# Outline

- Summary of game results
- A two-agent learning model
- Feedback loops and misinformation

# Two-agent newsvendor

- The teams:
  - » Find your team number so that you can identify yourself in the graphs that follow.
  - » "Trust policy" team assumes that central believes whatever the field asks for.

| |
|---|
| 1-David/Leo |
| 2-Luke/Hill |
| 3-Diana/Natasha |
| 4-Joe/Aria |
| 5-Mark/Kevin |
| 6-Alice/Chad |
| 7-Hannah/Julian |
| 8-Dan/Seth |
| 9-Brandon/Genna |
| 10-Yanran/Stephen |
| 11-Prakhar/Kevin |
| 12-Druv/Amy |
| 13-Jay/Charquia |
| 14-Kevin/Michael |
| 15-Adam/Chris |
| 16-Andrew/Akshay "DJ Shake" |
| 17-Danny/Lisa |
| 18-Eric/Nate |
| 19-Adriana/Jieming |
| 20-Christian/Kevin |
| 21-Philip/Michael |
| 22-Eileen/Christine |
| 23- |
| 24- |
| 25- |
| 26- |
| 27- |
| 28- |
| 29- |
| 30- |
| 31- |
| 32- |
| 33- |
| 34-J. Schoppe/C. Schoppe |
| 35-"Trust" Policy |

# Two-agent newsvendor

- Round 1: exogenous data

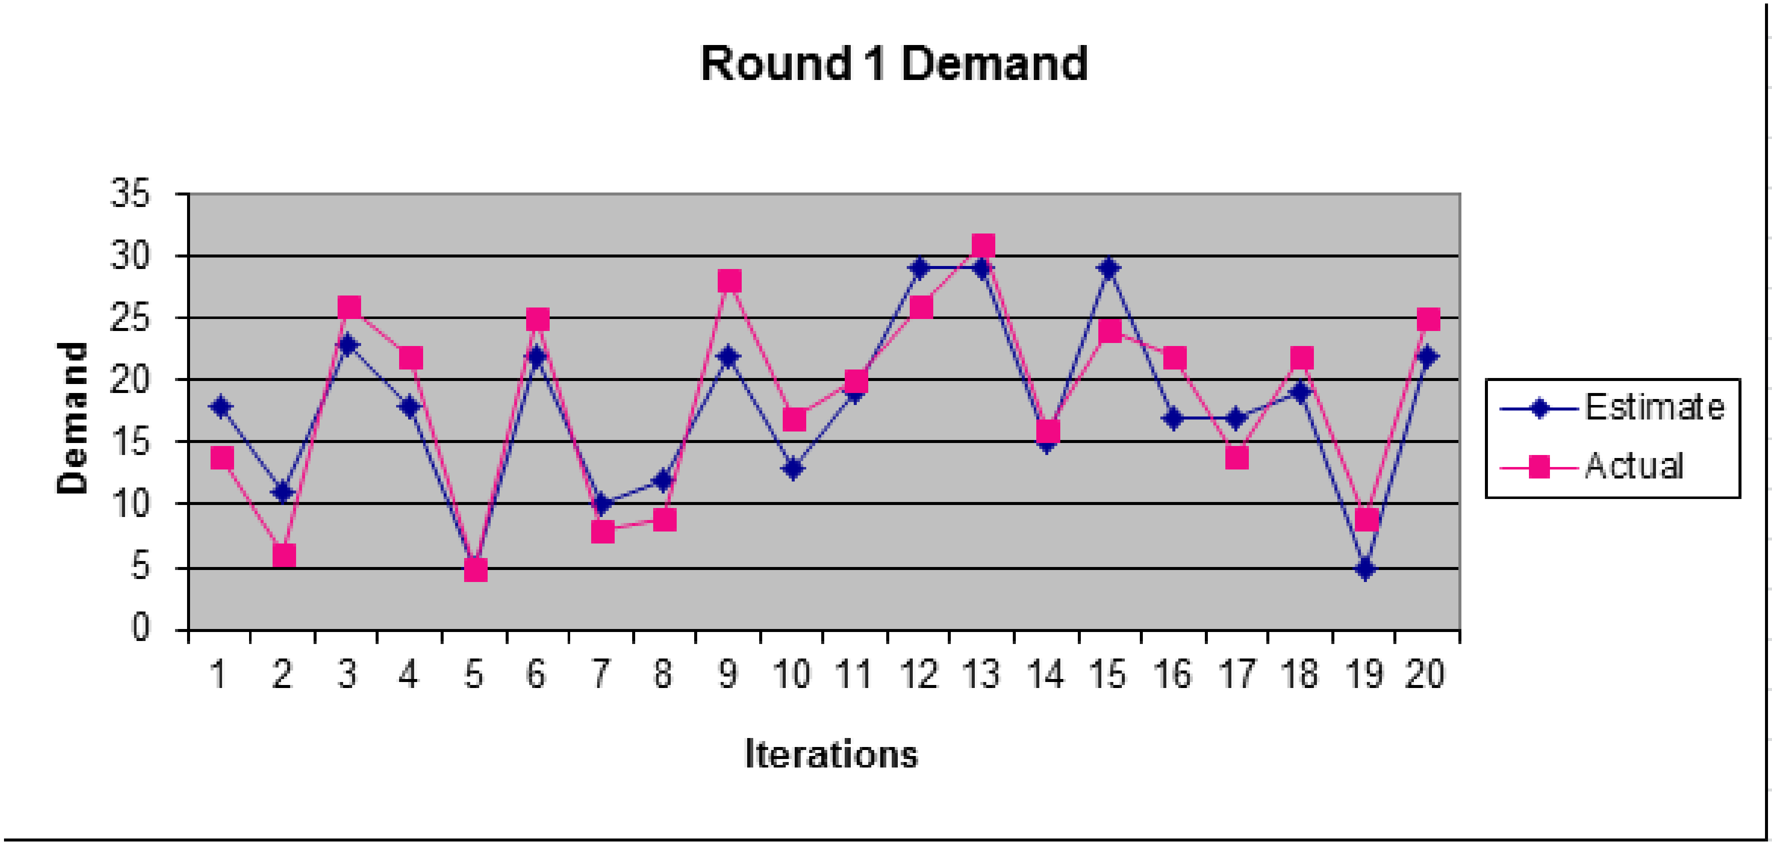

# Two-agent newsvendor

- Round 2: exogenous data

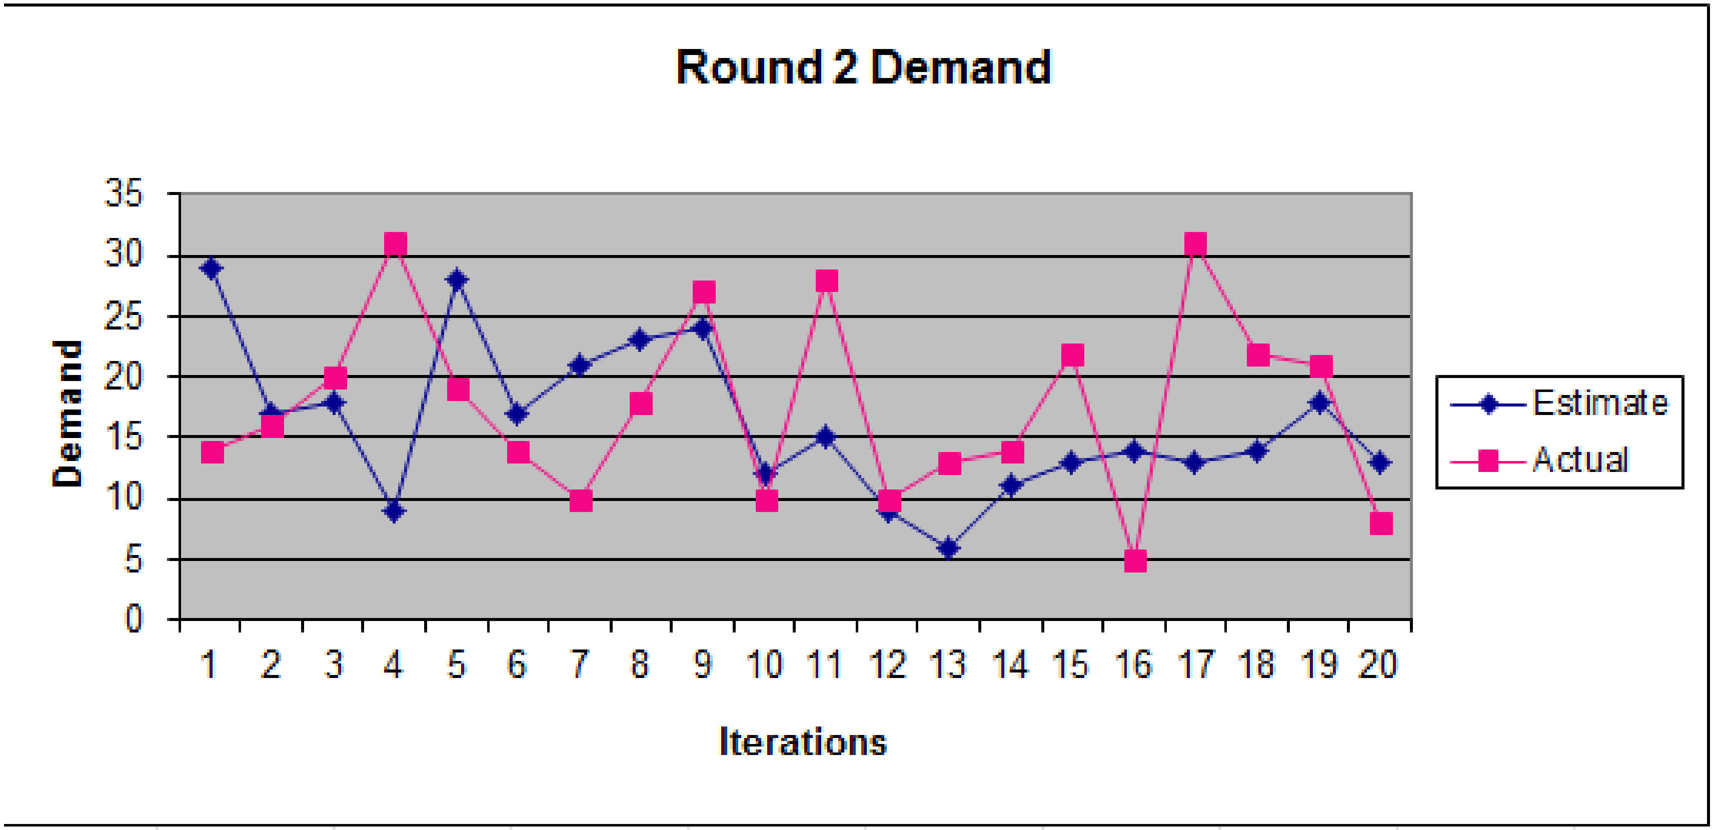


 

# Two-agent newsvendor

- The "trust model"
  - » Field orders exactly the initial estimate.
  - » Central command provides exactly what the field asks for.
- Who wins – class or trust policy?
  - » Who was better at beating the trust policy: field or central?
  - » Who was better: the teams (field + central) or trust policy?
  - » Switch to other presentation for answer.
  - » Discuss: What is the implication for management?

| **Round 1** | | | Costs | Central | Field |
|---|---|---|---|---|---|
| | | | Overage | 5 | 2 |
| | | | Underage | 5 | 10 |

| | Field | | Central | Actual | Costs | |
|---|---|---|---|---|---|---|
| | Initial estimate | Request from central | Give to Field | exogenous demand | Central c^o = 5 c^u=5 | Field c^o=2 c^u=10 |
| 1 | 24 | 24 | 24 | 27 | 15 | 30 |
| 2 | 36 | 36 | 36 | 41 | 25 | 50 |
| 3 | 39 | 39 | 39 | 39 | 0 | 0 |
| 4 | 19 | 19 | 19 | 17 | 10 | 4 |
| 5 | 23 | 23 | 23 | 27 | 20 | 40 |
| 6 | 21 | 21 | 21 | 18 | 15 | 6 |
| 7 | 32 | 32 | 32 | 32 | 0 | 0 |
| 8 | 16 | 16 | 16 | 12 | 20 | 8 |
| 9 | 11 | 11 | 11 | 16 | 25 | 50 |
| 10 | 21 | 21 | 21 | 13 | 40 | 16 |
| 11 | 39 | 39 | 39 | 39 | 0 | 0 |
| 12 | 33 | 33 | 33 | 33 | 0 | 0 |
| 13 | 17 | 17 | 17 | 14 | 15 | 6 |
| 14 | 37 | 37 | 37 | 44 | 35 | 70 |
| 15 | 34 | 34 | 34 | 36 | 10 | 20 |
| 16 | 34 | 34 | 34 | 38 | 20 | 40 |
| 17 | 17 | 17 | 17 | 11 | 30 | 12 |
| 18 | 34 | 34 | 34 | 28 | 30 | 12 |
| 19 | 6 | 6 | 6 | 12 | 30 | 60 |
| 20 | 35 | 35 | 35 | 40 | 25 | 50 |
| | 26.4 | | | 26.85 | 365 | 474 |

