Distributionally Robust Linear Quadratic Control: Supplementary Material

Anonymous Author(s)
Affiliation
Address
email

The supplementary material is structured as follows. Appendix §A presents the well-known solution to the classic LQG problem using dynamic programming and Kalman Filter estimation. Appendix §B provides the definitions of the stacked system matrices utilized in the compact formulation (5) of the distributionally robust LQG problem. Appendix §C contains the proofs of the formal statements in the main text and provides additional technical results. Appendix §D derives the SDP reformulation of the dual problem (11). Appendix §E, finally, elaborates on the bisection algorithm used for solving the linearization oracle of the Frank-Wolfe algorithm.

A. Solution of the LQG Problem

The classic LQG problem can be solved efficiently via dynamic programming; see, e.g., [3]. That is, the unique optimal control inputs satisfy $u_t^\star = K_t \hat{x}_t$ for every $t \in [T-1]$, where $K_t \in \mathbb{R}^{n\times n}$ is the optimal feedback gain matrix, and $\hat{x}_t = \mathbb{E}_{\mathbb{P}}[x_t | y_0, \ldots, y_t]$ is the minimum mean-squared-error estimator of x_t given the observation history up to time t. Thanks to the celebrated separation principle, K_t can be computed by pretending that the system is deterministic and allows for perfect state observations, and $\hat{x}_t$ can be computed while ignoring the control problem.

To compute K_t, one first solves the deterministic LQR problem corresponding to the LQG problem at hand. Its value function $x_t^\top P_t x_t$ at time t is quadratic in x_t, and P_t obeys the backward recursion

$$P_t = A_t^\top P_{t+1} A_t + Q_t - A_t^\top P_{t+1} B_t (R_t + B_t^\top P_{t+1} B_t)^{-1} B_t^\top P_{t+1} A_t \quad \forall t \in [T-1] \tag{A.1a}$$

initialized by $P_T = Q_T$. The optimal feedback gain matrix K_t can then be computed from P_{t+1} as

$$K_t = -(R_t + B_t^\top P_{t+1} B_t)^{-1} B_t^\top P_{t+1} A_t \quad \forall t \in [T-1]. \tag{A.1b}$$

Since x_t and $(y_0, \ldots, y_t)$ are jointly normally distributed, the minimum mean-squared-error estimator $\hat{x}_t$ can be calculated directly using the formula for the mean of a conditional normal distribution. Alternatively, however, one can use the Kalman filter to compute $\hat{x}_t$ recursively, which is significantly more insightful and efficient. The Kalman filter also recursively computes the covariance matrix Σ_t of x_t conditional on $y_0, \ldots, y_t$ and the covariance matrix $\Sigma_{t+1|t}$ of x_{t+1} conditional on $y_0, \ldots, y_t$ evaluated under $\mathbb{P}$. Specifically, these covariance matrices obey the forward recursion

$$\left.\begin{aligned} \Sigma_t &= \Sigma_{t|t-1} - \Sigma_{t|t-1} C_t^\top (C_t \Sigma_{t|t-1} C_t^\top + V_t)^{-1} C_t \Sigma_{t|t-1} \\ \Sigma_{t+1|t} &= A_t \Sigma_t A_t^\top + W_t \end{aligned}\right\} \forall t \in [T-1] \tag{A.2}$$

initialized by $\Sigma_{0|-1} = X_0$. Using $\Sigma_{t|t-1}$, we then define the Kalman filter gain as

$$L_t = \Sigma_t C_t^\top V_t^{-1} \quad \forall t \in [T-1]$$

which allows us to compute the minimum mean-squared-error estimator via the forward recursion

$$\hat{x}_{t+1} = A_t \hat{x}_t + B_t u_t + L_{t+1} (y_{t+1} - C_{t+1}(A_t \hat{x}_t + B_t u_t)) \quad \forall t \in [T-1]$$

initialized by $\hat{x}_0 = L_0 y_0$. One can also show that the optimal value of the LQG problem amounts to

$$\sum_{t=0}^{T-1} \mathrm{Tr}((Q_t - P_t)\Sigma_t) + \sum_{t=1}^{T} \mathrm{Tr}(P_t(A_{t-1}\Sigma_{t-1}A_{t-1}^\top + W_{t-1})) + \mathrm{Tr}(P_0 X_0). \tag{A.3}$$

B. Definitions of Stacked System Matrices

The stacked system matrices appearing in the distributionally robust LQG problem (5) are defined as follows. First, the stacked state and input cost matrices $Q \in \mathbb{S}^{n(T+1)}$ and $R \in \mathbb{S}^{mT}$ are set to

$$Q = \begin{bmatrix} Q_0 & & & \\ & Q_1 & & \\ & & \ddots & \\ & & & Q_T \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} R_0 & & & \\ & R_1 & & \\ & & \ddots & \\ & & & R_{T-1} \end{bmatrix},$$

respectively. Similarly, the stacked matrices appearing in the linear dynamics and the measurement equations $C \in \mathbb{R}^{pT \times n(T+1)}$, $G \in \mathbb{R}^{n(T+1)\times n(T+1)}$ and $H \in \mathbb{R}^{n(T+1)\times mT}$ are defined as

