Susara van den Heever, PhD
6 October 2015

Decision Optimization: the Key Differentiator for Next-Generation Analytics

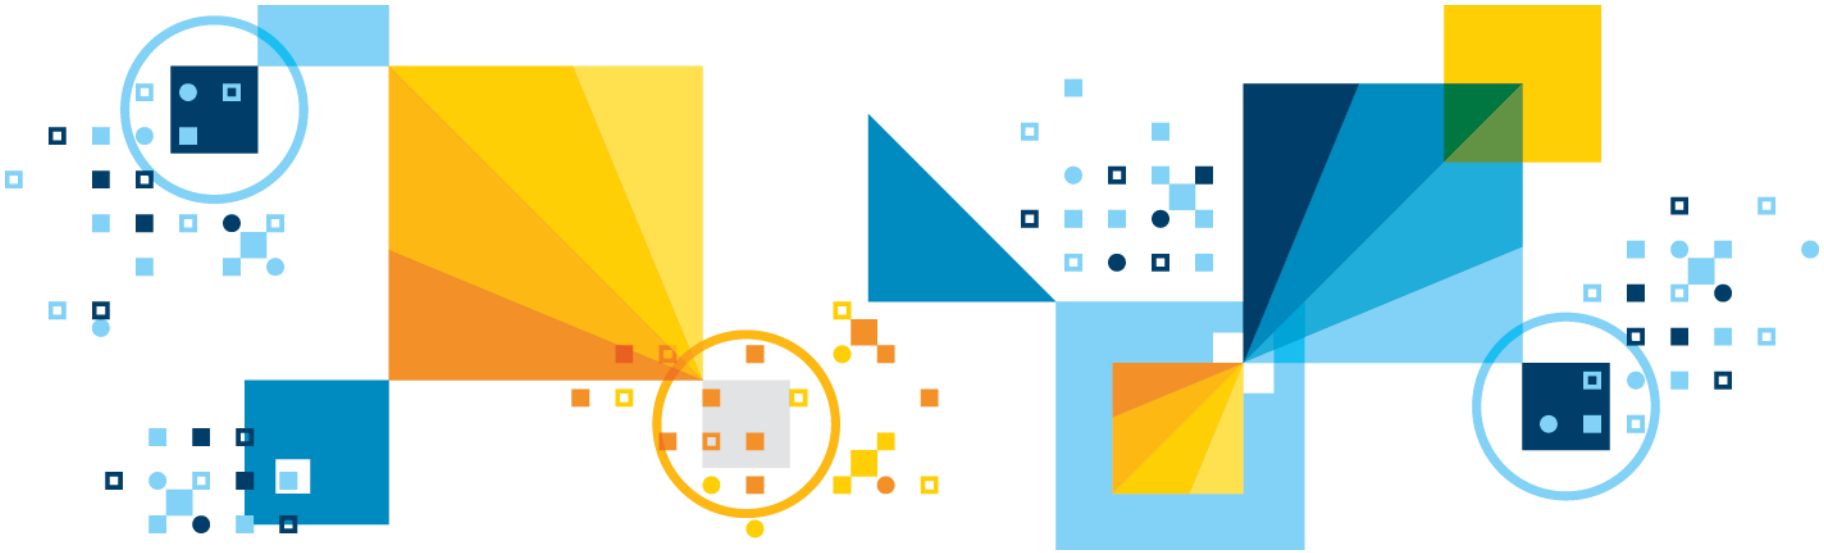

Important Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

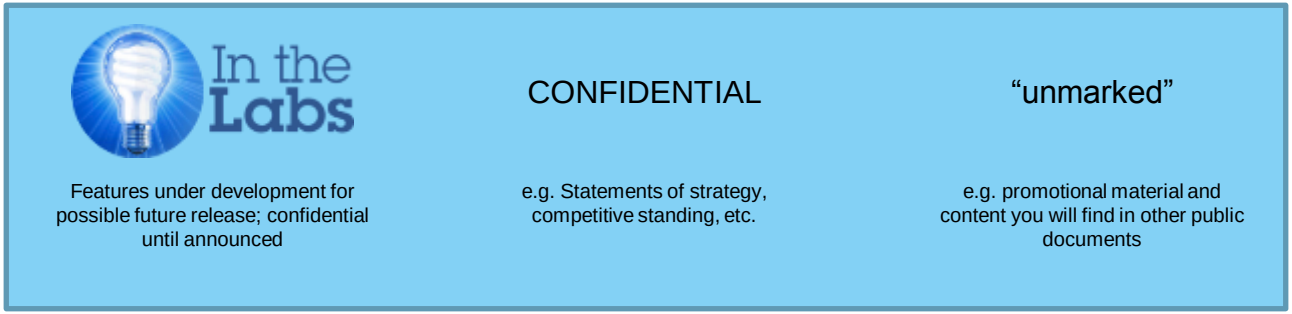

Overview

- Introduction to IBM Advanced Analytics
- How Decision Optimization on Cloud is changing the Prescriptive Analytics landscape
- The importance of uncertainty: the Uncertainty Toolkit
- Cognitive insights: a vision for consumable Watson Analytics-like optimization

Overview

- Introduction to IBM Advanced Analytics
- How Decision Optimization on Cloud is changing the Prescriptive Analytics landscape
- The importance of uncertainty: the Uncertainty Toolkit
- Cognitive insights: a vision for consumable Watson Analytics-like optimization

IBM Advanced Analytics

The most comprehensive portfolio of data and analytics offerings

Analytics Platform

Descriptive, Predictive, & Prescriptive Analytics

Software & Services

OpenSource Initiatives

Analytics / Industry Solutions

Pre-built solutions based on Analytics Platform, e.g. Predictive Maintenance & Quality (PMQ), Predictive Customer Intelligence (PCI), Energy Platform, and more

Watson

Cognitive base

Ability to learn, enhance, scale and accelerate, human expertise

Internet of Things

A platform designed for data-rich solutions that are secure and intelligent

Insight Cloud Services

Industry apps and solutions built in partnership with The Weather Company and Twitter that deliver new insights

Cloud Data Services

Cloud-based experience that enables developers to use the right platform-as-a-service tools for mobile and web app development

Unified Advanced Analytics Platform

Differentiated Analytic Solutions Made Possible By Breadth, Depth and Scale

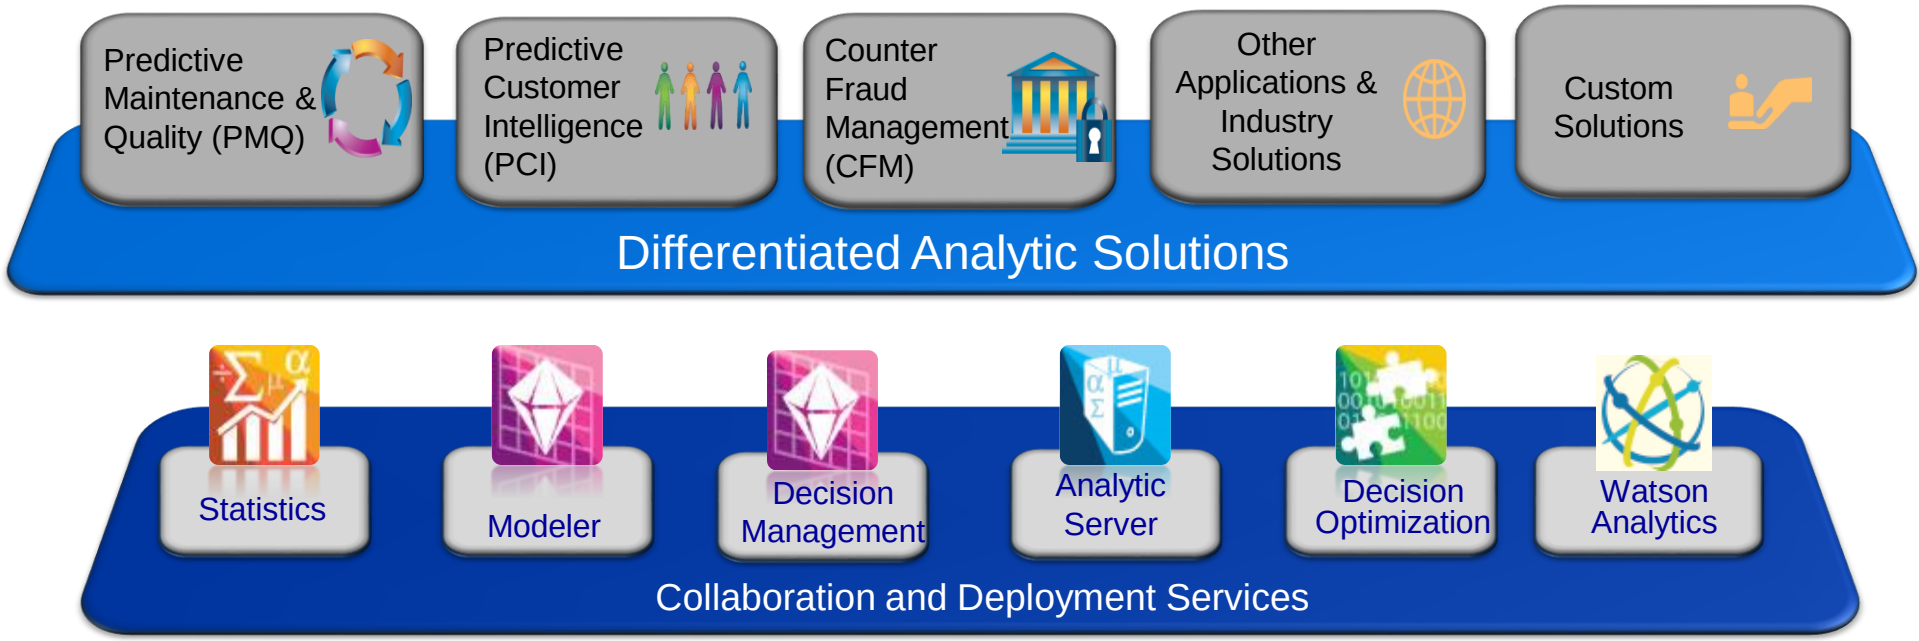

ALL Data
(Structured, Unstructured, Streaming)

IBM Analytics Platform: End-to-End Analytic Capabilities

Data Preparation	Analytics at Scale	Insight to Action
• Data models • Data connectors • Data Wrangling	• Predictive models • Statistical Analysis • Decision Optimization • Analytic models	• Real-Time Scoring • Optimized Decisions • APIs & services • Dashboards / Interactive apps

ALL Decisions
(People, Systems, Strategic, Operational, Real-Time)

Business Analytics and Optimization

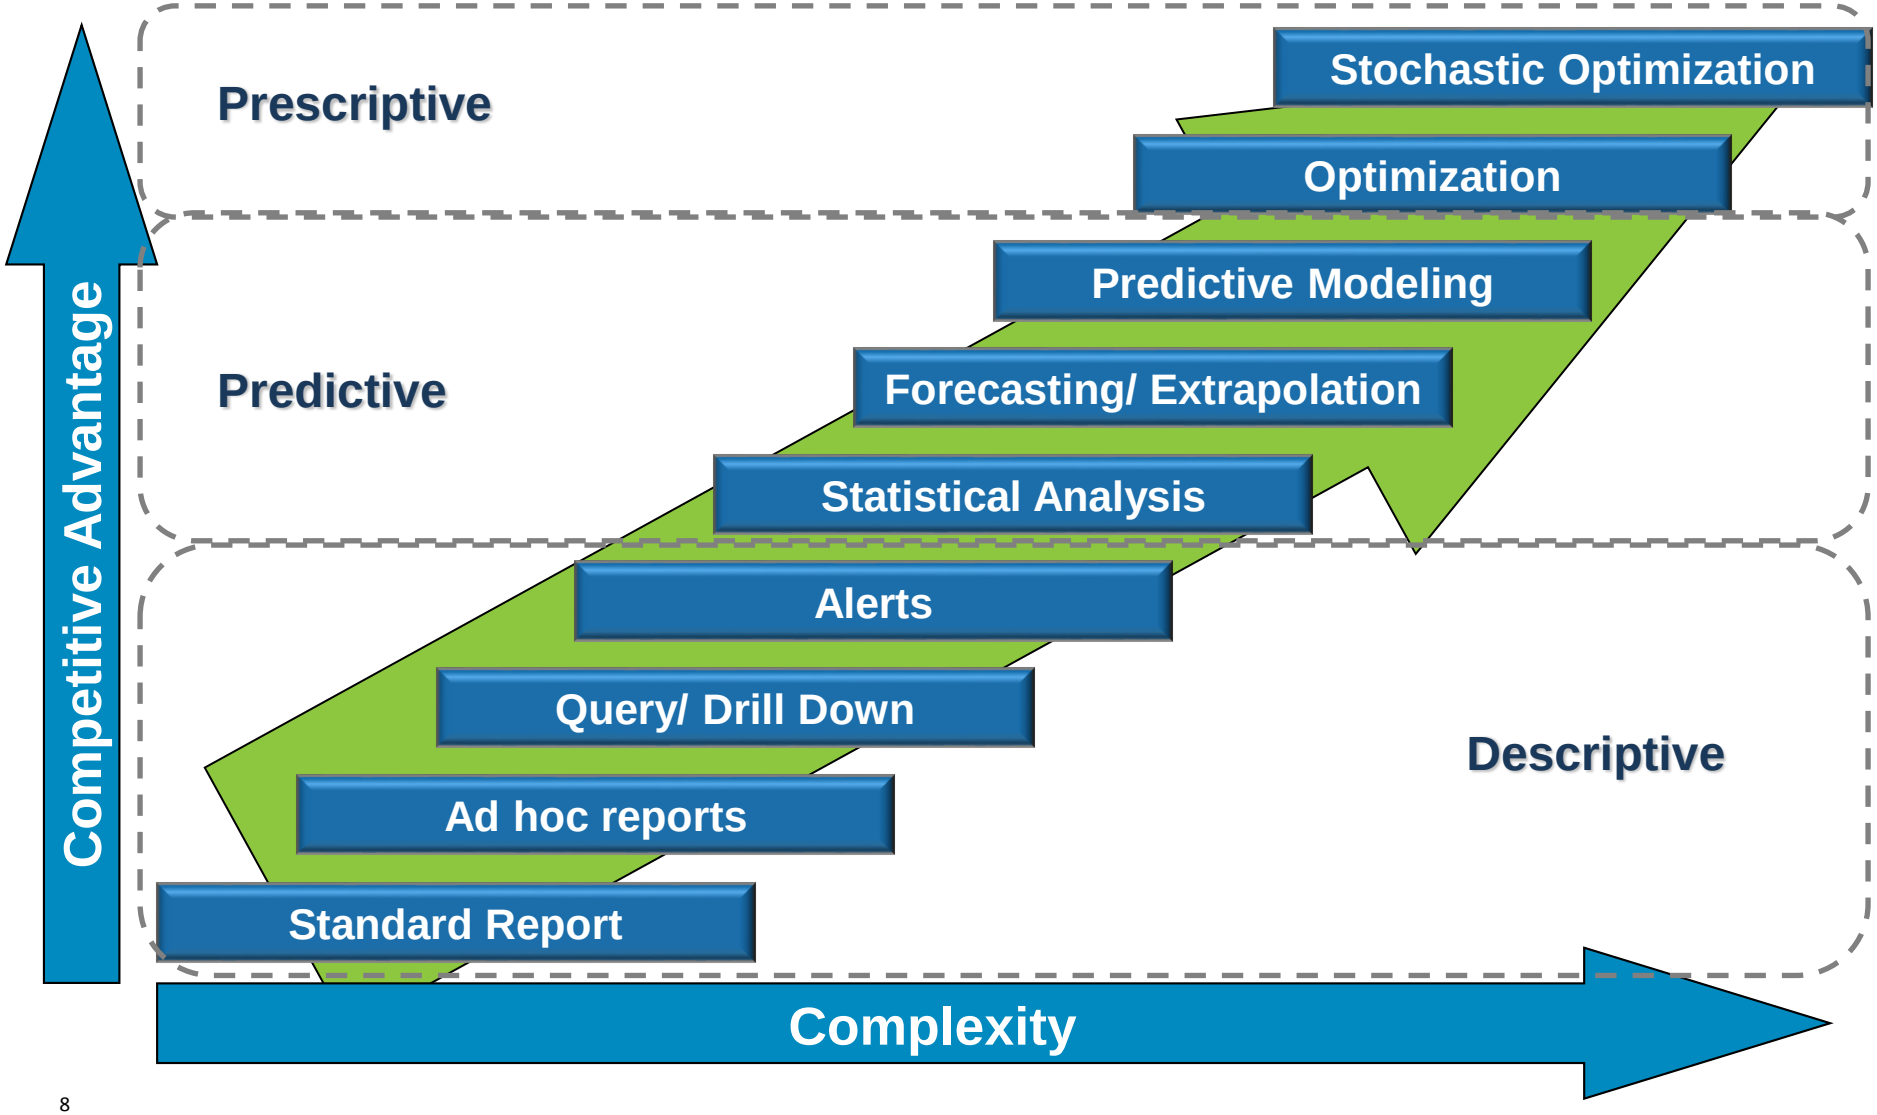

From the book "Competing on Analytics" by Thomas H. Davenport , Jeanne G. Harris

8

IBM's Decision Optimization Offering

CPLEX Optimization Studio (COS)
Optimization Engine

Model complex business problems.
Solve with IBM CPLEX Optimizer.
Prescribe precise and logical decisions.

