varrho\zeta_\infty}\Phi^\infty_{\tau,\vartheta}], \tag{30}$$

where $\mathbb{E}_x$ conditions on $X_0 = x$ and $e_{L,U}$ is the time-independent version of Equation (2). Exit is almost surely finite in this model. The stationary value is $V(x) = \sup_\tau \inf_\vartheta J_x(\tau, \vartheta) = \inf_\vartheta \sup_\tau J_x(\tau, \vartheta)$ when these quantities agree. For smooth w, write

$$\mathcal{G}w = \frac{\sigma^2}{2}w'', \qquad \mathcal{H}w = (\mathcal{G} - \varrho)w. \tag{31}$$

The ordered stationary variational inequality (VI), imposed on the interior of I with $V(\underline{x}) = L(\underline{x})$ and $V(\bar{x}) = U(\bar{x})$, is

$$\max\{L - V, \min\{U - V, \mathcal{H}V\}\} = 0. \tag{32}$$

Thus $\mathcal{H}V \leq 0$ in Max's contact region, $\mathcal{H}V = 0$ in continuation, and $\mathcal{H}V \geq 0$ in Min's contact region.

The known base obstacles $\ell_0, u_0 \in C^2([\underline{x}, \bar{x}])$ are modified by unknown constants

$$L(x) = \ell_0(x) + \alpha, \qquad U(x) = u_0(x) + \beta. \tag{33}$$

The learner observes labelled thresholds $\underline{x} < a < b < \bar{x}$: a belongs to Max's contact and b to Min's contact.

We use the strict disjoint-contact implementation class of Definition 3 in the appendix: the pasted value is smooth fit, lies strictly between the obstacles in continuation, and its two first-contact times form a saddle.

The owner labels are essential under globally ordered contracts: a single maximizer offered both $L \leq U$ at the same stopping time weakly prefers U on every event, and therefore cannot generate the lower-action contact in (41). Crossing action payoffs, including those of Wu (2018), define a different contract class.

B.2 Closed-form inverse map

For a proposed pair (a, b), smooth fit fixes the slope of the continuation value at both contacts. Solving the resulting two-boundary ODE gives that value, and value matching then reveals α and

β. The derivation below is explicit, but its role is simple: it turns observed contacts into a candidate reward pair before the global gate decides whether that pair is valid.

Let

$$\kappa = \frac{\sqrt{2\varrho}}{\sigma}, \quad p_a = \ell_0'(a), \quad q_b = u_0'(b), \quad z = \kappa(b-a). \tag{34}$$

The continuation ODE $\mathcal{H}v = 0$ and smooth fit $v'(a) = p_a$, $v'(b) = q_b$ give

$$v_{a,b}(x) = C_{a,b}\cosh\big(\kappa(x-a)\big) + \frac{p_a}{\kappa}\sinh\big(\kappa(x-a)\big), \tag{35}$$

$$C_{a,b} = \frac{q_b - p_a\cosh z}{\kappa\sinh z}. \tag{36}$$

Value matching at a and b forces the candidate transfers

$$\alpha(a,b) = \frac{q_b - p_a\cosh z}{\kappa\sinh z} - \ell_0(a), \tag{37}$$

$$\beta(a,b) = \frac{q_b\cosh z - p_a}{\kappa\sinh z} - u_0(b). \tag{38}$$

The global obstacle inequalities below determine exactly when these candidate transfers implement the prescribed stopping pair.

Theorem 5 (Strict stationary two-owner implementation)**.** *Form L, U from* (33) *using* (37)–(38)*, and define*

$$V_{a,b}(x) = \begin{cases} L(x), & x \le a, \\ v_{a,b}(x), & a < x < b, \\ U(x), & x \ge b. \end{cases} \tag{39}$$

Then the labelled thresholds a, b are implemented in the strict disjoint-contact class of Definition 3 if and only if

(i) $L < U$ *on* $[\underline{x}, \bar{x}]$*;*

(ii) $L < v_{a,b} < U$ *on* (a, b)*;*

(iii) $(\mathcal{G} - \varrho)L \le 0$ *on* $[\underline{x}, a]$*; and*

(iv) $(\mathcal{G} - \varrho)U \ge 0$ *on* $[b, \bar{x}]$.

When these conditions hold, (37)–(38) *are the unique constant transfers implementing the thresholds. The first entrance times*

$$\tau_a = \inf\{t \ge 0 : X_t \le a\}, \qquad \vartheta_b = \inf\{t \ge 0 : X_t \ge b\} \tag{40}$$

form a saddle, and $V_{a,b}$ is the game value.

Why the conditions are exact. Smooth fit fixes the Neumann solution, and value matching fixes both transfers. The four global conditions then verify the first-contact saddle through the general-diffusion sandwich proved in Appendix B. Thus the chart is necessary, sufficient, and unique within the declared class.

Write $\mathcal{D} = \{(a,b) : \underline{x} < a < b < \bar{x}\}$ and let $\mathcal{I} \subset \mathcal{D}$ contain exactly the pairs satisfying conditions (i)–(iv). Its interior below is relative to $\mathcal{D}$. Reported transfers require interval verification of all four conditions; the endpoint regularization is given in Appendix B.

The stationary proof reuses the same logic in a form that exposes every algebraic step. We first prove a general-diffusion verification lemma, then derive the Brownian Neumann solution and both transfer traces, and finally show that the four displayed gates are simultaneously necessary and sufficient. This order keeps the closed-form inversion separate from the stochastic saddle argument.

Stationary conventions and implementation class. The Brownian motion in Section 5 carries its completed natural filtration, and admissible strategies are all $[0,\infty]$-valued stopping times. A player stop precedes a later exit, a simultaneous pre-exit stop pays the lower obstacle, and exit settlement is immediate when the process starts at an endpoint.

Definition 3 (Strict disjoint-contact implementation class)**.** *A labelled pair (a, b) is implemented in the strict disjoint-contact class when its value V is C^1 on the compact interval, piecewise C^2 away from a, b, and*

$$\begin{aligned} &V = L < U, \quad x \in [\underline{x}, a], \\ &L < V < U, \quad x \in (a, b), \\ &L < V = U, \quad x \in [b, \bar{x}]. \end{aligned} \tag{41}$$

For every initial $x \in [\underline{x}, \bar{x}]$, the first entrance times into the two displayed contact regions form a saddle. Every stationary if-and-only-if statement is relative to this class.

Lemma 2 (General-diffusion verification)**.** *For the stationary diffusion $g = (A, B)$ of Equation* (15)*, let $L, U \in C^2(I)$ and let V be bounded, C^1, and piecewise C^2 at a, b, with bounded derivative. Suppose $L \le V \le U$, $V = L$ on $[\underline{x}, a]$, $V = U$ on $[b, \bar{x}]$, and $\mathcal{H}_g V \le 0$ on $(\underline{x}, b) \setminus \{a\}$ while $\mathcal{H}_g V \ge 0$ on $(a, \bar{x}) \setminus \{b\}$. Then, for every pair of stopping times τ, ϑ and every initial state,*

$$J_x(\tau, \vartheta_b) \le V(x) \le J_x(\tau_a, \vartheta). \tag{42}$$

In particular, the first entrance times τ_a and ϑ_b form a saddle for the general diffusion.

Proof of Lemma 2. Step 1 (upper saddle inequality). For arbitrary τ, set $\theta_n = \tau \wedge \vartheta_b \wedge \eta \wedge n$. Generalized Itô–Tanaka applies to the C^1, piecewise-C^2 function V; the local-time coefficients at a, b are the jumps of V' and hence vanish. Before ϑ_b, the drift of $e^{-\varrho t} V(X_t)$ is $e^{-\varrho t} \mathcal{H}_g V(X_t)\, dt \le 0$. Moreover, its stopped stochastic integral is a true martingale because

$$\mathbb{E}_x \int_0^{\theta_n} e^{-2\varrho s} 2A(X_s) |V'(X_s)|^2 ds \le \frac{\|A\|_\infty \|V'\|_\infty^2}{\varrho} < \infty. \tag{43}$$

Consequently, $\mathbb{E}_x[e^{-\varrho \theta_n} V(X_{\theta_n})] \le V(x)$. Uniform nondegeneracy on the compact interval makes the killed diffusion exit almost surely, so dominated convergence passes to $\tau \wedge \vartheta_b \wedge \eta$. At Max's stop $L \le V$; at Min's first upper contact $U = V$; and at exit the matching endpoint contract equals V. The tie convention also selects $L \le V$. Therefore $J_x(\tau, \vartheta_b) \le V(x)$.

Step 2 (lower saddle inequality). For arbitrary ϑ, repeat the argument with $\theta_n' = \tau_a \wedge \vartheta \wedge \eta \wedge n$. Before τ_a the drift is nonnegative, while at Min's stop $U \ge V$, at Max's first lower contact $L = V$, and at exit the matching endpoint contract equals V. Hence $V(x) \le J_x(\tau_a, \vartheta)$. Taking $\tau = \tau_a$ and $\vartheta = \vartheta_b$ proves the saddle sandwich and the lemma. □

Proof of Theorem 5. Step 1 (inverse chart). We first derive the inverse chart and then verify the saddle. Every solution of $\frac{\sigma^2}{2} v'' - \varrho v = 0$ on (a, b) is

$$v(x) = c_1 \cosh\big(\kappa(x - a)\big) + c_2 \sinh\big(\kappa(x - a)\big), \qquad c_1, c_2 \in \mathbb{R}. \tag{44}$$

Smooth fit at a gives $c_2 = p_a / \kappa$. Smooth fit at b then gives

$$q_b = \kappa c_1 \sinh z + p_a \cosh z, \tag{45}$$

so $c_1 = C_{a,b}$ in (36). Value matching $v(a) = L(a)$ yields $\alpha = c_1 - \ell_0(a)$, which is (37). At the upper boundary,

$$v(b) = \frac{q_b - p_a \cosh z}{\kappa \sinh z} \cosh z + \frac{p_a}{\kappa} \sinh z \tag{46}$$

$$= \frac{q_b \cosh z - p_a}{\kappa \sinh z}, \tag{47}$$

where $\cosh^2 z - \sinh^2 z = 1$ was used. Matching $v(b) = U(b)$ gives (38).

Step 2 (necessity and uniqueness). In any implementation in Definition 3, the strict partition gives (i)–(ii). The dynamic programming principle on the continuation interval gives $\mathcal{H} V = 0$ there. More explicitly, for $x < a$ let θ_ε be ε stopped on exiting a neighborhood contained in the lower contact region. Against Min's upper first-hit rule, Max's optimal immediate stop implies

$$L(x) \ge \mathbb{E}_x[e^{-\varrho \theta_\varepsilon} L(X_{\theta_\varepsilon})],$$

so Dynkin's formula, division by $\mathbb{E}_x[\theta_\varepsilon]$, and $\varepsilon \downarrow 0$ give $\mathcal{H}L(x) \le 0$. In the upper contact interior, Min's symmetric delayed-stop deviation gives $U(x) \le \mathbb{E}_x[e^{-\varrho\theta_\varepsilon}U(X_{\theta_\varepsilon})]$ and hence $\mathcal{H}U(x) \ge 0$. Continuity extends the signs to a, b, proving (iii)–(iv). The ODE, smooth fit inherited from the C^1 class, and value matching above also prove that its constant transfers must be (37)–(38).

Step 3 (sufficiency). For sufficiency, assume (i)–(iv). The candidate (39) is C^1 by value matching and smooth fit, is piecewise C^2, lies between its ordered obstacles, and satisfies

$$(\mathcal{G} - \varrho)V_{a,b} \begin{cases} \le 0, & x < a, \\ = 0, & a < x < b, \\ \ge 0, & x > b. \end{cases} \tag{48}$$

It therefore satisfies the hypotheses of Lemma 2. That lemma supplies the saddle inequalities for every admissible stopping-time pair and proves that $V_{a,b}$ is the value. Conditions (i)–(ii) make its three contact statements exactly the strict partition in Definition 3. This proves sufficiency; the preceding matching argument proves uniqueness and completes the theorem. □

Interval implementation gate. Analytic fixtures may be checked symbolically. Otherwise, for obstacles with validated interval extensions through second order, adaptive subdivision encloses the obstacle order, the two residuals, and

$$g_-(x) = \frac{v_{a,b}(x) - L(x)}{(x-a)^2}, \qquad g_+(x) = \frac{U(x) - v_{a,b}(x)}{(b-x)^2}, \tag{49}$$

for $a < x < b$, extended by $g_-(a) = -(\mathcal{H}L)(a)/\sigma^2$ and $g_+(b) = (\mathcal{H}U)(b)/\sigma^2$ at their matching contacts. Positive lower bounds certify strict continuation gaps. A proved violation and an unresolved enclosure both return ABSTAIN, with distinct internal reasons. Strict positive quotient bounds are a sufficient numerical certificate; their absence need not violate the exact weak-residual theorem.

C PROOFS FOR THE RESPONSE CHART AND WITNESS

For the proofs in this appendix, write $g = (A, B)$, $\mathcal{H}_g = \mathcal{H}_{A,B}$, $J_{a,b} = \mathsf{J}_{a,b}$ and $R = \mathcal{C}$. Set $r_a = \partial_a\beta$, $r_b = \partial_b\alpha$, $p = y_a(b)$ and $q = y_b(a)$. For any fixed $x_0 \in I$, define the integrating factor $\omega(x) = \exp(\int_{x_0}^x B(r)/A(r)\,dr)$.

The response analysis begins by differentiating the two-point boundary problem at one endpoint at a time. Two influence solutions propagate those endpoint changes across the continuation interval. An integrating-factor Wronskian then links the cross responses, after which the residual signs provide the P-matrix and global-univalence conclusions. The final witness gives an explicit regular branch with nonzero cross-owner responses.

Proof of Proposition 1. Step 1 (endpoint differentiation). Standard parameter dependence for linear two-point ODEs makes the chart C^1. Fix a and differentiate $v_{a,b}$ with respect to b. The derivative $w = \partial_b v_{a,b}$ solves $\mathcal{H}_g w = 0$ and

$$w'(a) = 0, \qquad w'(b) = u_0''(b) - v_{a,b}''(b) = \frac{\mathcal{H}_g U(b)}{A(b)}. \tag{50}$$

The last equality uses $v(b) = U(b)$, $v'(b) = U'(b)$, and $\mathcal{H}_g v(b) = 0$. Since $w = w(a)y_a$ and $w(a) = \partial_b\alpha$, the first identity in (16) follows. Smooth fit also removes the moving-endpoint term:

$$\partial_b\beta = w(b) + v'(b) - u_0'(b) = w(b) = (\partial_b\alpha)y_a(b). \tag{51}$$

Similarly, with b fixed, $r = \partial_a v_{a,b}$ satisfies

$$r'(b) = 0, \qquad r'(a) = \ell_0''(a) - v_{a,b}''(a) = \frac{\mathcal{H}_g L(a)}{A(a)}. \tag{52}$$

Because $r = r(b)y_b$ and $r(b) = \partial_a\beta$, this gives the second identity. The lower-endpoint chain rule and smooth fit give

$$\partial_a\alpha = r(a) + v'(a) - \ell_0'(a) = r(a) = (\partial_a\beta)y_b(a). \tag{53}$$

Step 2 (positivity and reciprocity). For completeness, $y_a'(b) > 0$ and $y_b'(a) < 0$. Indeed, $(\omega y')' = \varrho\omega y/A$. Starting from $y_a(a) = 1, y_a'(a) = 0$, the integral identity and a first-zero argument show that y_a stays positive and strictly increases to the right, so $y_a(b) > 1$; reversing the direction shows that $y_b(a) > 1$ and $y_b'(a) < 0$. Hence $K_+, K_- > 0$. For the Wronskian $\mathcal{W} = y_a y_b' - y_a' y_b$, Abel's identity gives $(\omega\mathcal{W})' = 0$. Evaluating at the two ends yields

$$\omega(a)y_b'(a) = -\omega(b)y_a'(b), \tag{54}$$

which is equivalent to $\omega(b)K_- = \omega(a)K_+$.