# Two-agent newsvendor

- Round 1: Field and Central

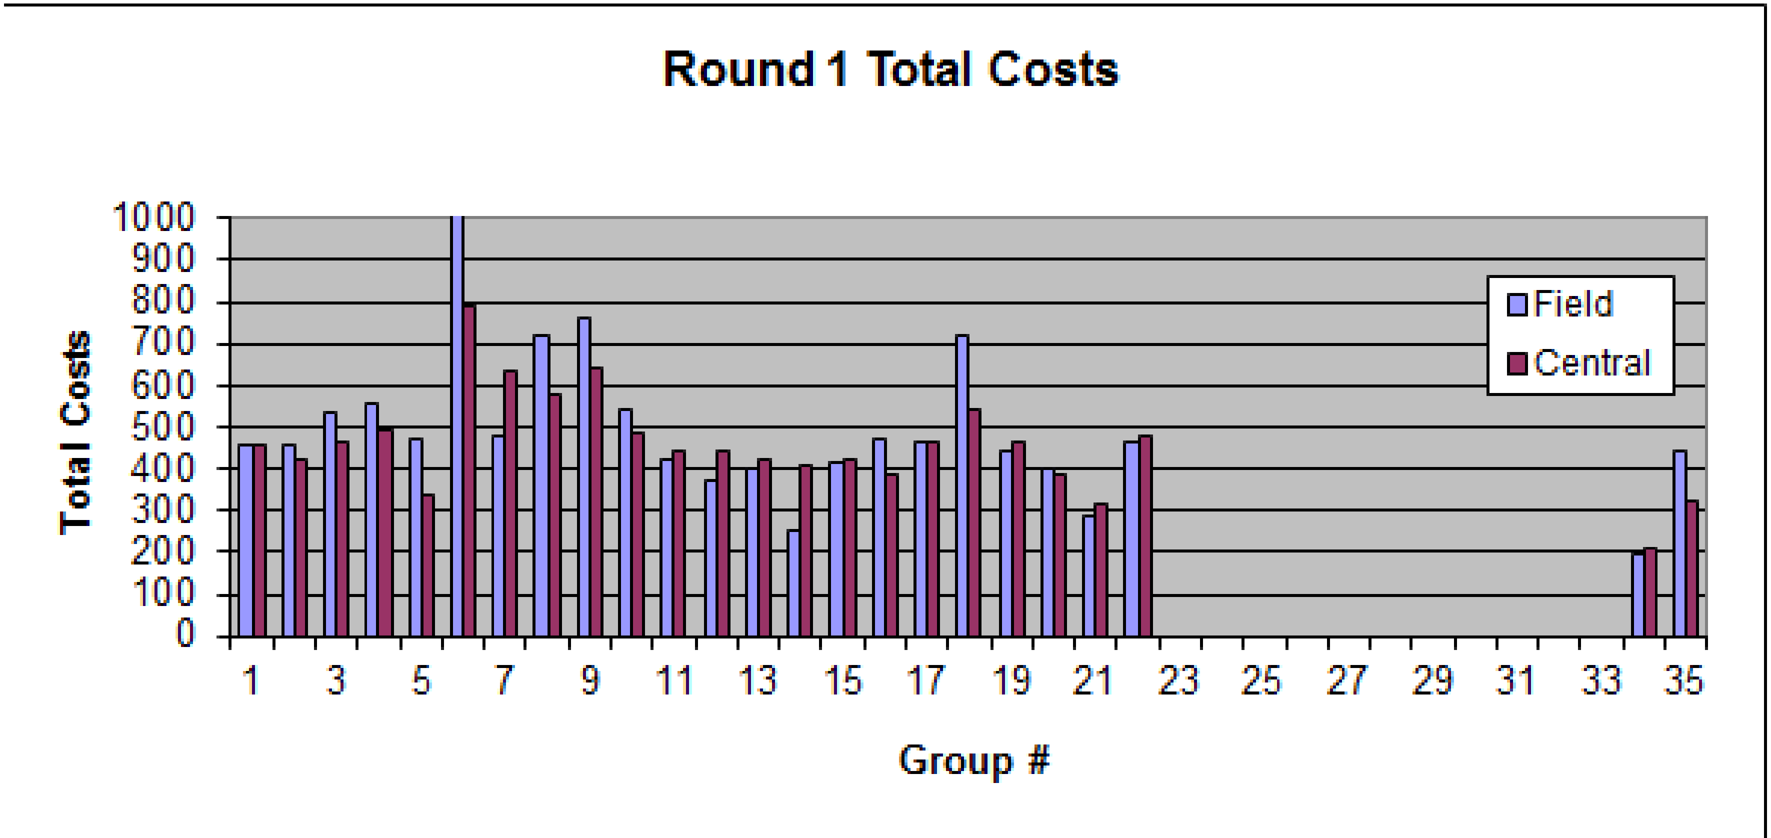

# Two-agent newsvendor

- Round 2: Field and Central

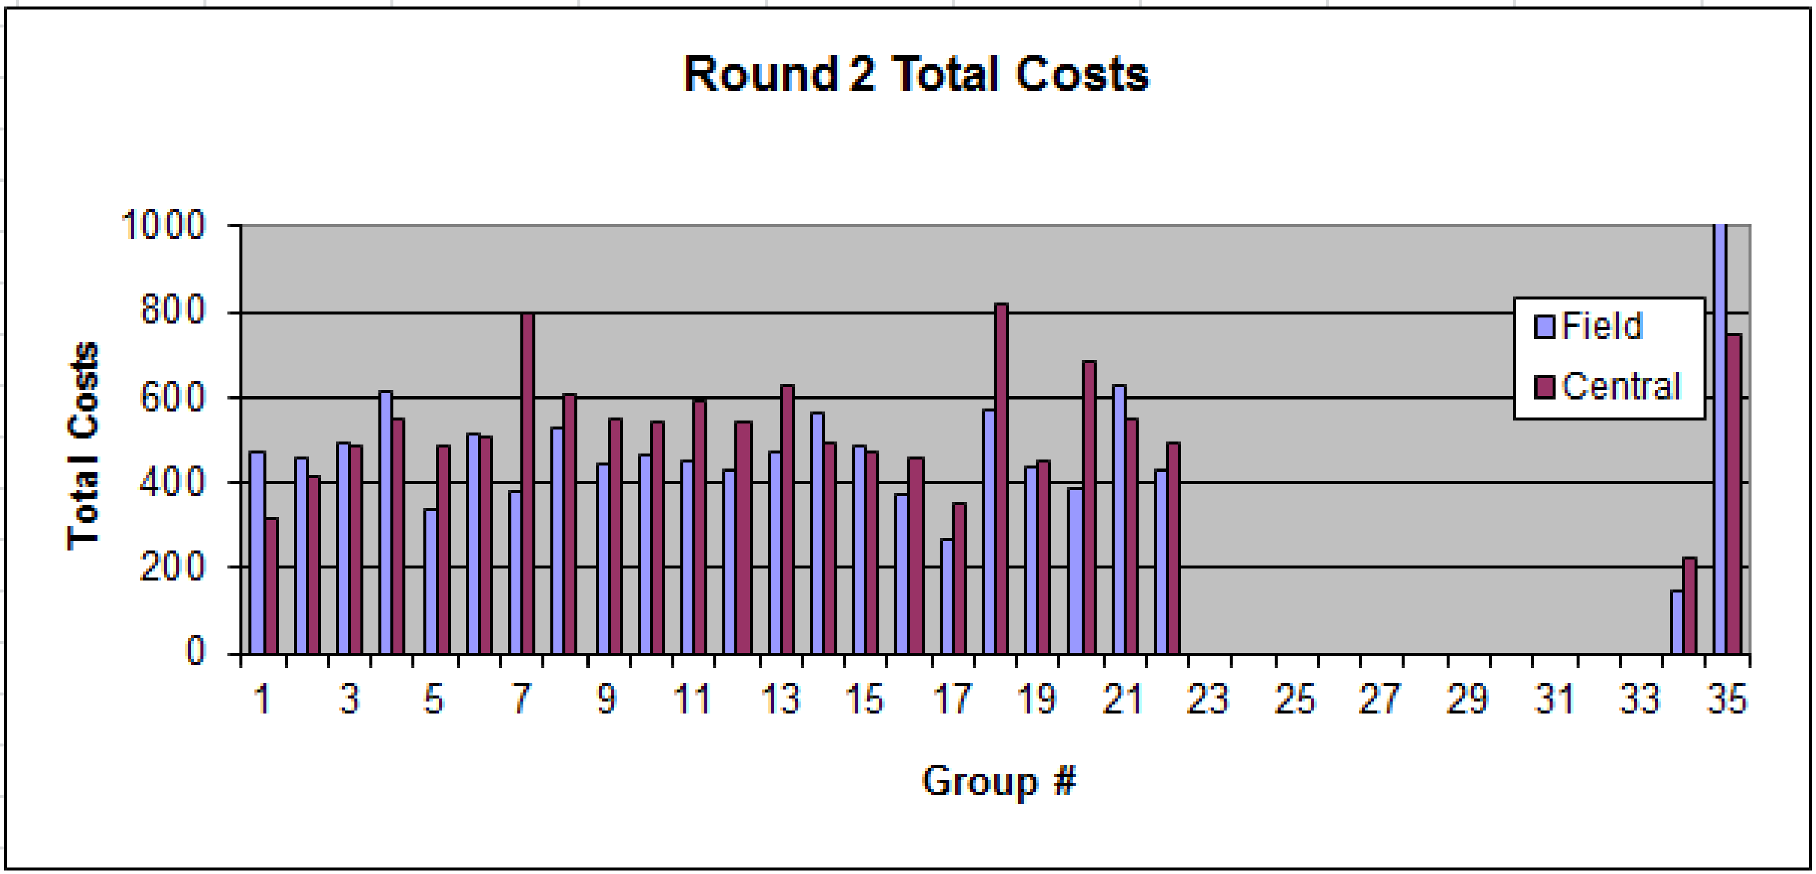

# Two-agent newsvendor

- Group totals: Field and Central

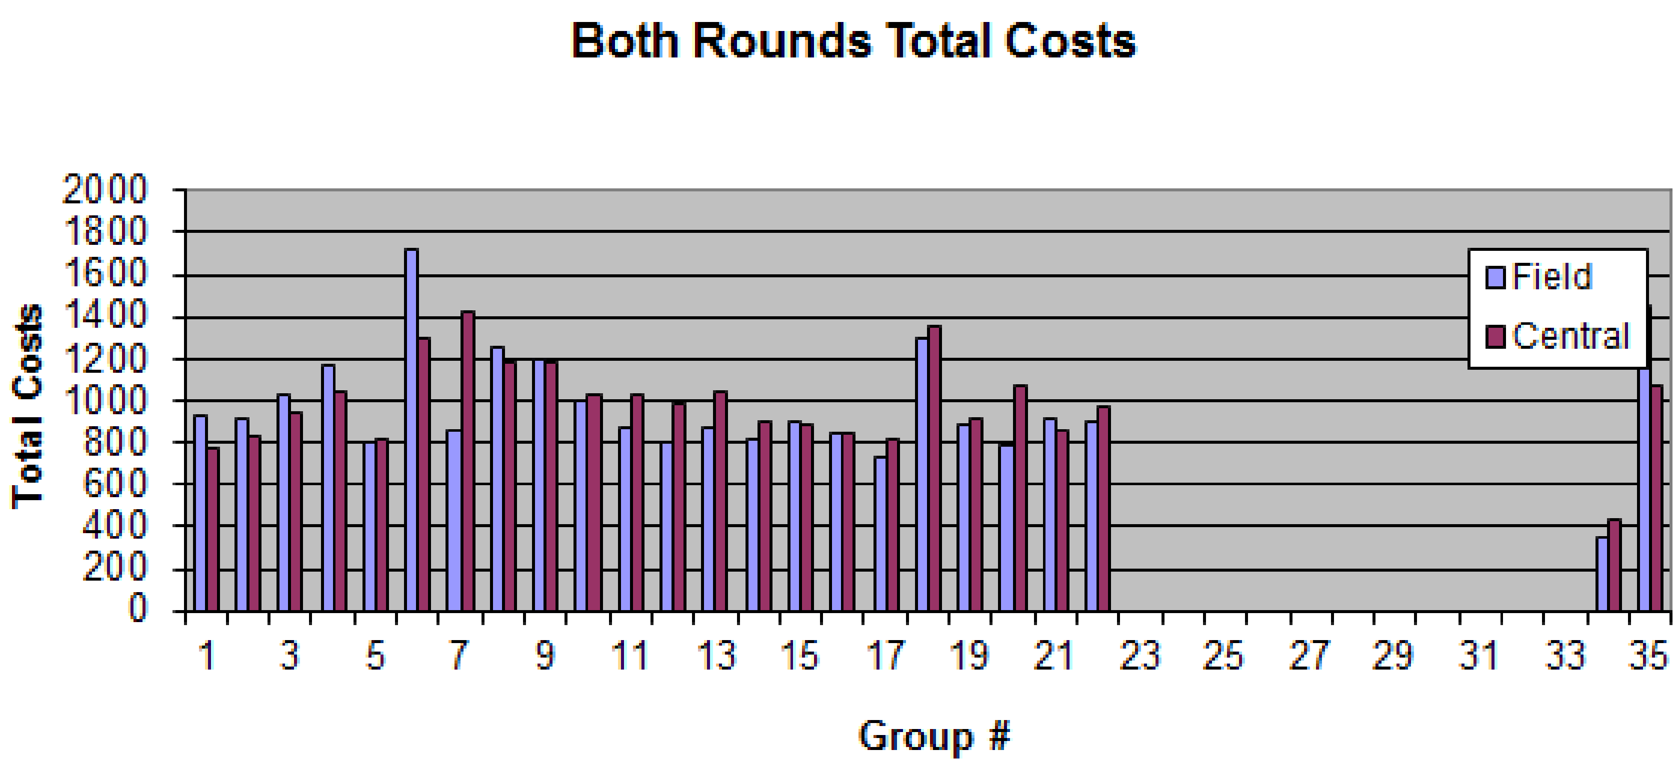

# Two-agent newsvendor

- Group totals: Rounds 1 and 2

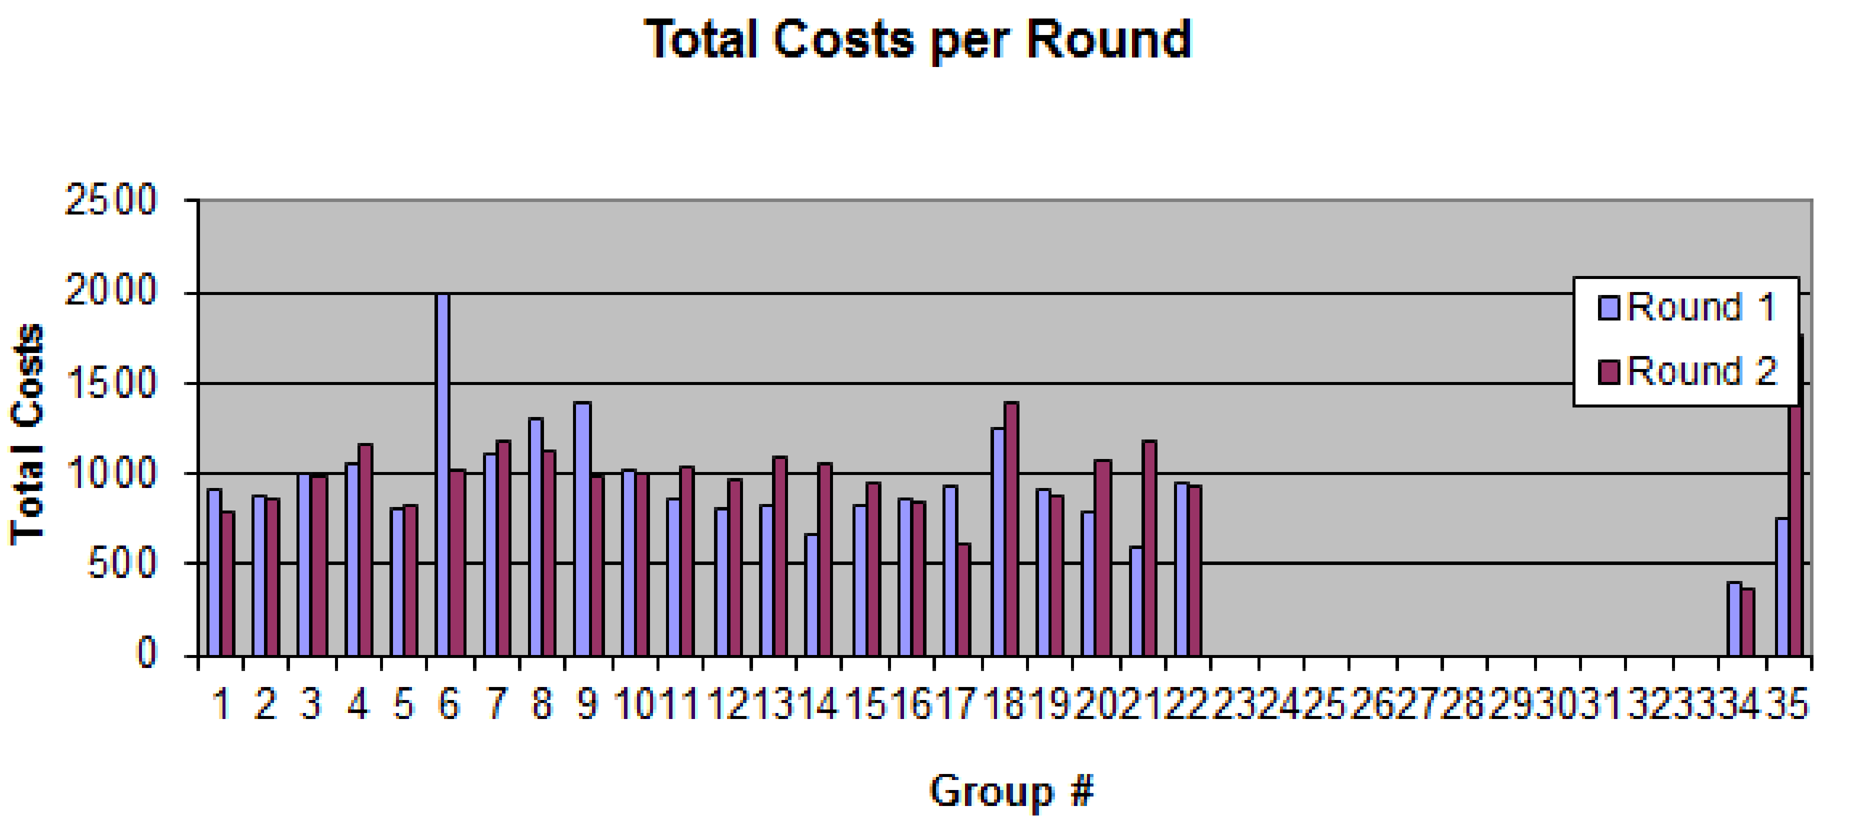