Decision Optimization on Cloud
SaaS Delivery

Prescriptive analytics as a service.
No install, no setup.
Embed in other applications.

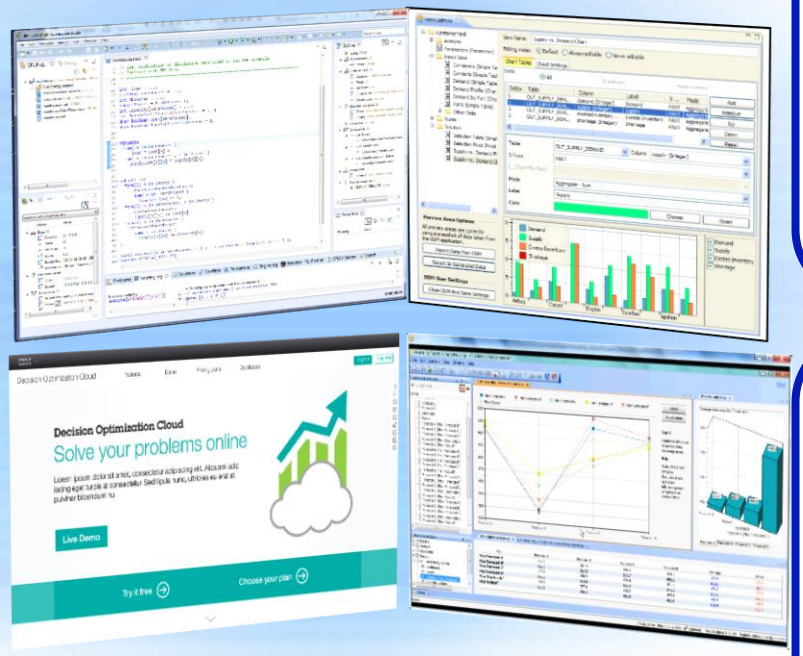

Decision Optimization Center (DOC)
Development & Deployment Platform

Build optimization solutions.
Includes data analysis & visualization, scenario management, collaborative planning, and what-if analysis.

Advanced Tools for
Optimization under Uncertainty

Compare multiple plans, scenarios, KPIs.
Understand trade-offs e.g. costs vs robustness.
Improve solution stability/robustness.

Overview

- Introduction to IBM Advanced Analytics
- How Decision Optimization on Cloud is changing the Prescriptive Analytics landscape
- The importance of uncertainty: the Uncertainty Toolkit
- Cognitive insights: a vision for consumable Watson Analytics-like optimization

IBM Decision Optimization on Cloud (DOcloud)

(30-day free trial https://ibm.co/docloudtrial)

Flexible configuration:
virtual machine
or bare metal servers,
cores and RAM options

Flexible pricing:
pay-as-you go,
committed hours,
reserved capacity

DropSolve

Drag-and-drop to solve models on cloud

DOcloud API

Embed DOcloud in any application

DOcloud Community

Quick access to OR, IT, and cloud experts

How does DOcloud fit with COS and DOC?

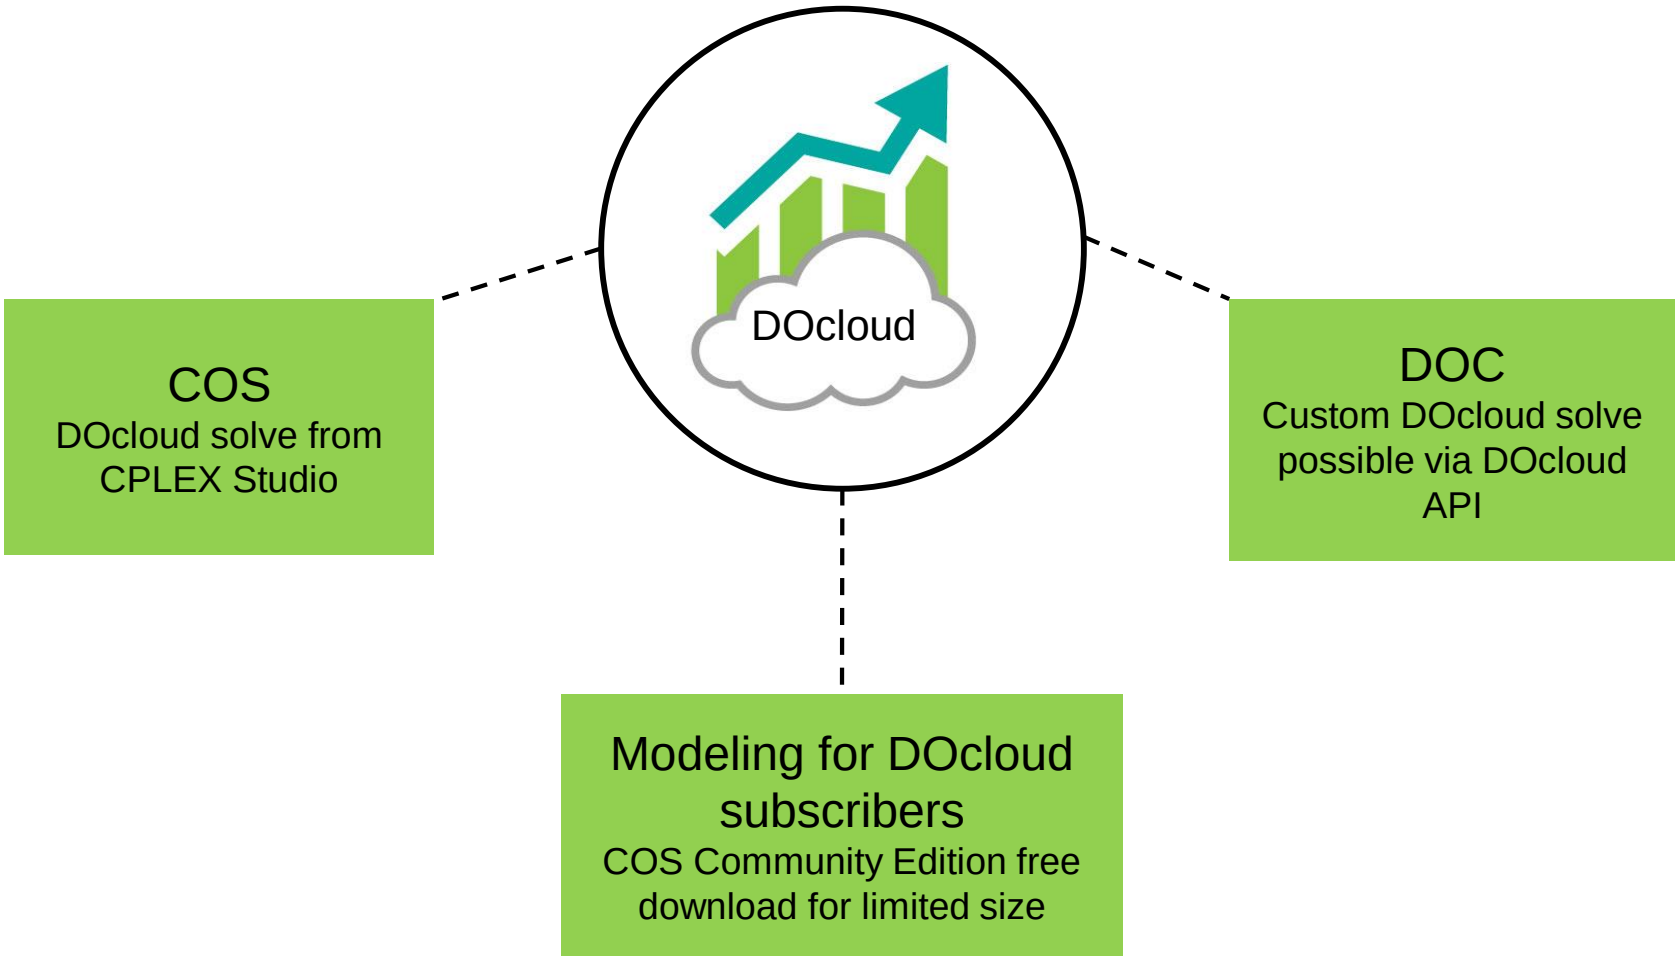

DOcloud personas

Oscar
OR Expert

I want to build models & solve them faster, with more capacity

Denise
Developer

I want to integrate Oscar's model into the end-to-end solution

Liam
Line of Business

I want to understand how this impacts the bottom line, and save on software license costs

Ethan
Analytics
Entrepreneur

I want to create an optimization-based application and sell it

Typical pain points

Oscar
OR Expert

I need more solve capacity

Denise
Developer

Why can't Oscar set this cloud server up himself? I'm swamped.

Liam
Line of Business

We only solve this huge problem once a year – I cannot invest in a server just for that.

Ethan
Analytics
Entrepreneur

I want to provide a CPLEX-based application to my customers, without requiring them to buy a license.

The six user experiences

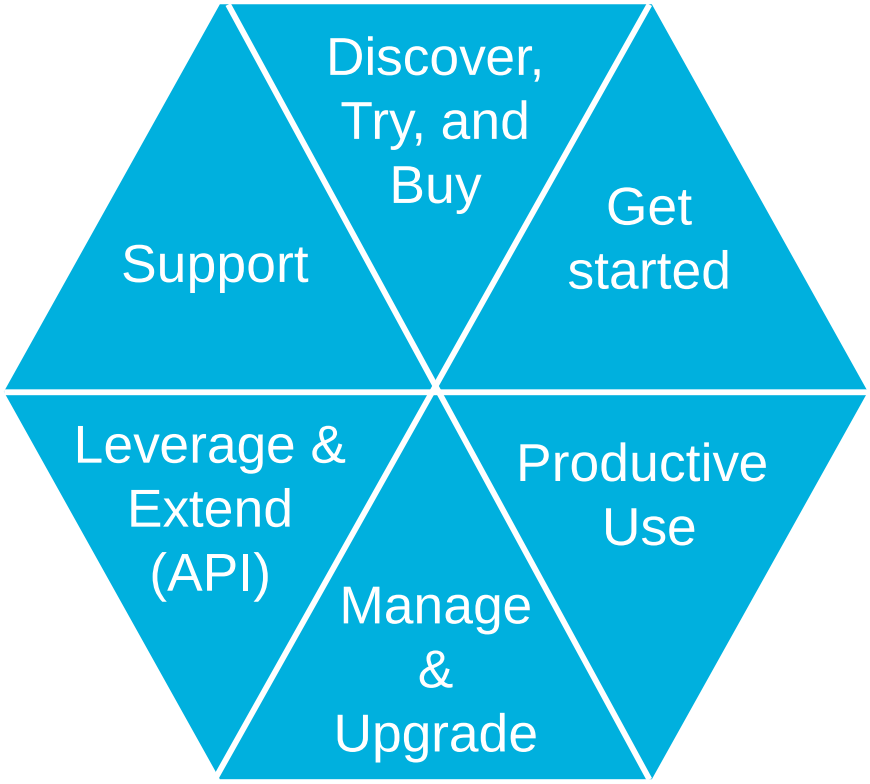

How does Decision Optimization on Cloud change the user experience?

How does DOcloud change the user experience?

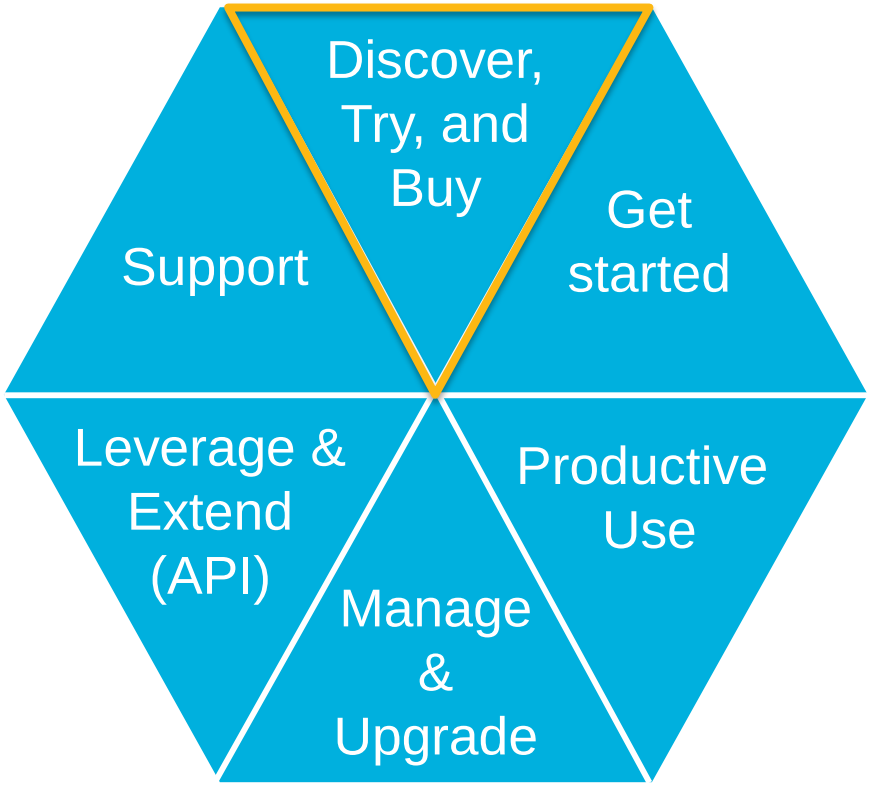

Self-service: buy online with a credit card

Or…if you prefer, you can still talk to an IBM representative as before

Free trial: https://ibm.biz/trydocloud

No installation or download required!

How does DOcloud change the user experience?

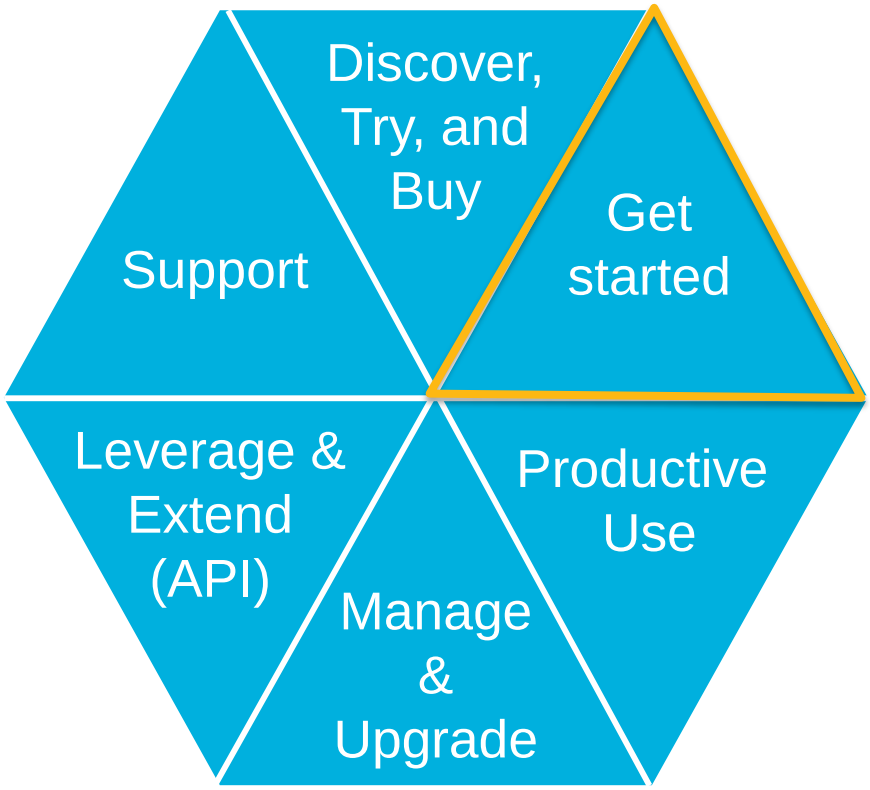

DropSolve: easy drag-and-drop interface to solve models on cloud

Online documentation & samples

How does DOcloud change the user experience?