Step 3 (Jacobian factorization). Equations (51) and (53) give the two columns of Equation (17). Expanding the determinant gives its second identity; the final factor is strictly positive, which proves the stated equivalence for a differentiable local inverse. Finally, a verified lower contact has $\mathcal{H}_g L(a) \le 0$ and a verified upper contact has $\mathcal{H}_g U(b) \ge 0$, proving the asserted cross-derivative signs and locating an exact zero residual on the singular first-order-chart locus. □

Proof of Theorem 3. Step 1 (monotonicity and diagonal envelopes). Proposition 1 gives $r_a, r_b \ge 0$ on a gate-verified rectangle, while its first-zero argument gives $p, q > 1$. Its column factorization is exactly the matrix in (17). The endpoint-differentiation solutions used above also give, at every $x \in [a, b]$,

$$\partial_a v_{a,b}(x) = r_a y_b(x), \qquad \partial_b v_{a,b}(x) = r_b y_a(x). \tag{55}$$

Both influence solutions are positive on $[a, b]$, proving the stated value monotonicity. Since R is a rectangle, integrate these nonnegative coordinate derivatives one coordinate at a time from c_ℓ to (a, b) and then to c_u. This proves (18); the two diagonal contacts belong to R, so the componentwise extrema are attained.

Step 2 (conditioning and global univalence). On a residual-regular rectangle, Equation (16) makes $r_a, r_b > 0$. Thus $\det J = r_a r_b(pq - 1) > 0$, and direct inversion gives

$$J_{a,b}^{-1} = \frac{1}{r_a r_b(pq-1)} \begin{pmatrix} r_b p & -r_b \\ -r_a & r_a q \end{pmatrix}, \tag{56}$$

$$\left\| J_{a,b}^{-1} \right\|_\infty = \max\left\{ \frac{p+1}{r_a(pq-1)}, \frac{q+1}{r_b(pq-1)} \right\}. \tag{57}$$

In particular,

$$\frac{\partial a}{\partial \beta} = -\frac{1}{r_a(pq-1)}, \qquad \frac{\partial b}{\partial \alpha} = -\frac{1}{r_b(pq-1)}. \tag{58}$$

Continuity and compactness make its supremum finite. Both diagonal entries of J and its determinant are positive, so every principal minor is positive. Residual strictness and compactness preserve these signs on an open rectangle containing R. The Gale–Nikaido global-univalence theorem applies there, and restriction makes F injective on the closed rectangle R (Gale & Nikaidô, 1965). Finally, along any continuously differentiable transfer path $z(t) \in F(R)$, set $c(t) = F^{-1}(z(t))$. Dots here denote derivatives with respect to the path parameter t; the chain rule gives $\dot{c} = J_c^{-1}\dot{z}$. Thus, with $\kappa_R = \sup_{c\in R} \left\| J_c^{-1} \right\|_\infty$, $\|\dot{c}(t)\|_\infty \le \kappa_R \|\dot{z}(t)\|_\infty$. □

Corollary 1 (Known-generator selective coverage)**.** *Fix a known generator, a true implemented pair $c^\star = (a^\star, b^\star)$, and $\delta \in (0, 1)$. Let $\widehat{R}$ be a random contact rectangle with $\Pr\{(a^\star, b^\star) \in \widehat{R}\} \ge 1 - \delta$. A procedure that reports the envelopes in Equation* (18) *only after verifying all of $\widehat{R}$, and otherwise abstains, satisfies*

$$\Pr\{\textsc{Abstain} \textit{ or both transfer and common-core value envelopes cover their truths}\} \ge 1-\delta. \tag{59}$$

Proof. On the event $(a^\star, b^\star) \in \widehat{R}$, either the procedure abstains or Theorem 3 contains the true transfers and every true common-core value. The complement of this event has probability at most δ. □

For the Brownian model in Section 5, $A = \sigma^2/2$, $B = 0$, and $K_+ = K_- = 1/(\kappa \sinh z)$. Hence

$$\partial_b \alpha = \frac{2\mathcal{H}U(b)}{\sigma^2 \kappa \sinh z}, \qquad \partial_a \beta = -\frac{2\mathcal{H}L(a)}{\sigma^2 \kappa \sinh z}. \tag{60}$$

Expanding gives

$$\partial_b \alpha = \frac{u_0''(b)}{\kappa \sinh z} + \frac{\ell_0'(a) - u_0'(b)\cosh z}{\sinh^2 z}, \tag{61}$$

$$\partial_a \beta = -\frac{\ell_0''(a)}{\kappa \sinh z} + \frac{u_0'(b) - \ell_0'(a)\cosh z}{\sinh^2 z}. \tag{62}$$

For (61), differentiate

$$\alpha + \ell_0(a) = \frac{q_b - p_a \cosh z}{\kappa \sinh z} \tag{63}$$

with respect to b, using $\partial_b q_b = u_0''(b)$ and $\partial_b z = \kappa$. Cancelling $\cosh^2 z - \sinh^2 z = 1$ yields

$$\partial_b \alpha = \frac{u_0''(b)}{\kappa \sinh z} + \frac{p_a - q_b \cosh z}{\sinh^2 z}. \tag{64}$$

The same calculation on (38), now using $\partial_a p_a = \ell_0''(a)$ and $\partial_a z = -\kappa$, gives (62). Equations (56)–(58) then give the Brownian inverse response on every residual-regular verified branch.

Corollary 2 (A globally certified strategic witness)**.** *Take* $[\underline{x}, \bar{x}] = [-2, 2]$, $\sigma = \sqrt{2}$, $\varrho = 1$, $a = -1$, $b = 1$, *and any* $c > 0$, *with*

$$\ell_0(x) = \cosh x - c(x+1)^2, \qquad u_0(x) = \cosh x + c(x-1)^2. \tag{65}$$

Then $\alpha = \beta = 0$, $V(x) = \cosh x$ *on* $(-1, 1)$, *all four conditions of Theorem 5 hold, and* $(-1, 1) \in \operatorname{int}_{\mathcal{D}} \mathcal{I}$. *Moreover,*

$$J_{-1,1} = \frac{2c}{\sinh 2} \begin{pmatrix} \cosh 2 & 1 \\ 1 & \cosh 2 \end{pmatrix}, \quad \det J_{-1,1} = 4c^2, \quad \frac{\partial a}{\partial \beta} = \frac{\partial b}{\partial \alpha} = -\frac{1}{2c \sinh 2} \neq 0. \tag{66}$$

For Corollary 2, $\kappa = 1$, $p_a = \ell_0'(-1) = -\sinh 1$, and $q_b = u_0'(1) = \sinh 1$. Hyperbolic addition identities in (37)–(38) give $\alpha = \beta = 0$ and $v(x) = \cosh x$. The obstacle order and continuation gaps are

$$U(x) - L(x) = 2c(x^2 + 1) > 0, \tag{67}$$

$$v(x) - L(x) = c(x+1)^2, \tag{68}$$

$$U(x) - v(x) = c(x-1)^2. \tag{69}$$

Since $\mathcal{G} = d^2/dx^2$ for $\sigma = \sqrt{2}$,

$$(\mathcal{G} - 1)\ell_0(x) = c\big((x+1)^2 - 2\big) \le -c, \qquad x \in [-2, -1], \tag{70}$$

$$(\mathcal{G} - 1)u_0(x) = c\big(2 - (x-1)^2\big) \ge c, \qquad x \in [1, 2]. \tag{71}$$

Thus all conditions of Theorem 5 hold. Direct differentiation of both entries of the inverse chart at $(-1, 1)$ gives (66); its determinant is

$$\frac{4c^2}{\sinh^2 2}(\cosh^2 2 - 1) = 4c^2. \tag{72}$$

Inserting the off-diagonal entry $2c/\sinh 2$ into (58) gives the response displayed in the corollary. To make the interior claim explicit, extend the two quotient gaps in (49) continuously to their matching contacts by Taylor's theorem. At $(-1, 1)$ both extensions equal c throughout the continuation interval. The obstacle-order margin is at least $2c$, and the two contact-residual margins are at least c. The chart, residuals, and extended gap quotients depend continuously on (a, b, x) over the relevant compact sets. They therefore retain their strict signs throughout a threshold neighborhood, so $(-1, 1) \in \operatorname{int}_{\mathcal{D}} \mathcal{I}$. Hence the displayed derivatives describe a locally verified equilibrium branch, completing the proof.

Analytic sufficient gates used in the figures. The witness plots use $I = [-2, 2]$, $\rho = 1$ and $c = 0.3$. At $\sigma^\star = \sqrt{2}$, write $v = \cosh x + w$. Then $w'' = w$, $w'(a) = p = -2c(a+1)$ and $w'(b) = q = 2c(b-1)$, so

$$w(x) = \frac{q\cosh(x-a) - p\cosh(b-x)}{\sinh(b-a)}, \qquad \|w\|_{[a,b]} \le M_{a,b} := \max\{|p|, |q|\}\coth((b-a)/2).$$

The bound follows by maximizing $\cosh(x-a) + \cosh(b-x)$ at an endpoint. The margins plotted in Figure 1(c) are

$$m_o = 2c + \beta - \alpha, \qquad m_g = c - M_{a,b}/2,$$
$$m_- = \alpha + 2c - c\max\{1, (a+1)^2\}, \quad m_+ = 2c - \beta - c\max\{1, (b-1)^2\}.$$

Positive margins imply the complete stationary gate: $U - L \ge m_o$, $-\mathcal{H}L \ge m_-$ on $[-2, a]$, and $\mathcal{H}U \ge m_+$ on $[b, 2]$. Also $(v-L)''$ and $(U-v)''$ are at least $2m_g$, so value matching and smooth fit give $v - L \ge m_g(x-a)^2$ and $U - v \ge m_g(b-x)^2$ on $[a, b]$. The figure colours $\min\{m_o, m_g, m_-, m_+\}$; its zero contour is a sufficient verification boundary.

For Figure 4, let $a \in [-1-h, -1+h]$, $b \in [1-h, 1+h]$ and $\sigma/\sigma^\star \in [1-\epsilon, 1+\epsilon]$, with $0 \le h, \epsilon < 1$. Put $A_- = (1-\epsilon)^2$, $\delta_A = 2\epsilon + \epsilon^2$ and

$$M = 2ch(1+\epsilon)\coth\left(\frac{1-h}{1+\epsilon}\right) + \delta_A\cosh(1+h).$$

Indeed, for $A = \sigma^2/2$, the equation for $w = v - \cosh x$ is $Aw'' - w = -(A-1)\cosh x$. Split w into its homogeneous Neumann solution and a zero-Neumann solution of this forcing. The preceding hyperbolic bound controls the first by the first term of M; the maximum principle controls the second by $\delta_A\cosh(1+h)$. Thus $|w| \le M$ uniformly on the complete contact–volatility product. Valid uniform lower margins are

$$\underline{m}_o = 2c - 2M - 2ch^2,$$
$$\underline{m}_g = c - \tfrac{1}{2}\{M/A_- + (A_-^{-1} - 1)\cosh(1+h)\},$$
$$\underline{m}_r = 2cA_- - c - \delta_A\cosh 2 - M.$$

The first bounds $U - L$, using $\alpha = w(a) + c(a+1)^2$ and $\beta = w(b) - c(b-1)^2$. The second bounds both squared-distance gap quotients because $|v'' - \cosh x| \le M/A_- + (A_-^{-1} - 1)\cosh(1+h)$. The third bounds both stopping residuals, using $|x \pm 1| \le 1$ on their respective stopping intervals and $\cosh x \le \cosh 2$. At $(h, \epsilon) = (0.15, 0.01)$, outward interval evaluation gives lower margins above 0.2518, 0.1970 and 0.0451, respectively.

Figure 4(c) returns only when all three lower margins exceed 10^{-10}. For each returned product, Theorem 3 reduces the transfer envelope to the two diagonal contact corners. Outward interval evaluation of Equation (14) over volatility subintervals encloses their entire images: panel (b) uses 128 subintervals and panel (c) uses 16. The displayed width is the larger of the two outer transfer widths. The five filled images in panel (b) illustrate the geometry; the enclosing box covers all intermediate volatilities.

Figure 3 evaluates the chart, the cross response, and the Jacobian determinant on the witness family for contact pairs with $b - a \ge 0.08$.

Figure 4 illustrates Theorems 3 and 4 on the same family: the transfer envelope is read off the two diagonal corners of a contact rectangle, dynamics uncertainty widens it into the joint set, and the stated sufficient gate separates returned outer enclosures from abstentions.

D NUISANCE-AWARE CONTACT-BY-DYNAMICS TOMOGRAPHY

This appendix builds the joint certificate from its probabilistic and analytic components. We first construct a finite-sample generator set from complete stopped records. The likelihood construction below uses the stated Gaussian discrete transition experiment; Q5 instead evaluates its scalar Brownian specialization with independent volatility calibration. A fixed-generator bridge then connects every retained model to a verified Dynkin game. The generator-response theorem describes the local nuisance geometry, while closed gap quotients make the strict gates compact. We next derive the contact rectangle, prove measurable selection over the full product, assemble selective coverage, and finish by showing that the returned sets contract under identification.

