

FRA: Financial Risk Analytics for IT Outsourcing Contracts

Sinem Güven¹ Shu Tao¹ Sherif Goma²

¹ IBM T. J. Watson Research Center, Yorktown Heights, NY, 10598, USA

² IBM Business Performance Services, Somers, NY 10589, USA



IBM Research

© 2013 S. Güven, S. Tao and S. Goma

This experience session track paper describes the Financial Risk Analytics (FRA) tool, which is designed to predict and quantify IT Strategic Outsourcing (SO) contract risks in order to help with contract signing and pricing decisions.

FRA enables Quality Analysts and Risk Managers to learn about and proactively manage potential risks before they materialize, while also providing guidance to contract Pricers to include the necessary cost contingencies into pricing considerations, in case of a high risk contract.

Outline

- Introduction
 - Motivation
- FRA (**F**inancial **R**isk **A**alytics)
 - Approach
- FRA: Technical Contributions
 - Contract Similarity
 - Gross Profit Margin Prediction
 - Risk Prediction and Quantification
- FRA Pilots
 - FRA Tool
 - Pilot I: Quantitative
 - Pilot II: Qualitative
- Conclusions and Future Work

© 2013 S. Güven, S. Tao and S. Goma

In the next few sections, we provide an overview of FRA along with its technical contributions, and report on the results of our pilots.

Motivation

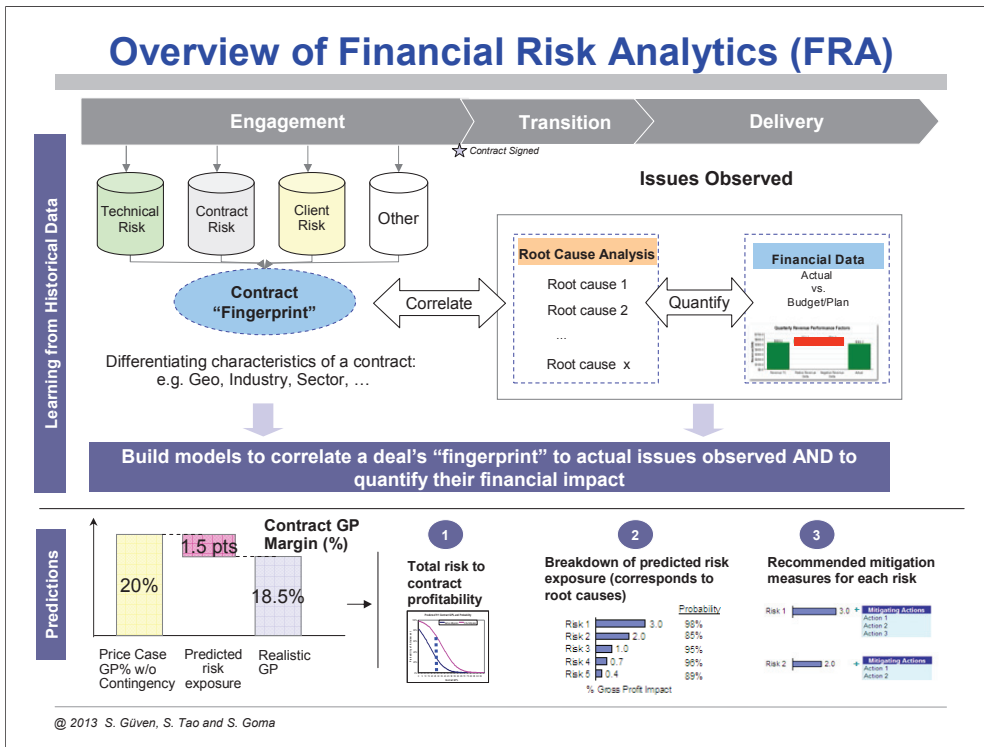
- Contract risk prediction and quantification is one of the major pain points IT service providers are facing today
- Service providers need to know about the potential risks for a given new opportunity ahead of contract signing in order to:
 - make educated decisions about whether to undertake the IT operations of a potential client
 - be proactive about mitigation planning if they are willing to take on a risky contract
 - price the contracts accordingly to cover for risks that cannot be mitigated
- Existing risk management processes have limitations
 - Service providers often need to decide on a contract with limited access to the client's IT environment without thoroughly understanding potential risks
 - Too many risks can be identified at engagement, too little resources to manage them all
 - Lack of a quantitative approach to objectively evaluate risk impact and prioritize risk management tasks
- The Financial Risk Analytics (FRA) tool is designed to use risk prediction models to predict potential contract risks and their impact

@ 2013 S. Güven, S. Tao and S. Goma

IT outsourcing allows clients to offload the operation of their IT systems and processes to a specialized service provider, so they can focus on their core business functions. As such, service providers strive to provide uninterrupted, high quality delivery of service to achieve high client satisfaction, while maintaining continuous profitability of the contract at the same time.