# Two-agent newsvendor

- Individual totals: Players 1 and 2

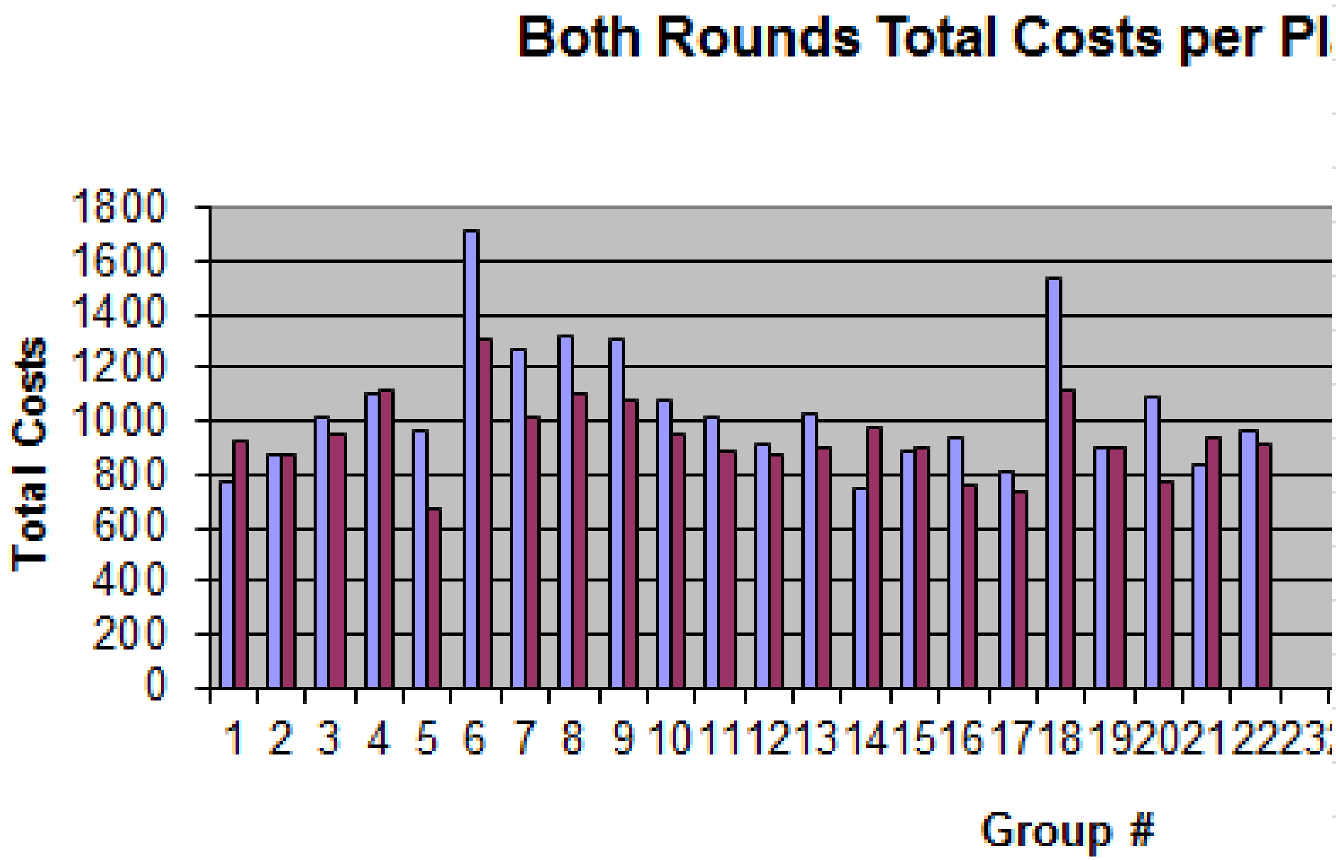


1-David/Leo
2-Luke/Hill
3-Diana/Natasha
4-Joe/Aria
5-Mark/Kevin
6-Alice/Chad
7-Hannah/Julian
8-Dan/Seth
9-Brandon/Genna
10-Yanran/Stephen
11-Prakhar/Kevin
12-Druv/Amy
13-Jay/Charquia
14-Kevin/Michael
15-Adam/Chris
16-Andrew/Akshay "DJ Shake"
17-Danny/Lisa
18-Eric/Nate
19-Adriana/Jieming
20-Christian/Kevin
21-Philip/Michael
22-Eileen/Christine