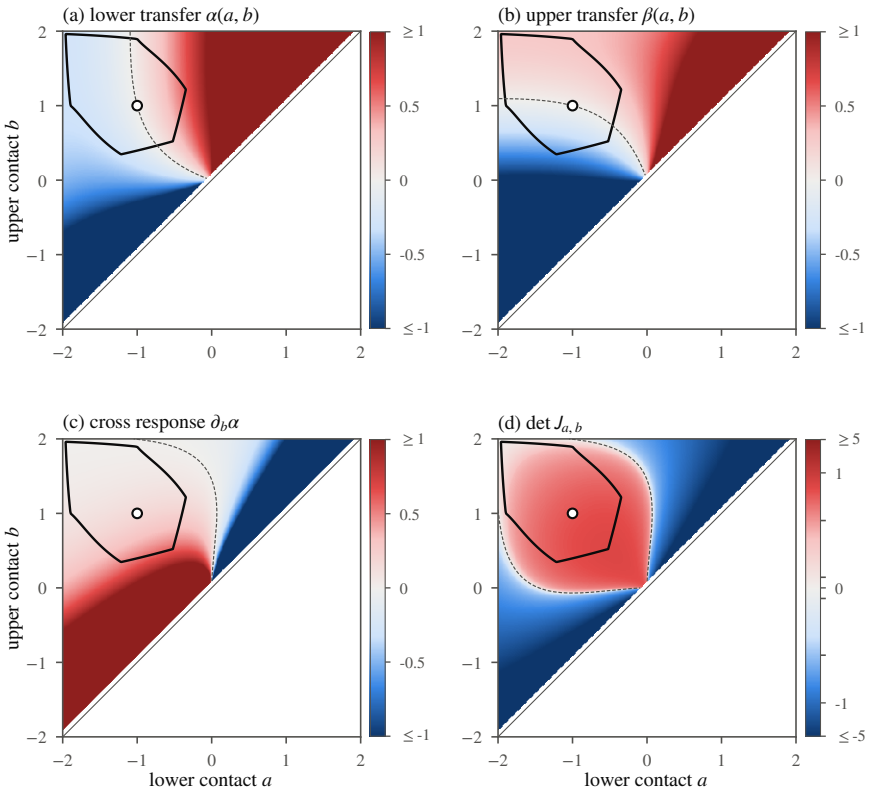


Figure 3: The stationary chart of Section 5 on the witness family ($\sigma = \sqrt{2}$, $\varrho = 1$, $c = 0.3$). (a,b) Both transfers. Because $\cosh x$ is $\mathcal{H}$-harmonic when $\sigma = \sqrt{2}$ and $\varrho = 1$, only the quadratic parts of ℓ_0 and u_0 enter the chart, and their mirror antisymmetry gives $\beta(a, b) = -\alpha(-b, -a)$. (c) The cross response $\partial_b \alpha$ of Proposition 1, nonnegative throughout the certified set and changing sign only outside it. (d) $\det \mathrm{J}_{a,b}$, positive throughout the certified set as required for the P-matrix conclusion of Theorem 3. Thick contours bound the sufficient verification region of Figure 1(c); dashed contours are the zero sets of the displayed quantity; the white marker is the witness.

D.1 Model, joint law, and complete theorem conditions

Let $I = [\underline{x}, \bar{x}]$ and consider the compact six-parameter family

$$\begin{aligned} A_\xi(x) &= q_0 + q_1 \sin(\pi x/2) + q_2 \cos(\pi x), \\ B_\xi(x) &= r_0 + r_1 \sin(\pi x/2) + r_2 x, \qquad \xi = (q, r) \in \Xi_0. \end{aligned} \tag{73}$$

The functions A_ξ, B_ξ are C^1 in x, continuous in ξ, and $0 < \lambda \le A_\xi \le \Lambda$ uniformly; every candidate killed diffusion is well posed. Calibration uses the exact conditional experiment

$$\Delta X_j = B_\xi(X_j)h + \sqrt{2A_\xi(X_j)h}\, Z_j, \qquad Z_j \mid \mathcal{F}_j \sim \mathcal{N}(0, 1). \tag{74}$$

The split-likelihood construction below applies universal inference (Wasserman et al., 2020) to complete records from this experiment.

All probabilities in this section are taken under one joint law $\mathbb{P}^\star$ for the calibration records, contact records, and any auxiliary randomization. Independence between the calibration and contact sigma-fields is not assumed. Let $\mathfrak{R}$ be the measurable space of compact ordered contact rectangles and adjoin an isolated sentinel †. The contact stage returns $\widehat{R}_c \in \mathfrak{R} \sqcup \{\dagger\}$, where $\widehat{R}_c = \dagger$ means upstream abstention; the downstream output is then the sentinel $\perp$. Membership $c^\star \in \dagger$ is defined to be false, and $d_H(\dagger, \{c^\star\}) = +\infty$, so

$$\mathcal{C}_c = \{\widehat{R}_c = \dagger\} \,\dot\cup\, \{\widehat{R}_c \in \mathfrak{R},\ c^\star \in \widehat{R}_c\} \tag{75}$$

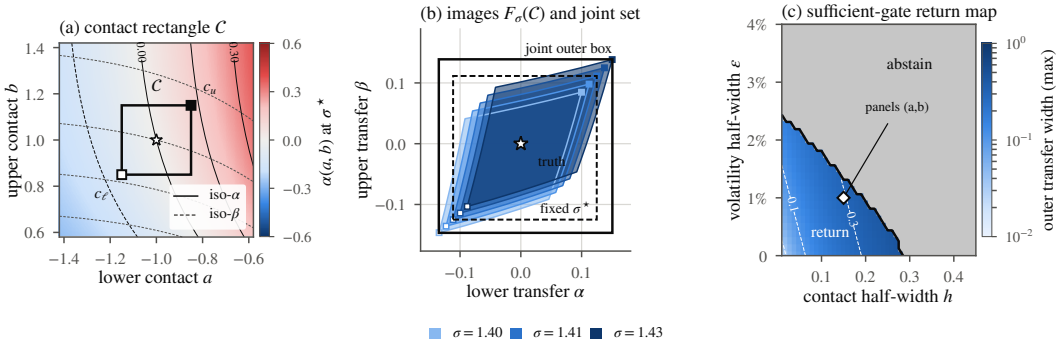


Figure 4: Certified outer sets on the witness family for $\mathcal{C} = [-1.15, -0.85] \times [0.85, 1.15]$ and $\mathcal{Q} = [0.99\sigma^\star, 1.01\sigma^\star]$. (a) Both transfers increase in both contacts on a gate-verified rectangle, so their extremes sit at the diagonal corners c_ℓ and c_u (Theorem 3). (b) Images $F_\sigma(\mathcal{C})$ for five volatilities in $\mathcal{Q}$; the dashed box is the fixed-$\sigma^\star$ envelope and the solid box an interval outer enclosure of the joint transfer image. (c) Maximum of the two outer transfer widths for products passing the analytic sufficient gate in Appendix C; grey cells abstain.

is a measurable disjoint union. On $\{\widehat{R}_c \in \mathfrak{R}\}$, write $R = \widehat{R}_c$.

For $c = (a, b)$, let $v_{\xi,c}$ solve $\mathcal{H}_\xi v = 0$ with $v'(a) = \ell_0'(a)$ and $v'(b) = u_0'(b)$, where $\mathcal{H}_\xi = A_\xi \partial_{xx} + B_\xi \partial_x - \varrho$, and set

$$F_\xi(c) = \begin{pmatrix} v_{\xi,c}(a) - \ell_0(a) \\ v_{\xi,c}(b) - u_0(b) \end{pmatrix}. \tag{76}$$

Write $L_{\xi,c} = \ell_0 + F_{\xi,1}(c)$ and $U_{\xi,c} = u_0 + F_{\xi,2}(c)$.

The complete conditions for Theorem 4 are as follows. Let $\ell_0, u_0 \in C^3(I)$, $\varrho > 0$, $\xi^\star \in \Xi_0$, and let $c^\star = (a^\star, b^\star)$ belong to the strict disjoint-contact implementation class. Define $(\alpha^\star, \beta^\star) = F_{\xi^\star}(c^\star)$ and let $V^\star$ denote its pasted game value. For $0 < \delta < 1$, $0 \le \delta_c < 1$, and $\delta + \delta_c \le 1$, require

$$\mathbb{P}^\star\{\xi^\star \in \widehat{\Xi}_\delta\} \ge 1 - \delta, \qquad \mathbb{P}^\star(\mathcal{C}_c) \ge 1 - \delta_c. \tag{77}$$

On every nonsentinel record, Q is a measurable finite union of closed boxes satisfying $\widehat{\Xi}_\delta \subseteq Q \subseteq \Xi_0$, while $R = [a_\ell, a_u] \times [b_\ell, b_u]$ has measurable endpoints and

$$\underline{x} < a_\ell \le a_u < b_\ell \le b_u < \bar{x}. \tag{78}$$

The procedure reports only after obtaining strictly positive validated lower bounds for obstacle order, both closed continuation quotients, both contact-region residuals, ellipticity, contact separation, the Neumann denominator, and common-core containment uniformly on $Q \times R$. Every reported output encloses the complete retained product, and a game condition never removes a likelihood-retained dynamics box. All other outcomes map to $\perp$, identified with ABSTAIN in the main text.

Corollary 3 (Hard-support and simultaneous validity)**.** *Suppose the monitored contact states obey a known hard error bound and carry exact latent first-hit owner and timing labels as specified in Lemma 7. Then $\delta_c = 0$ in Theorem 4. For any fixed collection of K games with fixed levels $\delta_1, \dots, \delta_K$,*

$$\mathbb{P}^\star\{\textit{every game abstains or covers both transfers and its common-core value}\} \ge 1 - \sum_{k=1}^{K} \delta_k. \tag{79}$$

No cross-game independence is required.

D.2 SPLIT LIKELIHOOD ON COMPLETE STOPPED RECORDS

Index trajectories before any calibration increment is generated. Fix disjoint pilot and confirmation index sets I_0, I_1 independently of the realized paths; no trajectory is split across the two sets. Every reported confirmation prefix is likewise a fixed prefix of trajectory identifiers, not a likelihood- or

coverage-selected prefix. Write D_0 for the pilot records. For confirmation trajectory i, condition on its parameter-free initial state and design variables $\mathcal{G}_i$, let $N_i \le N_{\max} < \infty$, and define

$$\mathcal{F}_{i,n} = \sigma(D_0, \mathcal{G}_i, \Delta X_{i,0}, \dots, \Delta X_{i,n-1}), \qquad 0 \le n \le N_i. \tag{80}$$

Then $X_{i,n}$ is $\mathcal{F}_{i,n}$-measurable. Let $\tau_i \le N_i$ be a stopping or killing time that is the same nonanticipating functional of the canonical recorded history under every candidate parameter. The complete stopped record is

$$O_i = (X_{i,0}, \tau_i, \Delta X_{i,0}, \dots, \Delta X_{i,\tau_i-1}). \tag{81}$$

In particular, exit-after-proposal killing retains the exit-causing increment.

Under Equation (74), set

$$p_\xi(y \mid x) = \frac{1}{\sqrt{4\pi h A_\xi(x)}} \exp\left[-\frac{(y - hB_\xi(x))^2}{4hA_\xi(x)}\right], \qquad L_i(\xi) = \prod_{j=0}^{\tau_i-1} p_\xi(\Delta X_{i,j} \mid X_{i,j}). \tag{82}$$

The indicator that the common stopping functional produced τ_i, and any parameter-free design density, are part of the stopped-record density but cancel from its likelihood ratios.

Let $\widetilde{\xi}_m$, w_m, $m = 1, \dots, M_0$, be pilot-measurable anchors and weights, with $\widetilde{\xi}_m \in \Xi_0$ almost surely, $w_m \ge 0$, and $\sum_m w_m = 1$. Their number, displacement rule, optimizer, optimizer randomness, and weights are all $\sigma(D_0)$-measurable; none may use a confirmation increment. Define

$$L_1(\xi) = \prod_{i \in I_1} L_i(\xi), \qquad E_1(\xi) = \sum_{m=1}^{M_0} w_m \frac{L_1(\widetilde{\xi}_m)}{L_1(\xi)}, \qquad \widehat{\Xi}_\delta = \{\xi \in \Xi_0 : E_1(\xi) \le \delta^{-1}\}. \tag{83}$$

Proposition 2 (Split-likelihood generator set)**.** *Suppose Equation* (74) *holds, trajectory membership in I_0, I_1 is fixed independently of the increments, initial states and design kernels are parameter independent or conditioned upon, and distinct confirmation trajectories are conditionally independent. At any bounded, candidate-common stopping or killing times described above,*

$$\mathbb{P}^\star\{\xi^\star \in \widehat{\Xi}_\delta\} \ge 1 - \delta \qquad \textit{for every } \xi^\star \in \Xi_0. \tag{84}$$

If the final increment is observed only through a coarsening, the conclusion continues to hold when the last point density is replaced by the candidate-dependent probability or density of that observed coarsening.

Proof. Step 1 (one stopped trajectory). Condition on D_0 and all parameter-free design variables, so every anchor and weight is fixed. For anchor m, define

$$R^m_{i,n} = \prod_{j=0}^{(n \wedge \tau_i)-1} \frac{p_{\widetilde{\xi}_m}(\Delta X_{i,j} \mid X_{i,j})}{p_{\xi^\star}(\Delta X_{i,j} \mid X_{i,j})}, \qquad R^m_{i,0} = 1. \tag{85}$$

The event $\{\tau_i > n\}$ belongs to $\mathcal{F}_{i,n}$. Uniform ellipticity gives common positive support, while normalization of the conditional Gaussian density gives

$$\mathbb{E}_{\xi^\star}[R^m_{i,n+1} \mid \mathcal{F}_{i,n}] = R^m_{i,n}\left(\mathbf{1}\{\tau_i \le n\} + \mathbf{1}\{\tau_i > n\} \int \frac{p_{\widetilde{\xi}_m}(y \mid X_{i,n})}{p_{\xi^\star}(y \mid X_{i,n})} p_{\xi^\star}(y \mid X_{i,n})\, dy\right) \tag{86}$$

$$= R^m_{i,n}. \tag{87}$$

Thus $R^m_{i,n}$ is a nonnegative martingale. Bounded optional stopping gives $\mathbb{E}_{\xi^\star} R^m_{i,\tau_i} = 1$. This step uses the complete record in Equation (81); deleting the exit-causing increment deletes its likelihood factor and generally destroys the identity.

Step 2 (confirmation product). Conditional independence across confirmation trajectories and iterated expectation give

$$\mathbb{E}_{\xi^\star}\left[\frac{L_1(\widetilde{\xi}_m)}{L_1(\xi^\star)}\middle| D_0\right] = 1. \tag{88}$$

If a terminal increment is coarsened, its induced conditional observation density is normalized, so the same one-step calculation applies in the coarsened filtration.

Step 3 (mixture and inversion). Pilot measurability and Equation (88) yield $\mathbb{E}_{\xi^\star}[E_1(\xi^\star) \mid D_0] = \sum_m w_m = 1$. Markov's inequality therefore gives $\Pr_{\xi^\star}\{E_1(\xi^\star) > 1/\delta\} \le \delta$, which is Equation (84). This is the split-likelihood argument of universal inference (Wasserman et al., 2020), applied to the complete stopped-record law. □

The proposition applies to the exact experiment in Equation (74). Finite-step observations of a continuously evolved state-dependent SDE instead use their exact transition density or a validated likelihood-ratio enclosure.