In practice, a significant number of new service contracts underperform financially when compared to the original budget and plan. This is because service providers often need to make a decision about whether to undertake a contract without having proper access to the client's IT environment to understand potential risks. During the Engagement phase prior to contract signing, clients are often reluctant to reveal critical or precise information about their IT operations as there is no guarantee that the service provider they are negotiating with would eventually be the one who takes over their operations. Another reason for poor financial performance in the early stages of a contract is often the lack of a quantitative approach to objectively evaluate risk impact. Even if risks are known ahead of time, it may not be possible to quantify their impact, which makes it difficult to put price contingencies in contracts should the service provider decide to take on a risky contract. Previous research on impact quantification [1, 2] mostly focused on high level IT risks and associated costs rather than quantifying contract risks at a fine level of granularity.

We present the Financial Risk Analytics (FRA) tool, which is designed to address these issues. FRA includes predictive models built from historical contract data and observed risks, and provides insights on potential risks and their impact for a given new opportunity ahead of contract signing. FRA, thus, enables service providers to make educated decisions about whether to undertake the IT operations of the potential client, or proactively mitigate risks if they are willing to take on a risky contract. Finally, service providers can also use FRA insights to adjust contract price according to the predicted risk impact, if risk mitigation is not feasible.



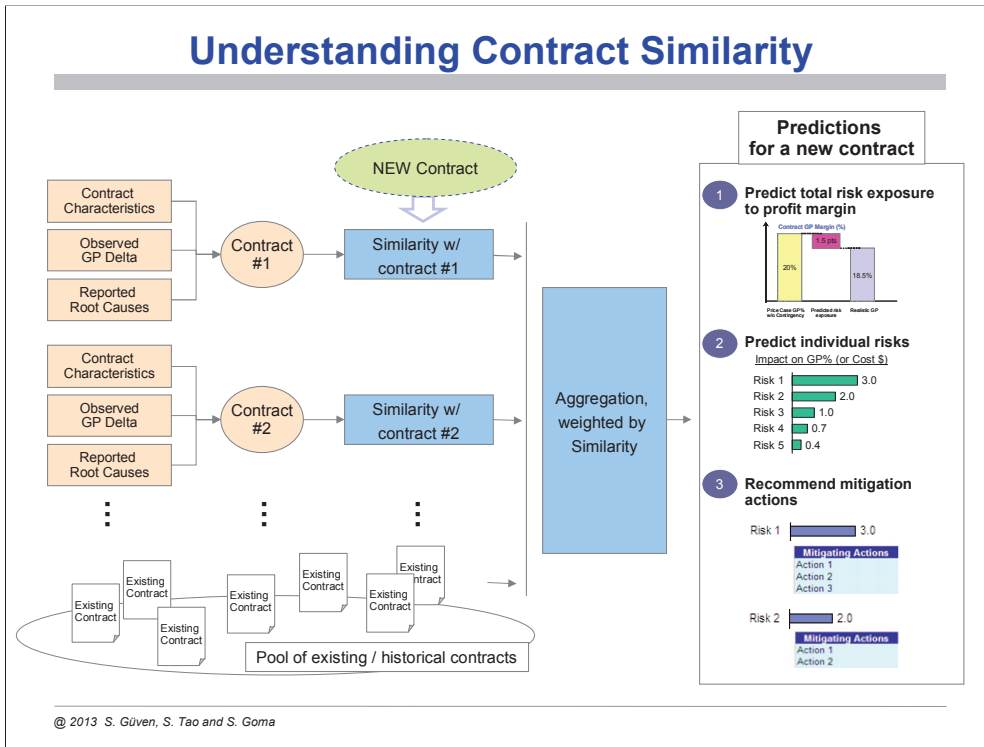
The Financial Risk Analytics (FRA) tool provides predictive models to shed light into potential contract risks and quantify their impact, while also recommending mitigation actions for proactive risk management.

The predictive model needs at least three types of training data from historical contracts:

- 1) Various risk assessments from historical contracts (such as Technical, Contract and Client risk assessments) as well as differentiating characteristics of contracts (such as geography, industry, etc), which altogether are termed *contract fingerprint*. Each feature in a contract fingerprint can be converted into a numerical or categorical value.
- 2) *Risks* (or root causes) observed for these historical contracts during contract reviews during transition or delivery.
- 3) The *financial performance* of the historical contracts, namely the actual performance compared to the original plan.

Trained with the above data, the predictive model predicts for a new contract, based on its fingerprint:

- 1) Whether it is likely to meet the profit target, and if not, miss by how much;
- 2) Potential risks along with their probability (likelihood of happening) and financial impact;
- 3) Mitigation actions recommended for proactive management of the predicted risks.



The core of FRA's predictive model is the similarity measure between contracts. That is, our prediction is based on the measurement of similarity between a given new contract and a set of historical contracts, based on their fingerprints.

