Intermediate Summary

♣ **So far,**

♠ **We have defined the notion of** *uncertain conic problem* **– a family $\mathcal{P}$ of instances**

$$\min_x \left\{ c^T x : \begin{array}{rcl} A_1[\zeta]x + b_1[\zeta] & \in & \mathbf{K}_1 \\ & \cdots\cdots\cdots & \\ A_m[\zeta]x + b_m[\zeta] & \in & \mathbf{K}_m \end{array} \right\}, \zeta \in \mathcal{Z}$$

with the *data* $A_1, b_1, ..., A_m, b_m$ **affinely parameterized by the** *perturbation* ζ **running through a given** *perturbation set* $\mathcal{Z}$**. Here $\mathbf{K}_i$ are simple cones, specifically,** *nonnegative rays* **(uncertain LO), or** *Lorentz/Semidefinite* **cones (uncertain CQO/SDO), and $\mathcal{Z}$ w.l.o.g. can be assumed convex and closed.**

♠ **We associated with uncertain problem $\mathcal{P}$ its** *Robust Counterpart* **– the** *semi-infinite* **convex problem**

$$\min_x \left\{ c^T x : \left. \begin{array}{c} A_1[\zeta]x + b_1[\zeta] \in \mathbf{K}_1 \\ \cdots\cdots\cdots \\ A_m[\zeta]x + b_m[\zeta] \in \mathbf{K}_m \end{array} \right\} \forall \zeta \in \mathcal{Z} \right\} \qquad \text{(RC)}$$

and treated the optimal solution of (RC) as the best "uncertainty-immunized" solution to the uncertain problem of interest.

$$\min_x \left\{ c^T x : \begin{array}{c} A_1[\zeta]x + b_1[\zeta] \in \mathbf{K}_1 \\ \dots\dots\dots \\ A_m[\zeta]x + b_m[\zeta] \in \mathbf{K}_m \end{array} \right\} \forall \zeta \in \mathcal{Z} \right\} \qquad \text{(RC)}$$

♠ **The major theoretical issue we focused on was the one of** *computational tractability* **of the RC. We have seen that this crucial property**

— *always takes place in uncertain LO*, **provided that the perturbation set $\mathcal{Z}$ is computationally tractable,**

— **takes place in the case of** *scenario uncertainty* $\mathcal{Z} = \mathrm{Conv}\{\zeta^1, ..., \zeta^N\}$,

— *sometimes* **takes place in Uncertain CQO/SDO.**

In the case of Uncertain CQO/SDO, we have listed known "solvable cases" (same as "nearly solvable" ones – those where the RC admits a tight safe tractable approximation), and we have seen that these (nearly) solvable cases cover a reasonably wide variety of interesting and important applications.

• I believe building tractable reformulations/tight safe tractable approximations of the RC (or, which is the same, of semi-infinite conic constraints) is a rich, challenging and nontrivial research area.

Challenges: Adjustable Robust Optimization

♣ **Aside of applications of the RO methodology in various subject areas, an important venue of the RO-related research is** *extending the RO methodology beyond the scope of the RC approach as presented so far.*

The most important in this respect is, we believe, passing to *Adjustable Robust Optimization, where the decision variables are allowed to "adjust themselves"*, **to come extent,** *to the true values of the uncertain data.*

♠ **One of the central assumptions which led us to the notion of Robust Counterpart reads:**

> **A.1.** All decision variables in uncertain problem represent "here and now" decisions; they should be assigned specific numerical values as a result of solving the problem *before* the actual data "reveals itself."

While being adequate to many decision making situations, A.1 is not a "universal truth."

♠ **In some cases,** *not all decision variables represent "here and now" decisions.* **In dynamical decision making some of the variables represent "wait and see" decisions and as such can depend on the portion of the true data which "reveals itself" before the moment when the decision is being made.**
Example: In an inventory affected by uncertain demand, there are no reasons to specify all replenishment orders in advance; the true time to specify the replenishment order of period t is the beginning of this period, and thus we can allow this order to depend on the actual demands in periods $1, ..., t-1$.
♠ **Usually,** *not all decision variables represent actual decisions;* **there exist also "analysis" (or slack) variables which do not represent decisions at all and are used to convert the problem into a desired form, e.g., one of a LO problem.** *Since the analysis variables do not represent actual decisions, why not to allow them to depend on the entire true data?*

Example: The convex constraint $\sum_i |a_i^T x - b_i| \leq \tau$ can be represented by a system of linear constraints

$$-y_i \leq a_i^T x - b_i \leq y_i, \sum_i y_i \leq \tau. \qquad (*)$$

When the data a_i, b_i are uncertain and x_j represent "here and now" decisions and thus should be assigned values independent of the true data, *there are absolutely no reasons to impose the same restriction on the slack variables y.* **To see the difference,**

- **The "true" RC of the uncertain constraint** $\sum_i |a_i^T[\zeta]x - b_i[\zeta]| \leq \tau$**,** $\zeta \in \mathcal{Z}$ **is**

$$\sum_i |a_i^T[\zeta]x - b_i[\zeta]| \leq \tau \ \forall \zeta \in \mathcal{Z},$$

and the "true" robust feasible set is

$$\{x : \forall \zeta \in \mathcal{Z} : \exists y : -y_i \leq a_i^T[\zeta]x - b_i[\zeta] \leq y_i, \sum_i y_i \leq \tau\} \qquad (1)$$