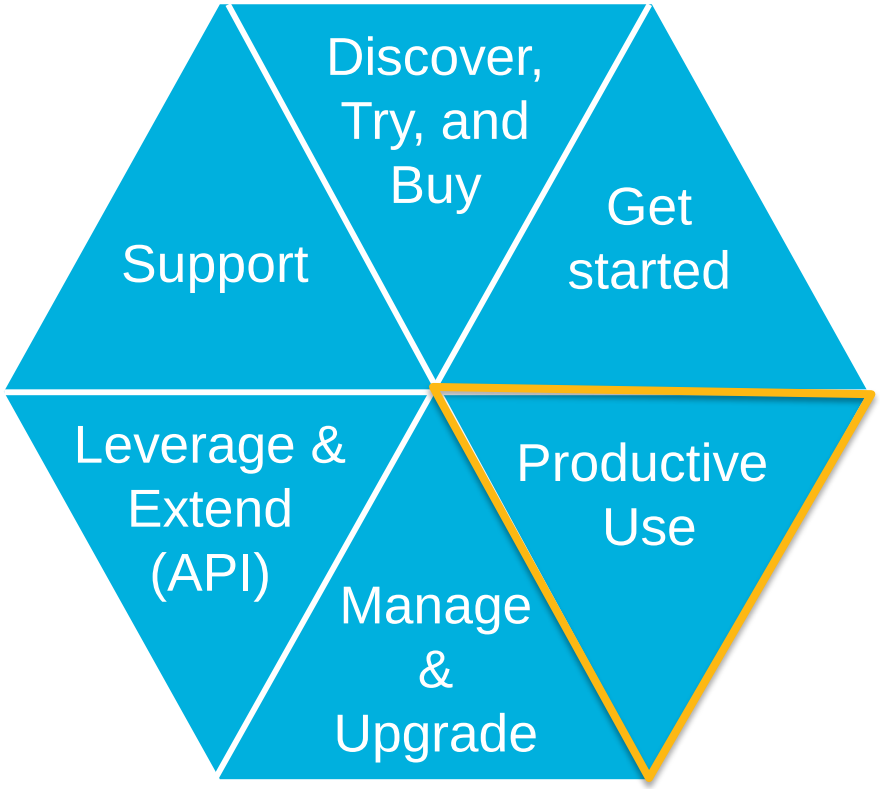

MP and CP on cloud, with choice of VM or bare metal servers on IBM SoftLayer

Pay for what you need – scale up or down according to demand

Infrequent-use excess capacity available on cloud, e.g. annual planning: no investment in additional onsite capacity required

Reduced total cost of ownership!

How does DOcloud change the user experience?

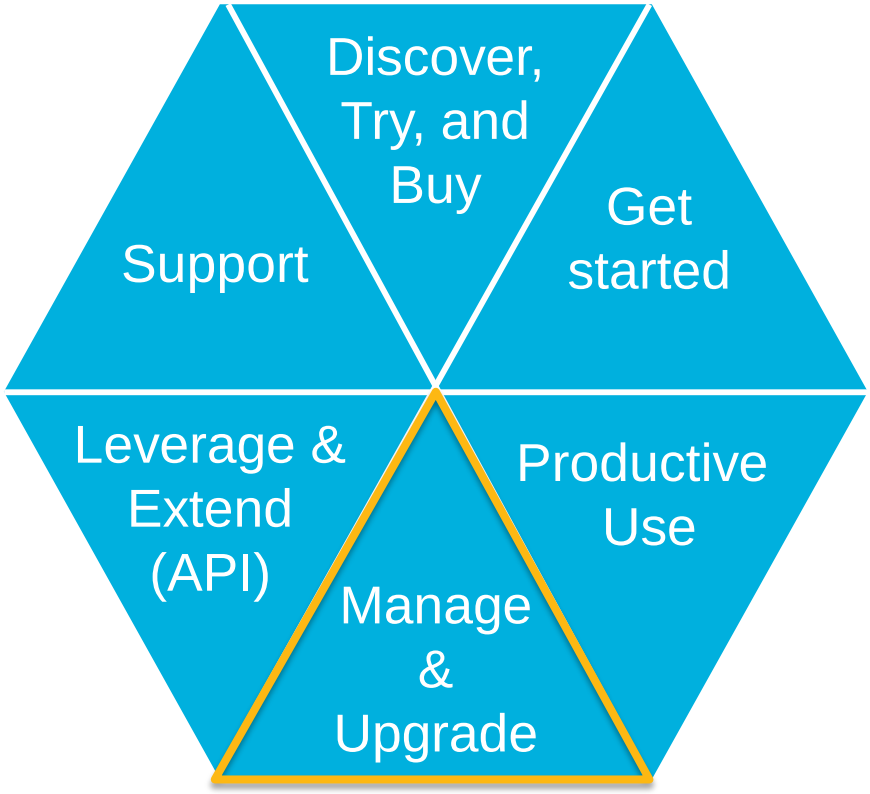

No installation, maintenance, upgrades – all managed by IBM

Continuous delivery means the latest version is always on the cloud!

How does DOcloud change the user experience?

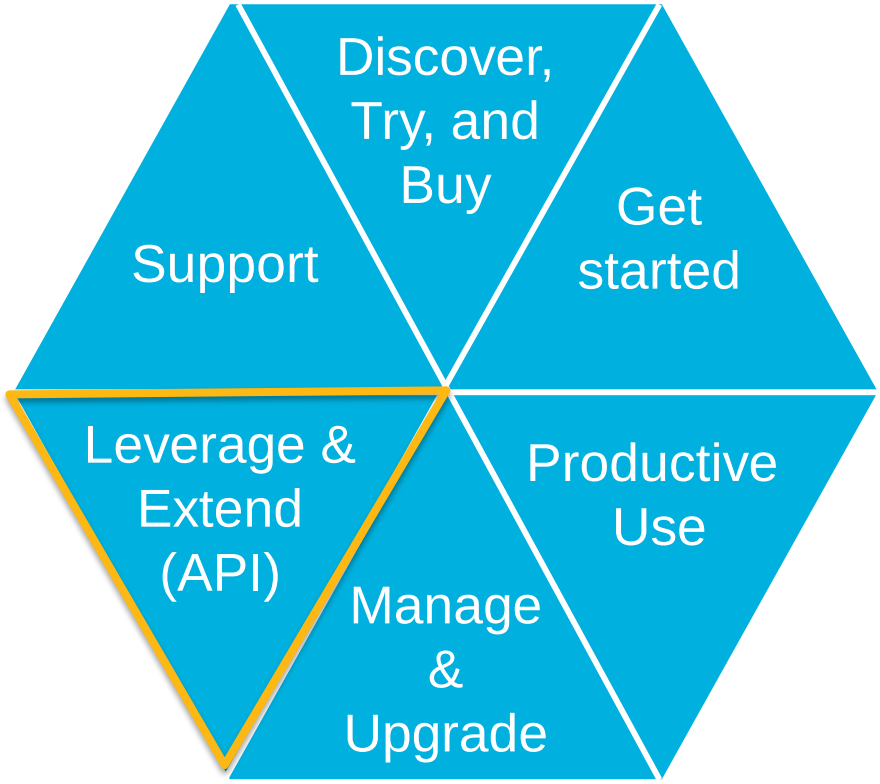

DOcloud API: embed CPLEX Optimizer on cloud in any application (on-prem or cloud)

Option to complement with other IBM Analytics services on cloud (e.g. SPSS scoring, BI)

Option to deploy optimization-based services/applications which embed DOcloud, on IBM Bluemix

How does DOcloud change the user experience?

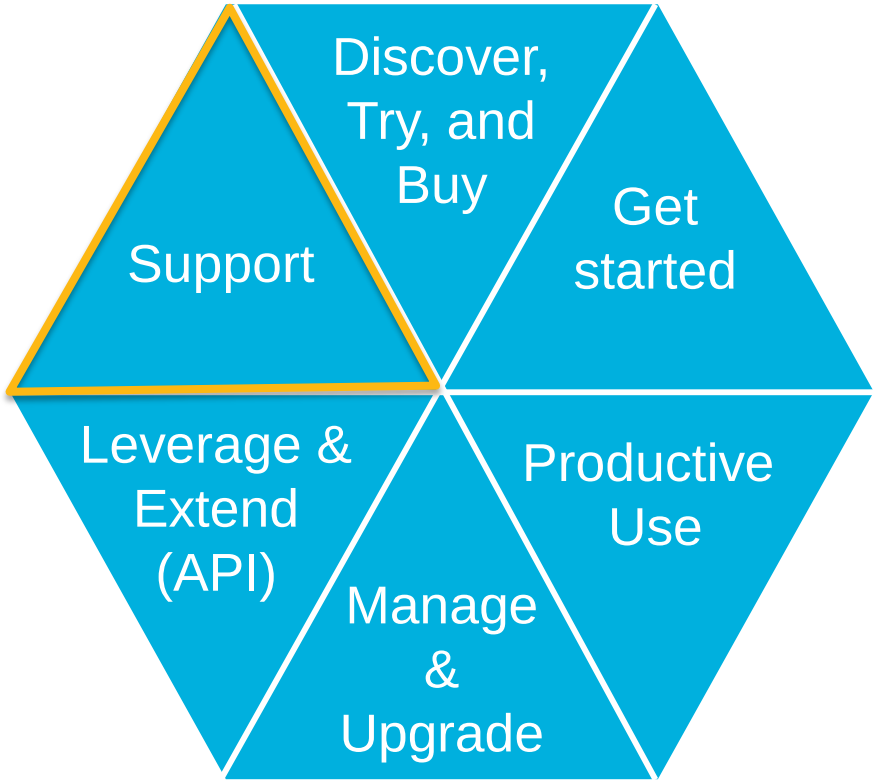

Same support channels as before

Access to wide range of OR, IT, and cloud experts within IBM

developerWorks community: https://developer.ibm.com/docloud/

Summary: Prescriptive Analytics on the Cloud

- Reduced IT spend
  - No license fees
  - No maintenance, upgrades
  - Reduced on-prem hardware spend
  - Reduced IT expertise to set up / maintain server environment
  - Pay-per-use options

- Increased consumability
  - Easy access (no installation / setup)
  - Flexible embedding via Rest APIs
  - Emphasis on communities / collaboration
  - Combine with other analytics and data services for end-to-end applications

Optimization combined with cloud & user-centric analytics platforms open up a new world to LoB, data scientists, OR & IT experts

Overview

- Introduction to IBM Advanced Analytics
- How Decision Optimization on Cloud is changing the Prescriptive Analytics landscape
- The importance of uncertainty: the Uncertainty Toolkit
- Cognitive insights: a vision for consumable Watson Analytics-like optimization

Uncertainty affects all industries…

Plans are out of date as soon as they are created.

No matter how prepared we are, each planning cycle is afflicted by future uncertainties - prices, demand, supply, weather effects…

We don't have a clear vision of possible future scenarios, and their effect on plans.

We get caught in a costly reactive cycle where we fix issues after the fact, instead of anticipating and planning for them.

The Uncertainty Toolkit

- Create plans to anticipate the future
- Hedge against uncertainty in future scenarios
  - Reduce reactive changes
  - Increase margins
  - Reduce response time
- Compare and evaluate alternative plans based on
  - KPIs
  - Risk-reward trade-offs across future scenarios

Uncertainty Toolkit origins

- 2013/14 Joint Program between IBM Research and Decision Optimization
- Goals
  - Increase customer solution resilience, reliability, and stability
  - Improve trust & understanding of optimization technology
- Approach
  - Leverage Decision Optimization & mathematical optimization to hedge against uncertainty (e.g. uncertain demand, task durations, prices, resource availability)
  - A user-friendly toolkit as plug-in to IBM Decision Optimization Center
- Cross-industry applications

5 steps Uncertainty Toolkit process

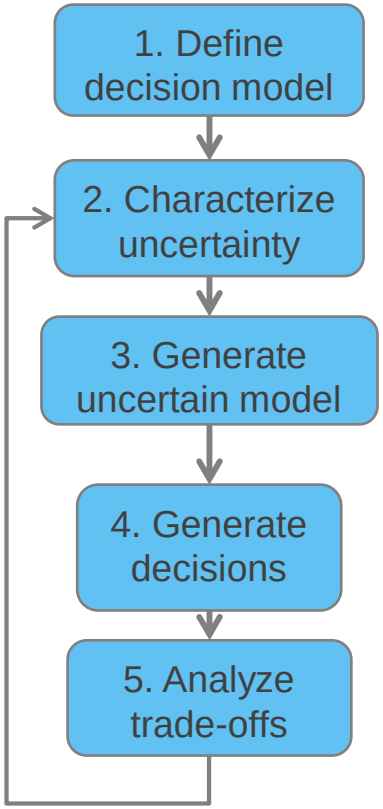

- Create optimization model with IBM CPLEX Studio
- Some modeling skill required, or existing assets
- Embed in IBM Decision Optimization Center

"Steve" the IT expert, & "Keith" the OR consultant

- OR consultant's "wizard": 7 screens
- Defines uncertainty, scenario generation, risk measures

- Built-in automated reformulation, based on steps 1 and 2
- No modeling knowledge required
- "Robustification" (make the original model robust to change)

- Business user's "wizard"
- Automated solution generation
- Automated scenario comparison

"Anne" the business user

- Built-in visual analytics
- Analyze KPI trade-offs across multiple plans & scenarios

27

Uncertainty Toolkit: automated reformulations

Robust / Stochastic approach	Applicable model types	Resulting model types	Uncertainty characterization	Restrictions
Single-stage penalty approach *(Mulvey et al., 1995)*	LP	LP (or QP)	Scenarios (finite)	No uncertain data in objective function
	MILP	MILP (or MIQP)		
Two-stage penalty approach *(Mulvey et al., 1995)*	LP	LP (or QP)	Scenarios (finite)	No uncertain data in objective function
	MILP	MILP (or MIQP)		
Multistage Stochastic (e.g. *King & Wallace, 2012*)	LP	LP	Scenarios (finite)	None
	MILP	MILP		
Safety margin approach with ellipsoidal uncertainty sets (*Ben-Tal & Nemirovski, 1999*)	LP	QCP	Range	No uncertain data in standalone parameters or equality constraints
	MILP	MIQCP		
Safety margin approach with polyhedral uncertainty sets (*Bertsimas & Sim, 2004*)	LP	LP	Range	No uncertain data in standalone parameters or equality constraints
	MILP	MILP		
Extreme Scenario approach (*Lee, 2014*)	LP	LP	Range	No uncertain data in variable coefficients
	MILP	MILP		
Distributionally robust reformulation (*Mevissen et al., 2013*)	LP	LP	Scenarios	Uncertainty in standalone parameters handled as penalty term in objective
	MILP	MILP		

Uncertainty Toolkit: automated reformulations

Robust / Stochastic approach	Applicable model types	Resulting model types	Uncertainty characterization	Restrictions
Single-stage penalty approach *(Mulvey et al., 1995)*	LP	LP (or QP)	Scenarios (finite)	No uncertain data in objective function
	MILP	MILP (or MIQP)		
Two-stage penalty approach *(Mulvey et al., 1995)*	LP	LP (or QP)	Scenarios (finite)	No uncertain data in objective function
	MILP	MILP (or MIQP)		
Multistage Stochastic (e.g. *King &*	LP	LP	Scenarios (finite)	None
Safety mar uncertainty (*Ben-Tal &*				lone
Safety mar uncertainty sets (*Bertsimas & Sim, 2004*)	MILP	MILP		lone parameters or equality constraints
Extreme Scenario approach (*Lee, 2014*)	LP	LP	Range	No uncertain data in variable coefficients
	MILP	MILP		
Distributionally robust reformulation (*Mevissen et al., 2013*)	LP	LP	Scenarios	Uncertainty in standalone parameters handled as penalty term in objective
	MILP	MILP		

Q: How do I know which of these methods to use?