$$C = \begin{bmatrix} C_0 & 0 & & & \\ & C_1 & 0 & & \\ & & \ddots & \ddots & \\ & & & C_{T-1} & 0 \end{bmatrix}, \quad G = \begin{bmatrix} A_0^0 & & & \\ A_0^1 & A_1^1 & & \\ \vdots & & \ddots & \\ A_0^T & A_1^T & \dots & A_T^T \end{bmatrix}$$

and

$$H = \begin{bmatrix} 0 & & & & \\ A_1^1 B_0 & 0 & & & \\ A_1^2 B_0 & A_2^2 B_1 & 0 & & \\ \vdots & & & \ddots & \\ \vdots & & & & 0 \\ A_1^T B_0 & A_2^T B_1 & \dots & \dots & A_T^T B_{T-1} \end{bmatrix},$$

respectively, where $A_s^t = \prod_{k=s}^{t-1} A_k$ for every $s < t$ and $A_s^t = I_n$ for $s = t$.

Using the stacked system matrices, we can now express the purified observation process η as a linear function of the exogenous uncertainties w and v that is *not* impacted by u; see also [2, 7]

Lemma B.1. *We have $\eta = Dw + v$, where $D = CG$.*

Proof of Lemma B.1. The purified observation process is defined as $\eta = y - \hat{y}$. Recall now that the observations of the original system satisfy $y = Cx + v$. Similarly, one readily verifies that the observations of the fictitious noise-free system satisfy $\hat{y} = C\hat{x}$. Thus, we have $\eta = C(x - \hat{x}) + v$. Next, recall that the state of the original system satisfies $x = Hu + Gw$, and note that the state of the fictitious noise-free system satisfies $\hat{x} = Hu$. Combining all of these linear equations finally shows that u cancels out and that $\eta = CGw + v = Dw + v$. □

C. Proofs

C.1. Additional Technical Results

It is well known that every causal controller that is linear in the original observations y can be reformulated as a causal controller that is linear in the purified observations η and vice versa [2, 7]. Perhaps surprisingly, however, the one-to-one transformation between the respective coefficients of y and η is *not* linear. To keep this paper self-contained, we review these insights in the next lemma.

Lemma C.1. *If $u = U\eta + q$ for some $U \in \mathcal{U}$ and $q \in \mathbb{R}^{pT}$, then $u = U'y + q'$ for $U' = (I + UCH)^{-1}U$ and $q' = (I + UCH)^{-1}q$. Conversely, if $u = U'y + q'$ for some $U' \in \mathcal{U}$ and $q' \in \mathbb{R}^{pT}$, then $u = U\eta + q$ for $U = (I - U'CH)^{-1}$ and $q = (I - U'CH)^{-1}q'$.*

Proof of Lemma C.1. If $u = U\eta + q$ for some $U \in \mathcal{U}$ and $q \in \mathbb{R}^{pT}$, then we have

$$u = U\eta + q = U(y - \hat{y}) + q = Uy - UC\hat{x} + q = Uy - UCHu + q,$$

where the second equality follows from the definition of η, the third equality holds because $y = Cx + v$, and the last equality exploits our earlier insight that $\hat{y} = C\hat{x}$. The last expression depends only on y and u. Solving for u yields $u = U'y + q'$, where $U' = (I + UCH)^{-1}U$ and $q' = (I + UCH)^{-1}q$.

Note that $(I + UCH)$ is indeed invertible because $I + UCH$ is a lower triangular matrix with all diagonal entries equal to one, ensuring a determinant of one.

Similarly, if $u = U'y + q'$ for some $U' \in \mathcal{U}$ and $q' \in \mathbb{R}^{pT}$, then we have

$$u = U'y + q' = U'(\eta + \hat{y}) + q' = U'\eta + U'C\hat{x} + q' = U'\eta + U'CHu + q'.$$

Solving for u yields $u = U\eta + q$, where $U = (I - U'CH)^{-1}U'$ and $q = (I - U'CH)^{-1}q'$. Note again that $(I - U'CH)$ is indeed invertible because $(I - U'CH)$ is a lower triangular matrix with all diagonal entries equal to one. □

C.2. Proofs of Section 3

Proof of Proposition 3.2. In problem (8), both u and x are linear in w and v, i.e., $u = q+UDw+Uv$ and $x = Hu + Gw = Hq + HUDw + HUv + Gw$. By substituting the linear representations of u and x into the objective function of problem (8), we obtain the following equivalent reformulation.

$$\begin{aligned}\min_{\substack{q\in\mathbb{R}^{pT}\\ U\in\mathcal{U}}} \max_{\mathbb{P}\in\mathcal{G}} \; & \mathbb{E}_{\mathbb{P}}\left[w^\top \left(D^\top U^\top (R + H^\top QH)UD + 2D^\top U^\top H^\top QG + G^\top QG\right) w\right] \\ & + \mathbb{E}_{\mathbb{P}}\left[v^\top \left(U^\top (R + H^\top QH)U\right) v\right] + q^\top (R + H^\top QH)q\end{aligned}$$