# Two-agent newsvendor

- Trust policy versus the class

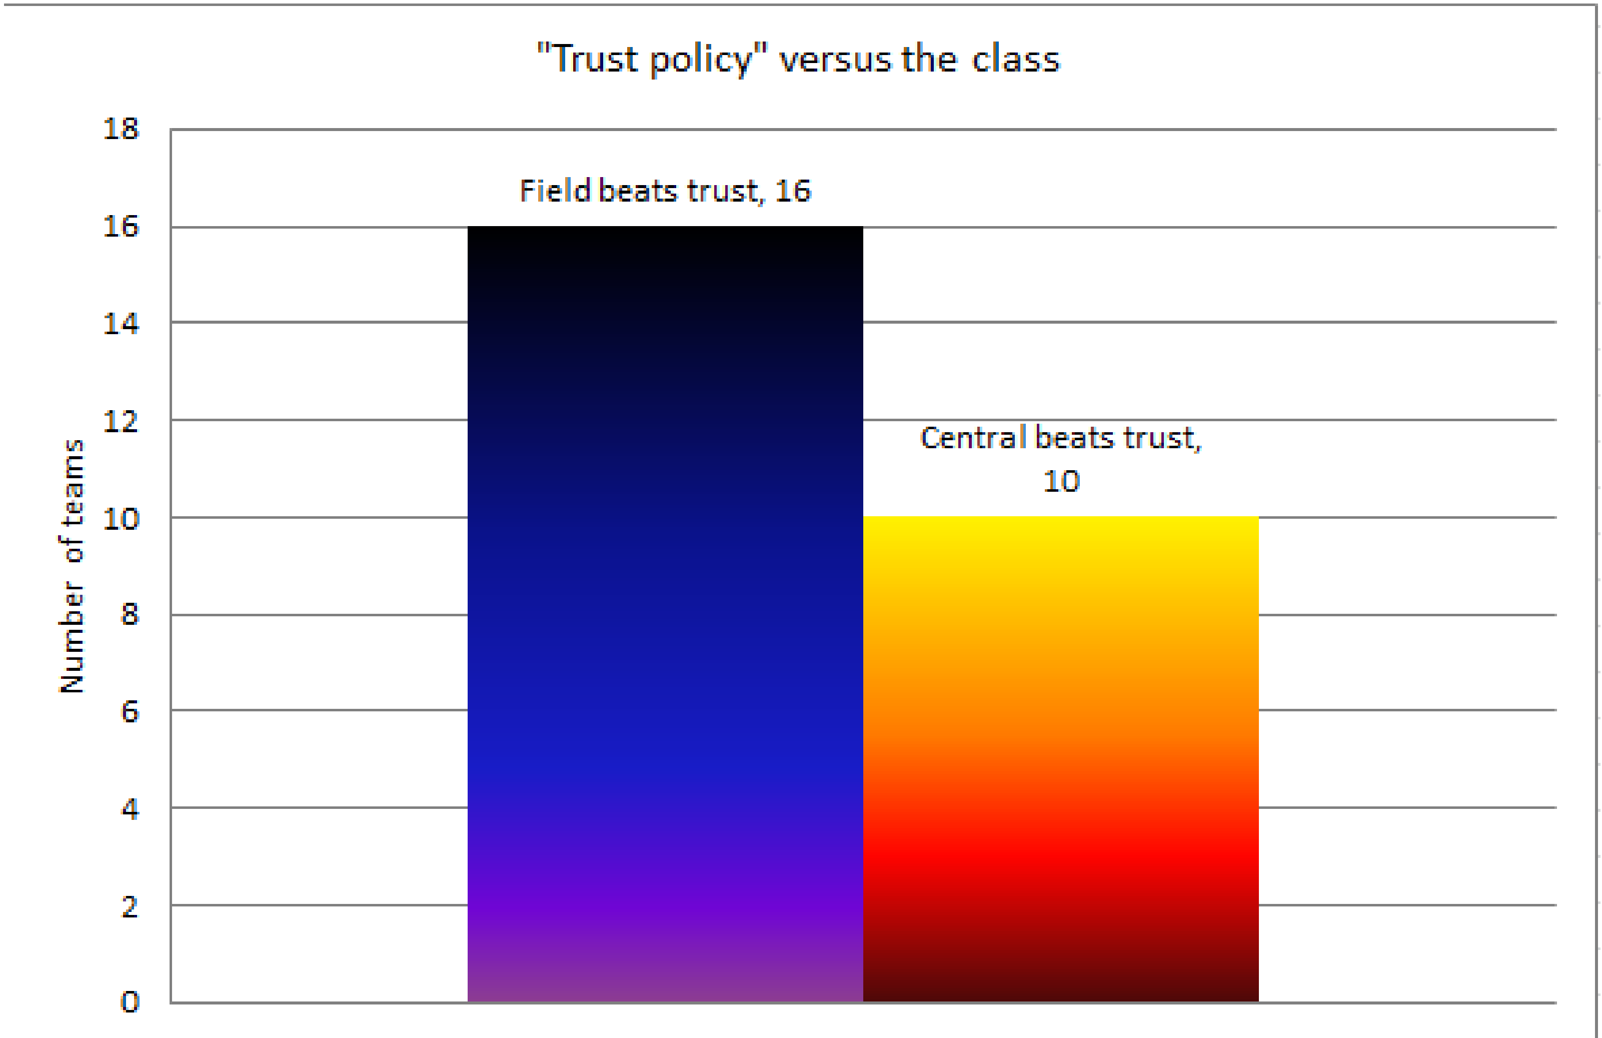

# Two-agent newsvendor

- Field+central as a team versus trust policy

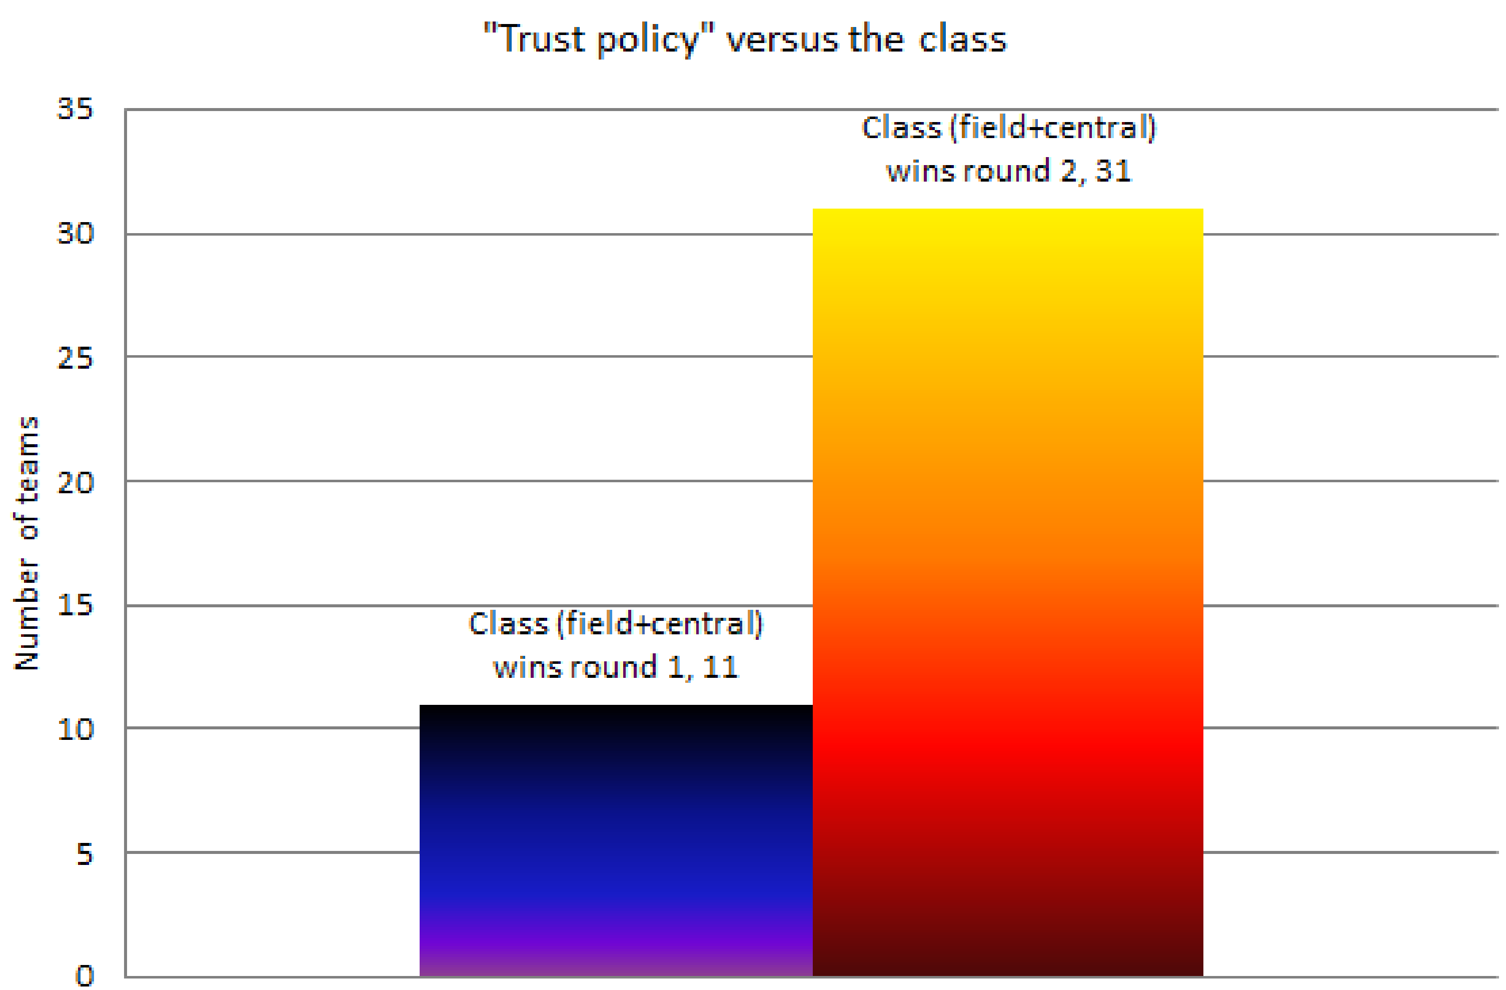

# Outline

- Summary of game results
- A two-agent learning model
- Feedback loops and misinformation

# 2-Player Decision-Making Model

Version 2.2

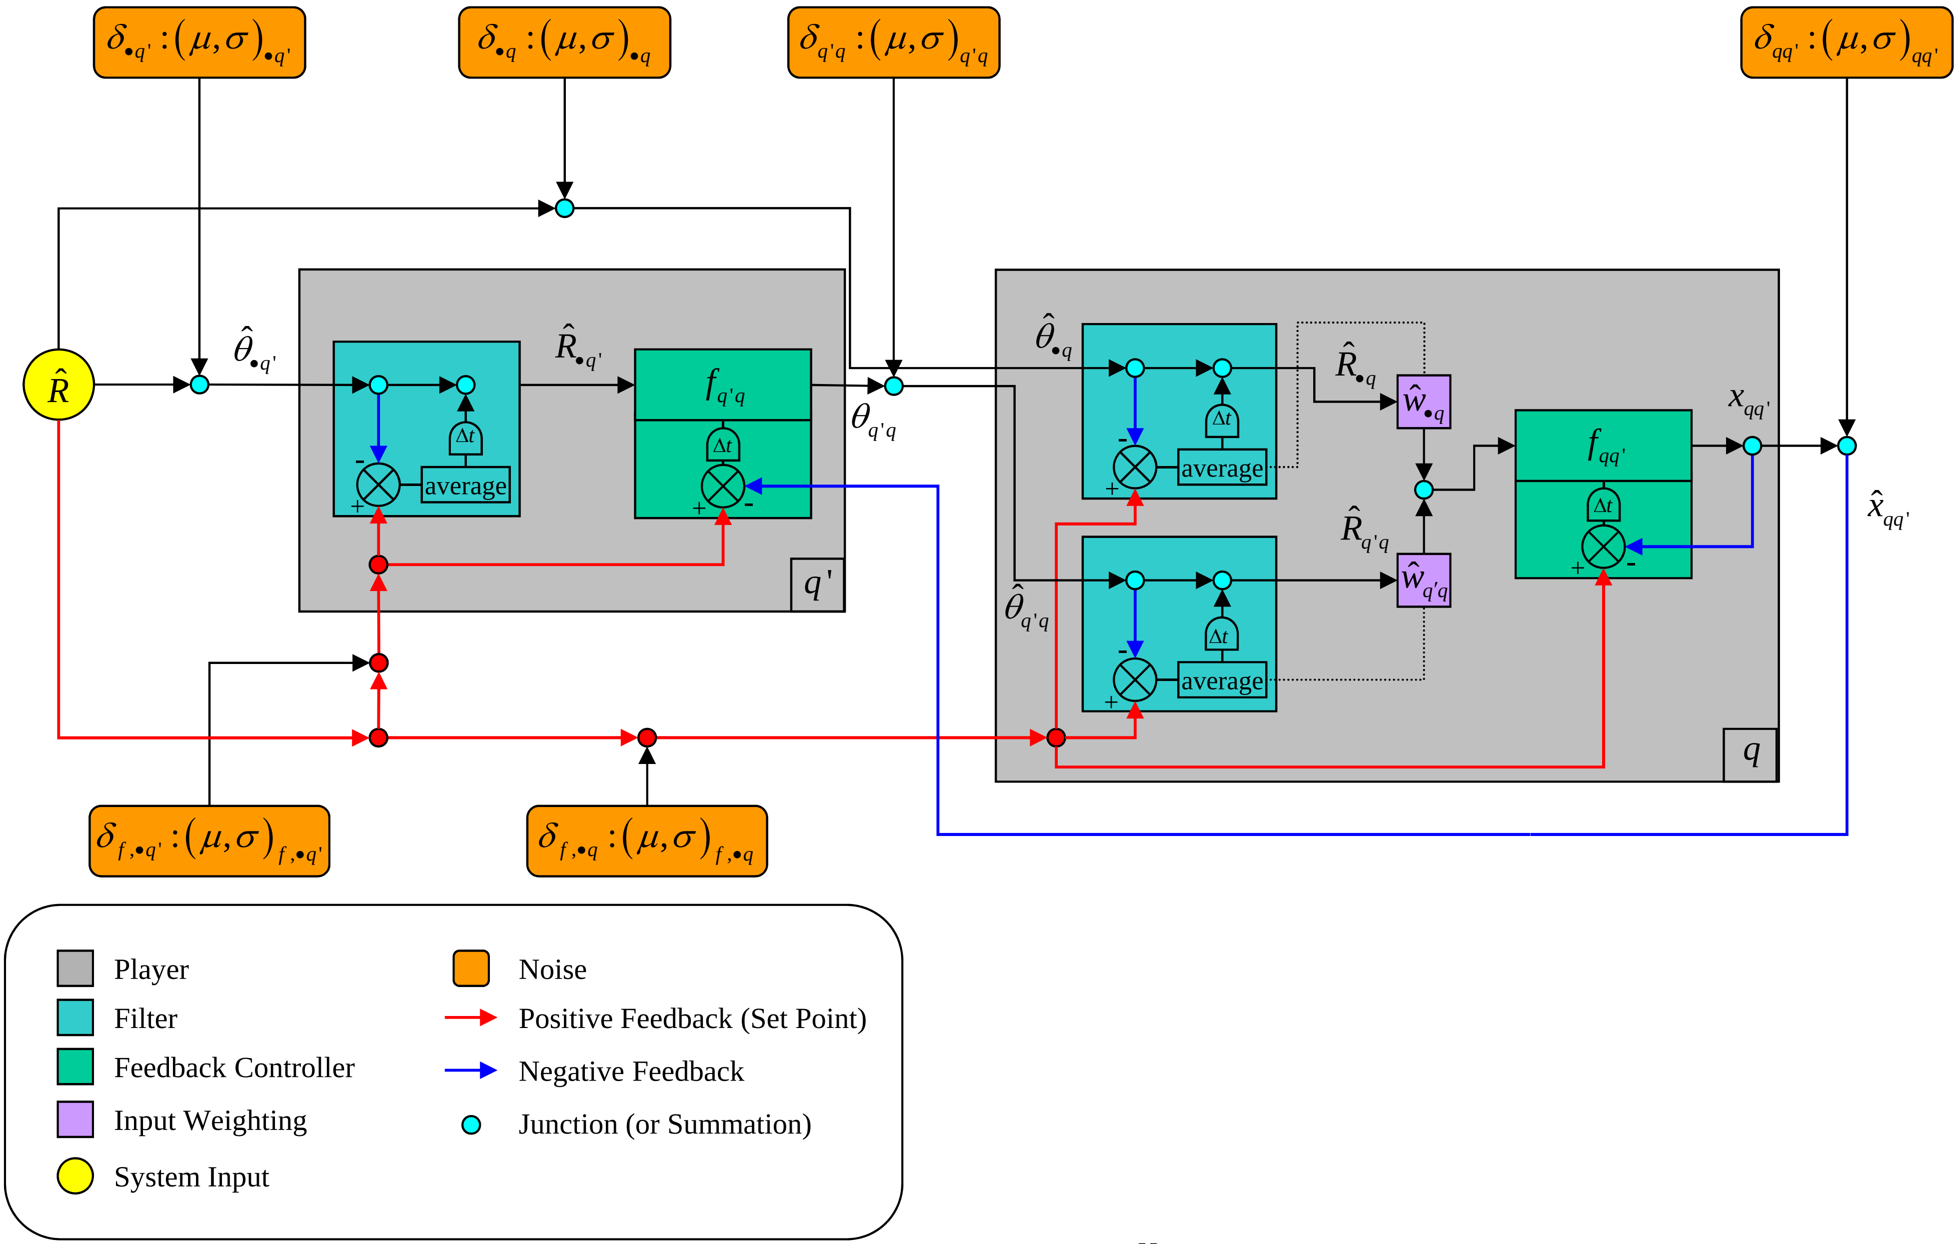

# 2-Player Decision-Making Model

Version 2.2

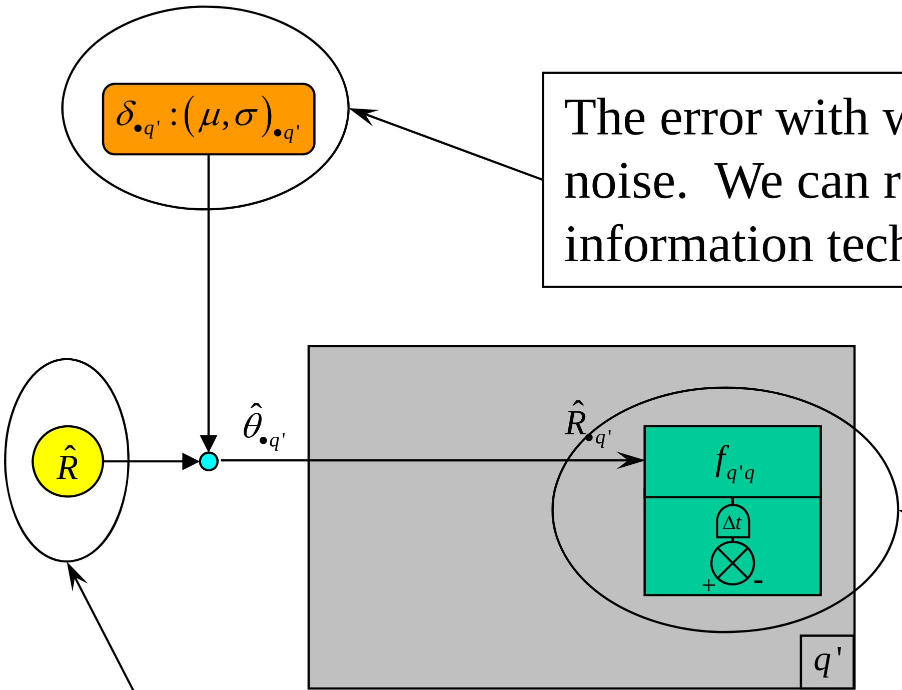


The error with which we perceive the need, with bias and noise. We can reduce these by making investments in information technology.

The "agent" that has to think and decide how much to order.

Exogenous requirement. $\hat{R}$ represents the real need which we cannot measure perfectly.