A: The Uncertainty Toolkit will decide automatically based on your input into the Consultant's Wizard

Uncertainty Toolkit Decision Tree (automated)

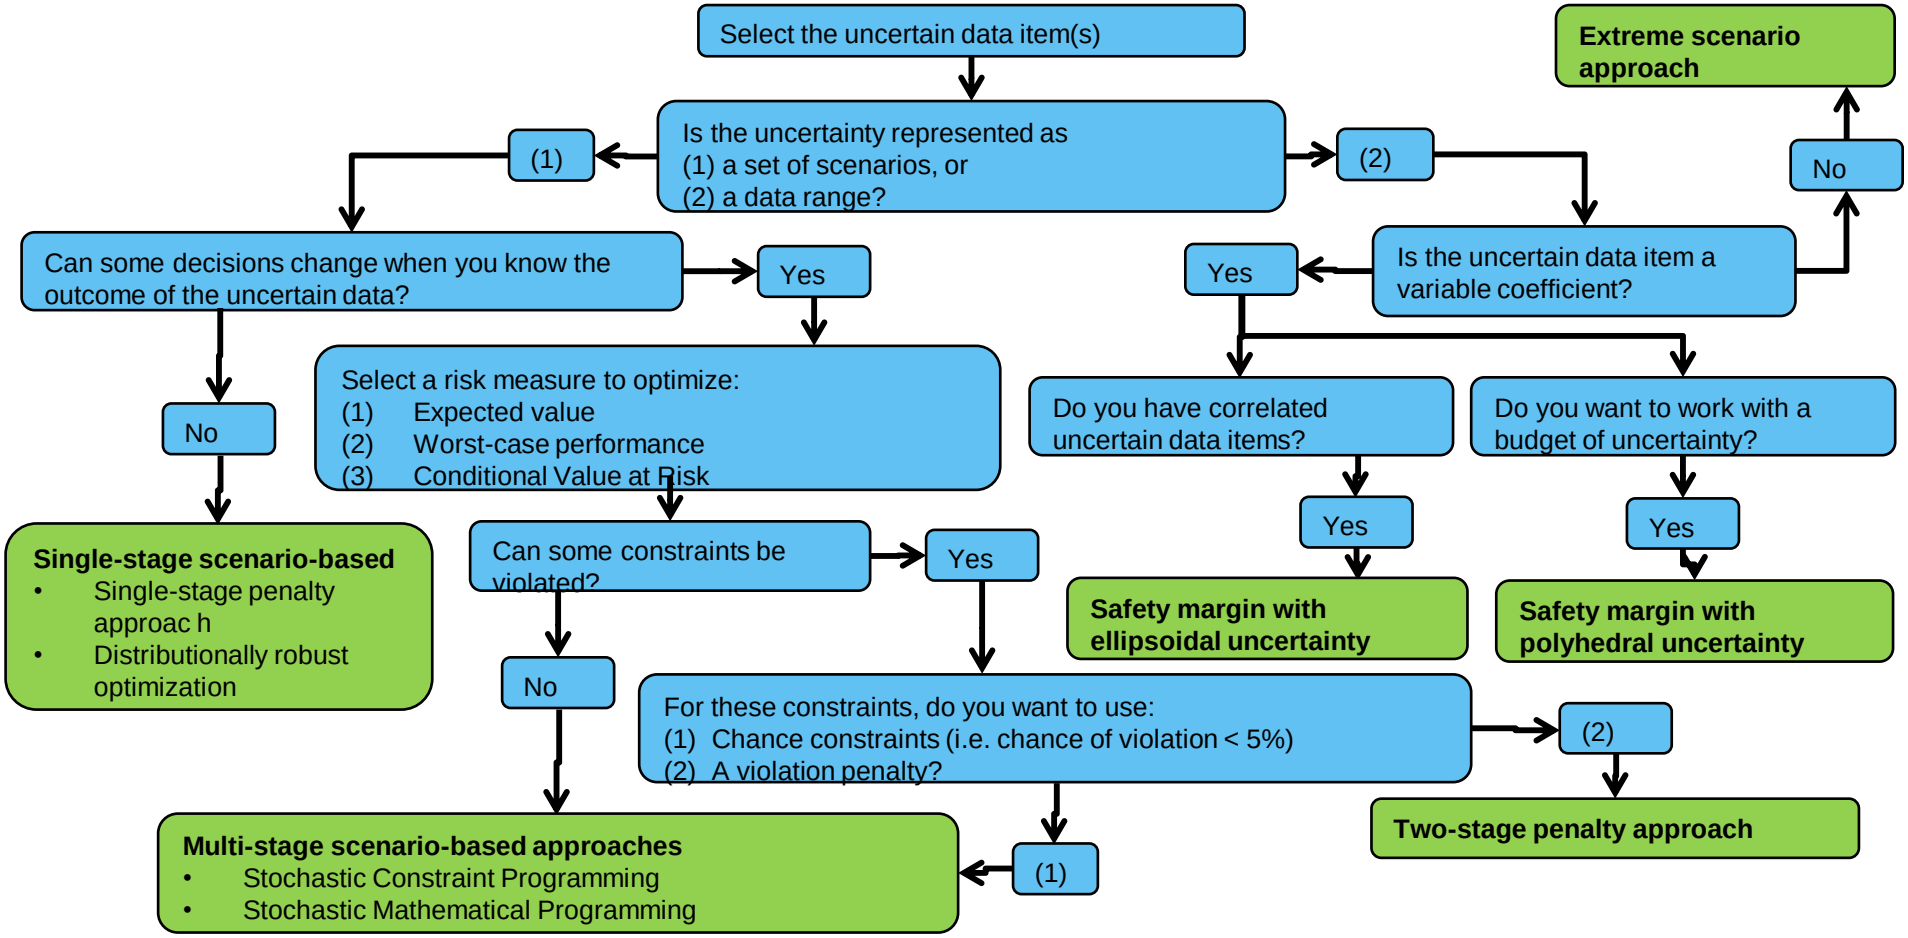

Uncertainty Toolkit Decision Tree (automated)

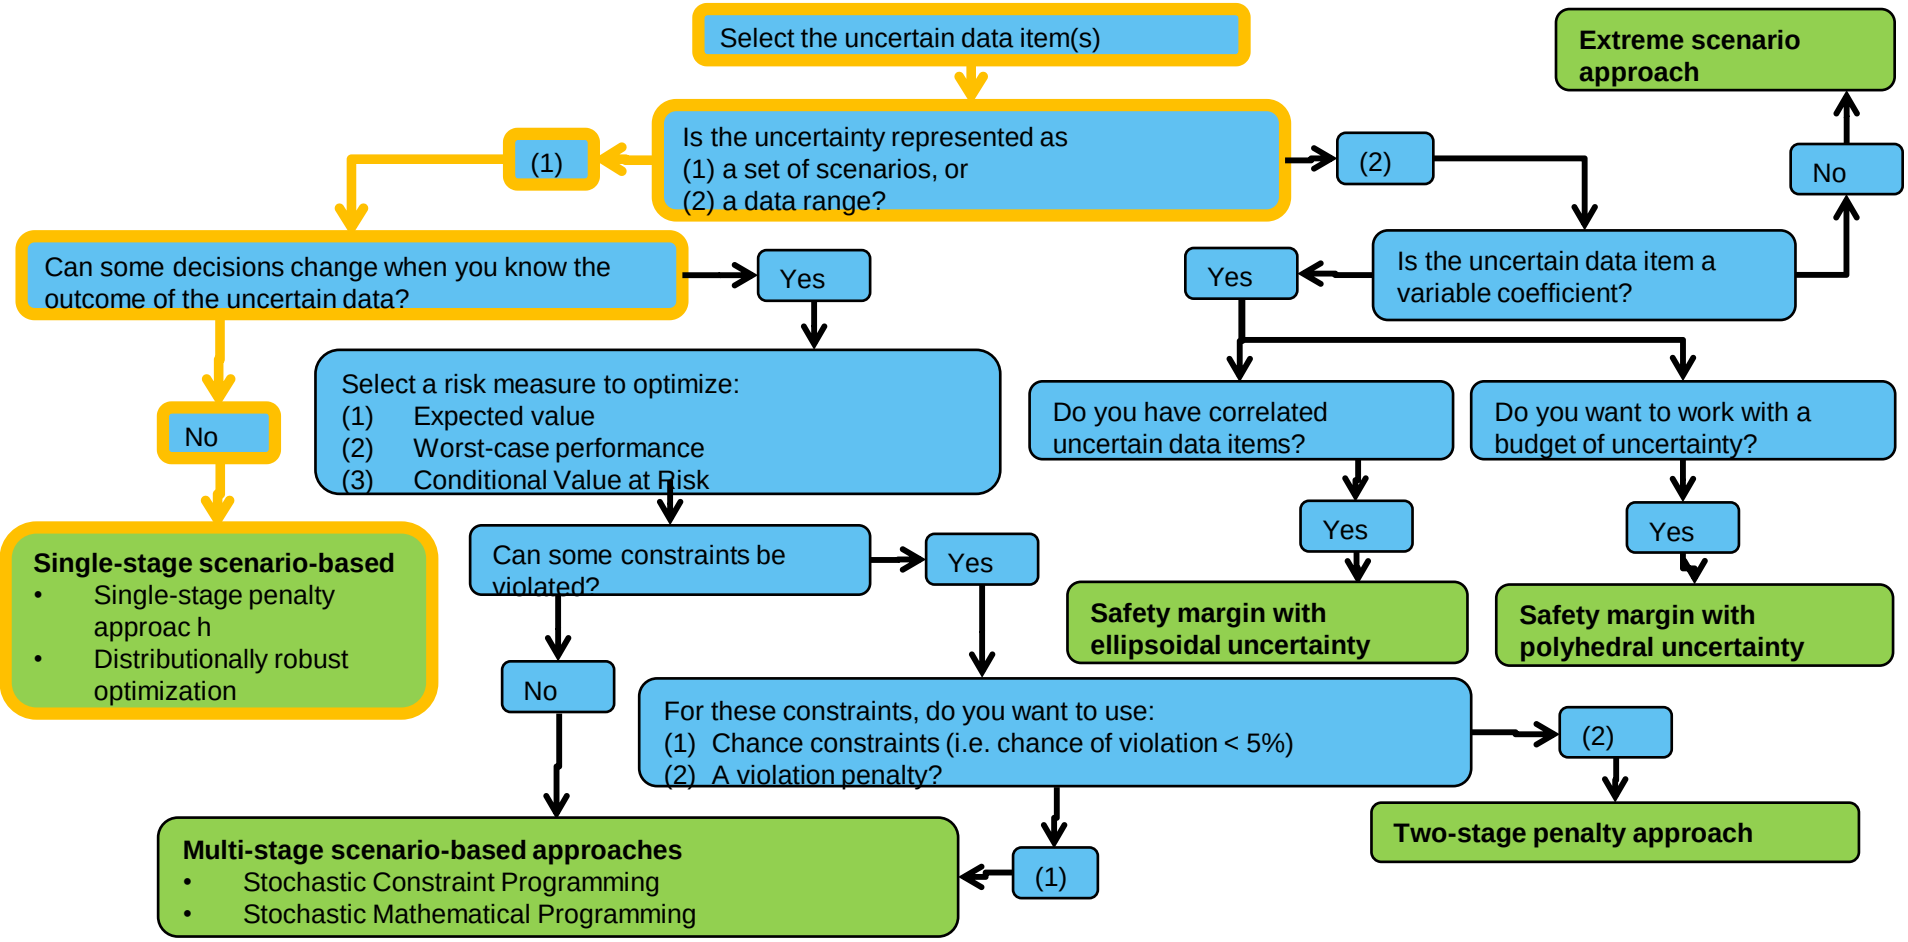

Example: Automated model reformulation for stochastic CP

```
18
19 dvar interval task[i in 1..n] size TaskDuration[i];
20 dvar sequence seq in task types all(i in 1..n) i;
21 dexpr int station[i in 1..n] = startOf(task[i]) div c;
22
23 minimize 1+max(i in 1..n) station[i];
24 subject to {
25   noOverlap(seq, Setups);
26   forall (p in Precedences)
27     endBeforeStart(task[p.pred], task[p.succ]);
28   forall (i in 1..n)
29     station[i] == (endOf(task[i]) - 1) div c;
30 };
31
```

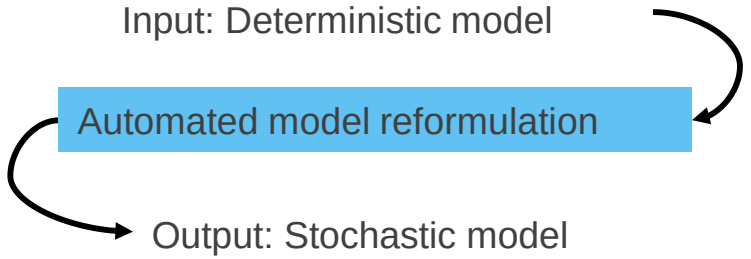


```
39
40 dvar interval task[i in 1..n][s in Scenarios] size TaskDuration[i][s];
41 dvar sequence seq[s in Scenarios] in all(i in 1..n) task[i][s] types all(i in 1..n) i;
42
43 dexpr int station[i in 1..n][s in Scenarios] = startOf(task[i][s]) div c;
44
45 minimize sum(s in Scenarios) Probability[s]*(1 + max(i in 1..n) station[i][s]);
46 subject to {
47
48   forall (s in Scenarios) {
49     noOverlap(seq[s], Setups);
50     forall (p in Precedences)
51       endBeforeStart(task[p.pred][s], task[p.succ][s]);
52     forall (i in 1..n)
53       station[i][s] == (endOf(task[i][s]) - 1) div c;
54     }
55
56   forall (s in 1..(S-1), i in 1..n) {
57     typeOfNext(seq[s], task[i][s], -1) == typeOfNext(seq[s+1], task[i][s+1], -1);
58   }
59 };
60
61
```

IBM Confidential

Case study: Water treatment/distribution energy cost reduction

- Big picture: Cork County Council must reduce energy consumption by 20% by 2020
- 95% of this utility's water-related energy costs due to pump operations
- New dynamic energy pricing schemes leverage renewables (wind energy)
- Trade-off: Cleaner energy at lower prices, but uncertainty in price due to
  - Wind uncertainty
  - Network outages
  - Other weather conditions
- Goal: Schedule pumps leveraging dynamic prices, while hedging against uncertainty in price prediction

Simplified network (illustration purposes)

Goal: Optimize pump schedules to minimize (uncertain) energy costs while meeting demand and respecting plant and network constraints*

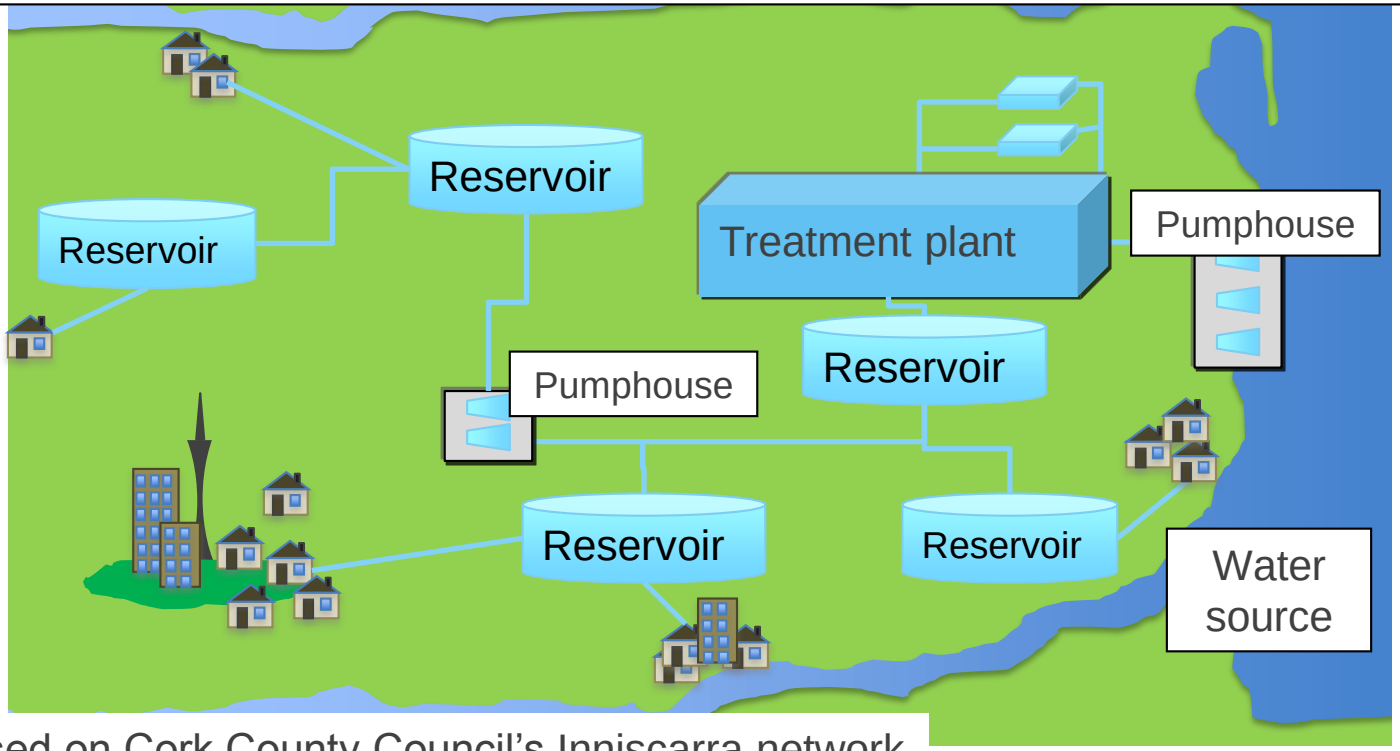

* Based on Cork County Council's Inniscarra network

Uncertainty in price prediction

- Forecasted (D-1) post ante price from supplier
  - Considers forecasted demand based on weather, special events, wind, etc.
- Actual (D+4) price charged 4 days after the event
  - Forecasted (D-1) and Settled price (D+4) can differ due to changes in predicted wind energy availability, weather, and unpredicted grid events