For any fixed $\mathbb{P} \in \mathcal{G}$, we can express the expectation in the objective function of the above problem in terms of the covariance matrices $W = \mathbb{E}_{\mathbb{P}}[ww^\top]$ and $V = \mathbb{E}_{\mathbb{P}}[vv^\top]$. Thus, the problem becomes

$$\begin{aligned}\min_{\substack{q\in\mathbb{R}^{pT}\\ U\in\mathcal{U}}} \max_{W,V,\mathbb{P}} \quad & \mathrm{Tr}\left(\left(D^\top U^\top (R+H^\top QH)UD+ 2G^\top QHUD +G^\top QG\right)W\right) \\ & + \mathrm{Tr}\left(\left(U^\top (R + H^\top QH)U\right)V\right)+q^\top (R+H^\top QH)q \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{G}, \;\; W = \mathbb{E}_{\mathbb{P}}[ww^\top], \;\; V = \mathbb{E}_{\mathbb{P}}[vv^\top].\end{aligned} \tag{A.4}$$

Recall now the definition of $\mathcal{G}$, and note that the requirements $\mathbb{G}(X_0, \hat{X}_0) \le \rho_{x_0}$, $\mathbb{G}(W_t, \hat{W}_t) \le \rho_{w_t}$ and $\mathbb{G}(V_t, \hat{V}_t) \le \rho_{v_t}$ are equivalent to the convex constraints $\mathbb{G}(X_0, \hat{X}_0)^2 \le \rho_{x_0}^2$, $\mathbb{G}(W_t, \hat{W}_t)^2 \le \rho_{w_t}^2$ and $\mathbb{G}(V_t, \hat{V}_t)^2 \le \rho_{v_t}^2$, respectively, for all $t \in [T - 1]$. The definition of $\mathcal{G}$ also implies that

$$W = \mathbb{E}_{\mathbb{P}}[ww^\top] = \mathrm{diag}(X_0, W_0, \ldots, W_{T-1}) \quad \text{and} \quad V = \mathbb{E}_{\mathbb{P}}[vv^\top] = \mathrm{diag}(V_0, \ldots, V_{T-1}).$$

Problem (A.4) thus constitutes a relaxation of problem (9). Indeed, the feasible set of the inner maximization problem in (A.4) is a subset of the feasible set of the inner maximization problem in (9). Moreover, for any W and V feasible in the inner maximization problem in (9), the distribution $\mathbb{P} = \mathbb{P}_{x_0} \times (\times_{t=0}^{T-1}\mathbb{P}_{w_t}) \times (\times_{t=0}^{T}\mathbb{P}_{v_t})$ defined through $\mathbb{P}_{x_0} = \mathcal{N}(0, X_0)$, $\mathbb{P}_{w_t} = \mathcal{N}(0, W_t)$ and $\mathbb{P}_{v_t} = \mathcal{N}(0, V_t)$, $t \in [T - 1]$, is feasible in the inner maximization problem in (A.4) with the same objective value. The relaxation is thus exact, and the optimal values of (8), (9) and (A.4) coincide. □

Proof of Proposition 3.4. Recall that the space $\mathcal{U}_y$ of all causal output-feedback controllers coincides with the space $\mathcal{U}_\eta$ of all causal *purified* output-feedback controllers. We can thus replace the feasible set $\mathcal{U}_\eta$ of the inner minimization problem in (10) with $\mathcal{U}_y$. Hence, for any fixed $\mathbb{P} \in \mathcal{W}_\mathcal{N}$, the inner minimization problem in (10) constitutes a classic LQG problem. By standard LQG theory [3], it is solved by a *linear* output-feedback controller of the form $u = U'y+q'$ for some $U' \in \mathcal{U}$ and $q' \in \mathbb{R}^{pT}$; see also Appendix §A. Lemma C.1 shows, however, that any linear output-feedback controller can be equivalently expressed as a linear *purified*-output feedback controller of the form $u = U\eta + q$ for some $U \in \mathcal{U}$ and $q \in \mathbb{R}^{pT}$. In summary, the above reasoning shows that the feasible set of the inner minimization problem in (10) can be reduced to the family of all linear purified-output feedback controllers without sacrificing optimality. Thus, problem (10) is equivalent to

$$\begin{aligned}\max_{\mathbb{P}\in\mathcal{W}_\mathcal{N}} \min_{q,U,x,u} \quad & \mathbb{E}_{\mathbb{P}}\left[u^\top Ru + x^\top Qx\right] \\ \text{s.t.} \quad & U \in \mathcal{U}, \;\; u = q + U\eta, \;\; x = Hu + Gw.\end{aligned}$$

Using a similar reasoning as in the proof of Proposition 3.2, we can now substitute the linear representations of u and x into the objective function and reformulate the above problem as

$$\begin{aligned}\max_{W,V,\mathbb{P}} \min_{\substack{q\in\mathbb{R}^{pT}\\ U\in\mathcal{U}}} \quad & \mathrm{Tr}\left(\left(D^\top U^\top (R+H^\top QH)UD+ 2G^\top QHUD +G^\top QG\right)W\right) \\ & + \mathrm{Tr}\left(\left(U^\top (R + H^\top QH)U\right)V\right)+q^\top (R+H^\top QH)q \\ \text{s.t.} \quad & \mathbb{P} \in \mathcal{W}_\mathcal{N}, \;\; W = \mathbb{E}_{\mathbb{P}}[ww^\top], \;\; V = \mathbb{E}_{\mathbb{P}}[vv^\top].\end{aligned}$$