D.3 Fixed-generator bridge and closed gates

For $c = (a, b)$ define the first-order continuation system

$$\frac{d}{dx}\begin{pmatrix} v \\ v' \end{pmatrix} = M_\xi(x)\begin{pmatrix} v \\ v' \end{pmatrix}, \qquad M_\xi(x) = \begin{pmatrix} 0 & 1 \\ \varrho/A_\xi(x) & -B_\xi(x)/A_\xi(x) \end{pmatrix}. \tag{89}$$

Let $\Phi_\xi(x, a)$ be its fundamental matrix, and abbreviate $p_a = \ell_0'(a)$ and $q_b = u_0'(b)$. Propagation of the two Neumann traces gives

$$v_{\xi,c}(a) = \frac{q_b - [\Phi_\xi(b, a)]_{22} p_a}{[\Phi_\xi(b, a)]_{21}}. \tag{90}$$

Lemma 3 (Fixed-generator bridge)**.** *Fix $\xi \in \Xi_0$ and $\underline{x} < a < b < \bar{x}$. The boundary problem*

$$\mathcal{H}_\xi v = 0 \text{ on } (a, b), \qquad v'(a) = \ell_0'(a), \quad v'(b) = u_0'(b) \tag{91}$$

has the unique solution in Equations (89)*–*(90)*. Value matching uniquely fixes the two constant transfers in Equation* (76)*. If the resulting obstacles and continuation value satisfy*

$$\begin{array}{ll} L_{\xi,c} < U_{\xi,c} & \text{on } I, \\ L_{\xi,c} < v_{\xi,c} < U_{\xi,c} & \text{on } (a, b), \\ \mathcal{H}_\xi L_{\xi,c} \le 0 & \text{on } [\underline{x}, a], \\ \mathcal{H}_\xi U_{\xi,c} \ge 0 & \text{on } [b, \bar{x}], \end{array} \tag{92}$$

then the pasted candidate is the game value and the two first-contact times form a saddle. These transfers are the unique constant transfers implementing (a, b) in Definition 3.

Proof. Step 1 (positive denominator and existence). Let y solve $\mathcal{H}_\xi y = 0$ with $y(a) = 1$ and $y'(a) = 0$, and set

$$\omega_\xi(x) = \exp\left(\int_{x_0}^{x} \frac{B_\xi(r)}{A_\xi(r)}\, dr\right). \tag{93}$$

Then $(\omega_\xi y')' = \varrho\omega_\xi y/A_\xi$. A first-zero argument shows that y remains positive and strictly increases to the right of a. Hence $[\Phi_\xi(b, a)]_{21} = y'(b) > 0$, and Equation (90) gives a solution with both prescribed derivative traces.

Step 2 (Neumann uniqueness). If w is the difference of two solutions, then $w'(a) = w'(b) = 0$ and $(\omega_\xi w')' = \varrho\omega_\xi w/A_\xi$. Multiplying by w and integrating gives

$$-\int_a^b \omega_\xi(x)(w'(x))^2\, dx = \varrho \int_a^b \frac{\omega_\xi(x)}{A_\xi(x)} w(x)^2\, dx. \tag{94}$$

The left side is nonpositive and the right side is nonnegative, so $w = 0$.

Step 3 (matching and verification). Continuity at the two contacts forces $\alpha = v(a) - \ell_0(a)$ and $\beta = v(b) - u_0(b)$. Under Equation (92), the pasted candidate is C^1, lies between the ordered obstacles, and has the generator signs required by Lemma 2. That lemma gives the first-contact saddle from every initial state. Conversely, any implementation in Definition 3 has Equation (91) in continuation and inherits the two derivative traces by smooth fit. Steps 1–2 and value matching therefore force the same continuation value and transfers in the declared class. □

D.4 GENERATOR RESPONSE AND ISO-TRANSFER GEOMETRY

The contact derivatives in Proposition 1 hold the generator fixed. We now differentiate in the complementary direction: the contacts remain fixed while drift and diffusion vary. The resulting formula is an exact infinite-dimensional sensitivity representation. Combining it with the contact Jacobian then gives the local geometry of curves along which the recovered transfers remain unchanged.

Fix a verified pair (ξ, c), abbreviate $v = v_{\xi,c}$, and write

$$\omega(x) = \exp\left(\int_{x_0}^{x} \frac{B_\xi(r)}{A_\xi(r)}\, dr\right). \tag{95}$$

Define separately

$$m(x) = \frac{\omega(x)}{A_\xi(x)}. \tag{96}$$

The differentiability statement below is made on an explicit function space. Let

$$\mathcal{X} = C(I) \times C(I) \tag{97}$$

with norm

$$\|(A, B)\|_{\mathcal{X}} = \max\{\|A\|_\infty, \|B\|_\infty\}. \tag{98}$$

The uniformly elliptic coefficient domain

$$\mathcal{U} = \left\{(A, B) \in \mathcal{X} : \inf_{x\in I} A(x) > 0\right\} \tag{99}$$

is open in $\mathcal{X}$. For $(A, B) \in \mathcal{U}$, write

$$\mathcal{H}_{A,B} = A\partial_{xx} + B\partial_x - \varrho. \tag{100}$$

Also let

$$\mathcal{Z} = \{w \in C^2([a, b]) : w'(a) = w'(b) = 0\} \tag{101}$$

with its usual C^2 norm.

Let the endpoint response kernels ψ_a, ψ_b be the unique solutions of

$$\mathcal{H}_\xi \psi_a = 0, \qquad \psi_a'(a) = \omega(a)^{-1}, \qquad \psi_a'(b) = 0, \tag{102}$$

$$\mathcal{H}_\xi \psi_b = 0, \qquad \psi_b'(a) = 0, \qquad \psi_b'(b) = -\omega(b)^{-1}. \tag{103}$$

Lemma 4 (Uniform Neumann resolvent and endpoint representers)**.** *There is a neighborhood $\mathcal{N} \subset \mathcal{U}$ of (A_ξ, B_ξ) and a finite constant $C_\mathcal{N}$ with the following properties.*

(i) For every $(A, B) \in \mathcal{N}$ and $g \in C([a, b])$, the homogeneous Neumann problem

$$\mathcal{H}_{A,B} w = g \tag{104}$$

with

$$w'(a) = 0 \qquad \text{and} \qquad w'(b) = 0 \tag{105}$$

has a unique solution $w \in \mathcal{Z}$, and

$$\|w\|_{C^2} \le C_\mathcal{N} \|g\|_\infty. \tag{106}$$

(ii) At the base coefficients, the endpoint evaluation functionals of the resolvent have the exact representations

$$w(a) = \int_a^b \psi_a(x) m(x) g(x)\, dx \tag{107}$$

and

$$w(b) = \int_a^b \psi_b(x) m(x) g(x)\, dx. \tag{108}$$

(iii) The solution operators

$$\mathcal{R}_{A,B} : g \longmapsto w \tag{109}$$

vary continuously in operator norm from $C([a,b])$ to $\mathcal{Z}$ over a possibly smaller neighborhood $\mathcal{N}$.

Proof. We begin at the base coefficient pair. Multiplication of $\mathcal{H}_\xi w = g$ by $m = \omega/A_\xi$ yields

$$(\omega w')' - \varrho m w = mg. \tag{110}$$

If $g = 0$, multiply this equation by w and integrate over $[a,b]$. The homogeneous derivative traces remove the boundary term and give

$$\int_a^b \omega(x)|w'(x)|^2\,dx + \varrho \int_a^b m(x)|w(x)|^2\,dx = 0. \tag{111}$$

Both weights are strictly positive and $\varrho > 0$, so w vanishes identically. Existence for continuous forcing follows from the standard two-point linear ODE construction. Thus

$$\mathcal{H}_{A_\xi,B_\xi} : \mathcal{Z} \longrightarrow C([a,b]) \tag{112}$$

is a bounded bijection. The bounded inverse theorem supplies a finite base resolvent norm.

For a perturbation $(\Delta A, \Delta B)$, define the bounded operator

$$\mathcal{E}_{\Delta A,\Delta B} w = \Delta A\, w'' + \Delta B\, w'. \tag{113}$$

Its norm from $\mathcal{Z}$ to $C([a,b])$ is bounded by a constant times $\|(\Delta A, \Delta B)\|_{\mathcal{X}}$. Choose the coefficient neighborhood so that

$$\left\|\mathcal{R}_{A_\xi,B_\xi}\mathcal{E}_{\Delta A,\Delta B}\right\| < 1. \tag{114}$$

The operator identity

$$\mathcal{H}_{A,B} = \mathcal{H}_{A_\xi,B_\xi}\left(I + \mathcal{R}_{A_\xi,B_\xi}\mathcal{E}_{\Delta A,\Delta B}\right) \tag{115}$$

then gives an inverse by a convergent Neumann series. Shrinking the neighborhood if necessary makes the resulting inverse norms uniformly bounded, proving part (i). The resolvent identity

$$\mathcal{R}_{A,B} - \mathcal{R}_{\widetilde{A},\widetilde{B}} = -\mathcal{R}_{A,B}\left(\mathcal{H}_{A,B} - \mathcal{H}_{\widetilde{A},\widetilde{B}}\right)\mathcal{R}_{\widetilde{A},\widetilde{B}} \tag{116}$$

and the uniform bound prove part (iii).

It remains to identify the endpoint functionals. For either $\psi \in \{\psi_a, \psi_b\}$, Green's identity applied to Equation (110) gives

$$\int_a^b \psi(x)m(x)g(x)\,dx = \left[\omega(x)\{\psi(x)w'(x) - w(x)\psi'(x)\}\right]_a^b. \tag{117}$$

For ψ_a, its two derivative traces reduce the right side to $w(a)$. For ψ_b, they reduce it to $w(b)$. This proves Equations (107) and (108). □

Lemma 5 (Joint coefficient–contact regularity)**.** *Let*

$$\mathcal{D} = \left\{((A,B),a,b) \in \mathcal{U} \times \mathbb{R}^2 : \underline{x} < a < b < \bar{x}\right\}. \tag{118}$$

For $((A,B),a,b) \in \mathcal{D}$, let $F_{A,B}(a,b)$ be the two-transfer chart obtained from the Neumann problem in Equation (91)*. The map*

$$((A,B),a,b) \longmapsto F_{A,B}(a,b) \tag{119}$$

is continuously Fréchet differentiable on $\mathcal{D}$. Its derivative in the coefficient coordinates is the response operator in Equation (131)*, and its derivative in the two contact coordinates is the Jacobian $J_{A,B,a,b}$ from Proposition 1.*

Proof. The point requiring care is that the Neumann problem lives on the moving interval $[a,b]$. We keep the coefficients on the fixed ambient interval I and work directly with the fundamental matrix, so no differentiability of the translation operator on $C(I)$ is needed.

Step 1 (the fundamental matrix as a joint map). For $(A, B) \in \mathcal{U}$, write

$$M_{A,B}(x) = \begin{pmatrix} 0 & 1 \\ \varrho/A(x) & -B(x)/A(x) \end{pmatrix}. \tag{120}$$

The coefficient map $(A, B) \mapsto M_{A,B}$ is continuously Fréchet differentiable from $\mathcal{U}$ into $C(I; \mathbb{R}^{2\times 2})$. In a direction $\Delta = (\Delta A, \Delta B)$, its derivative is

$$DM_{A,B}[\Delta](x) = \begin{pmatrix} 0 & 0 \\ -\varrho \Delta A(x)/A(x)^2 & B(x)\Delta A(x)/A(x)^2 - \Delta B(x)/A(x) \end{pmatrix}. \tag{121}$$

Let $\Phi_{A,B}(x, y)$ solve

$$\partial_x \Phi_{A,B}(x, y) = M_{A,B}(x)\Phi_{A,B}(x, y), \qquad \Phi_{A,B}(y, y) = I_2. \tag{122}$$

Picard iteration on the compact interval gives locally uniform bounds for Φ and its inverse. Subtracting the two Volterra equations, dividing by the perturbation size, and applying Gronwall's inequality yields the coefficient derivative

$$D_{(A,B)}\Phi_{A,B}(b, a)[\Delta] = \int_a^b \Phi_{A,B}(b, t) DM_{A,B}[\Delta](t)\Phi_{A,B}(t, a)\, dt. \tag{123}$$

The same estimate applied to the difference of two derivatives proves continuity in operator norm.

The semigroup identity for Φ gives the endpoint derivatives without moving the coefficients:

$$\partial_b \Phi_{A,B}(b, a) = M_{A,B}(b)\Phi_{A,B}(b, a) \tag{124}$$

and

$$\partial_a \Phi_{A,B}(b, a) = -\Phi_{A,B}(b, a)M_{A,B}(a). \tag{125}$$

Equations (123)– (125), together with their locally uniform remainder bounds, show that $((A, B), a, b) \mapsto \Phi_{A,B}(b, a)$ is jointly C^1 on $\mathcal{D}$.

Step 2 (the Neumann trace is a regular quotient). Put

$$p(a) = \ell_0'(a) \qquad \text{and} \qquad q(b) = u_0'(b). \tag{126}$$

The lower value trace that enforces the upper derivative condition is

$$\eta(A, B, a, b) = \frac{q(b) - [\Phi_{A,B}(b, a)]_{22} p(a)}{[\Phi_{A,B}(b, a)]_{21}}. \tag{127}$$

The denominator is strictly positive by the first-zero argument in Lemma 3. It remains bounded away from zero on a neighborhood of any base point by continuity. Since $\ell_0, u_0 \in C^3(I)$, Equation (127) is therefore a jointly C^1 quotient.

The upper value trace is

$$\zeta(A, B, a, b) = e_1^\top \Phi_{A,B}(b, a) \begin{pmatrix} \eta(A, B, a, b) \\ p(a) \end{pmatrix}, \qquad e_1 = (1, 0)^\top. \tag{128}$$

Thus the complete chart can be written as

$$F_{A,B}(a, b) = \begin{pmatrix} \eta(A, B, a, b) - \ell_0(a) \\ \zeta(A, B, a, b) - u_0(b) \end{pmatrix}. \tag{129}$$

This representation proves the asserted joint C^1 regularity.

Step 3 (identifying the two derivative blocks). With (a, b) fixed, differentiating the boundary problem in (A, B) gives the homogeneous-Neumann response equation

$$\mathcal{H}_{A,B} w = -\Delta A\, v'' - \Delta B\, v'. \tag{130}$$

The endpoint representers in Lemma 4 identify its two traces with $\mathsf{S}_{A,B,a,b}(\Delta A, \Delta B)$. With (A, B) fixed, ordinary differentiation of Equation (129) gives exactly the two contact columns in Proposition 1. These are the two blocks of the joint Fréchet derivative. □

Theorem 6 (Exact generator response and iso-transfer correction)**.** *Let $A_\xi, B_\xi \in C^1(I)$, let $A_\xi > 0$, and let $c = (a, b)$ be a residual-regular verified contact pair. View the ambient coefficient pair (A_ξ, B_ξ) on I as a point of the open Banach domain $\mathcal{U}$ in Equation* (99)*. For a direction $(\dot{A}, \dot{B}) \in \mathcal{X}$, define*

$$\mathsf{S}_{\xi,c}(\dot{A}, \dot{B}) = -\begin{pmatrix} \displaystyle\int_a^b \psi_a(x)m(x)\big(\dot{A}(x)v''(x) + \dot{B}(x)v'(x)\big)\,dx \\ \displaystyle\int_a^b \psi_b(x)m(x)\big(\dot{A}(x)v''(x) + \dot{B}(x)v'(x)\big)\,dx \end{pmatrix}. \tag{131}$$

Then:

(i) On a neighborhood of the base coefficient pair in $\mathcal{U}$, the boundary-to-transfer map from $\mathcal{X}$ to $\mathbb{R}^2$ is Fréchet differentiable. Its derivative is

$$D_{(A,B)}F_\xi(c)[\dot{A}, \dot{B}] = \mathsf{S}_{\xi,c}(\dot{A}, \dot{B}). \tag{132}$$

Equivalently, its derivative is obtained from one inhomogeneous Neumann solve, or from the two precomputed kernels in Equation (131)*.*

(ii) Let r_A, r_B be nonnegative continuous radii. Define the functional coefficient box

$$\mathfrak{D} = \left\{(\dot{A}, \dot{B}) \in \mathcal{X} : \begin{array}{l} |\dot{A}(x)| \le r_A(x), \\ |\dot{B}(x)| \le r_B(x) \end{array} \text{ for every } x \in [a, b]\right\}. \tag{133}$$

Let

$$\mathcal{K}_{\xi,c} = \mathrm{cl}_{\mathbb{R}^2}(\mathsf{S}_{\xi,c}(\mathfrak{D})) \subset \mathbb{R}^2. \tag{134}$$

The Euclidean closure is essential: a pointwise maximizing sign can be discontinuous even though the admissible coefficient directions are continuous. Then $\mathcal{K}_{\xi,c}$ is a compact, centrally symmetric convex set. For every $u = (u_1, u_2) \in \mathbb{R}^2$, its exact support function is

$$\begin{aligned} h_{\mathcal{K}_{\xi,c}}(u) &:= \max_{y \in \mathcal{K}_{\xi,c}} u^\top y \\ &= \sup_{(\dot{A},\dot{B}) \in \mathfrak{D}} u^\top \mathsf{S}_{\xi,c}(\dot{A}, \dot{B}) \\ &= \int_a^b m(x)|u_1\psi_a(x) + u_2\psi_b(x)|\big(r_A(x)|v''(x)| + r_B(x)|v'(x)|\big)\,dx. \end{aligned} \tag{135}$$

Thus the closed first-order image of a functional coefficient box is determined by one quadrature for each direction u.