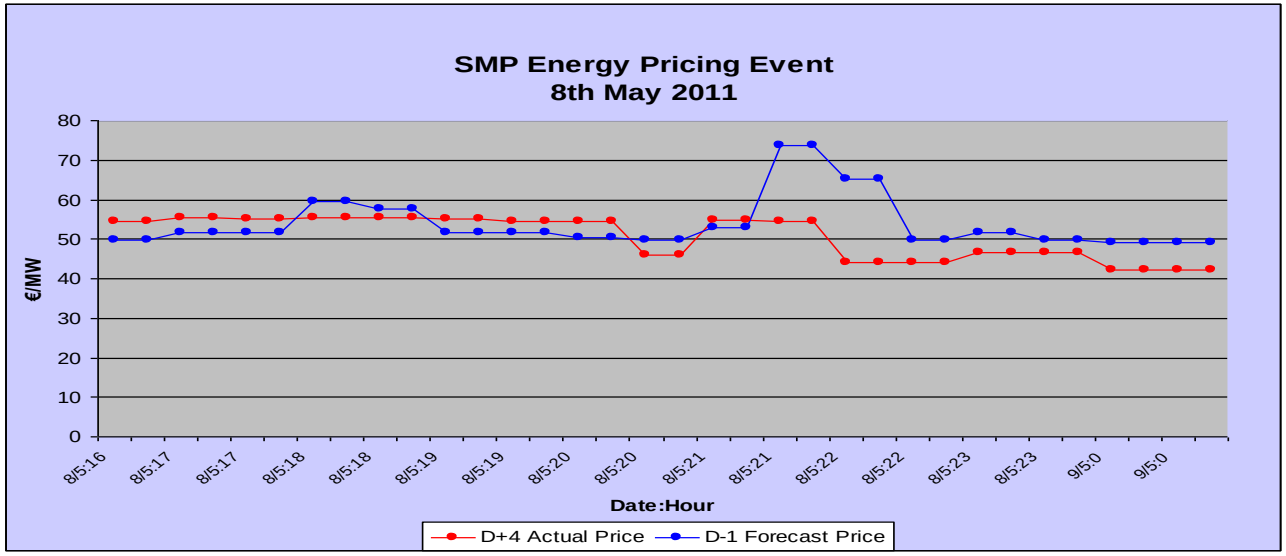

Question:
Should utility switch to a dynamic pricing scheme?
Step 1: Prove dynamic pricing benefits
Step 2: Prove optimization benefits
Step 3: Deal with uncertainty

Step 1: Define decision model

- Define objective, decisions, constraints (mathematical modeling skill required)
  - Objective: minimize energy costs from pump operations
  - Decisions: when to switch pumps on/off (decided every 30 minutes for 24 hours in advance)
  - Constraints: satisfy tank levels, pump operation rules, customer demand, network constraints
- Model using CPLEX Studio, assuming certain data (“deterministic” model)

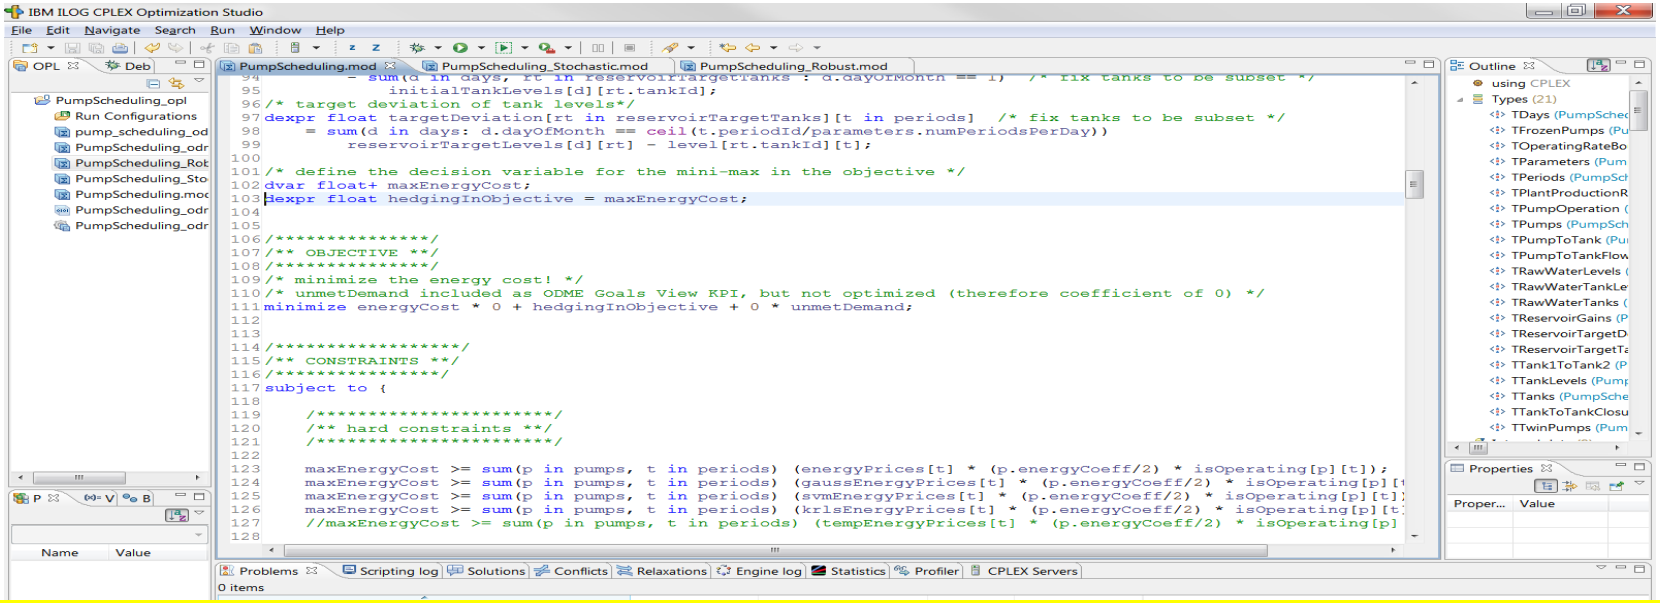

Note: When data is fairly certain, deterministic models are sufficient to provide significant benefit

Decision model in **Decision Optimization Center prototype**: Predicted energy cost savings from pump optimization and (known) dynamic energy prices

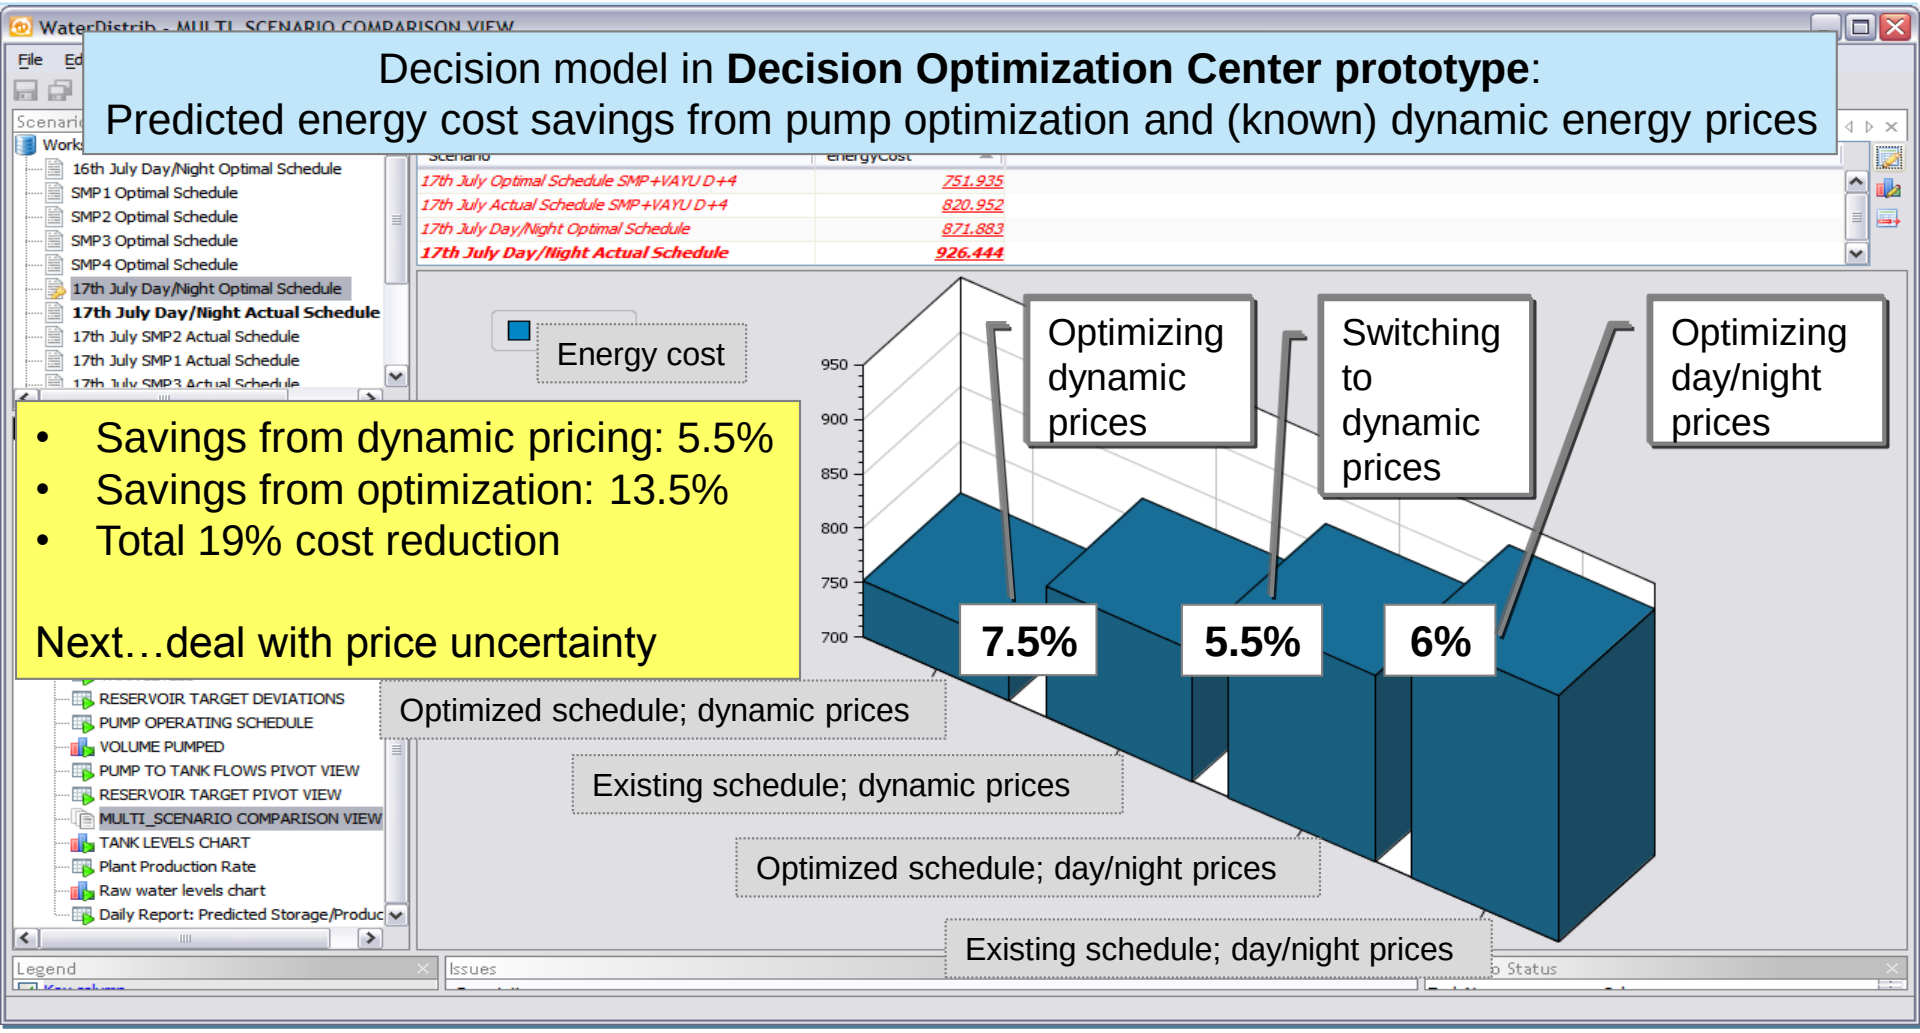

- Savings from dynamic pricing: 5.5%
- Savings from optimization: 13.5%
- Total 19% cost reduction

Next…deal with price uncertainty

Step 2: Characterize uncertainty

- Price scenarios, with likelihoods:
  - From energy provider
  - From IBM Research forecasts

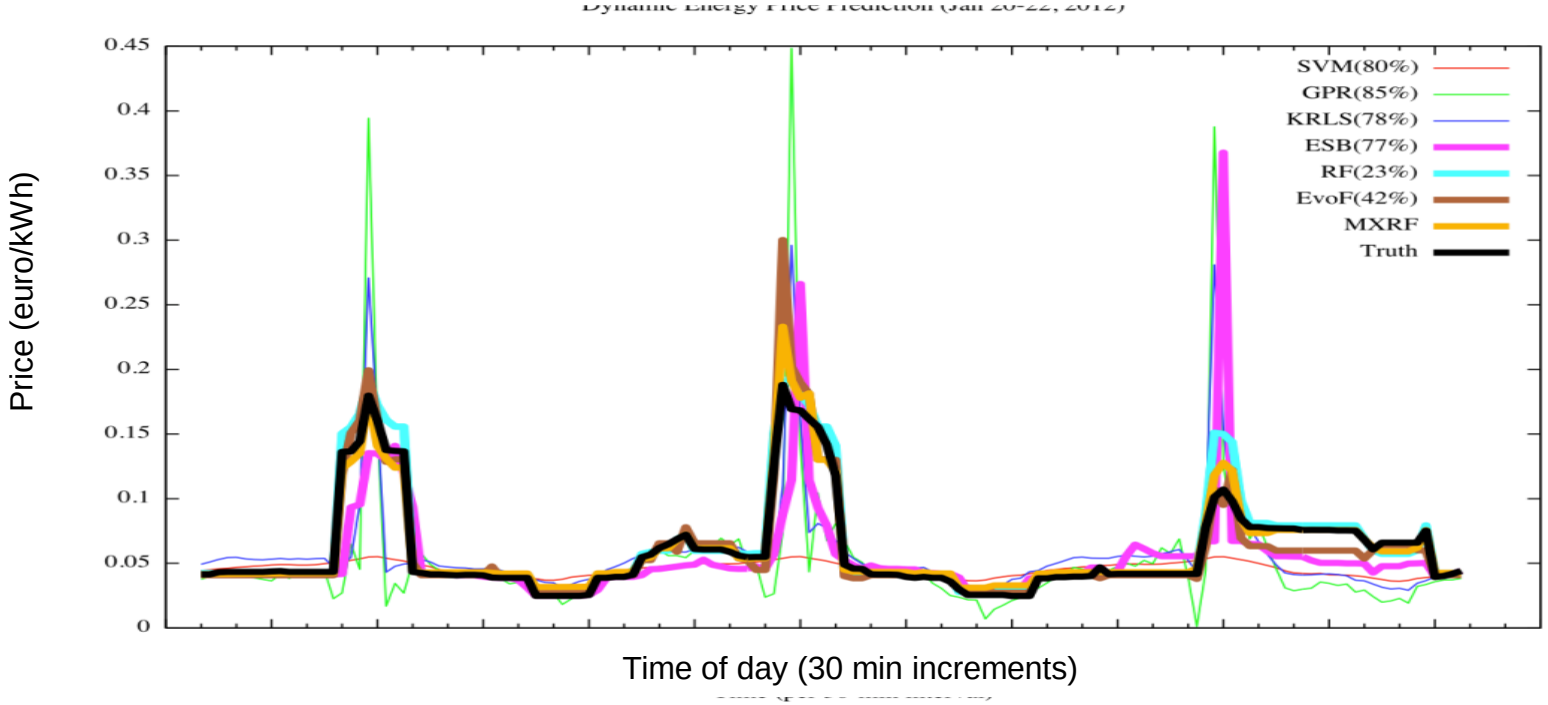

Step 2: Uncertainty Toolkit wizard for consultant input (e.g.)