As $\mathcal{W}_\mathcal{N}$ contains only *normal* distributions, Proposition 3.3 implies that $\mathbb{W}(\mathbb{P}_{x_0}, \hat{\mathbb{P}}_{x_0}) = \mathbb{G}(X_0, \hat{X}_0)$, $\mathbb{W}(\mathbb{P}_{w_t}, \hat{\mathbb{P}}_{w_t}) = \mathbb{G}(W_t, \hat{W}_t)$ and $\mathbb{W}(\mathbb{P}_{v_t}, \hat{\mathbb{P}}_{v_t}) = \mathbb{G}(V_t, \hat{V}_t)$ for all $t \in [T-1]$. We may thus replace the requirement $\mathbb{W}(\mathbb{P}_{x_0}, \hat{\mathbb{P}}_{x_0}) \leq \rho_{x_0}$ in the definition of $\mathcal{W}_\mathcal{N}$ by $\mathbb{G}(X_0, \hat{X}_0) \leq \rho_{x_0}$, which is equivalent to the convex constraint $\mathbb{G}(X_0, \hat{X}_0)^2 \leq \rho_{x_0}^2$. The conditions on the marginal distributions of w_t and v_t, $t \in [T-1]$, admit similar reformulations. The definition of $\mathcal{W}_\mathcal{N}$ also implies that

$$W = \mathbb{E}_\mathbb{P}[ww^\top] = \mathrm{diag}(X_0, W_0, \ldots, W_{T-1}) \quad \text{and} \quad V = \mathbb{E}_\mathbb{P}[vv^\top] = \mathrm{diag}(V_0, \ldots, V_{T-1}).$$

Thus, the feasible set of the outer maximization problem in (11) constitutes a relaxation of that in (10). One readily verifies that the relaxation is exact by using similar arguments as in the proof of Proposition 3.2. Thus, the claim follows. □

Proof of Theorem 3.5. By Proposition 3.2, $\bar{p}^\star$ coincides with the minimum of (9). Similarly, by Proposition 3.4 $\underline{d}^\star$ coincides with the maximum of (11). Note that problems (9) and (11) only differ by the order of minimization and maximization. Note also that $\mathcal{U}$ is convex and closed, $\mathcal{G}_W$ and $\mathcal{G}_V$ are convex and compact by virtue of [5, Lemma A.6], and the (identical) trace terms in (9) and (11) are bilinear in (W, V) and (U, q). The claim thus follows from Sion's minimax theorem [6]. □

C.3. Proofs of Section 4

Note that Proposition 4.1 is consistent with Corollary 3 because the optimal LQG controller corresponding to $\mathbb{P}^\star$ is linear in the past observations.

Proof of Proposition 4.1. By [5, Lemma A.3], the inner problem in (9) admits a maximizer $(W^\star, V^\star)$ with $X_0^\star \succeq \lambda_{\min}(\hat{X}_0)$ as well as $W_t^\star \succeq \lambda_{\min}(\hat{W}_t)$ and $V_t^\star \succeq \lambda_{\min}(\hat{V}_t)$ for all $t \in [T-1]$. Thus, the optimal value of problem (9) and its strong dual (11) does not change if we restrict $\mathcal{G}_W$ and $\mathcal{G}_V$ to $\mathcal{G}_W^+$ and $\mathcal{G}_V^+$, respectively. We may thus conclude that problem (11) has a maximizer $(W^\star, V^\star)$ with $V_t^\star \succeq \lambda_{\min}(\hat{V}_t) \succ 0$ for all $t \in [T-1]$. This in turn implies that problem (6) is solved by a normal distribution $\mathbb{P}^\star$ under which the covariance matrix of the observation noise v_t satisfies $V_t^\star \succ 0$ for every $t \in [T-1]$. As (5) and (6) are strong duals, the optimal solution $u^\star$ of problem (5) forms a Nash equilibrium with $\mathbb{P}^\star$, i.e., $u^\star$ is a best response to $\mathbb{P}^\star$ and thus solves the *classic* LQG problem corresponding to $\mathbb{P}^\star$. As $R_t \succ 0$ for every $t \in [T-1]$, this best response $u^\star$ is unique, and as $V_T^\star \succ 0$ for every $t \in [T-1]$, $u^\star$ is in fact the Kalman filter-based optimal output-feedback strategy corresponding to $\mathbb{P}^\star$ (which can be obtained using the techniques highlighted in Appendix §A). □