- **The RC of the uncertain system** $(*)$ **is**

$$-y_i \leq a_i^T[\zeta]x - b_i[\zeta] \leq y_i, \sum_i y_i \leq \tau \ \forall \zeta \in \mathcal{Z},$$

and the robust feasible set is

$$\{x : \exists y : \forall \zeta \in \mathcal{Z} : -y_i \leq a_i^T[\zeta]x - b_i[\zeta] \leq y_i, \sum_i y_i \leq \tau\} \qquad (2)$$

(2) **is smaller than** (1)**, and the difference can be dramatic:**

$$|x + \zeta| + |x - \zeta| \leq 2, \ \zeta \in [-1, 1] \Rightarrow \begin{cases} (1) = \{-1 \leq x \leq 1\} \\ (2) = \{0\} \end{cases}$$

$$\mathcal{P} = \left\{ \min_x \{ c^T[\zeta]x + d[\zeta] : \sum_{j=1}^n x_j A_j[\zeta] \leq b[\zeta] \} : \zeta \in \mathcal{Z} \right\}$$

Adjustable and Affinely Adjustable Robust Counterpart

♣ **In order to allow for the decision variables in $\mathcal{P}$ to "adjust themselves," to some extent, to the true values of the uncertain data, we could act as follows:**

• **We fix matrices P_j and allow the decision variable x_j to be an arbitrary function of the "portion" $P_j\zeta$ of the true data: $x_j = X_j(P_j\zeta)$**

• **We plug the** *decision rules* $X_j(P_j\zeta)$ **into $\mathcal{P}$ and require them to be** *robust feasible*:

$$\sum_{j=1}^n X_j(P_j\zeta)A_j[\zeta] \leq b[\zeta] \;\forall \zeta \in \mathcal{Z}$$

• **We associate with uncertain problem $\mathcal{P}$ its** *Adjustable Robust Counterpart*

$$\min_{t, X_j(\cdot)} \left\{ t : \begin{array}{l} \sum_j c_j[\zeta]X_j(P_j\zeta) + d[\zeta] \leq t \\ \sum_j X_j(P_j\zeta)A_j[\zeta] \leq b[\zeta] \end{array} \right\} \forall \zeta \in \mathcal{Z} \right\} \quad \text{(ARC)}$$

Note: When the decision rules $X_j(\cdot)$ are restricted to be constant, (ARC) recovers the usual RC of $\mathcal{P}$.

$$\mathcal{P} = \left\{ \min_x \{ c^T[\zeta]x + d[\zeta] : \sum_{j=1}^n x_j A_j[\zeta] \leq b[\zeta] \} : \zeta \in \mathcal{Z} \right\} \Rightarrow$$

$$\min_{t, X_j(\cdot)} \left\{ t : \begin{array}{l} \sum_j c_j[\zeta] X_j(P_j\zeta) + d[\zeta] \leq t \\ \sum_j X_j(P_j\zeta) A_j[\zeta] \leq b[\zeta] \end{array} \right\} \forall \zeta \in \mathcal{Z} \Bigg\} \quad \text{(ARC)}$$

♠ **While perfectly well suited to capture the adjustability, if any, of decision variables to the true data, (ARC) has a severe built-in drawback:** *it is a "genuine" infinite-dimensional problem, and as such is, in general, severely computationally intractable.* **It is unclear even how to represent candidate decision rules – which are functions of many variables! – in a computer.** Seemingly the *only* techniques allowing to handle ARC are offered by Dynamic Programming, and thus suffer from the "curse of dimensionality."

♠ **Remedy:** to restrict ourselves with *parametric* decision rules, specifically, with *affine* ones:

$$X_j(P_j\zeta) \equiv \xi_j + \eta_j^T P_j \zeta.$$

♡ **Restricted to affine decision rules, the ARC becomes a** *finite-dimensional* **semi-infinite problem**

$$\min_{t, \xi_j, \eta_j} \left\{ t : \begin{array}{l} \sum_j c_j[\zeta][\xi_j + \eta_j^T P_j \zeta] + d[\zeta] \leq t \\ \sum_j [\xi_j + \eta_j^T P_j \zeta] A_j[\zeta] \leq b[\zeta] \end{array} \right\} \forall \zeta \in \mathcal{Z} \Bigg\}$$

called the *Affinely Adjustable RC* **of the uncertain problem $\mathcal{P}$.**

$$\min_{t,\xi_j,\eta_j}\left\{t : \begin{array}{l}\sum_j c_j[\zeta][\xi_j+\eta_j^T P_j\zeta]+d[\zeta]\leq t\\ \sum_j[\xi_j+\eta_j^T P_j\zeta]A_j[\zeta]\leq b[\zeta]\end{array}\right\}\forall\zeta\in\mathcal{Z}\Bigg\} \qquad \text{(AARC)}$$

Definition: We say that $\mathcal{P}$ is *with fixed recourse*, if the coefficients $c_j[\zeta]$, $A_j[\zeta]$ of every adjustable (i.e., with $P_j \neq 0$) variable x_j are in fact certain.

Observation: Under fixed recourse, (AARC) is of the same structure as the RC of $\mathcal{P}$, specifically, is a problem with linear objective and *bi-affine* in $\zeta \in \mathcal{Z}$ and in (x,t) constraints.