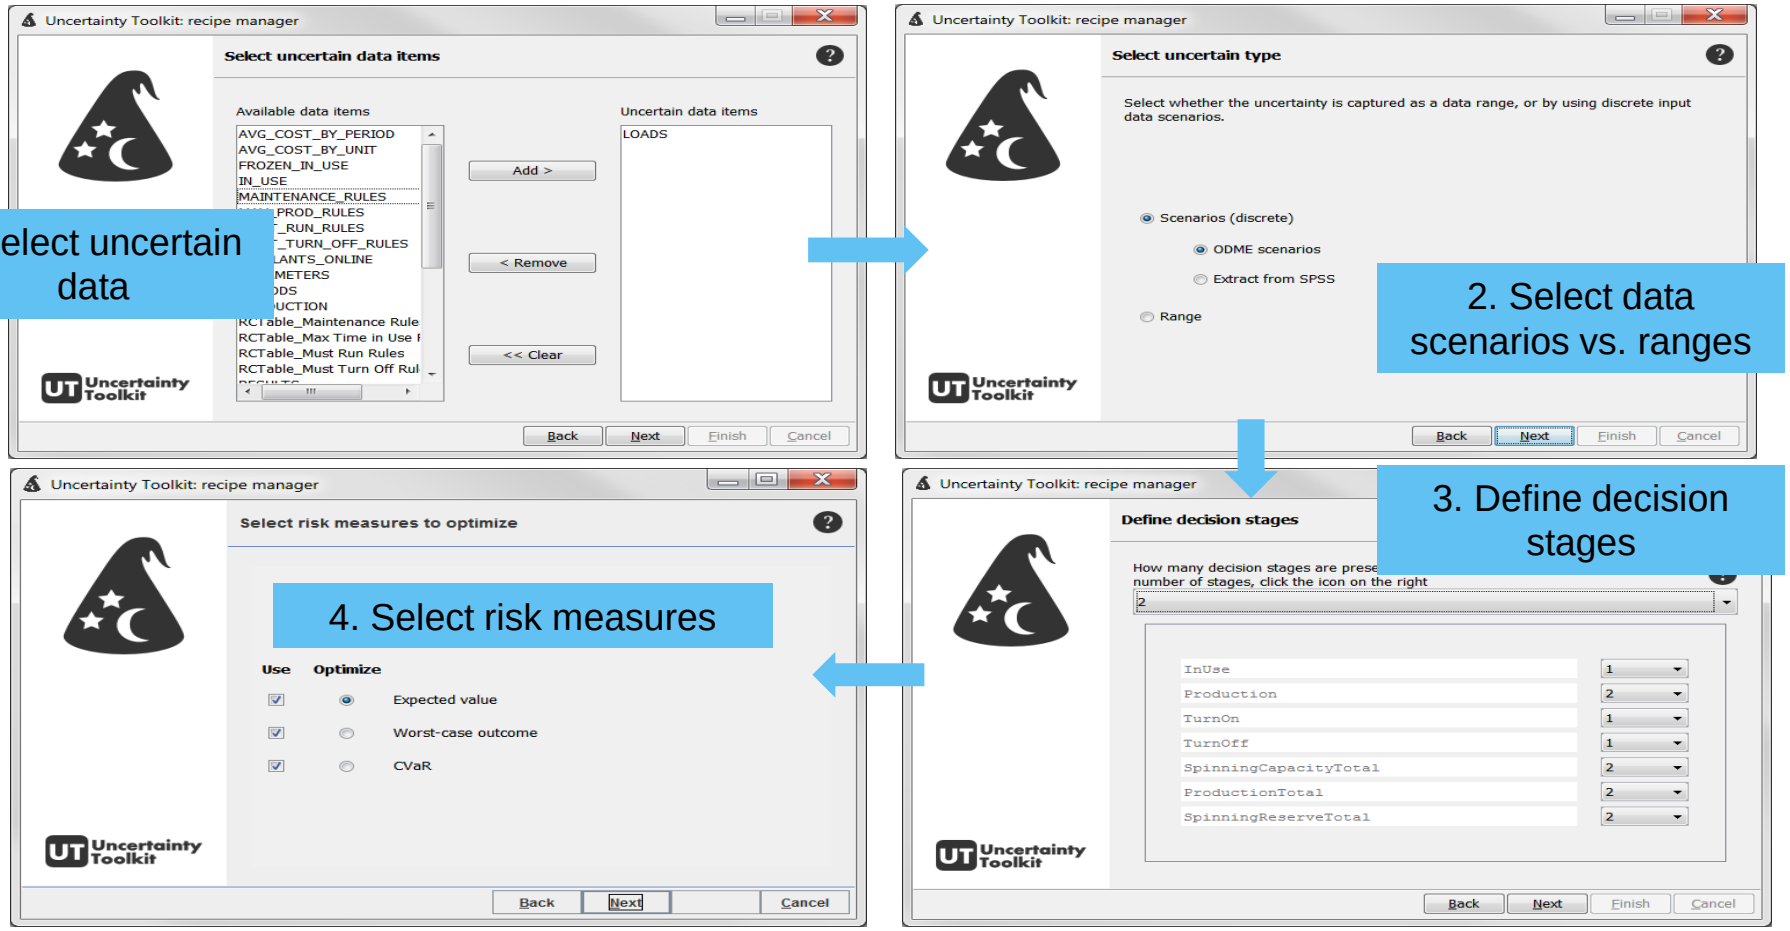

Step 3: Generate uncertain model

- Uncertainty Toolkit automatically generates the uncertain model(s) depending on choices in Steps 1 and 2

- Uncertain models are typically classified as
  - "Robust": hedging against worst case outcome(s)
  - "Stochastic": optimizing for expected outcome(s)
  - If choice unclear, use both & visualize trade-offs

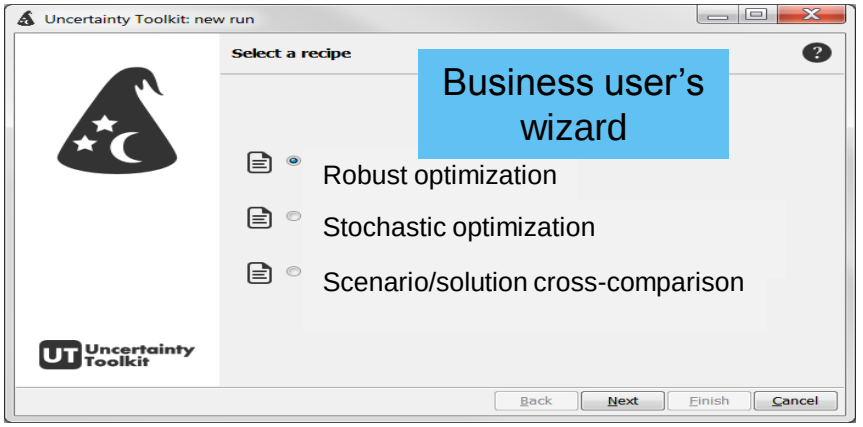

Step 4: Generate plans

- Uncertainty Toolkit generates multiple solutions (deterministic, robust, stochastic)
- Uncertainty Toolkit automatically does solution-scenario cross-comparison
  - What is the impact of change on each plan

Step 5: Analyze trade-offs

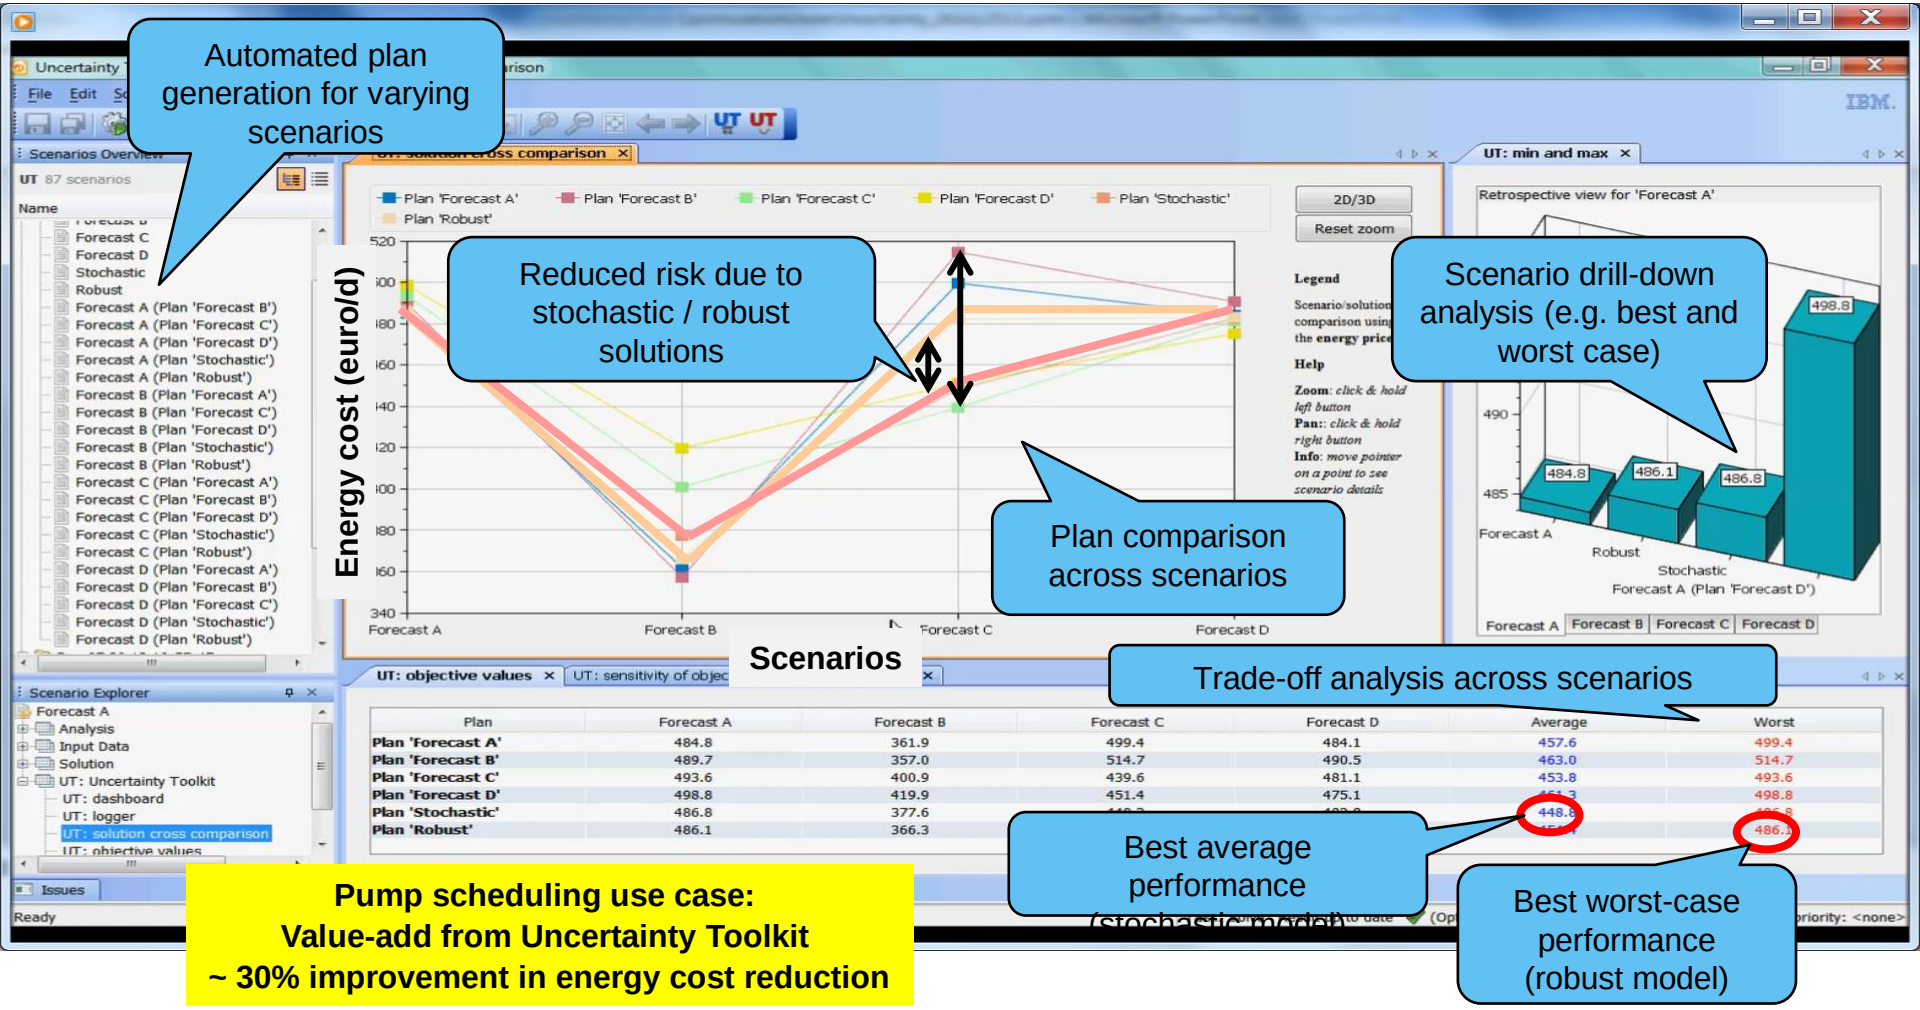

Example: pressure management in water distribution networks

- Problem: Leakage (non-revenue water) leads to 5 – 60% of treated water lost
- Existing solution:
  - 1% pressure reduction ~ 1% leakage reduction
  - Place and set pressure reducing valves to minimize leakage for given demand pattern (deterministic plans)
- New challenge: demand uncertain
  - Unexpected short large draw-offs by industrials
  - Variations depending on time of day / week
  - Deterministic plans not robust – often infeasible & sub-optimal when demand changes
- Uncertainty Toolkit creates robust plans
  - Stable (robust) valve settings / placement (no need for frequent changes as demand changes)
  - Leakage reduction across demand scenarios