Before proving Proposition 4.2, recall that $f(W, V)$ is called β-smooth for some $\beta > 0$ if

$$|\nabla f(W, V) - \nabla f(W', V')| \leq \beta \left(\|W - W'\|_F^2 + \|V - V'\|_F^2\right)^{\frac{1}{2}} \quad \forall\, W, W' \in \mathcal{G}_W^+,\ V, V' \in \mathcal{G}_V^+,$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Proof of Proposition 4.2. The function $f(W, V)$ is concave because the objective function of the inner minimization problem in (11) is linear (and hence concave) in W and V and because concavity is preserved under minimization. To prove that $f(W, V)$ is β-smooth, we first recall from Proposition 3.3 that it coincides with the optimal value of the inner minimization problem in (10). As $\mathcal{U}_\eta = \mathcal{U}_y$, $f(W, V)$ can thus be viewed as the optimal value of the classic LQG problem corresponding to the normal distribution $\mathbb{P}$ determined by the covariance matrices W and V. Hence, $f(W, V)$ coincides with (A.3), where Σ_t, for $t \in [T-1]$, is a function of (W, V) defined recursively through the Kalman filter equations (A.2). Note that all inverse matrices in (A.2) are well-defined because any $V \in \mathcal{G}_V^+$ is strictly positive definite. Therefore, Σ_t constitutes a proper rational function (that is, a ratio of two polyonmials with the polynomial in the denominator being strictly positive) for every $t \in [T-1]$. Thus, $f(W, V)$ is infinitely often continuously differentiable on a neighborhood of $\mathcal{G}_W^+ \times \mathcal{G}_V^+$.

As $f(W, V)$ is concave and (at least) twice continuously differentiable, it is β-smooth on $\mathcal{G}_W^+ \times \mathcal{G}_V^+$ if and only if the largest eigenvalue of the Hessian matrix of $-f(W, V)$ is bounded above by β throughout $\mathcal{G}_W^+ \times \mathcal{G}_V^+$. Also, the largest eigenvalue of the positive semidefinite Hessian matrix $\nabla^2(-f(W, V))$ coincides with the spectral norm of $\nabla^2 f(W, V)$. We may thus set

$$\beta = \sup_{W \in \mathcal{G}_W^+, V \in \mathcal{G}_V^+} \|\nabla^2 f(W, V)\|_2, \tag{A.5}$$

where $\|\cdot\|_2$ denotes the spectral norm. The supremum in the above maximization problem is finite and attained thanks to Weierstrass' theorem, which applies because $f(W, V)$ is twice continuously differentiable and the spectral norm is continuous, while the sets $\mathcal{G}_W^+$ and $\mathcal{G}_V^+$ are compact by virtue of [5, Lemma A.6]. This observation completes the proof. ☐

D. SDP Reformulation of the Dual Problem (11)

Instead of solving the dual problem (11) with the customized Frank-Wolfe algorithm of Section 4, it can be reformulated as an SDP amenable to off-the-shelf solvers. This reformulation is obtained by dualizing the inner minimization problem and by exploiting the following preliminary lemma.

Lemma D.1. *For any $\hat{Z} \in \mathbb{S}_+^d$ and $\rho_z \geq 0$, the set $\mathcal{G}_Z = \{Z \in \mathbb{S}_+^d : \mathbb{G}(Z, \hat{Z}) \leq \rho_z\}$ coincides with*

$$\left\{ Z \in \mathbb{S}_+^d : \exists E_z \in \mathbb{S}_+^d \text{ with } \operatorname{Tr}(Z + \hat{Z} - 2E_z) \leq \rho_z^2,\ \begin{bmatrix} \hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}} & E_z \\ E_z & I \end{bmatrix} \succeq 0 \right\}.$$

Proof of Lemma D.1. By Definition 2, we have

$$\mathcal{G}_Z = \{Z \in \mathbb{S}_+^d : \operatorname{Tr}(Z + \hat{Z} - 2(\hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}})^{\frac{1}{2}}) \leq \rho_z^2\}.$$

Next, introduce an auxiliary variable $E_z \in \mathbb{S}_+^d$ subject to the matrix inequality $E_z^2 \preceq (\hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}})$. By [1, Theorem 1], this inequality can be recast as $E_z \preceq (\hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}})^{\frac{1}{2}}$. Hence, we can reformulate the nonlinear matrix inequality in the above representation of $\mathcal{G}_Z$ as $\operatorname{Tr}(Z + \hat{Z} - 2E_z) \leq \rho_z^2$. A standard Schur complement argument reveals that the inequality $E_z^2 \preceq (\hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}})$ is also equivalent to

$$\begin{bmatrix} \hat{Z}^{\frac{1}{2}} Z \hat{Z}^{\frac{1}{2}} & E_z \\ E_z & I \end{bmatrix} \succeq 0.$$

The claim then follows by combining all of these insights. ☐

We are now ready to derive the desired SDP reformulation of problem (11).