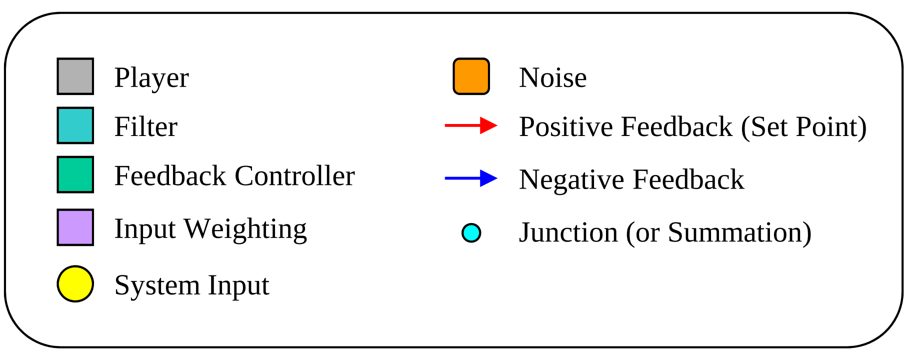


Our agent wants to solve:

$$\min_{x} \left\{ c^{o}_{q'} \left( x - \hat{R} \right)^{+} + c^{u}_{q'} \left( \hat{R} - x \right)^{+} \right\}$$

As in the past, he does not know exactly what is required, but does know if he requested too much or too little.

# 2-Player Decision-Making Model

Version 2.2

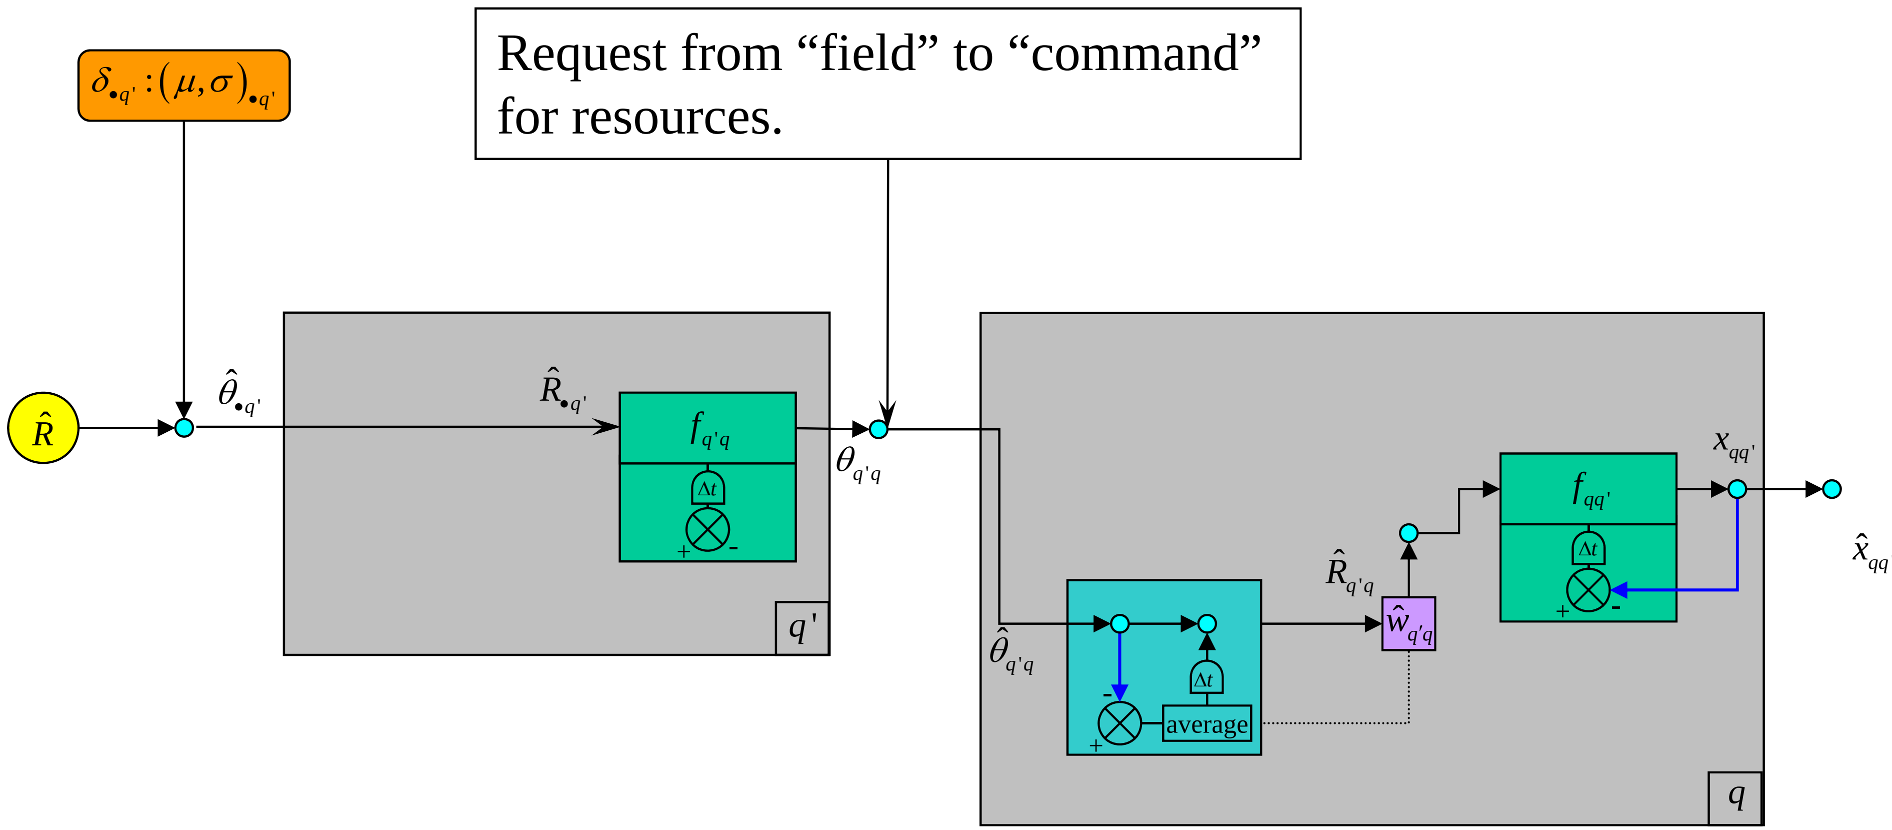


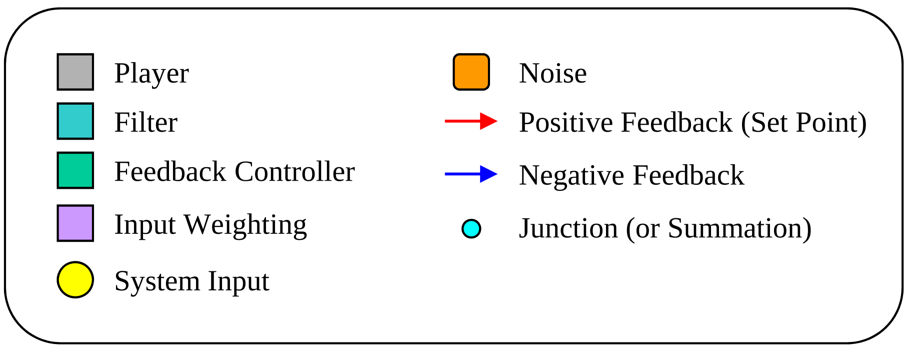

# Information exchange

The field agent $q'$ wants to solve:

$$\min_x \left\{ c^o_{q'} \left( x - \hat{R} \right)^+ + c^u_{q'} \left( \hat{R} - x \right)^+ \right\}$$

The central command $q$ wants to solve:

$$\min_x \left\{ c^o_{q} \left( x - \hat{R} \right)^+ + c^u_{q} \left( \hat{R} - x \right)^+ \right\}$$

Typically:

$$c^u_q < c^u_{q'}$$

because the field is usually much more risk averse.

*How does this affect the exchange of information between the agents?*

# 2-Player Decision-Making Model

Version 2.2

An independent source of information (usually with more noise).

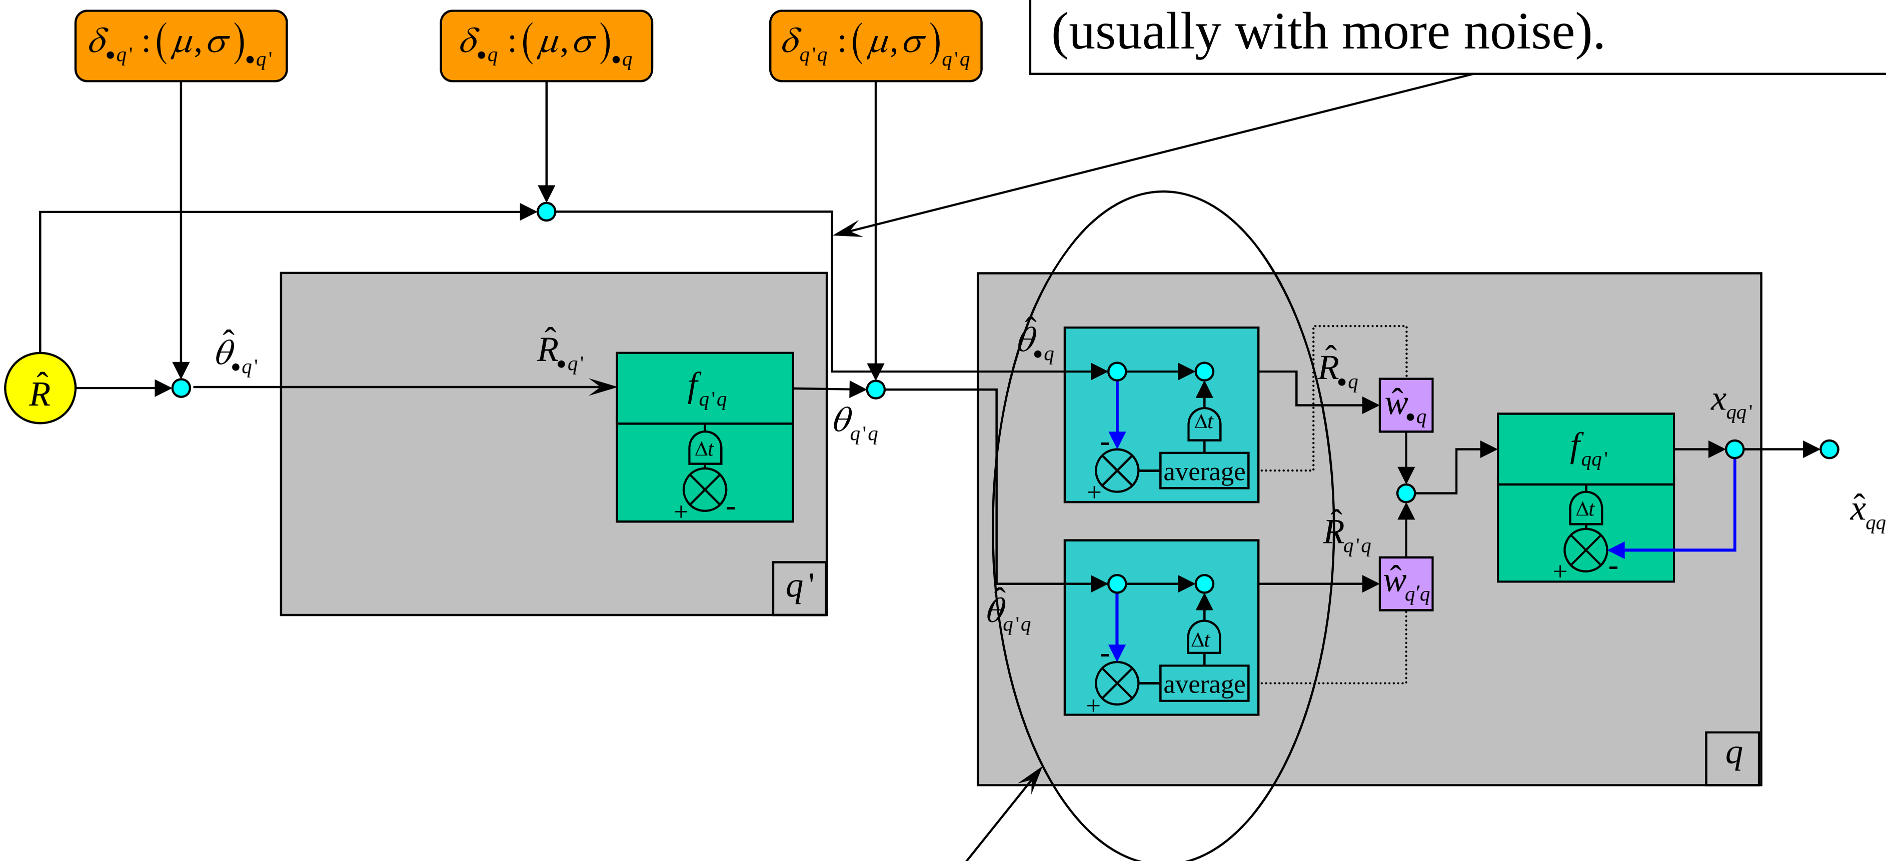


Our central command agent $q$ has to combine these two sources of information.

# Information exchange

Let:

$\hat{R}_{q'q}$ = Information in the form of a request for resources from $q'$ to $q$.

$\hat{R}_{\bullet q}$ = Information about the requirement from an independent source.

Each of these flows of information come with a bias $\mu$ and zero mean noise with variance $\sigma^2$. The 'central' combines these requests using a weighted average:

$$\hat{R}_q = \hat{w}_{q'q}\hat{R}_{q'q} + \hat{w}_{\bullet q}\hat{R}_{\bullet q}$$

The best weights are proportional to the inverse estimates of the variance of the quality of these two sources of information. Let:

$s^2_{q'q}$ = Estimate of the variance of the noise coming from $q'$.

$s^2_{\bullet q}$ = Estimate of the variance of the noise coming from the exogenous source.