We have arrived at the following

Theorem 1: *The AARC of an uncertain LO problem with fixed recourse is computationally tractable, provided the perturbation set $\mathcal{Z}$ is so.*

Theorem 2: *The AARC of an uncertain LO problem with non-fixed recourse and with $\cap$-ellipsoidal perturbation set $\mathcal{Z} = \mathcal{Z}_\rho = \{\zeta : \zeta^T Q_j \zeta \leq \rho^2, 1 \leq j \leq J\}$ $[Q_j \succeq 0, \sum_j Q_j \succ 0]$ admits a safe tractable approximation tight within the factor $O(1)\sqrt{\ln(J+1)}$. When $J=1$, the approximation is exact.*

Note: the conclusion of Theorem 2 remains valid when the decision rules for "fixed recourse" x_j (those which enter the problem solely with certain coefficients) are allowed to be quadratic in $P_j\zeta$, and the decision rules for the "non-fixed-recourse" variables are allowed to be affine in $P_j\zeta$.

♣ **How it works: Inventory Problem.** A single-product inventory comprised of a warehouse and I factories evolves over time horizon $1, ..., N$. The inventory is affected by uncertain demand $d = [d_1; ...; d_N]$ varying in a given domain D. No backlogged demand is allowed. Let

- x_t be the inventory level at the beginning of period t,
- w_{it} be the amount of product, ordered from and delivered by factory # i in period t.

Given the initial state x_1 of the inventory, bounds on the inventory levels and on the instant and cumulative replenishment orders, we want to minimize the worst, over the demand trajectories from D, overall ordering cost:

$$
\begin{array}{lll}
\min\limits_{C,x,w} & C & \text{[total ordering cost]} \\
\textbf{s.t.} & C \geq \sum_{t=1}^{N} \sum_{i=1}^{I} c_{it} w_{it} & \text{[cost description]} \\
& x_{t+1} = x_t + \sum_{i=1}^{I} w_{it} - d_t & \text{[state equations]} \\
& \underline{X} \leq x_t \leq \overline{X} & \text{[bounds on states]} \\
& 0 \leq w_{it} \leq W_{it} & \text{bounds on orders]} \\
& 0 \leq \sum_{\tau=1}^{N} w_{it} \leq \widehat{W}_i & \text{[bounds on accumulated orders]}
\end{array}
$$

♠ **Applying the AARC approach, we**

- **allow to our actual "wait and see decisions"** w_{it} **to depend** *affinely* **on** $P_t d \equiv [d_1; ...; d_{t-1}]$**:** $w_{it} = p_{it} + \sum_{\tau<t} q_{it}^{\tau} d_{\tau}$
- **allow to the "analysis variables"** x_2**,...,**x_{N+1} **to be arbitrary** *affine* **functions of** d**:** $x_t = \xi_t + \eta_t^T d$**;**
- **treat** C **as the only non-adjustable variable.**

We plug the decision rules into the model and require the constraints to be satisfied for all $d \in D$, thus ending up with the semi-infinite LO problem

$$
\begin{array}{ll}
\min\limits_{\substack{C,\xi_t\eta_t,\\ p_{it},q_{it}^\tau}} & C \\
\textbf{s.t.} & C \geq \sum_{t=1}^{N} \sum_{i=1}^{I} c_{it} \left[p_{it} + \sum_{\tau<t} q_{it}^\tau d_\tau\right] \\
& \xi_{t+1} + \eta_{t+1}^T d = \xi_t + \eta_t^T d + \sum_{i=1}^{I}[p_{it} + \sum_{\tau<t} q_{it}^\tau d_\tau] - d_t \\
& \underline{X} \leq \xi_t + \eta_t^T d \leq \overline{X} \\
& 0 \leq \sum_{i=1}^{I}[p_{it} + \sum_{\tau<t} q_{it}^\tau d_\tau] \leq W_{it} \\
& 0 \leq \sum_{t=1}^{N}[p_{it} + \sum_{\tau<t} q_{it}^\tau d_\tau] \leq \widehat{W}_i
\end{array}
$$

where the constraints should be satisfied for all $d \in D$.

• **Note:** We are in the situation of fixed recourse⇒ the AARC is tractable, provided that D is so.

• **Note:** We could handle easily a much more complicated problem (many products, additional components in the cost function, probabilistic constraints, etc., etc.) *All what matters is that the underlying problem is uncertain LO with uncertainty affecting the right hand sides of constraints only (and thus not affecting coefficients of adjustable variables).*

♣ Numerical illustration: $N = 48$, $I = 3$, D is a box:

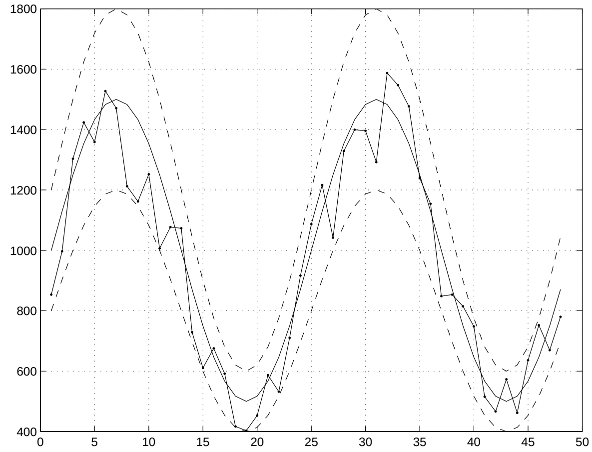

Demand box $0.8d^* \leq d \leq 1.2d^*$
and a sample demand trajectory

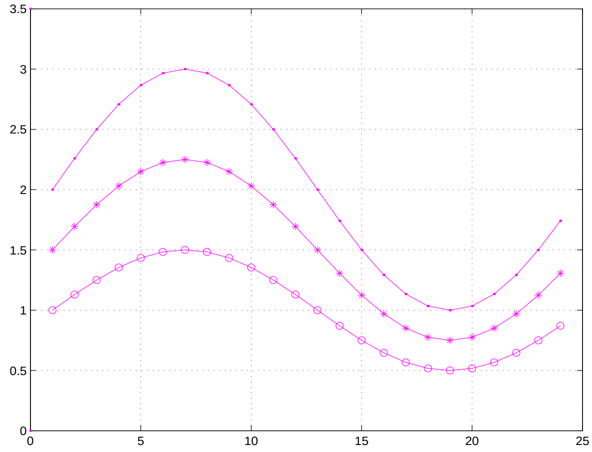

Ordering prices c_{it}
(periodic with period 24)

♡

We ran several hundreds of simulations and compared the average replenishment cost incurred by optimal *affine* **decision rules with** *utopian cost* **we could pay when** *knowing in advance* **the demand trajectory and** *optimizing accordingly our policy*. **This comparison (biased** *against* **affine decision rules) shows a surprisingly high quality of these rules:**

	AARC-based cost		Utopian cost	
Uncertainty	Mean	Std	Mean	Std
2.5%	33974	190	33878(-0.3%)	194
5%	34063	432	33864(-0.6%)	454
10%	34471	595	34009(-1.6%)	621
20%	35121	1458	33958(-3.4%)	1541

Note: With our setup, the RC is infeasible already at the 5% uncertainty level.

♡ **There is significant evidence that affine decision rules indeed work well in multi-stage Inventory problems.**

♣ **Recent (difficult!) result of Bertsimas, Iancu and Parrilo states that** *For a single-product Inventory with the cost function* $\sum_{t=1}^{N}[c_t w_t + h_t(x_t)]$ *(*$c_t > 0$*,* $h_t(\cdot)$ *are convex), state equations*

$$x_t = x_{t-1} + w_t - d_t$$

bounds $\underline{W}_t \leq w_t \leq \overline{W}_t$ *on replenishment orders and a box D in the role of the set of uncertain demand trajectories, the ARC is equivalent to the AARC, that is, the optimal, in terms of the worst-case management cost, decision rules can be chosen to be affine in the respective parts of the demand trajectory.*

While this theoretical result cannot be extended to a more general settings of the Inventory problem (say, it fails to be true when bounds on accumulated orders are added, and/or the box uncertainty set is replaced with a more general one), AARC, practically speaking, seems to be a good technique for *worst-case oriented* Inventory management.

Note: *When passing from minimizing the worst-case management cost to minimizing the average one, affine decision rules become by far non-optimal already for pretty simple Inventory models.*

♣ **Assume we want to solve a "restricted ARC" on an Uncertain LO problem with fixed recourse, that is, its ARC where the decision rules are restricted to reside in a given class (e.g., to be affine, or quadratic, or polynomial,... in their arguments).**

• When restricted to affine decision rules, the ARC becomes easy. *Is the affinity an actual restriction here?*

♠ **Assume that instead of affine decision rules** $x_j = \xi_j + \eta_j^T P_j \zeta$ **we indent to use rules from a general parametric family:**

$$x_j = \sum_{\ell} \eta_{j\ell} f_{j\ell}(P_j \zeta) \qquad (*)$$

• $f_{j\ell}(\cdot)$: **"basic functions"** • $\eta_{j\ell}$: **free parameters.**

♡Augmenting the perturbation ζ by the entries $\zeta_{j\ell} = f_{j\ell}(P_j \zeta)$, that is, extending ζ to the new perturbation vector

$$\widehat{\zeta}[\zeta] = [\zeta; \{f_{j\ell}(P_j \zeta)\}_{j,\ell}],$$

decision rules $(*)$ become affine in the new perturbation. Thus, for all practical purposes *all parametric decision rules can be thought of as affine ones.*

!!! Bottleneck: The AARC of an Uncertain LO with fixed recourse is easy due to both affinity of the decision rules *and the assumption* (which we always made) *that the perturbation set $\mathcal{Z}$ is tractable.* Passing from ζ to its nonlinear transform $\widehat{\zeta}[\zeta]$, the perturbation set becomes $\widehat{\mathcal{Z}} = \mathrm{Conv}\{\widehat{\zeta}[\zeta] : \zeta \in \mathcal{Z}\}$ and can easily lose tractability.