Proposition D.2. *If $\hat{V} \succ 0$, then problem* (11) *is equivalent to the SDP*

$$\begin{array}{ll}
\max & \operatorname{Tr}(G^\top QGW) - \operatorname{Tr}(F(R + H^\top QH)^{-1}) \\
\text{s.t.} & W \in \mathbb{S}_+^{n(T+1)},\ V \in \mathbb{S}_+^{pT},\ M \in \mathcal{M},\ F \in \mathbb{S}_+^{Tm} \\
& E_{x_0} \in \mathbb{S}_+^n,\ E_{w_t} \in \mathbb{S}_+^n,\ E_{v_t} \in \mathbb{S}_+^p\ \ \forall t \in [T-1] \\
& \operatorname{Tr}(W_0 + \hat{X}_0 - 2E_{x_0}) \leq \rho_{x_0}^2, \\
& \operatorname{Tr}(W_{t+1} + \hat{W}_t - 2E_{w_t}) \leq \rho_{w_t}^2,\ \operatorname{Tr}(V_t + \hat{V}_t - 2E_{v_t}) \leq \rho_{v_t}^2\ \ \forall t \in [T-1] \\
& \begin{bmatrix} \hat{X}_0^{\frac{1}{2}} X_0 \hat{X}_0^{\frac{1}{2}} & E_{x_0} \\ E_{x_0} & I_n \end{bmatrix} \succeq 0, \\
& \begin{bmatrix} \hat{W}_t^{\frac{1}{2}} W_{t+1} \hat{W}_t^{\frac{1}{2}} & E_{w_t} \\ E_{w_t} & I_n \end{bmatrix} \succeq 0,\ \begin{bmatrix} \hat{V}_t^{\frac{1}{2}} V_t \hat{V}_t^{\frac{1}{2}} & E_{v_t} \\ E_{v_t} & I_p \end{bmatrix} \succeq 0\ \ \forall t \in [T-1] \\
& \begin{bmatrix} F & H^\top QGWD^\top + M/2 \\ (H^\top QGWD^\top + M/2)^\top & DWD^\top + V \end{bmatrix} \succeq 0 \\
& W_0 \succeq \lambda_{\min}(\hat{X}_0) I,\quad W_{t+1} \succeq \lambda_{\min}(\hat{W}_t) I,\quad V_t \succeq \lambda_{\min}(\hat{V}_t) I\ \ \forall t \in [T-1].
\end{array} \tag{A.6}$$

Here, $\mathcal{M}$ denotes the set of all strictly upper block triangular matrices of the form

$$\begin{bmatrix} 0 & M_{1,2} & M_{1,3} & \dots & M_{1,T} \\ & 0 & M_{2,3} & & M_{2,T} \\ & & \ddots & & \vdots \\ & & & 0 & M_{T-1,T} \\ & & & & 0 \end{bmatrix} \in \mathbb{R}^{Tm \times Tp},$$

where $M_{t,s} \in \mathbb{R}^{m \times p}$ for every $t, s \in \mathbb{Z}$ with $1 \leq t < s \leq T$.

Proof of Proposition D.2. The proof relies on dualizing the inner minimization problem in (11). Note that strong duality holds because the primal problem is trivially feasible and involves only equality constraints, which implies that any feasible point is in fact a Slater point. In the following we use $M \in \mathcal{M}$ to denote the Lagrange multiplier of the constraint $U \in \mathcal{U}$, which requires all blocks of the matrix U above the main diagonal to vanish. The Lagrangian function of the inner minimization problem in (11) can therefore be represented as

$$\begin{aligned}\mathcal{L}(q, U, M) = \operatorname{Tr}\left(\big(D^\top U^\top (R + H^\top Q H) U D + G^\top Q G\big) W\right) + 2\operatorname{Tr}(G^\top Q H U D W)\\ + \operatorname{Tr}\left(\big(U^\top (R + H^\top Q H) U\big) V\right) + q^\top (R + H^\top Q H) q + \operatorname{Tr}(U M^\top).\end{aligned}$$

Recall now that $R \succ 0$ and $Q \succeq 0$, and thus $R + H^\top Q H \succ 0$. Consequently, $\mathcal{L}$ is minimized by $q^\star = 0$ for any fixed U and M. In addition, the partial gradient of $\mathcal{L}$ with respect U is given by

$$\frac{\partial \mathcal{L}}{\partial U} = 2(R + H^\top Q H) U D W D^\top + 2(R + H^\top Q H) U V + 2 H^\top Q G W D^\top + M.$$

Recall also that $V \in \mathcal{G}_V^+$ is strictly positive, which implies that $DWD^\top + V \succ 0$ is invertible. As we already know that $R + H^\top Q H \succ 0$ is invertible, as well, $\mathcal{L}$ is minimized by

$$U^\star = -(R + H^\top Q H)^{-1}\left(H^\top Q G W D^\top + M/2\right)(DWD^\top + V)^{-1}$$

for any fixed M. Substituting both $q^\star$ and $U^\star$ into $\mathcal{L}$ yields the dual objective function

$$\begin{aligned}g(M) = \mathcal{L}(q^\star, U^\star, M) = \operatorname{Tr}(G^\top Q G W)\\ - \operatorname{Tr}\left((R + H^\top Q H)^{-1}(H^\top Q G W D^\top + M/2)(DWD^\top + V)^{-1}(H^\top Q G W D^\top + M/2)^\top\right).\end{aligned}$$