Dublin's Chapelizod network: optimal robust valve placement

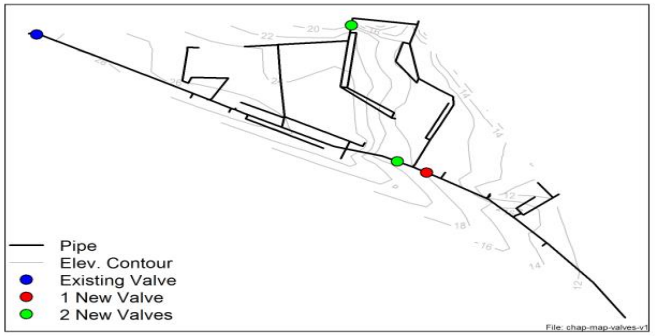

Benefits of Uncertainty Toolkit – pressure management use case

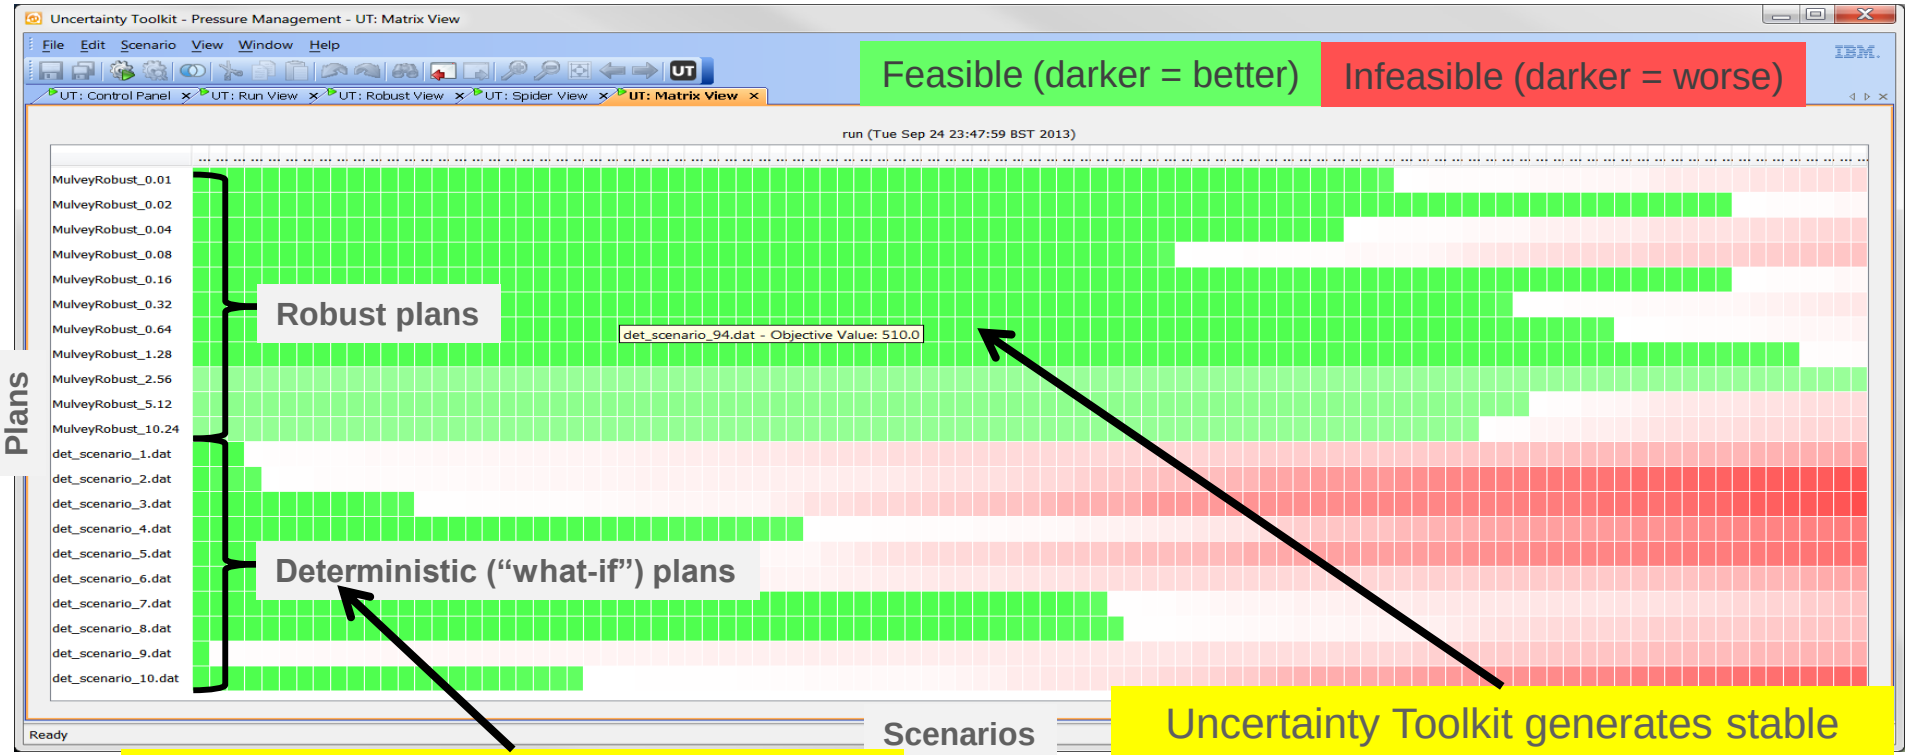

"What-if" plans could become unstable when data changes

Uncertainty Toolkit generates stable plans across scenarios, by hedging against uncertainty

Benefits of Uncertainty Toolkit – pressure management use case

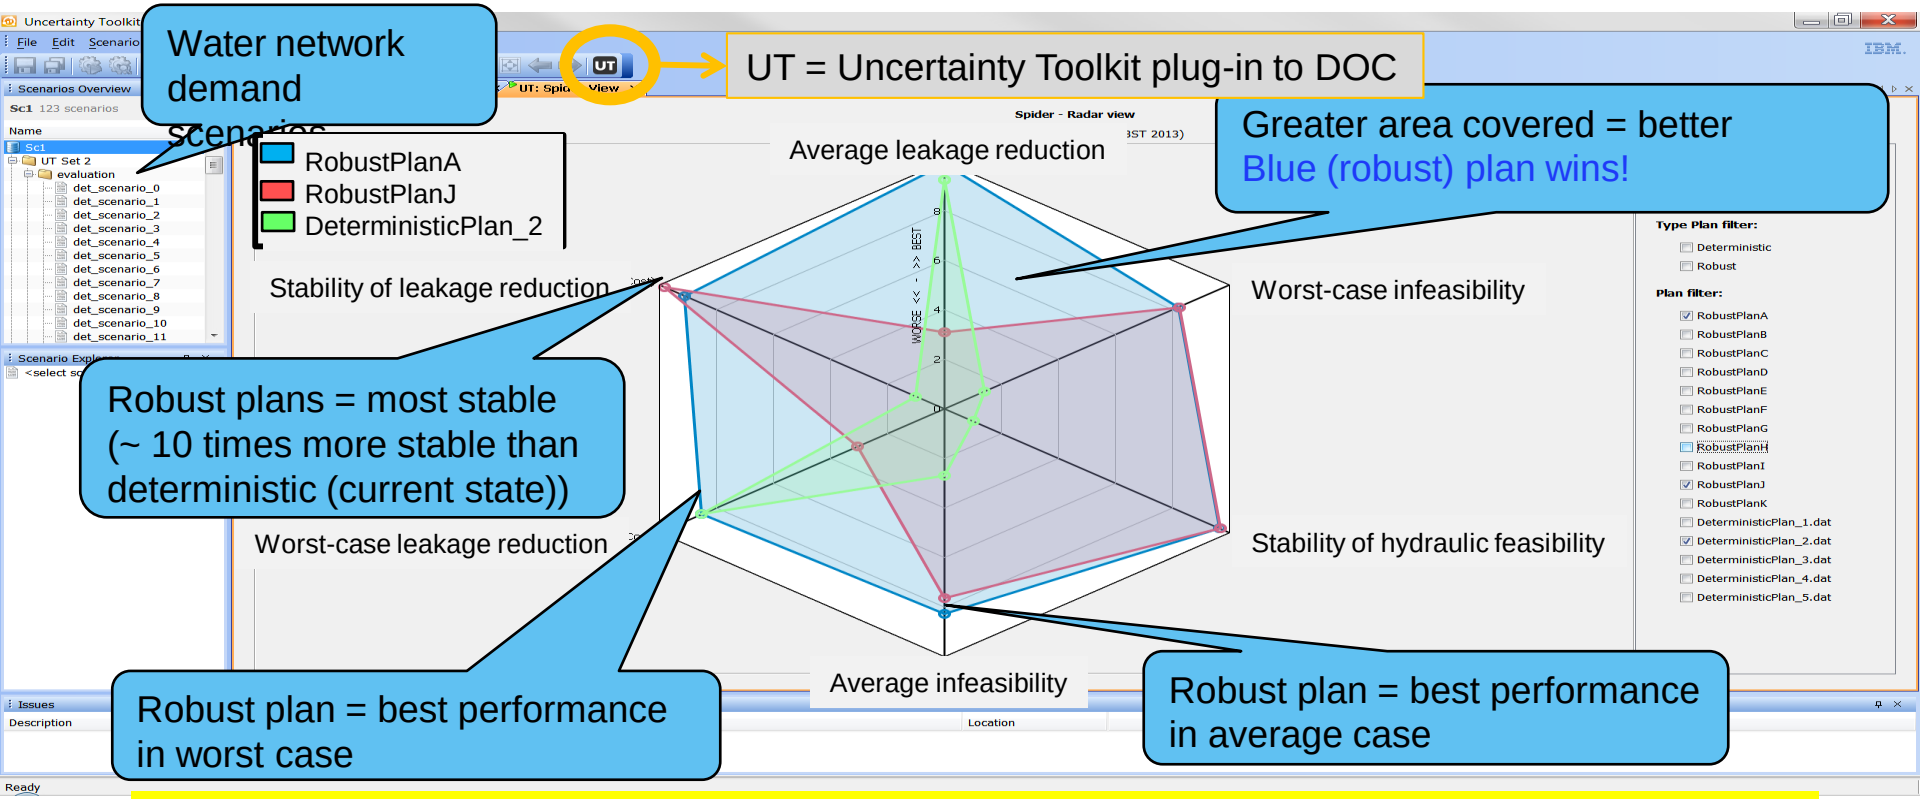

Pressure management use case:
Water network operational decisions 10 times more stable than current state, continue to perform well when data changes (i.e. “robust” plans)

Benefits of Uncertainty Toolkit – pressure management use case

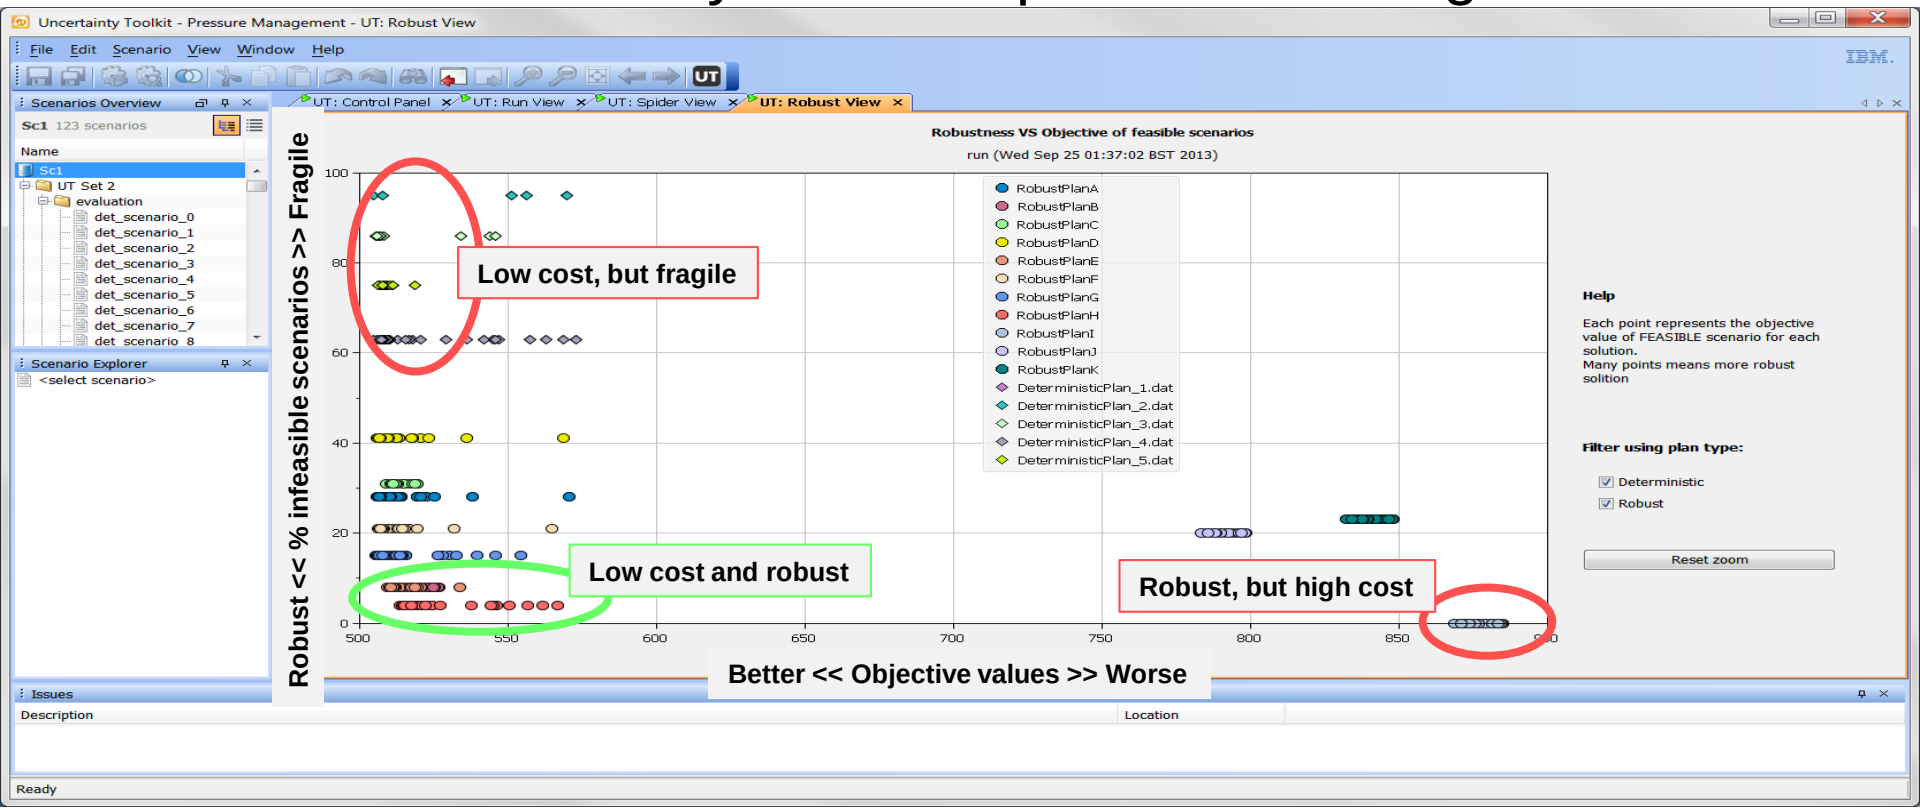

Visualization of trade-off: robustness vs. cost

Benefits of Uncertainty Toolkit – pressure management use case

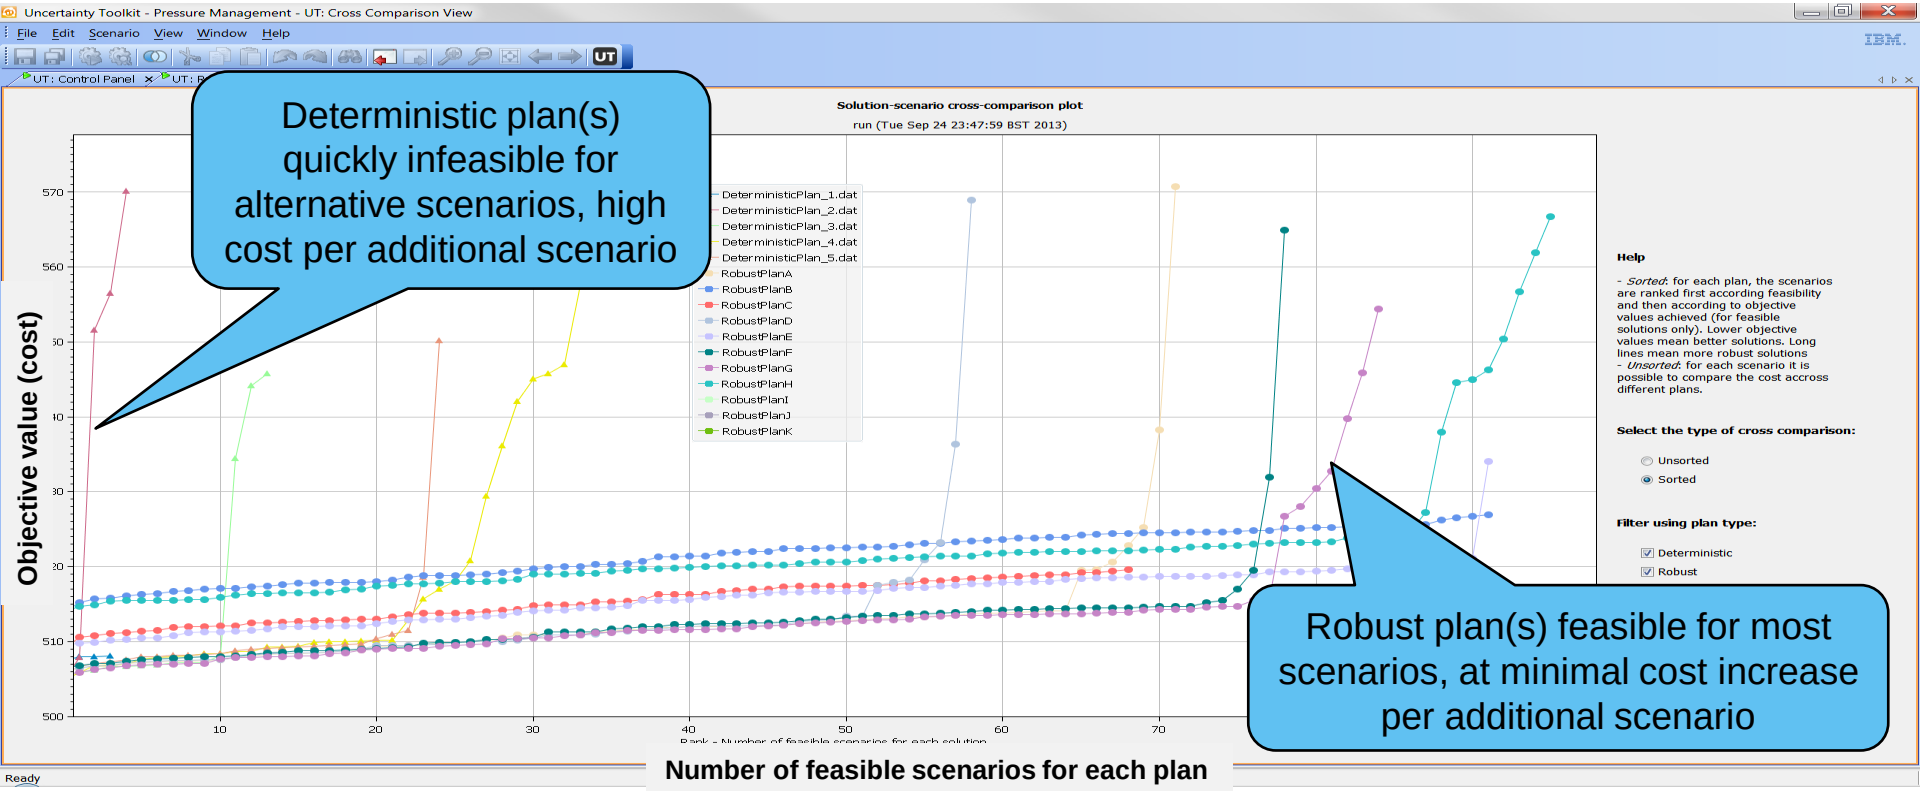

Effect on feasibility (robustness) by increasing cost

Completed use case demos

- Mulprod: Toy problem based on pasta production
- UnitCommitment: Extension of the Unit Commitment Demo distributed with DOC
- Replenish: Spare part replenishment scheduling
- Staff scheduling (hospital and airline)
- Power: Investment planning for power distribution networks
- Portfolio: Portfolio optimization
- Water network pressure management (network design and planning)
- Water network pump scheduling
- Task scheduling (only CP)
- Facility location

References

- A. J. King and S. W. Wallace. *Modeling with Stochastic Programming*. Springer, 2012.
- J. M. Mulvey, R. J. Vanderbei, and S. A. Zenios, “Robust Optimization of Large-Scale Systems,” *Operations Research*, vol. 43, no. 2, pp. 264-281, 1995.
- A. Ben-Tal and A. Nemirovski, “Robust solutions of uncertain linear programs,” *Operations research letters*, vol. 25, no. 1, pp. 1-13, 1999.
- D. Bertsimas and M. Sim, “The price of robustness,” *Operations Research*, vol. 52, no. 1, pp. 35-53, 2004.
- C. Lee, “Extreme scenario approach,” forthcoming paper, to be presented at the *International Federation of Operational Research Societies (IFORS) conference,* Barcelona, 2014.
- M. Mevissen, E. Ragnoli, and Y. Y. Jia, “Data-driven Distributionally Robust Polynomial Optimization,” in *Advances in Neural Information Processing Systems,* Nevada, United States, pp. 37-45, 2013.