Now set the weights to:

$$\hat{w}_{q'q} = \frac{\frac{1}{s^2_{q'q}}}{\frac{1}{s^2_{q'q}} + \frac{1}{s^2_{\bullet q}}} = \frac{s^2_{\bullet q}}{s^2_{q'q} + s^2_{\bullet q}} \qquad \hat{w}_{\bullet q} = \frac{\frac{1}{s^2_{\bullet q}}}{\frac{1}{s^2_{q'q}} + \frac{1}{s^2_{\bullet q}}} = \frac{s^2_{q'q}}{s^2_{q'q} + s^2_{\bullet q}}$$

# Information exchange

Take a close look at our information weights:

$$\hat{w}_{q'q} = \frac{s^2_{\bullet q}}{s^2_{q'q} + s^2_{\bullet q}} \qquad \hat{w}_{\bullet q} = \frac{s^2_{q'q}}{s^2_{q'q} + s^2_{\bullet q}}$$

The weight we put on the information from $q'$ increases with the variance of the information from the exogenous source.

Similarly, the weight we put on the information from the exogenous source increases with the variance of the information from agent $q'$.

This weighting makes intuitive sense. It is also the optimal weighting, producing a combined estimate with the lowest variance.

# 2-Player Decision-Making Model

Version 2.2

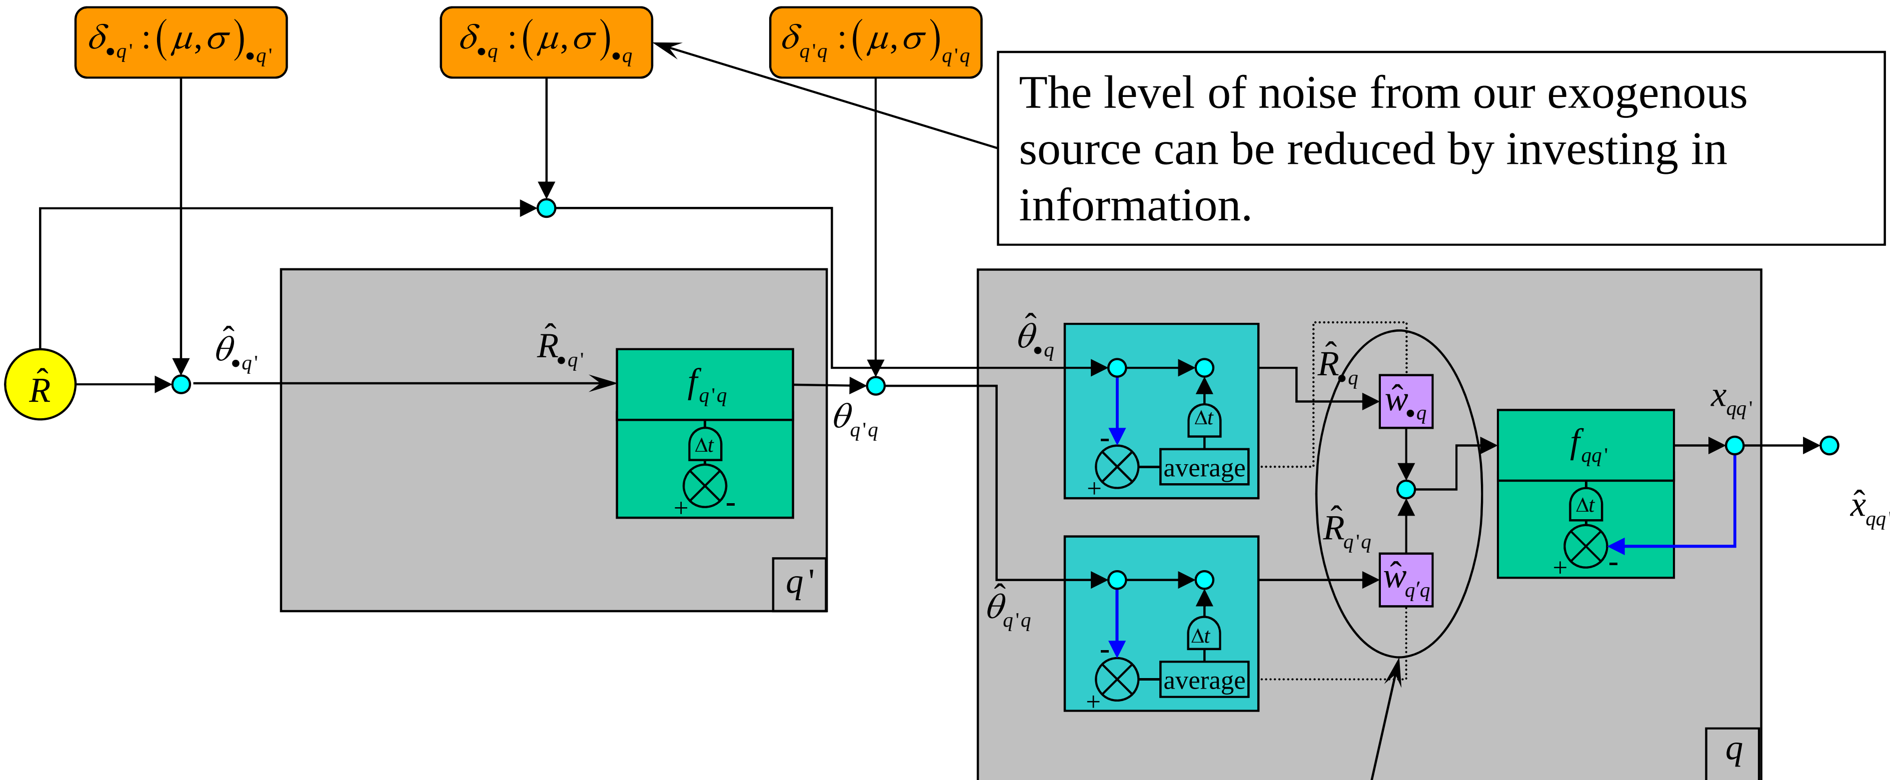


The level of noise from our exogenous source can be reduced by investing in information.

Weighting the two inputs provides the best estimate of the actual need.

***But how do we get the estimates of the variances?***

# 2-Player Decision-Making Model

Version 2.2

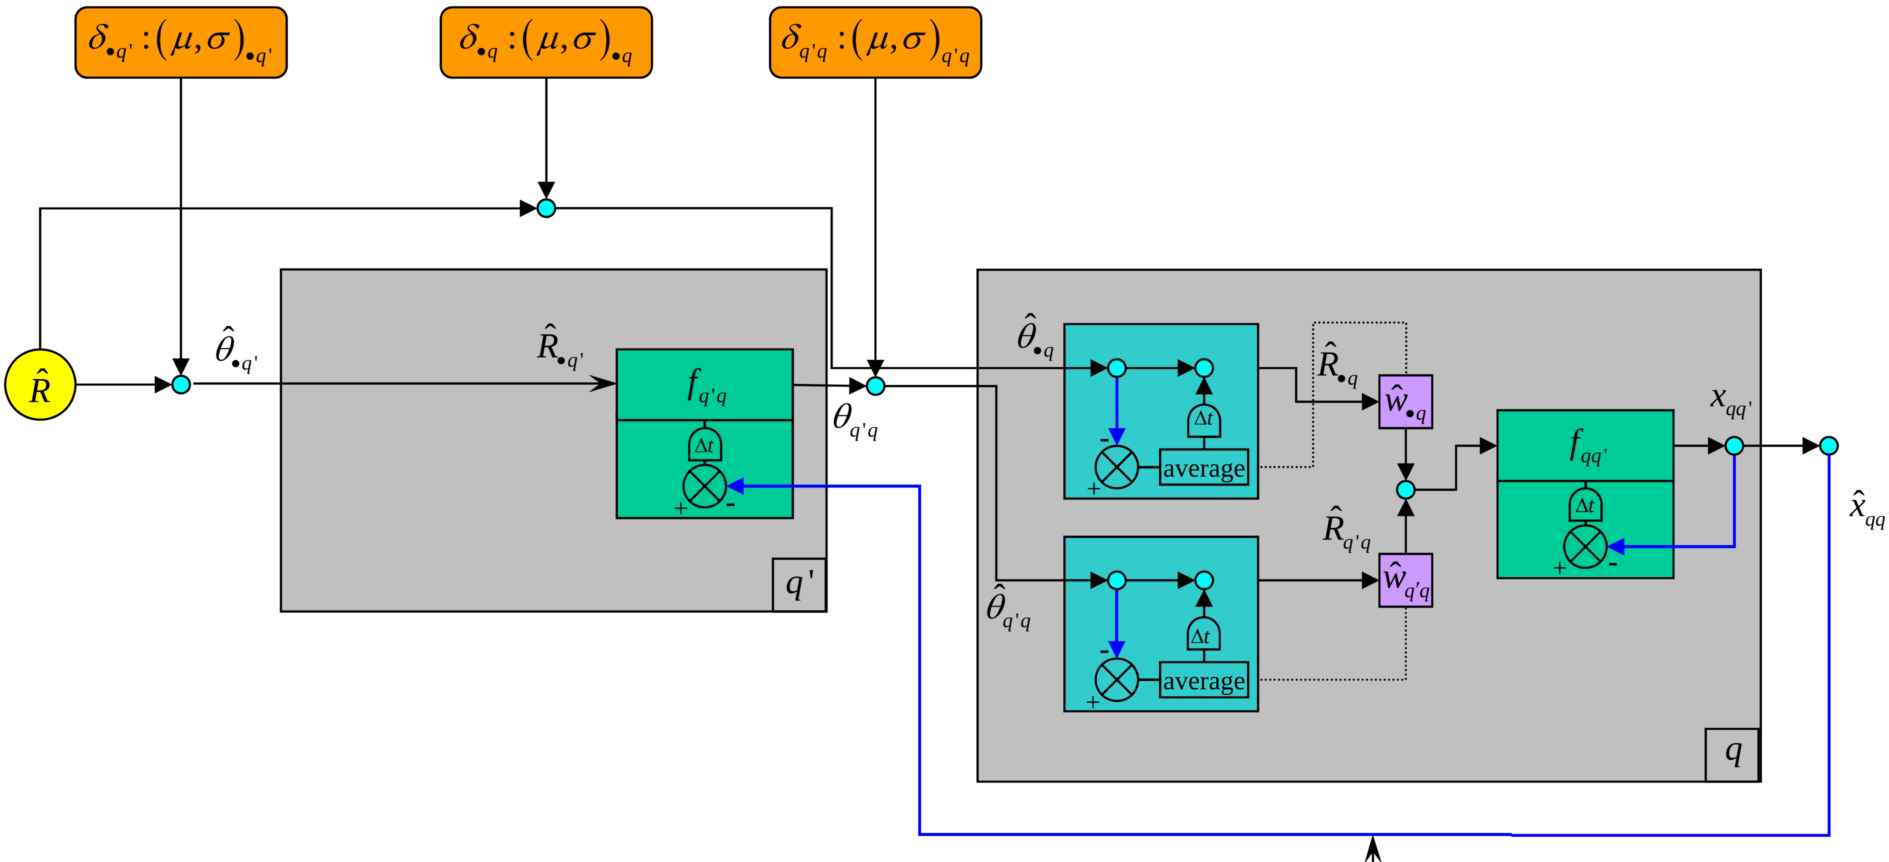


Resources granted by the "central command"