Good case: Quadratic decision rules, $\mathcal{Z}$ is an ellipsoid. Here $\widehat{\zeta}[\zeta] = \left[\begin{array}{c|c} & \zeta^T \\ \hline \zeta & \zeta\zeta^T \end{array}\right]$. Assuming $\mathcal{Z} = \{\zeta : \|\zeta\|_2 \leq 1\}$, $\widehat{\mathcal{Z}} := \mathrm{Conv}\{\widehat{\zeta}[\zeta] : \zeta \in \mathcal{Z}\}$ is computationally tractable:

$$\widehat{\mathcal{Z}} = \left\{ \left[\begin{array}{c|c} & \zeta^T \\ \hline \zeta & Z \end{array}\right] : \left[\begin{array}{c|c} 1 & \zeta^T \\ \hline \zeta & Z \end{array}\right] \succeq 0,\ \mathrm{Tr}(Z) \leq 1 \right\}$$

Semi-Good case: Quadratic decision rules, $\cap$-ellipsoidal uncertainty. Here

$$\mathcal{Z} = \{\zeta : \zeta^T Q_j \zeta \leq 1,\ 1 \leq j \leq J\}$$
$$[J > 1, Q_j \succeq 0, \textstyle\sum_j Q_j \succ 0]$$

and $\widehat{\zeta}[\cdot]$ is as above. Now the set

$$\widehat{\mathcal{Z}} = \mathrm{Conv}\{\widehat{\zeta}[\zeta] : \zeta \in \mathcal{Z}\}$$

can be intractable, but it admits an outer tractable approximation:

$$\widehat{\mathcal{Z}} \subset \widetilde{\mathcal{Z}} = \left\{ \left[\begin{array}{c|c} & \zeta^T \\ \hline \zeta & Z \end{array}\right] : \left[\begin{array}{c|c} 1 & \zeta^T \\ \hline \zeta & Z \end{array}\right] \succeq 0,\ \mathrm{Tr}(ZQ_j) \leq 1, 1 \leq j \leq J \right\}$$

which is tight within factor $\vartheta = O(1)\ln(J)$:

$$\vartheta^{-1}\widetilde{\mathcal{Z}} \subset \widehat{\mathcal{Z}} \subset \widetilde{\mathcal{Z}}.$$

Generic Application: Synthesis of Linear Controllers

♣ **Consider time-varying discrete time linear dynamical system**

$$\begin{array}{rcll} x_0 &=& z & \text{[initial state]} \\ x_{t+1} &=& A_t x_t + B_t u_t + R_t d_t & \left[\begin{array}{l}\text{state equations}\\ \bullet\ x_t\text{: state}\ \ \bullet\ u_t\text{: control}\\ \bullet\ d_t\text{: external disturbance}\end{array}\right] \\ y_t &=& C_t x_t + D_t d_t & \text{[observed output]} \end{array}$$

"closed" by *affine output-based control law*

$$u_t = g_t + \sum\nolimits_{\tau=0}^{t} G_t^\tau y_\tau. \qquad (*)$$

♠ **Given finite time horizon** $0 \leq t \leq N$**, we want to specify a control law** $(*)$ **which ensures that** *the state-control trajectory* $w = [x_0; ... x_{N+1}; u_0; ...; u_N]$ **satisfies given design specifications**

$$Aw \leq b \qquad (!)$$

robustly **w.r.t. the "perturbation"** $\zeta = [z; d_0; ...; d_N]$ **running through a given set** $\mathcal{Z}$**.**

Good news: by linearity of the system and the control law, the trajectory is affine in ζ.

$\Rightarrow$ The *Analysis problem: check whether a given control law* $(*)$ *robustly meets the design specifications* reduces to verifying whether a system of affine constraints on ζ is satisfied by all $\zeta \in \mathcal{Z}$. This is easy, provided $\mathcal{Z}$ is tractable.

$$\begin{array}{rcl} x_0 &=& z \\ x_{t+1} &=& A_t x_t + B_t u_t + R_t d_t \\ y_t &=& C_t x_t + D_t d_t \end{array} \qquad (S)$$

$$u_t = g_t + \sum\nolimits_{\tau=0}^{t} G_t^\tau y_\tau \qquad (*)$$

Bad news: the trajectory is highly nonlinear in the parameters $\gamma = \{g_t, G_t^\tau\}$ of the control law $(*)$

$\Rightarrow$ The *Synthesis problem: find control law* $(*)$, *if it exists, which robustly meets the design specifications* seems to be intractable.

Remedy: pass to affine purified-output-based control laws.

♠ **Consider, along with system** (S) **"closed" by some control law, its** *model*

$$\begin{array}{rcl} \widehat{x}_0 &=& 0 \\ \widehat{x}_{t+1} &=& A_t \widehat{x}_t + B_t u_t \\ \widehat{y}_t &=& C_t \widehat{x}_t \end{array} \qquad (M)$$

which we "feed" by the same controls u_t **as** (S). **We can run the model in an on-line fashion, and thus at time** t, **before the decision on** u_t **should be made, we have in our disposal** *purified output* $v_t = y_t - \widehat{y}_t$

Observation: *purified outputs are independent on the control law known in advance affine functions of* ζ.

Indeed, setting $\Delta_t = x_t - \widehat{x}_t$, we clearly have

$$v_t = C_t \Delta_t + D_t d_t, \ \Delta_0 = z, \ \Delta_{t+1} = A_t \Delta_t + R_t d_t.$$

System:	Model:
$$\begin{array}{rcl} x_0 &=& z \\ x_{t+1} &=& A_t x_t + B_t u_t + R_t d_t \\ y_t &=& C_t x_t + D_t d_t \end{array} \quad (S)$$	$$\begin{array}{rcl} \widehat{x}_0 &=& 0 \\ \widehat{x}_{t+1} &=& A_t \widehat{x}_t + B_t u_t \\ \widehat{y}_t &=& C_t \widehat{x}_t \end{array} \quad (M)$$
Purified outputs: $v_t = y_t - \widehat{y}_t$	
$u_t = \begin{cases} g_t + \sum_{\tau=0}^t G_t^\tau y_\tau & \text{[output-based affine law]} \quad (*) \\ h_t + \sum_{\tau=0}^t H_t^\tau v_t & \text{[purified-output-based affine law]} \quad (+) \end{cases}$	