(iii) Let $J_{\xi,c} = D_c F_\xi(c)$. There is a neighborhood of the base generator on which a unique contact map $c(A, B)$ preserves the two recovered transfers. Every corrected contact pair remains ordered and residual-regular after shrinking this neighborhood. If the base pair also has strictly positive compact gate minima in Equations (188)–(190)*, the corrected pairs remain verified. The derivative of the contact map is*

$$D_{(A,B)}c[\dot{A}, \dot{B}] = -J_{\xi,c}^{-1}\mathsf{S}_{\xi,c}(\dot{A}, \dot{B}). \tag{136}$$

Consequently, for simultaneous perturbations $h = (h_a, h_b)$ of the contacts and $(\dot{A}, \dot{B})$ of the generator,

$$F_{A+\varepsilon\dot{A},B+\varepsilon\dot{B}}(c + \varepsilon h) = F_\xi(c) + \varepsilon\big(J_{\xi,c}h + \mathsf{S}_{\xi,c}(\dot{A}, \dot{B})\big) + o(\varepsilon). \tag{137}$$

Proof. The proof separates the functional-analytic perturbation from the planar contact geometry. Lemma 4 provides a uniform inverse for the continuation problem. We use it first to obtain a quadratic remainder, then identify the two endpoint functionals, compute the support body, and finally apply the implicit-function theorem.

Step 1 (finite coefficient increments). Take an increment

$$\Delta = (\Delta A, \Delta B) \in \mathcal{X} \tag{138}$$

small enough that $(A_\xi + \Delta A, B_\xi + \Delta B) \in \mathcal{N}$, where $\mathcal{N}$ is the neighborhood from Lemma 4. Let v_Δ solve the continuation problem with the perturbed coefficients and the same two derivative traces as v. Set

$$z_\Delta = v_\Delta - v. \tag{139}$$

Subtracting the two continuation equations yields

$$\mathcal{H}_\xi z_\Delta = -\Delta A\, v_\Delta'' - \Delta B\, v_\Delta'. \tag{140}$$

The derivative boundary conditions cancel, so

$$z_\Delta'(a) = 0 \tag{141}$$

and

$$z_\Delta'(b) = 0. \tag{142}$$

The uniform resolvent bound implies

$$\|z_\Delta\|_{C^2} \le C_{\mathcal{N}} \|\Delta\|_{\mathcal{X}} \left(\|v\|_{C^2} + \|z_\Delta\|_{C^2}\right). \tag{143}$$

After shrinking the neighborhood once more, the last term can be absorbed into the left side. Consequently, there is a constant C_1 such that

$$\|z_\Delta\|_{C^2} \le C_1 \|\Delta\|_{\mathcal{X}}. \tag{144}$$

Step 2 (a quadratic Fréchet remainder). Let $w_\Delta \in \mathcal{Z}$ be the solution of

$$\mathcal{H}_\xi w_\Delta = -\Delta A\, v'' - \Delta B\, v'. \tag{145}$$

The map $\Delta \mapsto w_\Delta$ is bounded and linear. Define the remainder

$$r_\Delta = z_\Delta - w_\Delta. \tag{146}$$

Equations (140) and (145) give

$$\mathcal{H}_\xi r_\Delta = -\Delta A\, z_\Delta'' - \Delta B\, z_\Delta'. \tag{147}$$

Again r_Δ has homogeneous derivative traces. Combining the base resolvent estimate with Equation (144) gives

$$\|r_\Delta\|_{C^2} \le C_2 \|\Delta\|_{\mathcal{X}}^2 \tag{148}$$

for a finite local constant C_2. This explicit quadratic bound establishes Fréchet differentiability; it is stronger than a direction-by-direction linearization.

Step 3 (reading the derivative at both contacts). At fixed contacts and fixed obstacle bases, the transfer increment is

$$F_{A_\xi + \Delta A, B_\xi + \Delta B}(c) - F_\xi(c) = \begin{pmatrix} z_\Delta(a) \\ z_\Delta(b) \end{pmatrix}. \tag{149}$$

Apply the endpoint representers in Lemma 4 to the forcing in Equation (145). They give

$$\begin{pmatrix} w_\Delta(a) \\ w_\Delta(b) \end{pmatrix} = \mathsf{S}_{\xi,c}(\Delta A, \Delta B). \tag{150}$$

Endpoint evaluation is bounded by the C^2 norm, so Equation (148) controls both coordinates of the difference between Equations (149) and (150). This proves part (i).

Step 4 (the exact functional-box image). For $u = (u_1, u_2)$, define

$$\psi_u = u_1 \psi_a + u_2 \psi_b. \tag{151}$$

Linearity gives

$$u^\top \mathsf{S}_{\xi,c}(\dot{A}, \dot{B}) = -\int_a^b m(x) \psi_u(x) \left(\dot{A}(x) v''(x) + \dot{B}(x) v'(x)\right) dx. \tag{152}$$

The pointwise box constraints immediately upper bound this expression by the right side of Equation (135). To show that the bound is exact over continuous directions, put

$$\phi_A(x) = -m(x)\psi_u(x)v''(x) \tag{153}$$

and

$$\phi_B(x) = -m(x)\psi_u(x)v'(x). \tag{154}$$

For $\eta > 0$, the admissible continuous directions

$$\dot{A}_\eta(x) = r_A(x)\frac{\phi_A(x)}{\sqrt{\phi_A(x)^2 + \eta^2}} \tag{155}$$

and

$$\dot{B}_\eta(x) = r_B(x)\frac{\phi_B(x)}{\sqrt{\phi_B(x)^2 + \eta^2}} \tag{156}$$

approach the maximizing signs away from their zero sets. Dominated convergence shows that their objective values approach the pointwise upper bound. This proves the support identity in part (ii), including cases in which a discontinuous sign itself is not an admissible direction. The set $\mathfrak{D}$ is convex and centrally symmetric, and $\mathsf{S}_{\xi,c}$ is a bounded linear map. Its image is therefore convex, centrally symmetric, and bounded. Taking the Euclidean closure in Equation (134) preserves the first two properties and makes the set compact in $\mathbb{R}^2$; it does not change any support supremum. This proves all claims about $\mathcal{K}_{\xi,c}$ without assuming that the pointwise sign is itself continuous.

Step 5 (iso-transfer correction and persistence of verification). Residual regularity makes the two cross residuals strict. Proposition 1 and Theorem 3 give

$$\det J_{\xi,c} > 0. \tag{157}$$

Consider the joint residual map

$$\mathfrak{G}(A, B, c') = F_{A,B}(c') - F_{A_\xi,B_\xi}(c). \tag{158}$$

Lemma 5 makes this map continuously Fréchet differentiable near the base point and identifies its two derivative blocks. Its derivative in c' is $J_{\xi,c}$ and is invertible by Equation (157). The Banach-space implicit-function theorem therefore gives a unique local contact map satisfying

$$F_{A,B}(c(A, B)) = F_{A_\xi,B_\xi}(c). \tag{159}$$

Differentiating this identity gives

$$J_{\xi,c}\, D_{(A,B)}c + \mathsf{S}_{\xi,c} = 0, \tag{160}$$

which proves Equation (136).

Positive contact separation and the two strict contact residuals persist by continuity. Under the additional strict-margin hypothesis, the compact obstacle, continuation-quotient and stopping-region gate minima also remain positive after the neighborhood is reduced. This proves the stated persistence of verification. Finally, combining the coefficient derivative with the contact derivative yields Equation (137). □

The response operator also answers a natural inverse-design question: among all infinitesimal generator interventions that create a prescribed transfer change, which one has minimum quadratic size? Define the matrix-valued response density

$$\mathsf{G}(x) = -m(x)\begin{pmatrix} \psi_a(x)v''(x) & \psi_a(x)v'(x) \\ \psi_b(x)v''(x) & \psi_b(x)v'(x) \end{pmatrix} \tag{161}$$

and its 2×2 Gramian

$$\mathsf{M} = \int_a^b \mathsf{G}(x)\mathsf{G}(x)^\top\, dx. \tag{162}$$

Corollary 4 (Minimum-energy generator intervention)**.** *Suppose the continuation value v is not constant on $[a,b]$. Then M is positive definite. The generator-response operator extends uniquely from continuous directions to $L^2([a,b];\mathbb{R}^2)$ and has rank two. For any desired first-order transfer displacement $y \in \mathbb{R}^2$, the problem*

$$\inf_{\delta \in L^2([a,b];\mathbb{R}^2)} \left\{ \int_a^b \|\delta(x)\|_2^2 \, dx : \int_a^b \mathsf{G}(x)\delta(x) \, dx = y \right\} \tag{163}$$

has the unique solution

$$\delta_y^*(x) = \mathsf{G}(x)^\top \mathsf{M}^{-1} y. \tag{164}$$

The minimum energy is

$$\int_a^b \|\delta_y^*(x)\|_2^2 \, dx = y^\top \mathsf{M}^{-1} y. \tag{165}$$

Moreover, the infinitesimally transfer-invisible generator perturbations form the closed codimension-two subspace

$$\left\{ \delta \in L^2([a,b];\mathbb{R}^2) : \int_a^b \mathsf{G}(x)\delta(x) \, dx = 0 \right\}. \tag{166}$$

The optimizer in Equation (164) *is continuous, so the same value is obtained when the search is restricted to continuous coefficient directions.*

Proof. Continuity of G makes integration against it a bounded map on $L^2([a,b];\mathbb{R}^2)$, so the continuous-direction response has a unique extension to that space. The response identity can be written as

$$\mathsf{S}_{\xi,c}(\dot{A},\dot{B}) = \int_a^b \mathsf{G}(x) \begin{pmatrix} \dot{A}(x) \\ \dot{B}(x) \end{pmatrix} dx. \tag{167}$$

For $u \in \mathbb{R}^2$, let

$$\psi_u = u_1\psi_a + u_2\psi_b. \tag{168}$$

Then

$$u^\top \mathsf{M} u = \int_a^b m(x)^2 \psi_u(x)^2 \left(v''(x)^2 + v'(x)^2 \right) dx. \tag{169}$$

If $u \neq 0$, the boundary derivatives in Equations (102) and (103) prevent ψ_u from being identically zero. A nonzero solution of the second-order homogeneous ODE cannot vanish on an interval. Since v is nonconstant, v' is nonzero on some open subinterval. The integrand in Equation (169) is therefore positive on a set of positive measure. This proves $\mathsf{M} \succ 0$.

Equation (164) is feasible because

$$\int_a^b \mathsf{G}(x)\delta_y^*(x) \, dx = \mathsf{M}\mathsf{M}^{-1} y = y. \tag{170}$$

Any other feasible direction can be written as

$$\delta = \delta_y^* + h \tag{171}$$

with $\int_a^b \mathsf{G}h = 0$. The cross term vanishes:

$$\int_a^b \delta_y^*(x)^\top h(x) \, dx = y^\top \mathsf{M}^{-1} \int_a^b \mathsf{G}(x)h(x) \, dx = 0. \tag{172}$$

Pythagoras now proves uniqueness and Equation (165). Positive definiteness makes the response surjective onto $\mathbb{R}^2$, so its kernel is closed and has codimension two. Finally, G is continuous, and hence so is δ_y^*. □

Set $x_t = a + t(b-a)$. The two normalized gap quotients have the following closed extensions.

Lemma 6 (Closed continuation quotients)**.** *For (ξ, c) in any compact product with positive contact separation, define*

$$\bar{g}_-(\xi, c, t) = \begin{cases} \dfrac{v_{\xi,c}(x_t) - L_{\xi,c}(x_t)}{(x_t - a)^2}, & t > 0, \\ -\dfrac{\mathcal{H}_\xi L_{\xi,c}(a)}{2A_\xi(a)}, & t = 0, \end{cases} \tag{173}$$

$$\bar{g}_+(\xi, c, t) = \begin{cases} \dfrac{U_{\xi,c}(x_t) - v_{\xi,c}(x_t)}{(b - x_t)^2}, & t < 1, \\ \dfrac{\mathcal{H}_\xi U_{\xi,c}(b)}{2A_\xi(b)}, & t = 1. \end{cases} \tag{174}$$

Both functions are continuous on the closed domain $t \in [0, 1]$.

Proof. We prove the two endpoint extensions separately and then establish joint continuity over the compact parameter product.

Step 1 (the lower contact). Put

$$q_-(x) = v_{\xi,c}(x) - L_{\xi,c}(x). \tag{175}$$

Value matching gives

$$q_-(a) = 0. \tag{176}$$

Smooth fit gives

$$q_-'(a) = 0. \tag{177}$$

Because $\mathcal{H}_\xi v_{\xi,c} = 0$ in continuation, evaluation at the contact yields

$$A_\xi(a) q_-''(a) = -\mathcal{H}_\xi L_{\xi,c}(a). \tag{178}$$

For $h > 0$, Taylor's formula with integral remainder reads

$$q_-(a + h) = h^2 \int_0^1 (1 - s) q_-''(a + sh)\, ds. \tag{179}$$

After division by h^2 and passage to the limit, Equations (178) and (179) give

$$\lim_{h \downarrow 0} \frac{q_-(a + h)}{h^2} = -\frac{\mathcal{H}_\xi L_{\xi,c}(a)}{2A_\xi(a)}. \tag{180}$$

This is exactly the $t = 0$ branch in Equation (173).