The dual of the inner minimization problem in (11) is thus given by $\max_{M \in \mathcal{M}} g(M)$. To linearize the dual objective function, we next introduce an auxiliary variable $F \in \mathbb{S}_+^{mT}$ subject to the matrix inequality $F \succeq (H^\top Q G W D^\top + M/2)(DWD^\top + V)^{-1}(H^\top Q G W D^\top + M/2)^\top$. By using a standard Schur complement reformulation, we can then rewrite the dual problem as

$$\begin{array}{ll}\max & \operatorname{Tr}(G^\top Q G W) - \operatorname{Tr}((R + H^\top Q H)^{-1} F)\\ \text{s.t.} & M \in \mathcal{M},\ F \in \mathbb{S}_+^{mT}\\ & \begin{bmatrix} F & H^\top Q G W D^\top + M/2 \\ (H^\top Q G W D^\top + M/2)^\top & DWD^\top + V \end{bmatrix} \succeq 0.\end{array} \tag{A.7}$$

Next, by replacing the inner problem in (11) with its strong dual (A.7), we can reformulate (11) as

$$\begin{array}{ll}\max & \operatorname{Tr}(G^\top Q G W) - \operatorname{Tr}((R + H^\top Q H)^{-1} F)\\ \text{s.t.} & M \in \mathcal{M},\ F \in \mathbb{S}_+^{mT},\ W \in \mathbb{S}_+^{n(T+1)},\ V \in \mathbb{S}_+^{pT}\\ & \begin{bmatrix} F & H^\top Q G W D^\top + M/2 \\ (H^\top Q G W D^\top + M/2)^\top & DWD^\top + V \end{bmatrix} \succeq 0\\ & \mathbb{G}(X_0, \hat{X}_0)^2 \leq \rho_{x_0}^2,\ \mathbb{G}(W_t, \hat{W}_t) \leq \rho_{w_t}^2,\ \mathbb{G}(V_t, \hat{V}_t) \leq \rho_{v_t}^2 \quad \forall t \in [T-1].\end{array} \tag{A.8}$$

By Proposition 4.1, the inclusion of the constraints $X_0 \succeq \lambda_{\min}(\hat{X}_0) I$, $W_t \succeq \lambda_{\min}(\hat{W}_t) I$ and $V_t \succeq \lambda_{\min}(\hat{V}_t) I$ for all $t \in [T-1]$ has no effect on the solution to problem (A.8). In addition, by Lemma D.1, each (non-linear) Gelbrich constraint in (A.8) can be reformulated as an equivalent (linear) SDP constraint. Thus, problem (A.8) reduces to (A.6), and the claim follows. □

E. Bisection Algorithm for the Linearization Oracle

We now show that the direction-finding subproblem (14) can be solved efficiently via bisection. To this end, we first establish that (14) can be reduced to the solution of a univariate algebraic equation.

Proposition E.1 ([5, Proposition A.4 (iii)]). *If $\hat{Z} \in \mathbb{S}_{++}^d$, $\Gamma_Z \in \mathbb{S}_+^d$, $\Gamma_Z \neq 0$ and $\rho_z \in \mathbb{R}_{++}$, then*

$$\begin{array}{ll}\max & \langle \Gamma_Z, L - Z \rangle\\ \text{s.t.} & \mathbb{G}(L, \hat{Z}) \leq \rho_z\\ & L \succeq \lambda_{\min}(\hat{Z}) I\end{array} \tag{A.9}$$

is uniquely solved by $L^\star = (\gamma^\star)^2 (\gamma^\star I - \Gamma_Z)^{-1} \hat{Z} (\gamma^\star I - \Gamma_Z)^{-1}$, where $\gamma^\star$ is the unique solution of

$$\rho_z^2 - \langle \hat{Z}, (I - \gamma^\star (\gamma^\star I - \Gamma_Z)^{-1})^2 \rangle = 0 \tag{A.10}$$

in the interval $(\lambda_{\max}(\Gamma_Z), \infty)$.

In practice, we need to solve the algebraic equation (A.10) numerically. The numerical error in approximating $\gamma^\star$ should be contained to ensure that $L^\star$ approximates the exact maximizer of problem (A.9). The next proposition shows that, for any tolerance $\delta \in (0,1)$, a δ-approximate solution of (A.9) can be computed with an efficient bisection algorithm.

Proposition E.2 ([5, Theorem 6.4])**.** *For any fixed $\rho_z \in \mathbb{R}_{++}, \hat{Z} \in \mathbb{S}^d_{++}$ and $\Gamma_Z \in \mathbb{S}^d_+, \Gamma_Z \neq 0$, define $\mathcal{G}^+_Z = \{Z \in \mathbb{S}^d_+ : \mathbb{G}(Z,\hat{Z}) \le \rho_z, Z \succeq \lambda_{\min}(\hat{Z})\}$ as the feasible set of problem* (A.9)*, and let $Z \in \mathcal{G}^+_Z$ be any reference covariance matrix. Additionally, let $\delta \in (0,1)$ be the desired oracle precision, and define $\varphi(\gamma) = \gamma(\rho^2 + \langle \gamma(\gamma I - \Gamma_Z)^{-1} - I, \hat{Z}\rangle) - \langle Z, \Gamma_Z\rangle$ for any $\gamma > \lambda_{\max}(\Gamma_Z)$. Then, Algorithm A.1 returns in finite time a matrix $L^\delta_Z \in \mathbb{S}^d_+$ with the following properties. (i) Feasibility: $L^\delta_Z \in \mathcal{G}^+_Z$ (ii) δ-Suboptimality: $\langle L^\delta_Z - Z, \Gamma_Z\rangle \ge \delta \max_{L \in \mathcal{G}^+_Z} \langle \Gamma_Z, L - Z\rangle$.*