Facts:

♡ *Purified-output-based affine laws are equivalent to the output-based affine laws: every mapping $\zeta \to w$ which can be obtained when "closing" (S) by a law $(*)$, can be obtained by closing (S) by a law $(+)$, and vice versa.*

♡ *When (S) is closed by a purified-output-based affine control law $(+)$, the trajectory $w = W[\zeta, \eta]$ becomes bi-affine in ζ and in the parameters $\eta = \{h_t, H_t^\tau\}$ of the control law.*

♡ **As a result,** *Sticking to purified-output-based control laws make the Synthesis problem*

> *Given design specifications $Aw \le b$ on the state-control trajectory, find a control law, if one exists, which meets these specifications robustly w.r.t. $\zeta = [z; d_0; ...; d_N] \in \mathcal{Z}$*

becomes an efficiently solvable system of semi-infinite affinely perturbed linear constraints on η.

How it Works: Control of 3-Level Serial Inventory

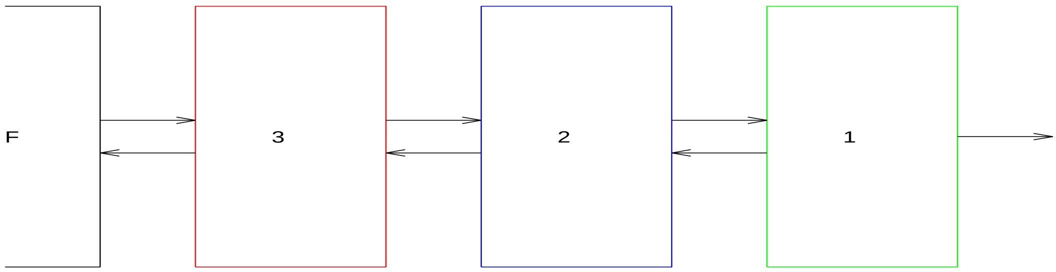

- Level 1 supplies extrnal demand
- Level 2 supplies Level 1
- Level 3 supplies Level 2 and is supplied from Factory
- There is 2-period delay in executing replenishment orders

The Inventory can be modeled as the 9-state LDS

$$
\begin{aligned}
x_1(t+1) &= x_1(t) + x_2(t) && -d_t \\
x_2(t+1) &= x_3(t) \\
x_3(t+1) &= && u_1(t) \\
x_4(t+1) &= x_4(t) + x_5(t) && -u_1(t) \\
x_5(t+1) &= x_6(t) \\
x_6(t+1) &= && u_2(t) \\
x_7(t+1) &= x_7(t) + x_8(t) && -u_2(t) \\
x_8(t+1) &= x_9(t) \\
x_9(t+1) &= && u_3(t) \\
\hline
y(t) &= x(t)
\end{aligned}
$$

- x_1, x_2, x_3 — **inventory levels**
- u_i — **replenishment orders** • d_t — **demands**

♣ **It is well known that serial inventories with delays suffer from** *bullwhip effect*: **variations in external demand result in** *much larger* **variations in the inventory levels, especially in the one closest to the factory, thus badly affecting the production.**
This is what happens with "naive" feedback:

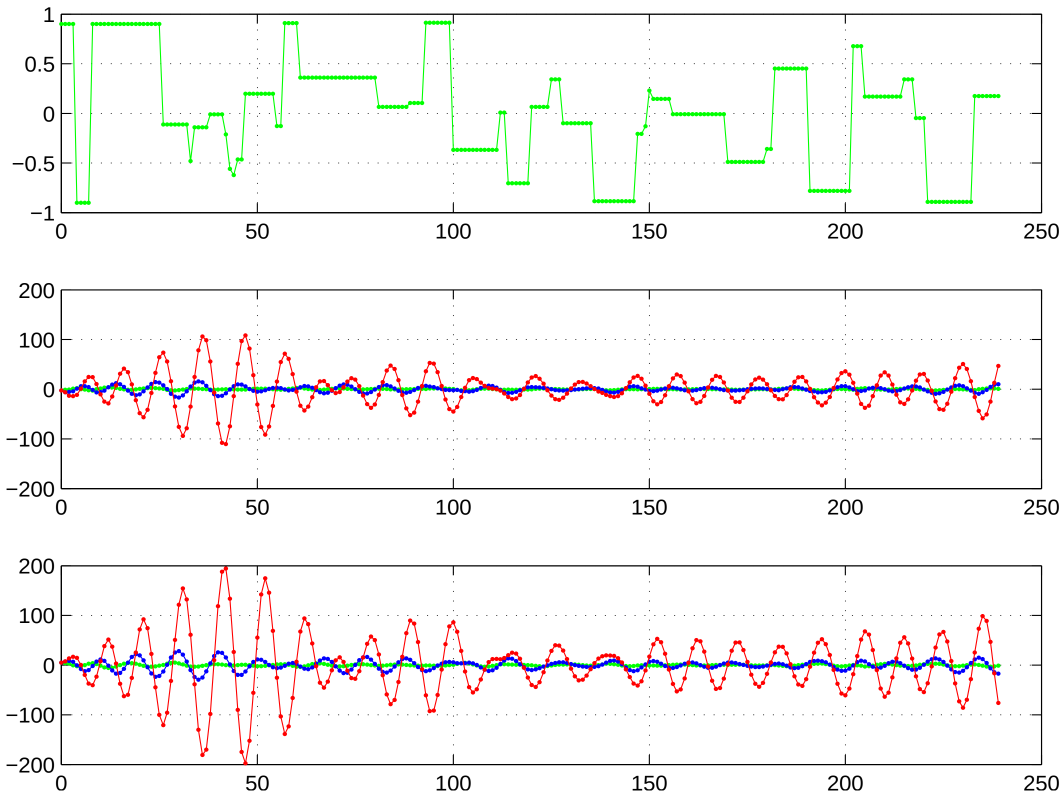

Bullwhip effect

Top: time-dependent demand varying in $[-1, 1]$
Middle: replenishment orders $u_1(t)$, $u_2(t)$, $u_3(t)$
Bottom: inventory levels (green: #1, blue: #2, red: #3)

Note: variations of the demand in the range $[-1, 1]$ result in huge (hundreds!) oscillations in the level #3 and in the replenishment orders.

♡ **To reduce the bullwhip effect, we can look for the best — with the largest decay rate as certified by Lyapunov Stability Certificate — linear feedback. With this control, the picture looks much better:**

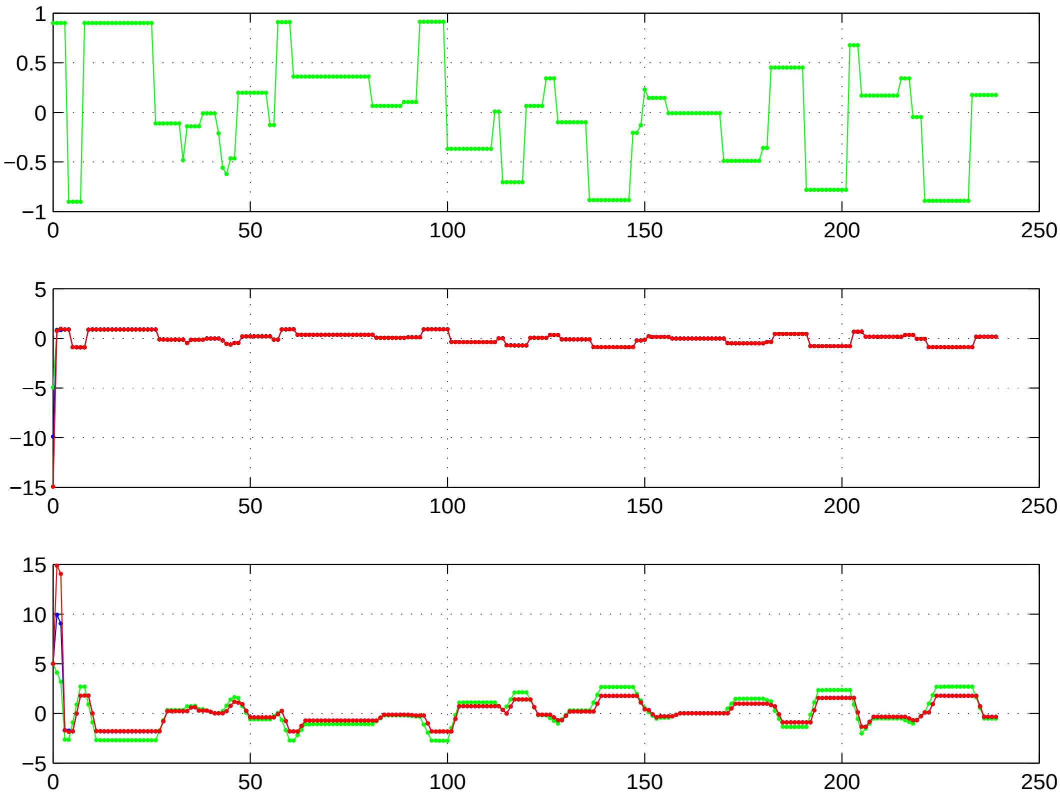

Good linear feedback

Top: time-dependent demand varying in $[-1, 1]$
Middle: replenishment orders $u_1(t)$, $u_2(t)$, $u_3(t)$
Bottom: inventory levels (green: #1, blue: #2, red: #3)

But: At the very beginning, we still have unpleasant jumps in the levels and replenishment orders.

♡ **To improve the behaviour of the process in the beginning, we can use purified-output-based affine control aimed at minimizing the initial jumps and converging to the above feedback control. This is what we get:**