Step 2 (the upper contact). Now put

$$q_+(x) = U_{\xi,c}(x) - v_{\xi,c}(x). \tag{181}$$

At b, value matching and smooth fit give separately

$$q_+(b) = 0 \tag{182}$$

and

$$q_+'(b) = 0. \tag{183}$$

The continuation equation gives

$$A_\xi(b) q_+''(b) = \mathcal{H}_\xi U_{\xi,c}(b). \tag{184}$$

For $h > 0$, the backward integral remainder is

$$q_+(b - h) = h^2 \int_0^1 (1 - s) q_+''(b - sh)\, ds. \tag{185}$$

Therefore

$$\lim_{h \downarrow 0} \frac{q_+(b - h)}{h^2} = \frac{\mathcal{H}_\xi U_{\xi,c}(b)}{2A_\xi(b)}, \tag{186}$$

which is the $t = 1$ branch in Equation (174).

Step 3 (joint continuity). Positive contact separation permits the common reparameterization

$$x_t = a + t(b - a). \tag{187}$$

Here t ranges over the closed unit interval. The coefficient family, obstacle bases, and derivative boundary data are continuous on the compact parameter product. The uniform Neumann-resolvent argument of Lemma 4, after rescaling to the fixed coordinate t, makes $v_{\xi,c}$ and its first two derivatives jointly continuous in (ξ, c, t). The integral-remainder formulas above converge uniformly at the two endpoints. They therefore paste the interior quotients to the displayed endpoint values continuously, proving the lemma. □

For compact Q and R, define the exact closed gate minima

$$G_{\mathrm{ord}} = \inf_{Q\times R} \inf_{x\in I} (U_{\xi,c}(x) - L_{\xi,c}(x)), \tag{188}$$

$$G_- = \inf_{Q\times R\times[0,1]} \bar{g}_-(\xi, c, t), \qquad G_+ = \inf_{Q\times R\times[0,1]} \bar{g}_+(\xi, c, t), \tag{189}$$

$$G_{\mathrm{low}} = \inf_{Q\times R} \inf_{x\in[\underline{x},a]} (-\mathcal{H}_\xi L_{\xi,c}(x)), \qquad G_{\mathrm{up}} = \inf_{Q\times R} \inf_{x\in[b,\bar{x}]} \mathcal{H}_\xi U_{\xi,c}(x). \tag{190}$$

The state domains in Equations (188)–(190) are compact after the affine transforms $x = \underline{x} + t(a - \underline{x})$ and $x = b + t(\bar{x} - b)$. The certificate requires strictly positive validated lower bounds for all five quantities.

D.5 Hard-bounded noisy contacts

For episode i, let m_i be its exact latent first-hit or administrative censoring index and let $d_i \in \{-, +, 0\}$ identify a lower-owner hit, upper-owner hit, or censoring. A hit path is in continuation at every $k < m_i$ and stops at m_i; a censored path remains in continuation through m_i. Let $\mathcal{C}$ collect exactly these continuation indices. Suppose the observed monitored states satisfy

$$Z_{i,k} = X_{i,k} + E_{i,k}, \qquad |E_{i,k}| \le \varepsilon. \tag{191}$$

The extrema below are formed only when both owner-index sets and $\mathcal{C}$ are nonempty; otherwise the contact procedure abstains before forming an extremum. On an admissible record, define

$$\begin{aligned} s_-^\varepsilon &= \max_{i:d_i=-} (Z_{i,m_i} - \varepsilon), \quad m_C^\varepsilon = \min_{(i,k)\in\mathcal{C}} (Z_{i,k} + \varepsilon), \\ M_C^\varepsilon &= \max_{(i,k)\in\mathcal{C}} (Z_{i,k} - \varepsilon), \qquad s_+^\varepsilon = \min_{i:d_i=+} (Z_{i,m_i} + \varepsilon). \end{aligned} \tag{192}$$

Lemma 7 (Hard-bounded noisy-contact inclusion)**.** *Conditional on the exact owner and timing labels, the coordinatewise existential projection of the latent label inequalities and Equation* (191) *is*

$$\mathcal{S}_Z^\varepsilon = \big([s_-^\varepsilon, m_C^\varepsilon) \times (M_C^\varepsilon, s_+^\varepsilon]\big) \cap \mathcal{D}. \tag{193}$$

This is the exact projection of the inequality-only observation model and an outer set for any model that additionally imposes dynamics. On every admissible realization, $c^\star \in \mathcal{S}_Z^\varepsilon$.

If both owners are observed and

$$\underline{x} < s_-^\varepsilon < m_C^\varepsilon < M_C^\varepsilon < s_+^\varepsilon < \bar{x}, \tag{194}$$

then

$$R_\varepsilon = [s_-^\varepsilon, m_C^\varepsilon] \times [M_C^\varepsilon, s_+^\varepsilon] \tag{195}$$

is a compact outer rectangle containing $c^\star$. Missing owners, ambiguous owner or timing labels, a failed hard-bound validity flag, an empty projection, invalid ordering, failed domain or separation gates, or failed common-core containment returns ABSTAIN.

Proof. Step 1 (lower contact). A lower stopping state satisfies $X_{i,m_i} \le a$. Its error interval contains such a state if and only if $Z_{i,m_i} - \varepsilon \le a$. Intersecting over lower hits gives $s_-^\varepsilon \le a$. A continuation

state satisfies $X_{i,k} > a$, which is feasible in its error interval if and only if $a < Z_{i,k} + \varepsilon$. Intersecting over $\mathcal{C}$ gives $a < m_{\mathcal{C}}^{\varepsilon}$.

Step 2 (upper contact). An upper stopping state satisfies $X_{i,m_i} \geq b$, which is feasible if and only if $b \leq Z_{i,m_i} + \varepsilon$. A continuation state satisfies $X_{i,k} < b$, which is feasible if and only if $Z_{i,k} - \varepsilon < b$. Intersecting gives $M_{\mathcal{C}}^{\varepsilon} < b \leq s_{+}^{\varepsilon}$.

Step 3 (projection and inclusion). The four inequalities are coordinatewise sufficient because each latent monitored state may be chosen independently within its observed error interval. Intersecting with $\mathcal{D}$ proves Equation (193), including its open continuation and closed stopping endpoints. For the realized latent states, $|E_{i,k}| \leq \varepsilon$ places $(a^{\star}, b^{\star})$ inside all four inequalities, proving pathwise inclusion. Under Equation (194), closure is the ordered compact rectangle in Equation (195) for the manuscript's basic ordered domain $\mathcal{D}$. Any additional registered domain or minimum-separation restriction is certified separately; its failure returns abstention rather than the unqualified hull. □

Proposition 3 (Finite-grid contact certificate)**.** *Suppose a Brownian trajectory is monitored at spacing Δ and each retained stopping record contains its last continuation state and first observed stopping state. On the event $\mathcal{O}$ that both owners are observed, let $(\widehat{a}, \widehat{b})$ be the midpoint of the owner-only rectangle $I_- \times I_+$ in Equation* (210)*. Then, pathwise,*

$$a \in I_-, \qquad b \in I_+, \qquad |\widehat{a} - a| \leq \tfrac{1}{2}|I_-|, \qquad |\widehat{b} - b| \leq \tfrac{1}{2}|I_+|. \tag{196}$$

If the normalized process has increments $X_{n+1} - X_n = \sqrt{2\Delta}\, Z_n$, where the Z_n are independent standard normal variables, and the complete dataset generates at most K increments, then for every $\delta \in (0, 1)$,

$$\Pr\Big(\mathcal{O} \cap \Big\{\max\{|\widehat{a} - a|, |\widehat{b} - b|\} > \sqrt{\Delta \log(2K/\delta)}\Big\}\Big) \leq \delta. \tag{197}$$

Whenever the full-path rectangle R_{X} is defined, its midpoint satisfies the same conclusions and its coordinate widths are no larger.

Proof. Every retained lower crossing brackets a between its stopping and preceding states; every upper crossing brackets b analogously. Intersecting valid brackets proves Equation (196). Each resulting interval is no wider than any contributing crossing increment. Moreover,

$$\Pr\Big\{|\sqrt{2\Delta}Z| > 2\sqrt{\Delta \log(2K/\delta)}\Big\} \leq \frac{\delta}{K}. \tag{198}$$

A union bound over at most K generated increments gives Equation (197). The full-path rectangle intersects the owner-only constraints with additional continuation inequalities, so its two coordinate widths cannot increase. □

D.6 Measurability and strict selection

Proposition 4 (Measurable strict selector)**.** *Suppose the likelihood stage returns a finite union Q of closed interval boxes with measurable endpoints and certifies $\widehat{\Xi}_{\delta} \subseteq Q$. Suppose the contact stage returns either abstention or a rectangle R with measurable endpoints satisfying Equation* (78)*. Retain every unresolved likelihood box. Return a certificate only if a finite directed-rounding algorithm establishes strictly positive lower bounds for Equations* (188)–(190) *and for ellipticity, the Neumann denominator, domain containment, contact separation, and common-core containment on the complete $Q \times R$.*

The algorithm returns ABSTAIN *on upstream contact abstention; an empty confidence set or empty/unresolved outer cover; an exhausted interval budget; any nonpositive or unresolved gate; or any output enclosure not proved to contain the complete retained product. A box may be removed only by a validated likelihood exclusion and never because a game gate fails. Then the output, including abstention, is measurable, and every nonabstaining product satisfies the hypotheses of Lemma 3.*

Proof. Step 1 (continuity on compact products). Uniform ellipticity and positive contact separation make the coefficient matrix in Equation (89) continuous on a compact domain. Continuous dependence for linear ODEs, together with a positive denominator, makes $(\xi, c, x) \mapsto v_{\xi,c}(x)$ and

$(\xi, c) \mapsto F_\xi(c)$ continuous. Lemma 6 makes the two quotient gates continuous on closed domains; the other gate integrands are continuous directly.

Step 2 (measurable extrema). The sets Q and R are measurable compact-valued correspondences. The measurable maximum theorem makes every exact gate extremum and every exact endpoint in Equations (19)–(20) measurable. Equivalently, the stipulated finite interval computation uses only measurable comparisons and outward-rounded arithmetic and therefore returns measurable endpoints and a measurable abstention flag.

Step 3 (strict semantics). On return, positivity of G_{ord} gives strict obstacle order, positivity of G_- and G_+ gives both strict continuation gaps, and positivity of G_{low} and G_{up} gives the required weak generator signs with strict margins. The remaining structural gates make the chart and common core well defined. Hence Lemma 3 applies to every retained pair. If any step is unresolved, the stipulated output is abstention rather than a selected subproduct, completing the proof. □

Proof of Theorem 4. Step 1 (uniform saddle validity). On return, Proposition 4 and Lemma 3 give the unique strict-class transfers and a first-contact saddle for every $(\xi, c) \in Q \times R$.

Step 2 (contact reduction at fixed dynamics). Fix $\xi \in Q$. The product gate verifies the whole contact rectangle for this generator. Proposition 1 gives

$$\partial_b \alpha_\xi = K_{+,\xi} \frac{\mathcal{H}_\xi U_{\xi,c}(b)}{A_\xi(b)} \geq 0, \qquad \partial_a \beta_\xi = -K_{-,\xi} \frac{\mathcal{H}_\xi L_{\xi,c}(a)}{A_\xi(a)} \geq 0, \tag{199}$$

where both influence factors are positive. Its column factorization makes the two direct derivatives nonnegative as well. For every $x \in [a_u, b_\ell]$, the same endpoint-differentiation solutions give nonnegative derivatives of $v_{\xi,c}(x)$ with respect to both contacts. Integrating one coordinate at a time across R yields

$$F_\xi(c_\ell) \preceq F_\xi(c) \preceq F_\xi(c_u), \qquad v_{\xi,c_\ell}(x) \leq v_{\xi,c}(x) \leq v_{\xi,c_u}(x). \tag{200}$$

Step 3 (complete dynamics envelope). No order is imposed on ξ. Taking the infimum and supremum of Equation (200) over every $\xi \in Q$ gives Equations (19)–(20). Compactness and continuity make the exact extrema finite; validated outer enclosures remain conservative.

Step 4 (joint selective coverage). Let

$$\mathcal{A} = \{\xi^\star \in \widehat{\Xi}_\delta\}, \qquad \mathcal{B} = \mathcal{C}_c. \tag{201}$$

On $\mathcal{A} \cap \mathcal{B}$, the sentinel $\widehat{R}_c = \dagger$ maps to final abstention. An unavailable dynamics cover or unresolved uniform condition also maps to abstention. On return, $\widehat{\Xi}_\delta \subseteq Q$, so the other part of the disjoint contact event gives $(\xi^\star, c^\star) \in Q \times R$; Step 3 then contains the true transfers and the complete true common-core value. Proposition 2, Equation (77), and a union bound under $\mathbb{P}^\star$ give $\mathbb{P}^\star(\mathcal{A} \cap \mathcal{B}) \geq 1 - \delta - \delta_c$, proving Equation (21). The selection incurs no additional probability loss because it never deletes a retained confidence box in response to a game gate. □

Proof of Corollary 3. Lemma 7 contains the true contacts on every record satisfying the hard bound and exact label semantics. Records failing a structural requirement abstain. Hence $\delta_c = 0$ in Equation (21). Apply the resulting per-game statement and take a union bound over the K dynamics-confidence failures to obtain Equation (79); no cross-game independence is used. □

D.7 Nonvacuity under identification

We state sufficient conditions for eventual return and shrinking bands. Let $0 < \delta_n < 1$ satisfy $\log(1/\delta_n) = o(n)$. Let ℓ_n be the confirmation stopped-record log likelihood and suppose there is a continuous function $K : \Xi_0 \times \Xi_0 \to [0, \infty)$ such that

$$\sup_{\xi \in \Xi_0} \left| \frac{\ell_n(\xi) - \ell_n(\xi^\star)}{n} + K(\xi^\star, \xi) \right| \xrightarrow{p} 0, \qquad \inf_{\|\xi - \xi^\star\| \geq e} K(\xi^\star, \xi) > 0 \quad \text{for every } e > 0. \tag{202}$$

At least one pilot anchor $\widetilde{\xi}_{m_n}$ satisfies $\widetilde{\xi}_{m_n} \to \xi^\star$ in probability and $w_{m_n} \geq w_0 > 0$. Let Q_n be a nonempty outer cover satisfying

$$\widehat{\Xi}_{\delta_n} \subseteq Q_n, \qquad \sup_{q \in Q_n} \operatorname{dist}(q, \widehat{\Xi}_{\delta_n}) \xrightarrow{p} 0. \tag{203}$$

Let $\widehat{R}_{c,n} \in \mathfrak{R} \sqcup \{\dagger\}$ denote the contact output and, on nonsentinel records, write $R_n = \widehat{R}_{c,n}$. Suppose

$$\mathbb{P}^\star\{\widehat{R}_{c,n} = \dagger\} \longrightarrow 0, \qquad d_H(\widehat{R}_{c,n}, \{c^\star\}) \xrightarrow{p} 0. \tag{204}$$

Finally, suppose all exact gate and structural minima at $(\xi^\star, c^\star)$ are strictly positive. The interval implementation uses a fixed deterministic refinement and budget sequence for which gate-enclosure excess and output-enclosure excess converge uniformly to zero on every compact product lying inside fixed positive margins.