# 2-Player Decision-Making Model

Version 2.2

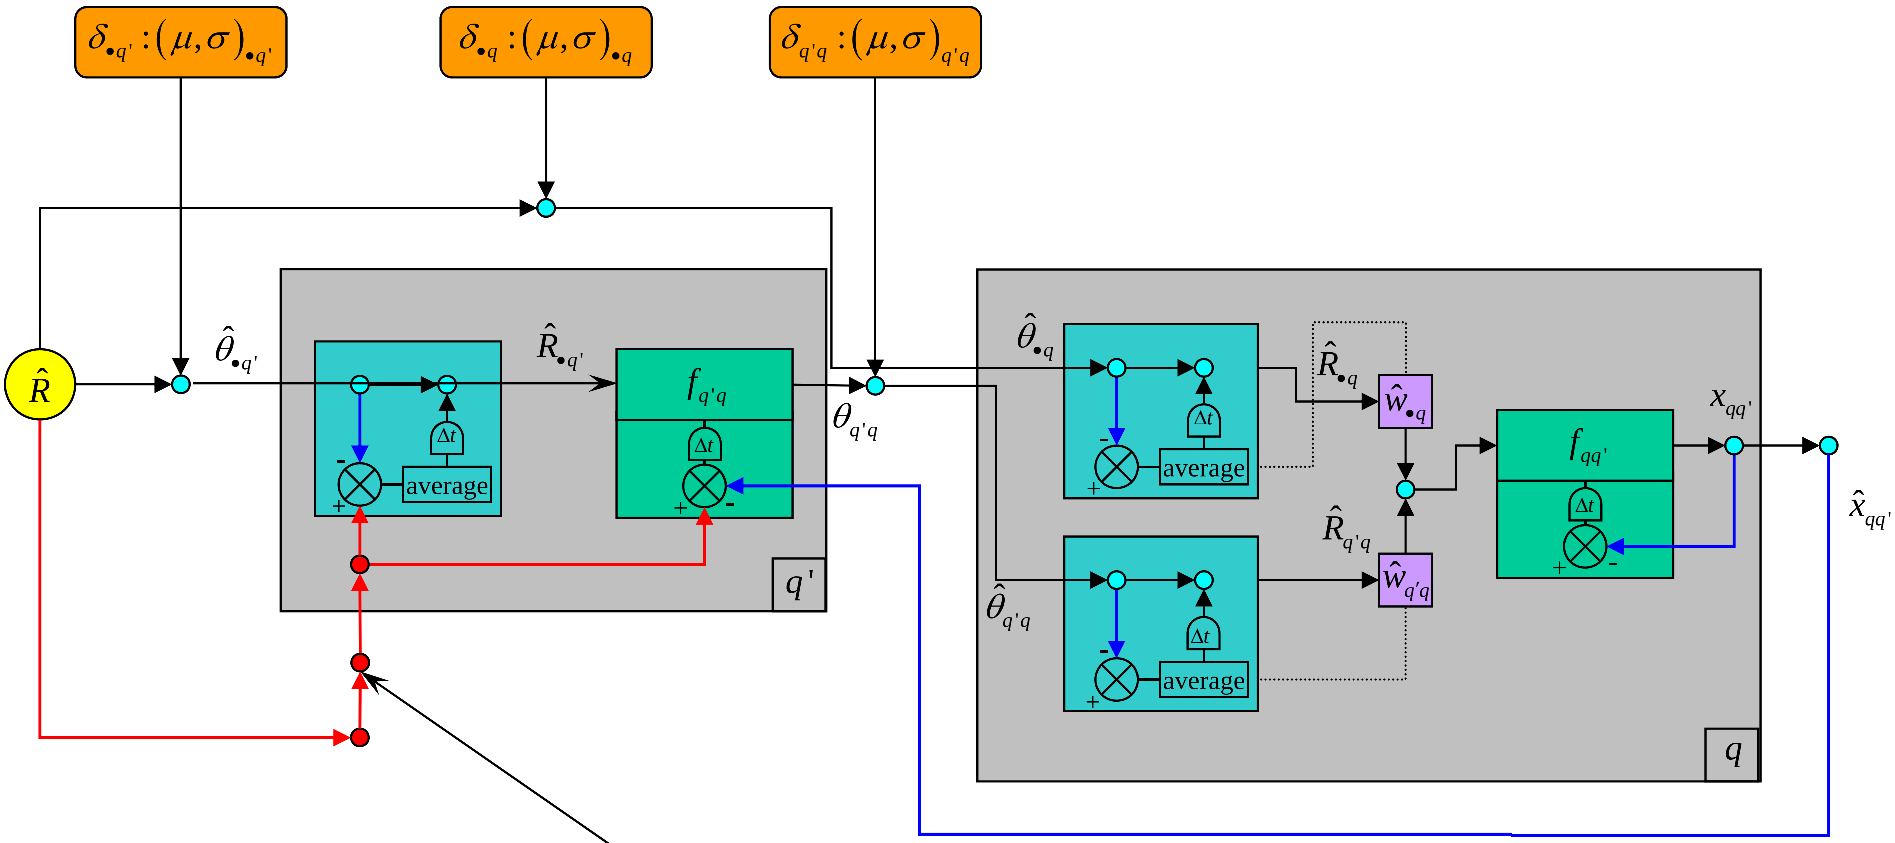


Here we "learn" what happened after the fact. We can use this information to learn about biases and noise, and use this for future calculations.

# Information exchange

Estimating biases:

$$\hat{\mu}^{k+1}_{\bullet q'} = \hat{\alpha}_{\bullet q'}(\hat{\theta}^{k}_{\bullet q'} - \hat{R}^{k}) + (1-\hat{\alpha}_{\bullet q'})\hat{\mu}^{k}_{\bullet q'}, \qquad \hat{\mu}^{0}_{\bullet q'} = 0$$

Estimating variances (wherever applicable):

$$\hat{\delta}^{2^{k+1}}_{\bullet q'} = \hat{\alpha}_{\bullet q'}(\hat{\theta}^{k}_{\bullet q'} - \hat{R}^{k})^2 + (1-\hat{\alpha}_{\bullet q'})\hat{\delta}^{2^{k}}_{\bullet q'}, \qquad \hat{\delta}^{2^{0}}_{\bullet q'} = 0$$

$$\hat{\sigma}^{2^{k+1}}_{\bullet q'} = \hat{\delta}^{2^{k+1}}_{\bullet q'} - (\hat{\mu}^{k+1}_{\bullet q'})^2$$

$$s^{2^{k+1}}_{\bullet q'} = \hat{\alpha}_{\bullet q'}\hat{\sigma}^{2^{k+1}}_{\bullet q'} + (1-\hat{\alpha}_{\bullet q'})s^{2^{k}}_{\bullet q'}, \qquad \text{for a given } s^{2^{0}}_{\bullet q'}$$

# 2-Player Decision-Making Model

Version 2.2

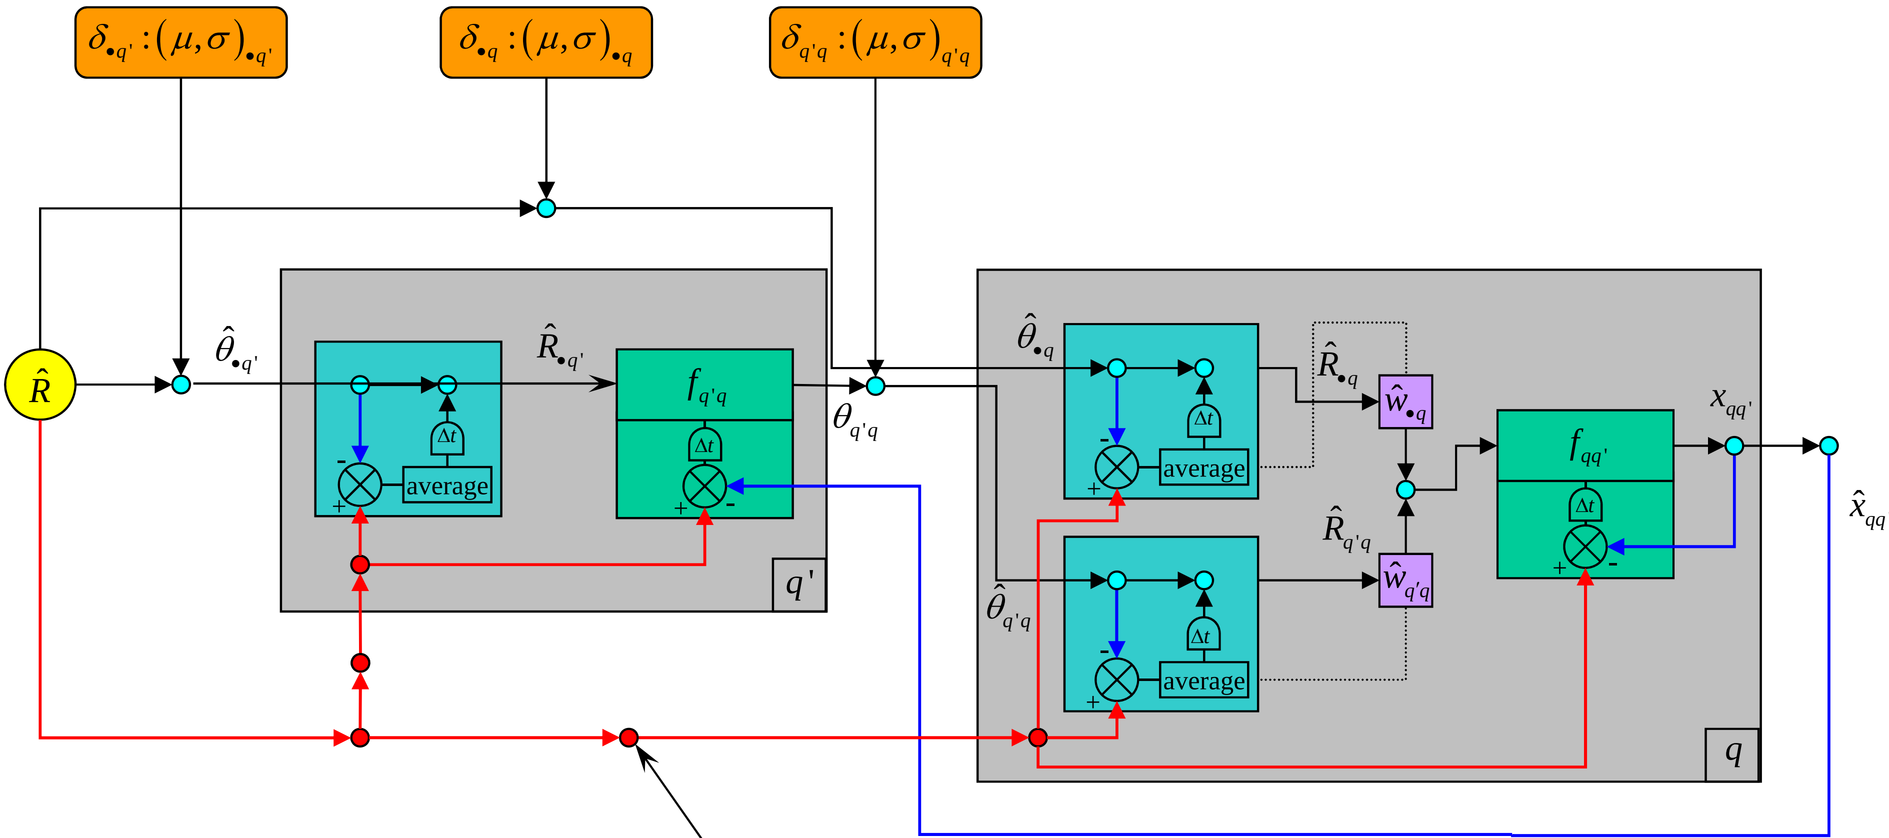


The central command also learns what happens after the fact, and is thus able to estimate biases, variances, etc.

# 2-Player Decision-Making Model

Version 2.2

There may be communication errors from the central command back to the field.

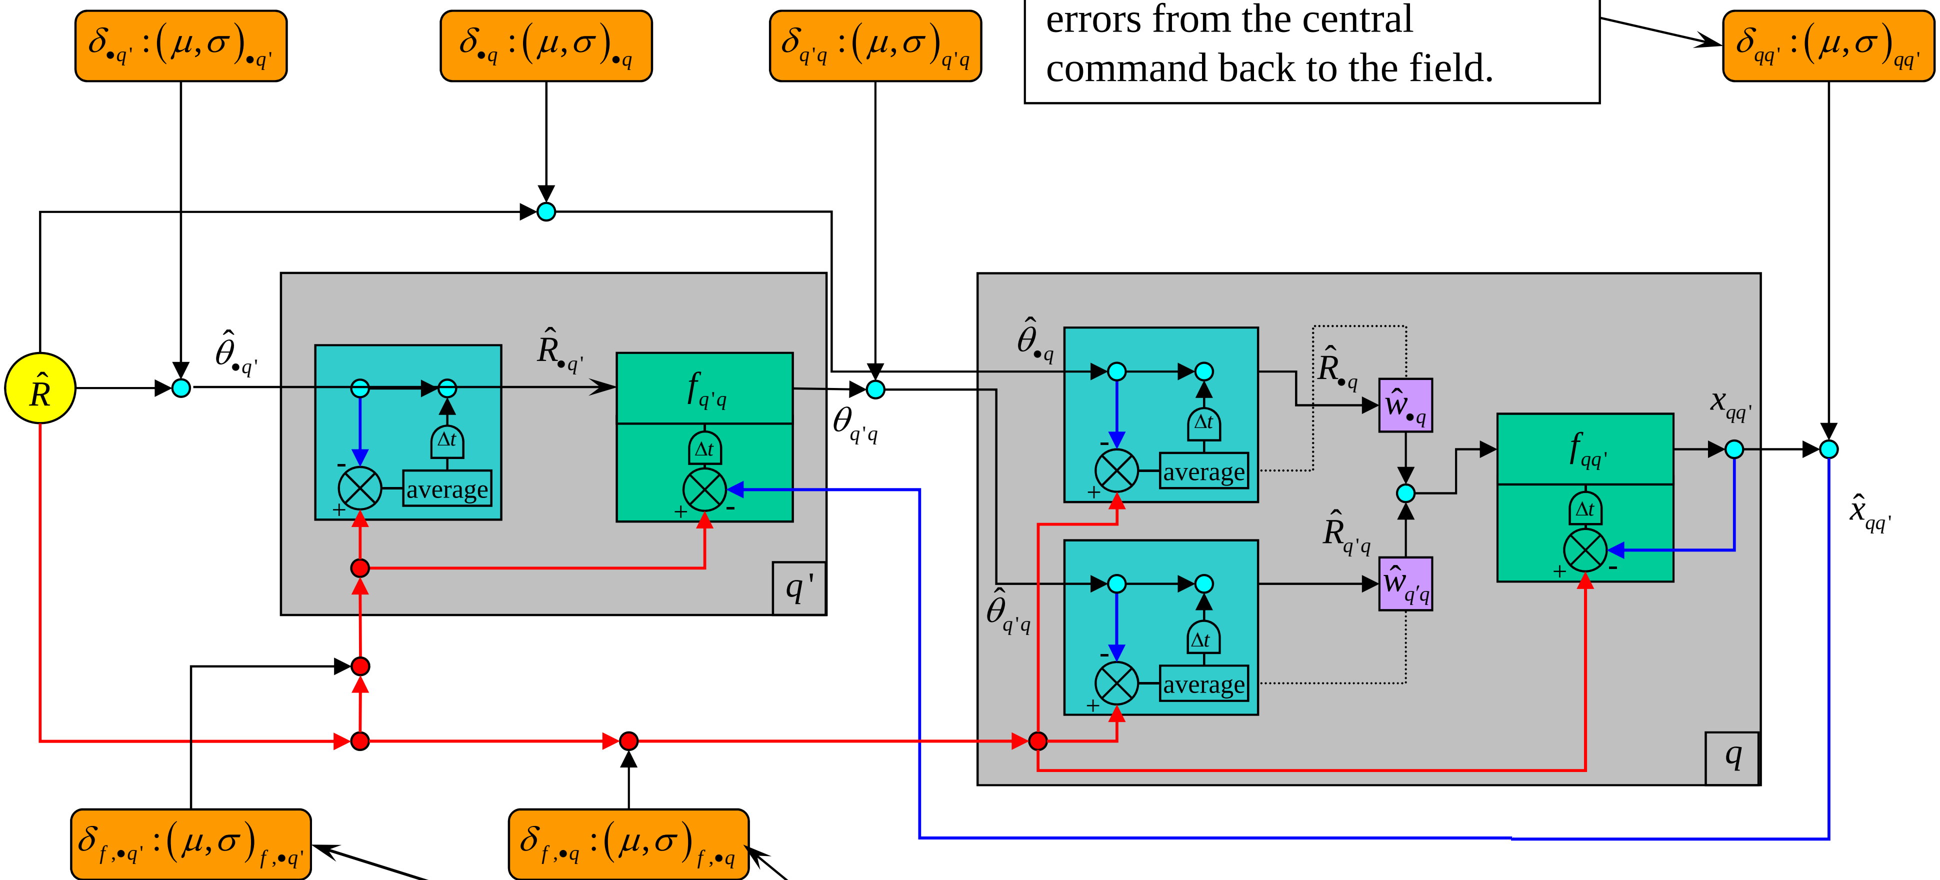


But we may not learn exactly what was needed, and each agent has different information channels with different levels of noise.