Overview

- Introduction to IBM Advanced Analytics
- How Decision Optimization on Cloud is changing the Prescriptive Analytics landscape
- The importance of uncertainty: the Uncertainty Toolkit
- Cognitive insights: a vision for consumable Watson Analytics-like optimization

IBM Watson Analytics: connect stakeholders, simplify analytics

Analytics in the hands of a broad range of users

Easier data access and refinement

Cloud for agility and speed

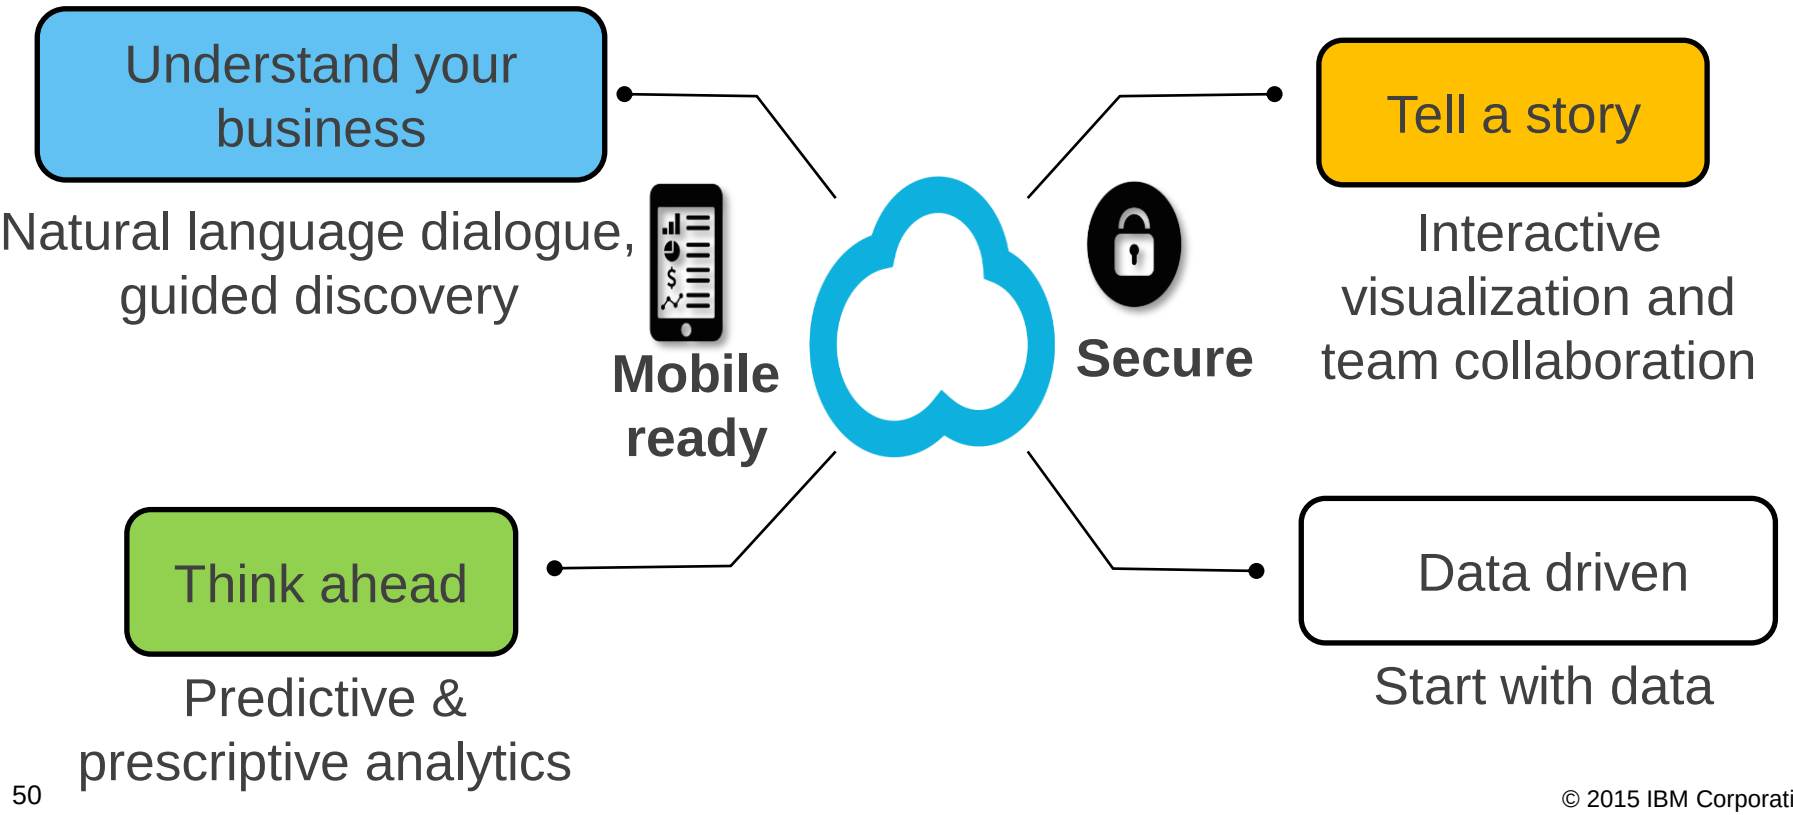

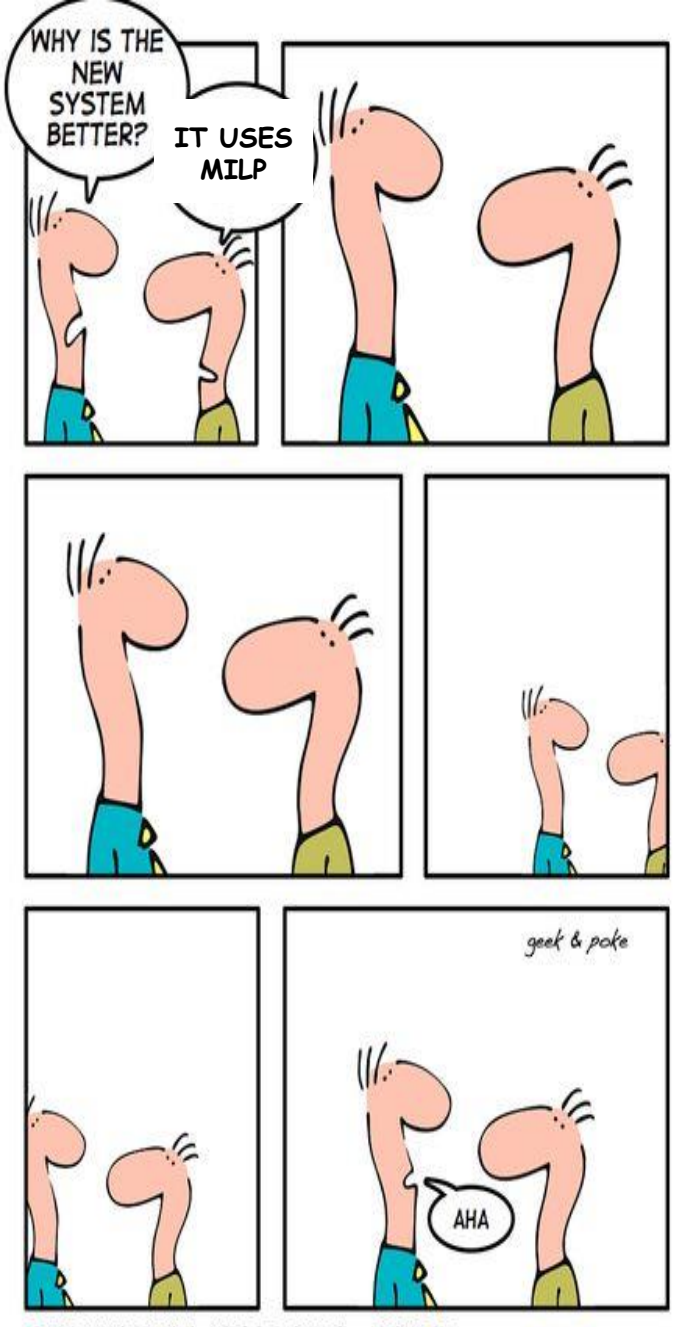

TO BE SUCCESSFUL YOU HAVE TO CONVINCE THE BUSINESS

How can Watson Analytics make mathematical optimization more consumable?

Bridget's story

Bridget Hart
Sales Manager

What is the revenue breakdown by region?

What are the predicted sales per region?

How should I create territories to maximize quota achievement fairly?

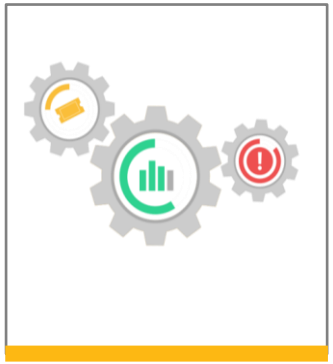

Descriptive Analytics

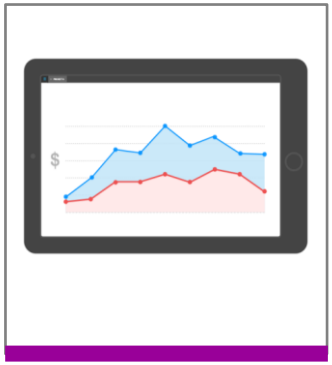

Predictive Analytics

Prescriptive Analytics

Bridget's story

Bridget Hart
Sales Manager

What is the revenue breakdown by region?

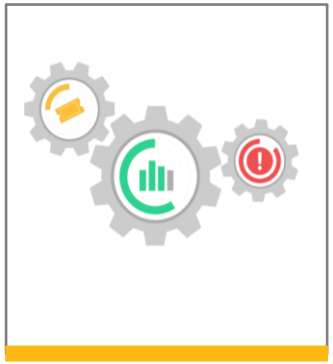

Descriptive Analytics

What are the predicted sales per region?

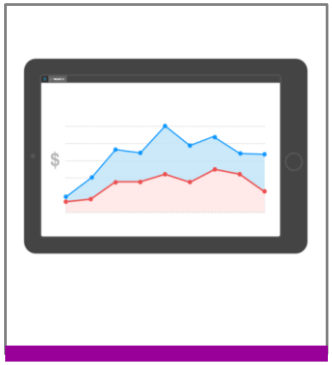

Predictive Analytics

Today's demo

How should I create territories to maximize quota achievement fairly?

Prescriptive Analytics

IBM Watson Analytics and Decision Optimization Demo

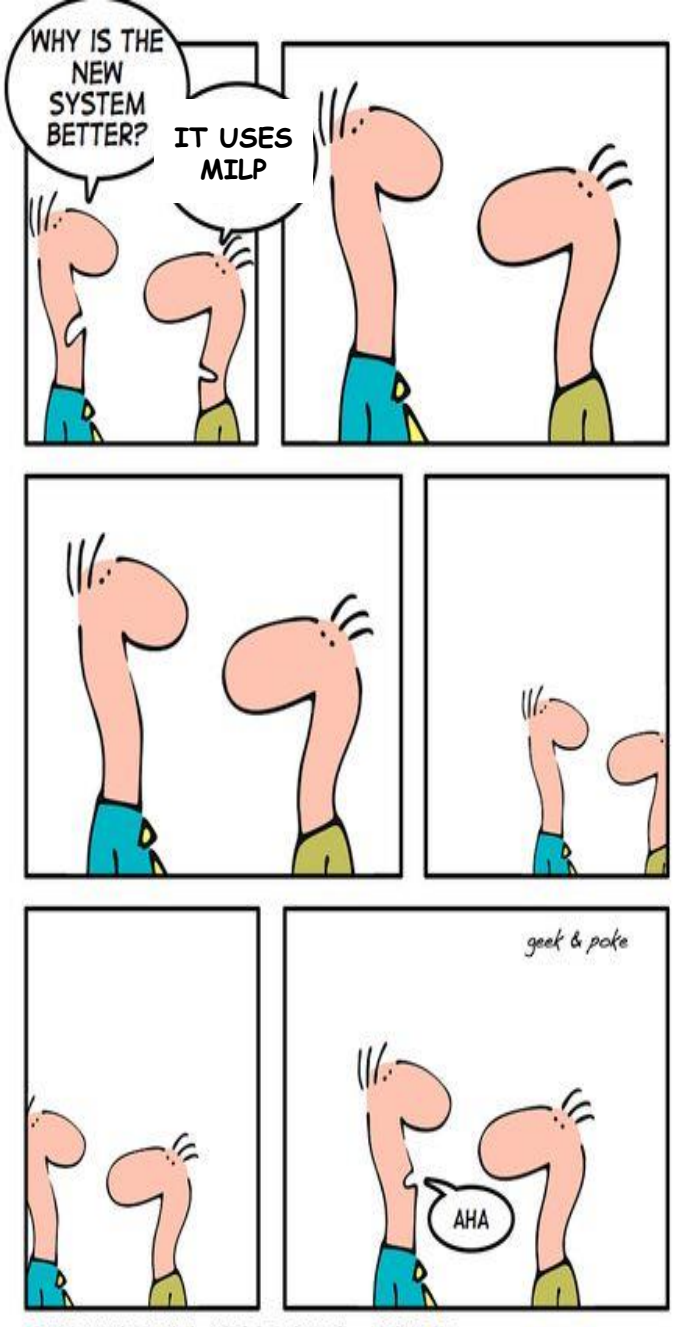

TO BE SUCCESSFUL YOU HAVE TO CONVINCE THE BUSINESS

How can Watson Analytics make mathematical optimization more consumable?

Tell a story

- Build a story flow
- LoB-focused visualization
- Natural language dialogue
- Mobile, team, community sharing
- Collaboration for end-to-end analytics

Leverage templates & services, e.g.

- Models
- Industry templates
- Analytics services
- Visualization services

Resources

- DOcloud
  - Free 30-day trial: ibm.co/docloudtrial
  - DeveloperWorks community:https://ibm.biz/doclouddev
  - API tutorial: https://www.youtube.com/watch?v=AqhusPZWBbc

- CPLEX Optimization Studio Community Edition (free, limit 1000 vars, 1000 constraints, DOcloud enabled):

http://www-01.ibm.com/software/websphere/products/optimization/cplex-studio-community-edition/

- Twitter:

  #ibmoptimization

  #ibmanalytics

- Contact me: Susara van den Heever (svdheever@fr.ibm.com)