Under these conditions,

$$\mathbb{P}^\star\{\text{the procedure returns}\} \longrightarrow 1, \qquad \max_i \operatorname{diam}(\mathcal{T}_{i,n}) + \sup_{x \in [a_{u,n}, b_{\ell,n}]} \operatorname{diam}(\mathcal{V}_n(x)) \xrightarrow{p} 0. \tag{205}$$

Proof. Step 1 (the split-likelihood set contracts). Let

$$S_n = \sum_m w_m \exp(\ell_n(\widetilde{\xi}_m)). \tag{206}$$

The consistent anchor and Equation (202) give

$$\frac{\log S_n - \ell_n(\xi^\star)}{n} \geq \frac{\log w_0}{n} + \frac{\ell_n(\widetilde{\xi}_{m_n}) - \ell_n(\xi^\star)}{n} = o_p(1). \tag{207}$$

Conversely, $S_n \leq \sup_{\xi \in \Xi_0} \exp(\ell_n(\xi))$, and the uniform law with $K \geq 0$ gives the matching upper bound. Hence $(\log S_n - \ell_n(\xi^\star))/n = o_p(1)$.

Fix $e > 0$. Strict separation in Equation (202) places every ξ with $\|\xi - \xi^\star\| \geq e$ below $\ell_n(\xi^\star)$ by a positive multiple of n, uniformly with probability tending to one. Since $\log(1/\delta_n) = o(n)$, such a point cannot satisfy

$$\ell_n(\xi) \geq \log S_n + \log \delta_n, \tag{208}$$

which is the superlevel form of Equation (83). Therefore

$$\sup_{\xi \in \widehat{\Xi}_{\delta_n}} \|\xi - \xi^\star\| \xrightarrow{p} 0. \tag{209}$$

The confidence set is nonempty. Indeed, continuity on compact Ξ_0 gives a confirmation maximum-likelihood point $\widehat{\xi}_n^{\mathrm{mle}}$, and $S_n \leq \exp(\ell_n(\widehat{\xi}_n^{\mathrm{mle}}))$, so $E_1(\widehat{\xi}_n^{\mathrm{mle}}) \leq 1 < 1/\delta_n$. Equation (203) now implies $d_H(Q_n, \{\xi^\star\}) \to 0$ in probability.

Step 2 (eventual verification). Equation (204) and Step 1 imply that $Q_n \times R_n$ enters every neighborhood of $(\xi^\star, c^\star)$ with probability tending to one. By Lemma 6, all exact gate minima are continuous on a structurally admissible compact neighborhood. Their strict values at the true pair therefore provide a smaller neighborhood with a common positive margin. The assumed convergent interval enclosures certify that neighborhood once n is large enough. Contact abstention vanishes by assumption, so the return probability tends to one.

Step 3 (envelope contraction). On the same neighborhood, positive contact separation and denominator make $F_\xi(c)$ and $v_{\xi,c}(x)$ uniformly continuous jointly in all arguments. The exact transfer image of $Q_n \times R_n$ consequently has diameter $o_p(1)$. The same statement holds uniformly in x on the common-core graph. The stipulated outward-enclosure excess is $o_p(1)$, proving Equation (205). □

E Experimental Protocol and Reproducibility

This section records the inputs and numerical decisions behind Section 7. The protocols specify the complete populations, observation streams, solver tolerances and final decision rules for each comparison. A failed global gate is recorded as an abstention rather than assigned a numerical error.

E.1 Relation to the closest stopping-design construction

The protocol-level distinction from Wu (2018) is summarized in Table 1. Wu's agent owns one stopping time and chooses one of two actions at that time. Our two players own different stopping rights and choose two stopping times. Both settings have a common waiting value. Separate opposing stopping rights impose the bilateral saddle inequalities here; their residual signs give the positive cross responses and diagonal-corner envelopes in Section 6.

E.2 Construction and response audits (Q1)

To specify the manufactured problem, write $m(t) = 0.18\sin(\pi t)$, $b_\pm(t) = m(t) \pm 0.8$, and $y = (x - b_-(t))/1.6$. The smooth continuation extended to all states is

$$v(t,x) = 1.1 + 0.15\cos(\pi y) + 0.08\sin(2\pi y)\cos(0.7\pi t) + 0.06t\,y(1-y).$$

The diffusion coefficient is $A(t,x) = 0.8 + 0.1\cos(0.7x + 0.4t)$, the drift is $\mu(t,x) = 0.12\sin(0.6x - 0.3t)$, and $f = -(\partial_t + A\partial_{xx} + \mu\partial_x - 0.4)v$. Choose $\alpha(t) = 0.12\sin(1.3\pi t) + 0.05t$ and $\beta(t) = -0.09\cos(0.8\pi t) + 0.04t$; set $\ell_0 = v - 0.18(x - b_-)^4 - \alpha$ and $u_0 = v + 0.23(x - b_+)^4 - \beta$. The horizon payoff equals the pasted L, v, U at $t = 1$. Consequently, the continuation gaps are the two positive quartics and all boundary data are known exactly. The uniform stopping-residual bracket is at least $12(0.7) - 4(0.12 + 0.18\pi)(1.38) - 0.4(1.38)^2 > 0$. This supplies the analytic gate; sampled residuals are separate numerical checks.

The moving-strip family uses horizon one, state interval $[-2, 2]$, discount 0.4, strip half-width 0.8, center amplitude 0.18, and lower/upper obstacle curvatures $0.18/0.23$. With n_x denoting the spatial grid size and n_t the number of time steps, finite differences use $(n_x, n_t) = (33, 384), (65, 384), (97, 384)$ for spatial refinement and $(97, 64), (97, 128), (97, 256)$ for temporal refinement. Independent Chebyshev collocation uses $(9, 256), (13, 256), (17, 256)$ and $(33, 64), (33, 128), (33, 256)$. The sampled gate uses 257 time points and 1,025 state points; the analytic certificate uses the manufactured identities and global bounds directly. Seven one-at-a-time perturbations target separation, both residual signs, continuation order, terminal consistency, and both corner conditions.

The aligned stationary VI suite uses the curvature parameter c of Equation (65), with $c \in \{0.05, 0.3, 1, 3\}$ and grids $65, 129, 257, 513$. Projected red–black and active-set policy iteration use residual tolerances 10^{-9} and 10^{-10}. The off-grid suite uses three asymmetric fixtures and grids $129, 257, 513$. Intervention tests use $c \in \{0.003, 0.01, 0.03, 0.1, 0.3, 1\}$ and centered step ratios $0.01, 0.003, 0.001$; $\sigma_{\min}(J) < 0.01$ triggers abstention before scoring, where $J = J_{a,b}$ and $\sigma_{\min}$ denotes its smallest singular value. The asymmetric Brownian fixtures use $(a, b, c) = (-0.973, 1.021, 0.1)$, $(-1.017, 0.986, 0.3)$ and $(-0.941, 1.033, 1.0)$ in the witness bases. The general-diffusion coefficients (A, B) are $(1, 0)$, $(1 + 0.2\cos(0.7x),\ 0.15\sin(0.5x) + 0.05x)$, and $(1.1 + 0.15\tanh(0.8x),\ -0.12 + 0.08\cos(0.6x))$, with $a = -1$, $b = 1$ and $\rho = 1$. The general-diffusion suite combines three coefficient fixtures, four ODE step counts, two integration schemes, three endpoint steps, and three curvatures.

E.3 Owner-labelled trajectory protocol (Q2–Q3)

For each candidate, independently draw $g \sim \mathrm{Unif}[0.24, 0.34]$, $s \sim \mathrm{Unif}[0.9, 1.2]$, $r_-, r_+ \sim \mathrm{Unif}[-0.02, 0.02]$, $q_-, q_+ \sim \mathrm{Unif}[-0.01, 0.01]$ and $d_-, d_+ \sim \mathrm{Unif}[-0.003, 0.003]$. The base payoffs are

$$\ell_0(x) = -g - s(1 + r_-)x + q_-x^2 + d_-x^3, \qquad u_0(x) = g - s(1 + r_+)x + q_+x^2 + d_+x^3.$$

Transfers, discount and volatility are independent uniform draws from the ranges below. The contact pair is recovered from the independently solved forward game and must have enclosed roots in $[-1.85, -0.55]$ and $[0.55, 1.85]$.

The candidate generator draws 192 independent games with cubic obstacles, transfers in $[-0.04, 0.04]$, discount in $[0.7, 1.2]$, and volatility in $[1.1, 1.5]$ on $[-2, 2]$. Qualification independently reconstructs the double-obstacle solution, encloses both contacts, and checks the complete stationary certificate. All 192 candidates qualify; the first 24 in candidate order are retained without replacement. Game generation, behavior trajectories and volatility calibration use independent random streams. Qualification is fixed before any trajectory outcome or recovery error is evaluated.

Table 1: Closest stopping-design comparison.

Object	Wu (2018)	This paper
Decision makers	one agent	two opposing players
Stopping rights	one stopping time	one stopping time per player
Terminal choice	one of two actions	owner-labelled lower/upper stop
Designed incentives	time–action transfers	two owner-specific transfers
Optimality notion	single-agent implementation	zero-sum saddle point
Global object	one-agent continuation value	bilateral double-obstacle value
Verification	agent deviation inequalities	bilateral obstacle and generator gates

Table 2: Numerical construction and response checks. Errors are maxima over the scored records unless marked "best."

Measurement	Result
Moving analytic certificate / sampled gate	PASS / PASS
Targeted gate perturbations detected	$7/7$
Best collocation value / transfer error	$1.905 \times 10^{-7}/1.799 \times 10^{-7}$
Aligned VI records converged	$32/32$
Off-grid VI residual / solver discrepancy	$5.873 \times 10^{-10}/1.090 \times 10^{-12}$
Intervention records / curvature abstentions	$60/1$
Maximum intervention relative error	6.164×10^{-6}
Exact-singular column norm / $\lvert \det J \rvert$	$2.220 \times 10^{-16}/1.382 \times 10^{-16}$
General-diffusion cross-law / reciprocity error	$1.453 \times 10^{-13}/5.219 \times 10^{-9}$

Each game receives 32,768 paths through horizon two. Starts cycle through $-0.5, 0, 0.5$. One finest Brownian array at step $0.04/64$ generates monitoring factors 16, 32, and 64 and prefixes 2,048, 8,192, 16,384, and 32,768. Trajectory generation and the 4,096-increment volatility calibration use independent random streams. The gate requires at least 32 unambiguous crossings per owner, nonempty ordered contact intervals, uniform separation at least 0.5, and query containment. The interval authority uses 70 decimal digits, 128 state subdivisions, series order 18, determinant floor 10^{-8}, strict margin floor 10^{-10}, and maximum transfer and value widths 0.5.

For lower-owner crossings let $(X_i^{\mathrm{pre}}, X_i^{\mathrm{stop}})$ be the last continuation/stopping pair, and let $(Y_i^{\mathrm{pre}}, Y_i^{\mathrm{stop}})$ denote the corresponding upper-owner pair. Each extremum below ranges over crossings of the indicated owner. The owner-only intervals are

$$I_- = \big[\max_i X_i^{\mathrm{stop}}, \min_i X_i^{\mathrm{pre}}\big], \quad I_+ = \big[\max_i Y_i^{\mathrm{pre}}, \min_i Y_i^{\mathrm{stop}}\big]. \tag{210}$$

For episode i, let $X_{i,k}$ be its state at monitoring index k, m_i its termination or administrative-censoring index, and $d_i \in \{-, +, 0\}$ its lower-hit, upper-hit, or censoring label. Let $\mathcal{C}$ collect pairs (i, k) with $k < m_i$ and also (i, m_i) when $d_i = 0$. With

$$\begin{aligned} s_- &= \max_{i:d_i=-} X_{i,m_i}, \quad m_C = \min_{(i,k)\in\mathcal{C}} X_{i,k}, \\ M_C &= \max_{(i,k)\in\mathcal{C}} X_{i,k}, \qquad s_+ = \min_{i:d_i=+} X_{i,m_i}, \end{aligned} \tag{211}$$

the full-path estimator is the closed hull $R_{\mathsf{X}} = [s_-, m_C] \times [M_C, s_+]$. It is the smallest closed rectangle containing every contact pair that reproduces all monitored labels, and $R_{\mathsf{X}} \subseteq I_- \times I_+$. Both constructions abstain when an owner is absent or their ordered common core is empty. The proof and the finite-grid localization bound are in Appendix D.5.

The owner-only and full-path contact sets use the same trajectories. With the planted true volatility held fixed, their transfer widths average 4.289×10^{-6} and 3.142×10^{-6}; the corresponding common-core value widths are 2.464×10^{-6} and 1.802×10^{-6}. Both known-volatility routes include all transfer and value targets. With the independently calibrated volatility set propagated, both routes also include all targets; mean transfer widths are $0.110577/0.110574$ and mean value widths are $0.259429/0.259424$ for owner-only/full-path (Table 3). The matched point-recovery comparisons are specified below.

Table 3: Q3: owner-only versus full-path recovery on 24 synthetic games at 32,768 trajectories per game. Fixed rows use the true volatility; joint rows propagate the same estimated-volatility set.

Width summary	Owner-only	Full-path	Reduction
Contact (10^{-6})	4.308	3.163	26.6%
Fixed transfer (10^{-6})	4.289	3.142	26.8%
Fixed value (10^{-6})	2.464	1.802	26.9%
Joint transfer	0.110577	0.110574	$< 0.003\%$
Joint value	0.259429	0.259424	$< 0.003\%$

E.4 Historical industry-return paths with planted contracts (Q4)

Scope and frozen data. This audit changes only the observation-stage increments in the fixed-generator trajectory experiment. The paths come from the value-weighted daily panel of the Fama–French 49 Industry Portfolios (Fama & French, 2026); the contracts, contact boxes, stopping labels, and generators remain planted. We obtained the official Data Library dataset on August 31, 2026; its documentation states that it was created using the 202606 CRSP database. It contains 26,274 dates from July 1, 1926 through June 30, 2026. Before outcome inspection, we fixed the data snapshot, transformation, games, decision rules, and numerical tolerances. The source data are copyrighted and have no stated open-data license.

Chronological construction. After mapping the two documented missing-value sentinels to missing, we divide percentage returns by 100 to obtain decimal simple returns R_{it} for industry i on date t. Valid dates through December 31, 1999 determine the sample mean $\widehat{\mu}_i$ and sample standard deviation $\widehat{s}_i$ of $\log(1 + R_{it})$. On the January 1, 2000–December 31, 2024 test window, dates common to all 49 industries form increments

$$Z_{it} = \frac{\log(1 + R_{it}) - \widehat{\mu}_i}{\widehat{s}_i\sqrt{252}}, \qquad \Delta X_{it}^{(g)} = \sigma_g Z_{it},$$

where σ_g is game g's pre-fixed volatility. The first 6,048 common days form 24 chronological, nonoverlapping 252-day blocks; crossing these blocks with 49 industries yields 1,176 paths per game. The remaining 241 days are dropped and no random numbers are used. Dependence across industries and blocks is retained, so the 24 contracts—not the 1,176 paths— index the descriptive summaries; they are not treated as independent draws.