# 2-Player Decision-Making Model

Version 2.2

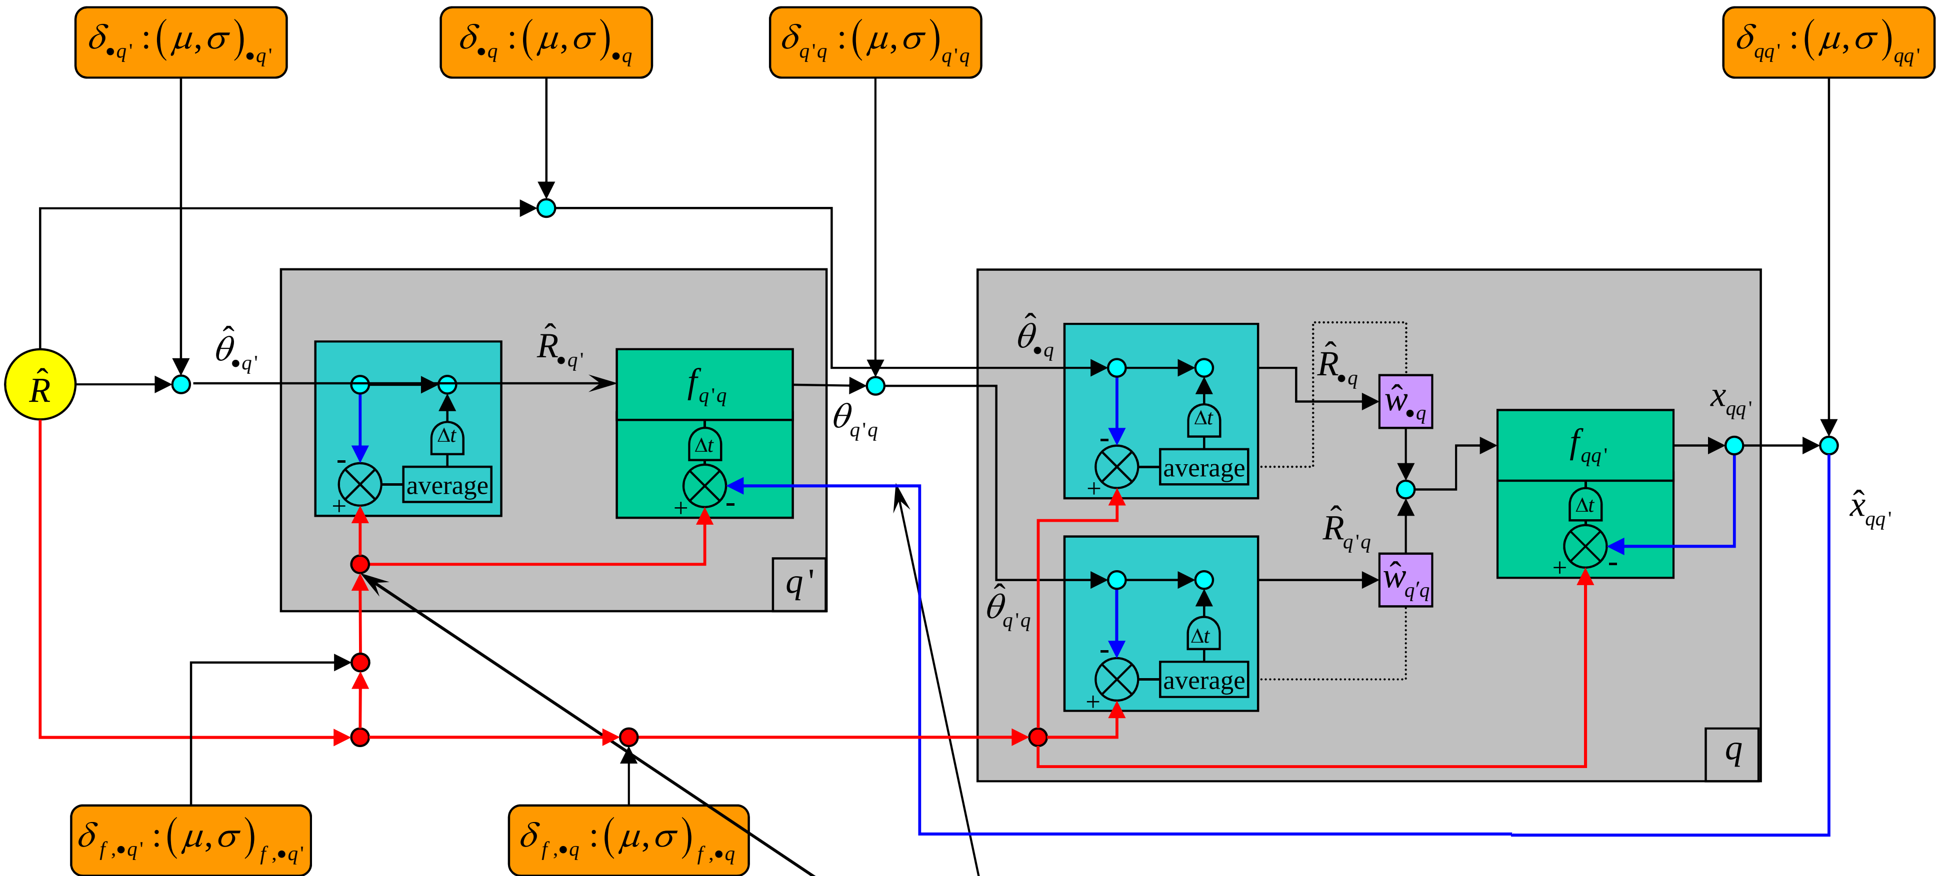


The agent looks at what was really needed (as best as he can measure it) and compare this to what the central command actually gave him. From this, he can identify biases in the behavior of the central command, and account for this in future requests.

# Information exchange

Given:

$$\tilde{R}^k = \hat{R}^k + \mu_{f,\bullet q'} + \varepsilon_{f,\bullet q'} \quad \text{where} \quad \varepsilon_{f,\bullet q'} \sim N(0, \sigma^2_{f,\bullet q'})$$

Estimating how much to order:

$$\theta^{k+1}_{q'q} = \hat{R}^{k+1}_{\bullet q'} + \beta^{k+1}_{q'q}$$

$$\beta^{k+1}_{q'q} = \beta^{k}_{q'q} - \gamma_{q'} \alpha^{k}_{q'q} \nabla F\left(\hat{x}^k_{qq'} - \tilde{R}^k, c^o_{q'}, c^u_{q'}\right)$$

where the stochastic gradient can be given by:

$$\nabla F\left(\hat{x}^k_{qq'} - \tilde{R}^k, c^o_{q'}, c^u_{q'}\right) = \begin{cases} c^o_{q'}, & \text{if } \hat{x}^k_{qq'} > \tilde{R}^k \\ -c^u_{q'}, & \text{if } \hat{x}^k_{qq'} \le \tilde{R}^k \end{cases}$$

# 2-Player Decision-Making Model

Version 2.2

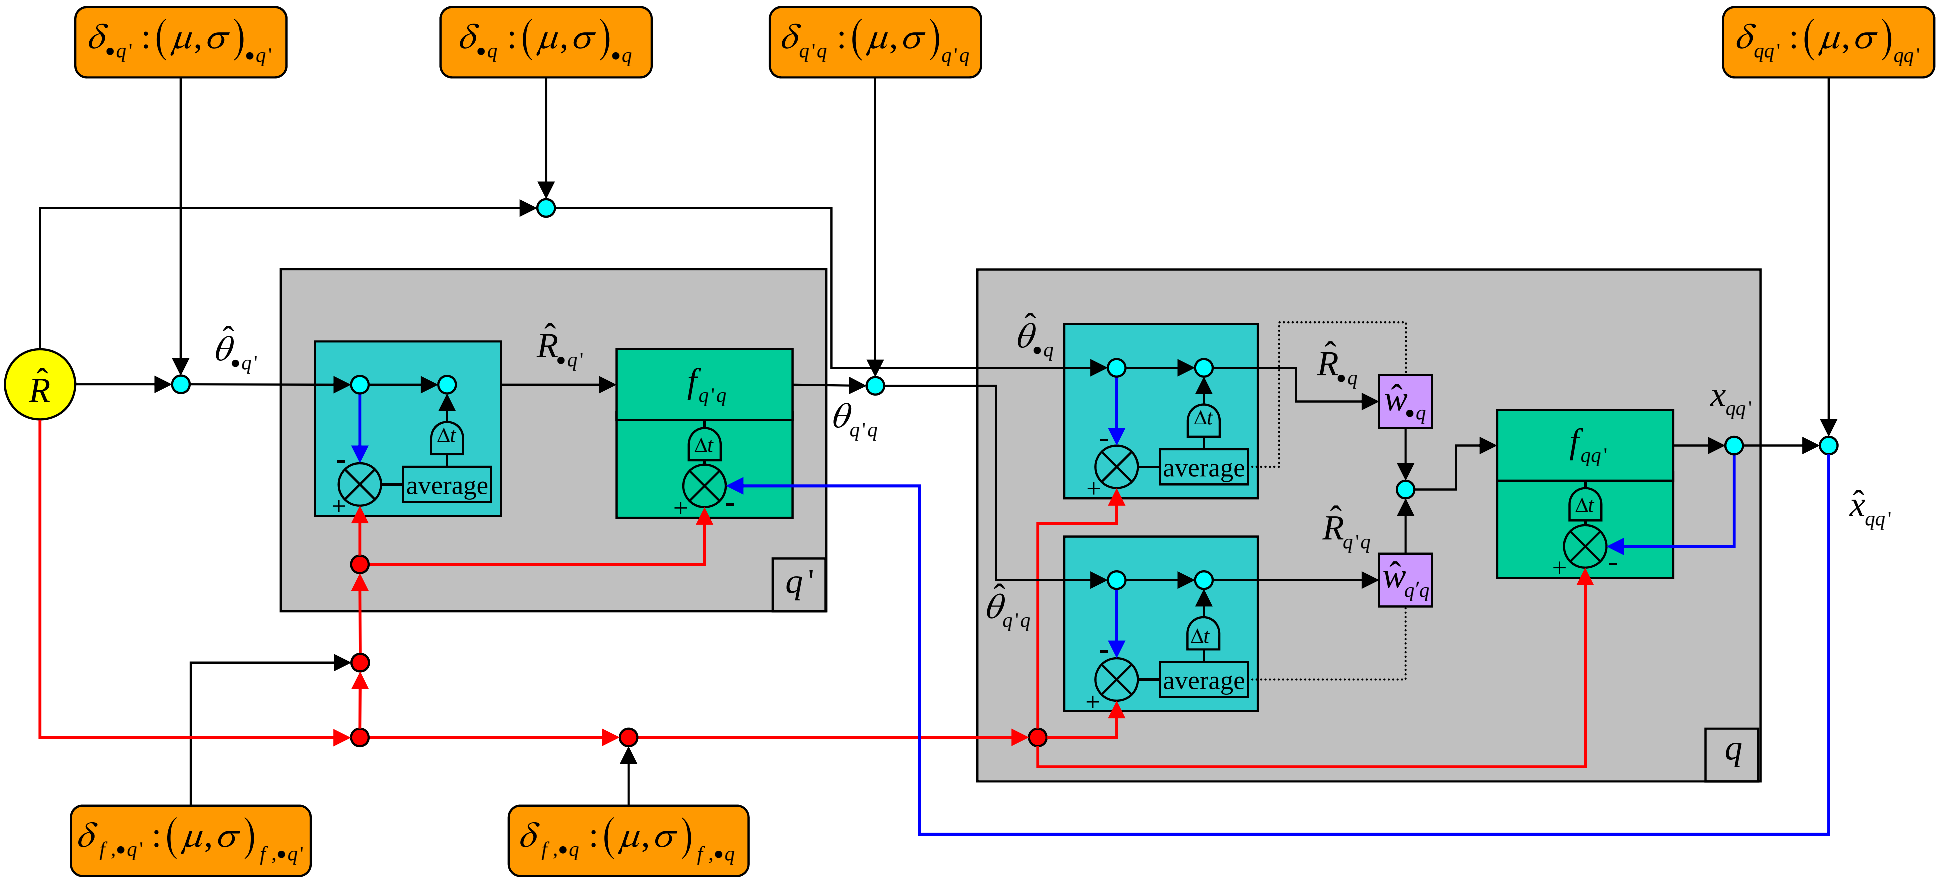


Player
Filter
Feedback Controller
Input Weighting
System Input
Noise
Positive Feedback (Set Point)
Negative Feedback
Junction (or Summation)

# Outline

- Summary of game results
- A two-agent learning model
- Feedback loops and misinformation

# Feedback loops and misinformation

- What are sources of misinformation in resource allocation problems?
  - » If we run out of resources, does this mean demand was too high?
  - » If we did not run out of resources, does this mean demand was too low?
  - » How do we estimate productivity? How do we estimate how many resources are required?

# The forms of misinformation

- Sources of misinformation in the feedback process:
  - » Bias from information hiding:
    - Censoring:
      - – Too few resources are allocated, all resources are used, and actual demand is hidden.
    - Inflation:
      - – Too many resources are allocated by central command.
      - – Field agent uses resources inefficiently to make it appear that the right amount was requested.
        - » Spend more time finishing a project
        - » Spend money unnecessarily – Ryder example, when management bought a lot of furniture the profits were reduced. Company did not want such a large bump in profits – would set unrealistic expectations.
    - Hiding
      - – The resources are there and available to be used, but the field agent does not allow them to be recorded.
        - » Hiding people: make them appear as if they are assigned to a project
        - » Hiding money: put it in a budget allocated to a project

# The forms of misinformation

- Sources of misinformation in the feedback process:
  - » Bias from information hiding (cont'd)
    - "Shrinkage"
      - – Resources are reduced due to theft, spoilage, breakage.
    - Rationing/hoarding:
      - – Too few resources are allocated, but resources are still left over. Agent has held resources for future potential uses, even if demand is not satisfied. So, it appears we have resources, but in reality we would have run out. We were "hiding" low priority demands.
  - » Measurement error
    - May be reduced through investment in information technology.
      - – Sensors on equipment; auditing accounts;