Algorithm A.1 Bisection algorithm to compute L^δ_Z

Input: nominal covariance matrix $\hat{Z} \in \mathbb{S}^d_{++}$, radius $\rho \in \mathbb{R}_{++}$,
reference covariance matrix $Z \in \mathcal{G}^+_Z$,
gradient matrix $\Gamma_Z \in \mathbb{S}^d_+$, $\Gamma_Z \neq 0$, precision $\delta \in (0,1)$,
dual objective function $\phi(\gamma)$ defined in Proposition E.2

1: set $\lambda_1 \leftarrow \lambda_{\max}(\Gamma_Z)$, and let p_1 be an eigenvector for λ_1
2: set $\underline{\gamma} \leftarrow \lambda_1(1 + (p_1^\top \hat{Z} p_1)^{\frac{1}{2}}/\rho)$ and $\overline{\gamma} \leftarrow \lambda_1(1 + \mathrm{Tr}(\hat{Z})^{\frac{1}{2}}/\rho)$
3: **repeat**
4: set $\tilde{\gamma} \leftarrow (\overline{\gamma} + \underline{\gamma})/2$ and $L \leftarrow (\tilde{\gamma})^2(\tilde{\gamma} I - \Gamma_Z)^{-1}\hat{Z}(\tilde{\gamma} I - \Gamma_Z)^{-1}$
5: **if** $\frac{\mathrm{d}\phi}{\mathrm{d}\gamma}(\tilde{\gamma}) < 0$ **then** set $\underline{\gamma} \leftarrow \tilde{\gamma}$ **else** $\overline{\gamma} \leftarrow \tilde{\gamma}$ **endif**
6: **until** $\frac{\mathrm{d}\phi}{\mathrm{d}\gamma}(\tilde{\gamma}) > 0$ and $\langle L - Z, \Gamma_Z\rangle \ge \delta\phi(\tilde{\gamma})$

Output: L

In summary, for any $Z \in \{X_0, W_0, \ldots, W_{T-1}, V_0, \ldots, V_{T-1}\}$, Algorithm A.1 computes a δ-approximate solutions to the direction-finding subproblem (14) with $\Gamma_Z = \nabla_Z f(W,V)$.

F. Additional Information on Experiments

Generation of Nominal Covariance Matrices. The nominal covariance matrices of the exogenous uncertainties are constructed randomly using the following procedure. For each exogenous uncertainty $z \in \{x_0, w_0, \ldots, w_{T-1}, v_0, \ldots, v_{T-1}\}$, we denote the dimension of z by d and sample a matrix $M_Z \in \mathbb{R}^{d\times d}$ from the uniform distribution on the hypercube $[0,1]^{d\times d}$. Next, we define $\Xi_Z \in \mathbb{R}^{d\times d}$ as the orthogonal matrix whose columns represent the orthonormal eigenvectors of the symmetric matrix $M_Z + M_Z^\top$. Finally, we set $\hat{Z} = \Xi_Z \Lambda_Z \Xi_Z^\top$, where Λ_Z is a diagonal matrix whose main diagonal is sampled uniformly from the interval $[1,2]^d$. The rationale for adopting this cumbersome procedure is to ensure that the covariance matrix $\hat{Z}$ is positive definite.

Optimality Gap. The optimality gap of the Frank-Wolfe algorithm visualized in Figure 1b is calculated as the sum of the surrogate optimality gaps $\langle L^\delta_Z - Z, \nabla_Z f(W,V)\rangle$ across all $Z \in \{X_0, W_0 \ldots, W_{T-1}, V_0, \ldots, V_{T-1}\}$. For more information on the surrogate optimality gaps see [4].

References

[1] Richard Bellman. Some inequalities for the square root of a positive definite matrix. *Linear Algebra and its Applications*, 1(3):321–324, 1968.

[2] Aharon Ben-Tal, Stephen Boyd, and Arkadi Nemirovski. Control of uncertainty-affected discrete time linear systems via convex programming. *Available at Optimization Online*, 2005.

[3] Dimitri Bertsekas. *Dynamic Programming and Optimal Control*, volume I. Athena Scientific, 2017.

[4] Martin Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, pages 427–435, 2013.

[5] Viet Anh Nguyen, Soroosh Shafieezadeh-Abadeh, Daniel Kuhn, and Peyman Mohajerin Esfahani. Bridging Bayesian and minimax mean square error estimation via Wasserstein distributionally robust optimization. *Mathematics of Operations Research*, 48(1):1–37, 2023.

[6] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.

[7] Joelle Skaf and Stephen P Boyd. Design of affine controllers via convex optimization. *IEEE Transactions on Automatic Control*, 55(11):2476–2487, 2010.