Games, observations, and controls. Before reading market outcomes, we froze the first 24 qualified contracts in candidate order from a pre-fixed synthetic-game candidate pool. A first event is a uniform crossing of a contract's complete certified lower or upper root box. Touching a box without a uniform crossing terminates that path as ambiguous; a path with neither a hit nor an ambiguity by day 252 is administratively censored. No ambiguous event occurs. Owner-only uses the final crossing brackets, whereas full-path also intersects all monitored continuation inequalities. Each resulting rectangle enters the same fixed-generator interval chart and global gate. Per game, four precondition-checked controls remove all hits, retain only eight contacts per owner, corrupt one owner label, or create an inconsistent crossing bracket.

Checks and interpretation. Both routes are eligible, pass, and include all planted contact, transfer, and query-value targets in 24/24 games. Full-path contact sets are subsets for all paired games, and its mean maximum contact, transfer, and value widths are 0.422, 1.004, and 1.337 in units of 10^{-3}, versus 0.737, 1.681, and 2.113 for owner-only. Across games, lower-owner, upper-owner, and censored counts range over 420–556, 179–288, and 337–577. All 96 controls execute and none passes. The three episode-start strata (2000–2007, 2008–2015, and 2016–2024) each use 392 paths and retain 24/24 passes under both routes. A separately implemented scalar audit reconstructs first events, contact rectangles, target containment, and claim aggregates. These are deterministic planted-target checks conditional on one correlated market panel: they do not invoke the Gaussian finite-grid bound, a generator confidence set, iid uncertainty, or recovery of natural investor preferences. Table 4 summarizes the width comparison.

Table 4: Q4: historical-increment stress test with planted contracts. Inclusion counts are contacts/transfers/query value (C/T/V); widths are scaled by 10^{-3}.

Route	Return	Inclusion (C/T/V)	Contact	Transfer	Value
Owner-only	24/24	24/24 / 24/24 / 24/24	0.737	1.681	2.113
Full-path	24/24	24/24 / 24/24 / 24/24	0.422	1.004	1.337
Reduction	—	—	42.7%	40.3%	36.7%

E.5 Adaptive forward-game calibration (Q2)

Both routes receive the same fitted volatility, full-path contact midpoints, discount and fixed cubic payoff coefficients. The adaptive baseline minimizes squared contact error in mesh units over $[-0.12, 0.12]^2$, starting from $(0, 0)$ independently at each of 513, 1,025 and 2,049 state points. Howard policy iteration solves the double-obstacle VI with banded linear systems; subsequent trial solves may use the previous VI field, clipped to the new obstacles. Contacts are smooth-fit roots of the discrete-harmonic interpolation of the solved continuation field. The inverse chart is not used to initialize or evaluate the baseline's objective.

Powell search uses relative objective tolerance 10^{-12}, parameter tolerance 10^{-9}, at most 60 iterations and 1,200 evaluations. Each VI solve has a 2,048-iteration cap and residual tolerance 10^{-8}. Success requires optimizer termination, VI convergence and contact RMSE at most 0.02 mesh cells. All 24 fits at each grid meet these conditions. The mean transfer distances to the chart are 3.6674×10^{-6}, 9.7685×10^{-7} and 2.5919×10^{-7}; mean solve counts are 238.3, 235.9 and 274.2. At the primary 2,049-state grid, maximum VI residual is 8.18×10^{-11}. The paired VI-minus-chart transfer-RMSE difference is 1.20×10^{-7}, with 95% interval $[2.04, 22.01] \times 10^{-8}$. Table 6 reports all games. The point-recovery and unilateral-gain intervals use two-sided Student-t calculations with 23 degrees of freedom over the 24 game-level values or paired differences.

Timing uses an AMD Ryzen 9 9950X3D CPU under Linux/WSL, float64, one pinned CPU and one BLAS/OpenMP thread. Each measurement covers the complete 24-game estimator batch, including solver construction and optimization on prepared inputs. After one warm-up batch, five measurements alternate method order; each short chart measurement averages 1,000 calls. Median times are 0.0244 ms and 35.50 s, with 6,581 VI solves in each timed baseline batch. Independent recomputation checks 144 forward fields, all timed outputs and complete replay of the 72 optimization histories. Table 5 summarizes the matched comparison.

Table 5: Q2: point recovery. Error is scaled by 10^{-3}; calls count forward-game solves per game.

Method	Fits	Transfer error	Mean calls
Direct chart	24/24	3.5604	0
Adaptive VI (2,049 states)	24/24	3.5605	274.2
Action IRL (1,025 states)	24/24	3.5373	1,856.3

E.6 Stopping-action game IRL (Q2)

We adapt the first-action expert–equilibrium margin of Wang & Klabjan (2018, Equation (8)) to the same two-transfer reward class. Exact tabular equilibrium solves replace neural policy approximation. For mesh spacing Δx, the nearest-neighbor Brownian chain has $\gamma = \sigma_{\text{fit}}^2/(\sigma_{\text{fit}}^2 + \varrho\Delta x^2)$ and continuation value $C_j = \gamma(V_{j-1} + V_{j+1})/2$. This discount equals the expected continuous-time discount over an exponential holding time of rate $\sigma_{\text{fit}}^2/\Delta x^2$. With rows and columns ordered as continue, stop, its local payoff matrix is

$$M_\theta(x) = \begin{pmatrix} C_\theta(x) & U_\theta(x) \\ L_\theta(x) & L_\theta(x) \end{pmatrix}, \qquad \theta = (\alpha, \beta). \tag{212}$$

This preserves Max's tie priority and owner-specific endpoint payments. Let f_θ^*, g_θ^* be equilibrium action distributions and V_θ the value. For an observed Max action a or Min action b, with e_a, e_b the corresponding unit vectors, the margins are

$$r_-(x, a) = V_\theta(x) - e_a^\top M_\theta(x) g_\theta^*(x), \qquad r_+(x, b) = f_\theta^*(x)^\top M_\theta(x) e_b - V_\theta(x). \tag{213}$$

Table 6: Per-game transfer RMSE for the direct chart, primary adaptive VI (2,049 states), and primary action IRL (1,025 states). All methods produce valid final fits for every game.

Game	Chart	Adaptive	Action	Game	Chart	Adaptive	Action
0	0.0023670	0.0023675	0.0017885	12	0.0056423	0.0056420	0.0072568
1	0.0036537	0.0036541	0.0025635	13	0.0066149	0.0066154	0.0047366
2	0.0065409	0.0065407	0.0083460	14	0.0028773	0.0028775	0.0011002
3	0.0006804	0.0006806	0.0006006	15	0.0034915	0.0034913	0.0048892
4	0.0011900	0.0011900	0.0017509	16	0.0031231	0.0031233	0.0026392
5	0.0012561	0.0012564	0.0009170	17	0.0099566	0.0099568	0.0094691
6	0.0004894	0.0004896	0.0019733	18	0.0014999	0.0015003	0.0018415
7	0.0033401	0.0033402	0.0024293	19	0.0015147	0.0015146	0.0024946
8	0.0021656	0.0021657	0.0018026	20	0.0012109	0.0012111	0.0011808
9	0.0014381	0.0014382	0.0013110	21	0.0035234	0.0035237	0.0026932
10	0.0022438	0.0022435	0.0045204	22	0.0130303	0.0130306	0.0121208
11	0.0060122	0.0060123	0.0037624	23	0.0015879	0.0015876	0.0027087

The objective averages these nonnegative margins over observed decisions. A continuation record supplies both continue actions; a Min stop also identifies Max's continue action. A Max stop leaves Min's action unobserved, so only r_- is included. Forced exits are settled at the endpoints.

The same 32,768 paths per game at monitoring step 0.000625 supply 693,394,926 observed decision records in total. Their state counts use 8,193 fixed bins, aggregated to 513 or 1,025 fit states. Payoff coefficients and the independent 4,096-increment volatility estimate match the chart. The two transfers lie in $[-0.12, 0.12]^2$. Three bounded Powell starts, $(0, 0)$, $(-0.08, 0.08)$ and $(0.08, -0.08)$, each allow 800 objective evaluations; a fixed Nelder–Mead refinement allows 600 more. The reward regularizer is zero, with scale fixed by the known payoff bases. Final convergence requires optimizer termination, VI residual at most 2×10^{-8} and an eight-direction local check at transfer radius 10^{-5}. All 24 final fits pass at both grids. At 1,025 states the mean cost, including all starts and checks, is 1,856.3 forward solves and 4.517 s per game on one CPU thread. Observation preparation takes 190.20 s for the complete population and is accounted for separately from the prepared-input fit times. Independent cold solves reproduce all 48 final action objectives by direct evaluation of Equation (212).

Mean transfer/contact RMSEs are $0.004321/0.002208$ at 513 states and $0.003537/0.001184$ at 1,025 states; the chart contact RMSE is 7.91×10^{-7}. For equilibrium evaluation, a separate 4,097-state solver fixes the observed midpoint policies and computes each player's unilateral best response. The maximum over states and owners is averaged across games. Under the common fitted dynamics, the action-minus-chart gain difference is 2.976×10^{-6} with paired 95% interval $[1.923, 4.030] \times 10^{-6}$. Under the planted true dynamics, the corresponding mean gains are 2.295×10^{-5} and 2.392×10^{-5}, with paired difference interval $[-6.688, 4.753] \times 10^{-6}$. These forward-solver diagnostics use a common numerical grid and a common policy-evaluation convention.

E.7 Independent volatility calibration (Q5)

This experiment conditions on the 24 existing full-path behavior records and their fixed known cubic payoff bases. For each game, four independent calibration streams contain 131,072 actual Gaussian increments at step $h = 0.000625$; prefixes of lengths $n = 512$, 8,192 and 131,072 form the paired budgets. The new streams are independent of the behavior records. The Brownian drift is known to be zero, and volatility is the single unknown dynamics parameter. With $S_n = \sum_{j=1}^{n} (\Delta X_j)^2 / h$, fit

$$\widehat{\sigma}_n = \sqrt{S_n/n}, \qquad \Sigma_n = \left[\sqrt{\frac{S_n}{\chi^2_{n,0.975}}}, \sqrt{\frac{S_n}{\chi^2_{n,0.025}}} \right], \tag{214}$$

where $\chi^2_{n,p}$ is the p quantile of the chi-square distribution. The pivot $S_n/\sigma^2 \sim \chi^2_n$ gives 95% coverage for each predefined interval. Quantiles and sum-of-squares calculations are enclosed with relative padding 10^{-10}, checked against a 70-digit Gamma CDF.

Both methods use the same observed contact rectangle R and point estimate. Joint propagation certifies $R \times \Sigma_n$; plug-in certifies $R \times \{\widehat{\sigma}_n\}$. The known payoff polynomials remain fixed throughout

both products. The inherited 70-decimal interval procedure uses two parameter subdivisions, 128 state subdivisions, series order 18, strict margin floor 10^{-10}, determinant floor 10^{-8}, and maximum transfer/value widths 0.5. Every budget retains all 96 calibration records. With this coarse two-way volatility subdivision, at 512 increments the joint products abstain at the prescribed gap/width gates; at the two larger budgets all joint products return and include both transfers and the query value. Plug-in returns on every record, with two-transfer inclusion $0/96$ at each budget. The main text reports the same certificate with the volatility interval split into 32 pieces, described below.

With the coarse subdivision, joint mean transfer/value widths are $0.079145/0.183564$ at 8,192 increments and $0.019796/0.045597$ at 131,072.

Volatility interval inclusion counts are $94/96$, $87/96$ and $92/96$; mean volatility widths are 0.15963, 0.03977 and 0.00993. Mean plug-in transfer widths are 9.540, 9.540 and 9.541 in units of 10^{-6}. Independent verification reconstructs every increment sample, recomputes inclusion and aggregation, and checks 1,728 corner/endpoint image evaluations with a separate exponential-basis solve.

Refining the volatility subdivision of the certificate. We rerun the same certificate on all 288 contact–volatility products, changing one numerical setting: the volatility interval is split into 32 equal pieces. Each piece uses the inherited two-way subdivision in both contacts and volatility, giving 256 parameter boxes per product instead of 8. A product returns only if every piece passes the implementation gates; the reported enclosure is the hull of the piecewise enclosures, and the width ceilings apply to that hull. The data, confidence intervals, payoff bases, gate tolerances and width ceilings remain fixed. Every refined enclosure lies inside its corresponding coarse enclosure. The refined certificate returns on all 288 products and contains both transfers and the query value in 96/96, 94/96, and 95/96 records across the three budgets (Table 7). At 512 increments, all 96 coarse abstentions have unresolved continuation-gap checks and exceed the value-width ceiling; all return with the finer subdivision.

Table 7: Q5 certificate with coarse and refined volatility subdivision on the same 288 products, and the uncertified delta-method comparator. Transfer width is the maximum of the two coordinate widths, averaged over returned records; value width is at $x_q = 0$. Inclusion requires both transfers and the query value. Time is CPU certificate latency per product.

Increments	Setting	Returned	Included	Transfer width	Value width	Time (s)
512	coarse (8 boxes)	$0/96$	—	—	—	1.43
512	refined (256 boxes)	$96/96$	$96/96$	0.058383	0.025825	58.46
512	delta method	$96/96$	$91/96$	0.050068	0.001701	—
8,192	coarse (8 boxes)	$96/96$	$96/96$	0.079145	0.183564	1.23
8,192	refined (256 boxes)	$96/96$	$94/96$	0.014609	0.006135	58.46
8,192	delta method	$96/96$	$87/96$	0.012522	0.000427	—
131,072	coarse (8 boxes)	$96/96$	$96/96$	0.019796	0.045597	1.21
131,072	refined (256 boxes)	$96/96$	$95/96$	0.003660	0.001539	58.03
131,072	delta method	$96/96$	$92/96$	0.003135	0.000108	—

For the refined outputs, increasing the budget from 8,192 to 131,072 increments changes mean transfer and value widths by -0.010949 and -0.004596, with paired 95% intervals $[-0.011512, -0.010385]$ and $[-0.004868, -0.004323]$. These are two-sided Student-t intervals with 23 degrees of freedom over 24 game-level differences, each averaging its four calibration repeats; budgets share nested prefixes. Verification times are medians per product on one CPU thread and measure the joint certificate, including image enclosures. The probability bound in Theorem 4 follows from the 95% volatility interval and pathwise contact inclusion; the table records empirical inclusion of the two transfers and the fixed query value.

The delta-method comparator uses the same estimate $\widehat{\sigma}_n$ and contact rectangle, with standard error $s_n = \widehat{\sigma}_n/\sqrt{2n}$. For each transfer coordinate and for $v_{\sigma,c}(0)$, write the scalar output as $g_\sigma(c)$. With $z_{0.975} = 1.9599639845$, it reports the coordinatewise normal interval

$$\big[g_{\widehat{\sigma}_n}(c_\ell) - z_{0.975}s_n|d(c_\ell)|, \quad g_{\widehat{\sigma}_n}(c_u) + z_{0.975}s_n|d(c_u)|\big],$$

where $d(c)$ is the centered finite difference in volatility with step $10^{-6}\widehat{\sigma}_n$. It returns on every record and checks no implementation gate. Joint inclusion of both transfers and the query value is

91/96, 87/96, and 92/96. The maximum of the two certified transfer widths is 1.15–1.20 times its delta-method counterpart on every record; Table 7 reports transfer and value widths separately.