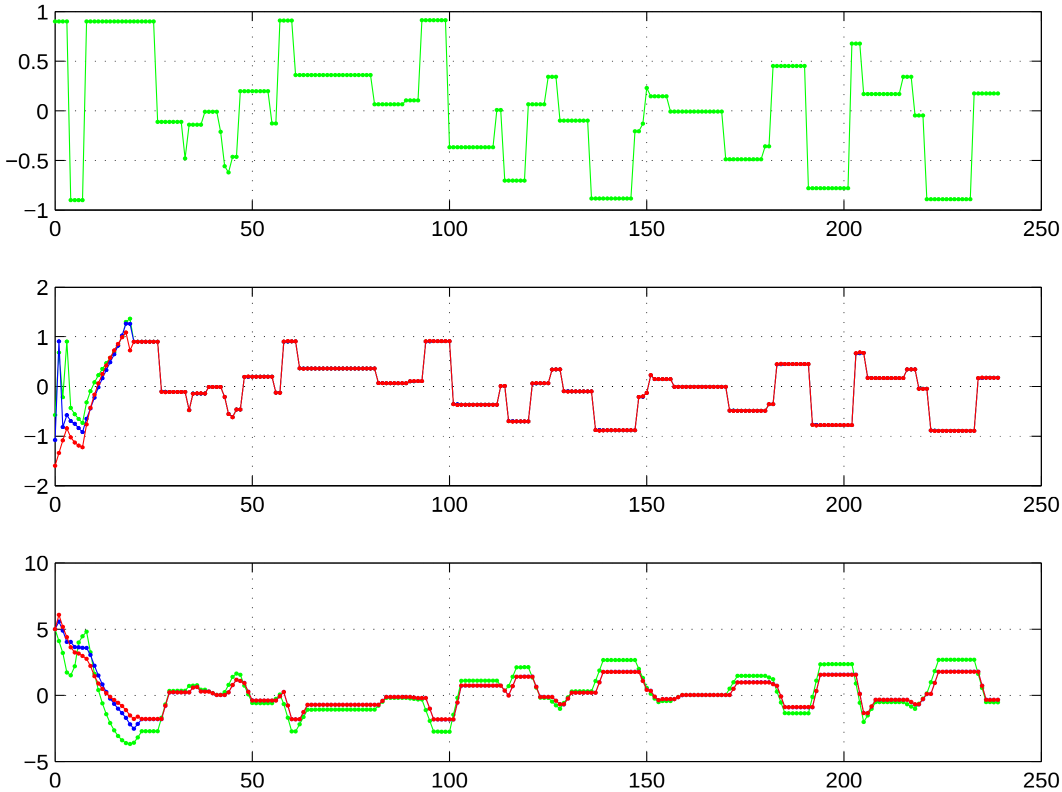

Combined p.o.b./feedback control

Top: time-dependent demand varying in $[-1, 1]$

Middle: replenishment orders $u_1(t)$, $u_2(t)$, $u_3(t)$

Bottom: inventory levels (green: #1, blue: #2, red: #3)

♡ **This is what we gain in the beginning, while loosing nothing in the long run:**

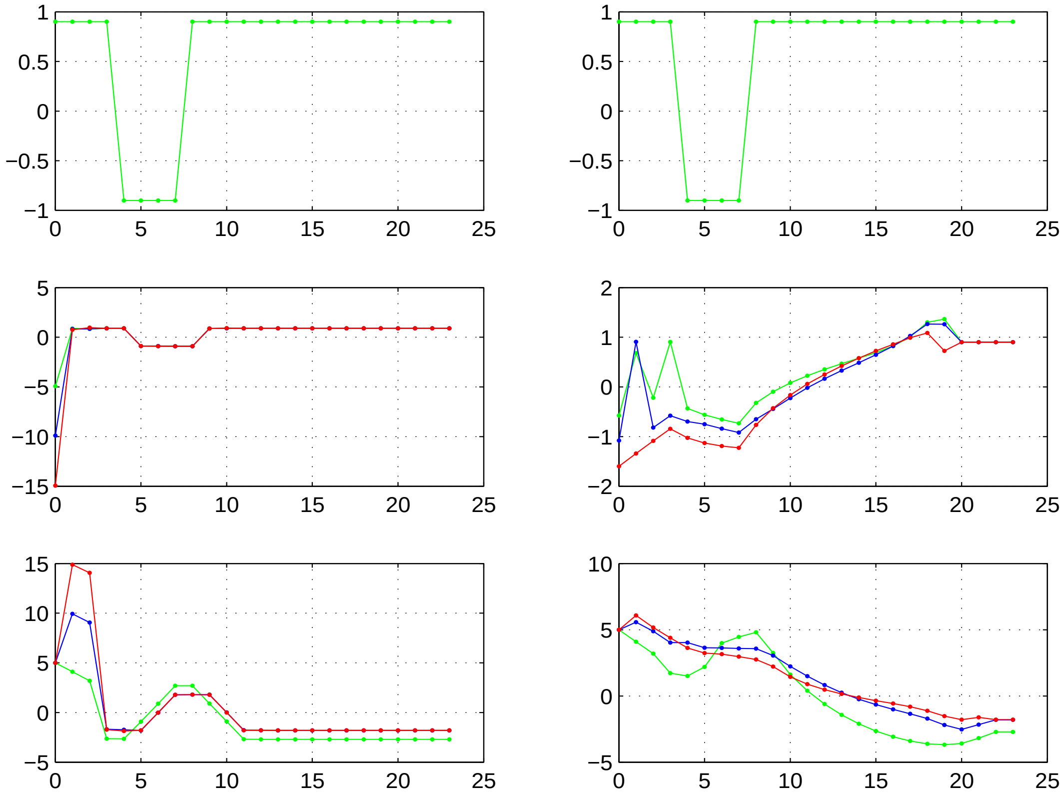

Pure feedback control (left)

vs.

combined p.o.b/feedback control (right)

Top: time-dependent demand varying in $[-1, 1]$

Middle: replenishment orders $u_1(t)$, $u_2(t)$, $u_3(t)$

Bottom: inventory levels (green: #1, blue: #2, red: #3)