The details of contract similarity measure will be defined on Slide 6. With this definition, FRA then provides:

- 1) Contract profitability prediction – (Slides 7 and 8)
- 2) Risk prediction – (Slide 9)
- 3) Recommended mitigation actions for each predicted risk – (Slide 11)

Determining Contract Similarity

For each pair of contracts (i,j) Compute $\text{dist}_r \rightarrow \text{Dist}(r_i, r_j)$	i
For each pair of contracts (i,j) Compute $\text{dist}_{\text{feature}} \rightarrow \text{Dist}(f_i, f_j)$	ii
For each feature f Pearson's Correlation $w_f = \text{Corr}(\text{dist}_{\text{feature}}, \text{dist}_r)$	iii
Given target opportunity i For each similar contract j For each feature f $\text{Dist}(i,j) = \sum_{f=1 \text{ to } N} (\text{Dist}(f_i, f_j) * w_f)$	iv
Given target opportunity i For each similar contract j $\text{Sim}(i,j) = 1 - \text{Dist}(i,j)$	v

© 2013 S. Güven, S. Tao and S. Goma

In our prediction model definition, two contracts are similar if they have similar contract fingerprints. In our data set, we have more than 300 features in a contract fingerprint, but not all features are equally important or useful for our risk predictions. In order to ensure that more significant features provide a greater contribution to the similarity measure, we need to assign higher weights to them. Since the goal of determining contract similarity is to predict risks, we assign weights to features based on their correlation with the actual similarity between a pair of contracts, in terms of their reported risks.

In order to calculate weights (w_f) for each feature f, we compute Pearson's Correlation [3] between risk distances (i) and feature distances (ii). The stronger the correlation, the higher weight will be assigned to feature f. Risk (root cause) distances (i) for all contracts are computed by comparing the risks for a pair of contracts at a time and calculating the difference between these risks, denoted by dist_r . Similarly, we calculate feature distances for all historical contracts in step (ii), denoted by $\text{dist}_{\text{feature}_j}$, by comparing the features for a pair of contracts at a time and calculating the difference between their values (iii). The resulting Pearson's Correlation coefficient, after normalization, is used as a weight (w_f) for each feature. Based on the weighted fingerprint (a vector of weighted features), we then calculate the Euclidian distance [4] between the target opportunity and each historical contract, as shown in step (iv).

The final step is to calculate contract similarity $\text{Sim}(i,j)$ between the target opportunity i and each historical contract j, as shown in step (v).

Predicting GP Delta

- Gross Profit (GP) Delta refers to contract profitability

$$GP\ Delta = GP\ Plan - GP\ Actual \quad \text{for a given contract}$$

- Approach I: Use an ordinal regression model
 1. Regress contract fingerprints (x_1 through x_N) against bucketed ranges of observed GP Deltas of historical contracts where optimal buckets ($bk_{1\text{ to }K}$) and their corresponding ranges [$r_{ai\text{ to }K}$, $r_{bi\text{ to }K}$] are determined based on the historical data distributions and expert input
 2. Given a new opportunity whose fingerprints are known, run the regression to yield:

<i>buckets</i>	$bk_1 = [r_{a1}, r_{b1}]$	$bk_2 = [r_{a2}, r_{b2}]$...	$bk_K = [r_{aK}, r_{bK}]$
<i>probabilities</i>	p_1	p_2	p_K

3. Calculate the Expected Value of GP Delta given the buckets and probabilities

$$E[GP\ Delta] = ((r_{a1} + r_{b1})/2) * p_1 + ((r_{a2} + r_{b2})/2) * p_2 + \dots + ((r_{aK} + r_{bK})/2) * p_K$$

@ 2013 S. Given, S. Tao and S. Goma

We measure contract profitability using Gross Profit (GP) Delta, which is determined by subtracting the planned GP (in percentage) from the actually observed GP (in percentage) for a given contract.

$$GP\ Delta = GP\ Plan - GP\ Actual$$

Our approach to predicting contract profitability is to build an ordinal regression model (step 1) by regressing the historical contract fingerprint features (x_1 through x_N in the above diagram) as the independent variables against several pre-defined buckets of observed GP Delta ranges from historical contracts (as the dependent variables) where the optimal range ($r_{ai\text{ to }K}$, $r_{bi\text{ to }K}$) of buckets are determined based on historical distributions and expert input.

Once the regression model is in place, given a new opportunity and its fingerprint, the regression model yields a set of (bucket, probability) pairs, as shown in step 2 above, that define the probability of the GP Delta prediction falling into a specific bucket. For example, the prediction could yield a 85% probability that the GP Delta will fall into bucket [0, 5] which would mean a positive GP Delta, indicating that the predicted profit margin is 0 to 5% higher than the plan.

Finally, the Expected Value for GP Delta is calculated by taking the mid-points of the ranges (assuming uniform distribution within the range) multiplied by the respective probabilities (p_i) of GP Delta falling in that bucket, as shown in step 3 above.

Predicting GP Delta

- Approach II: Merge regression and similarity models

$$GP\ Delta_{SR} = \sum_{i=1\ to\ N} (GP\ Delta_i * Similarity_i) / totalSimilarity(i,N)$$

where i is a similar contract within bucket range $[r_{a_i}, r_{b_i}]$ predicted by regression model
where contract similarity threshold = $x\%$

While the regression model (see Approach I) provides a decent prediction on GP Delta range when tested against our historical contracts (see Pilot I on Slide 12), we further incorporate contract similarity to provide more fine-grained prediction on GP Delta.

In this extended model, we treat the result of the regression model as an indicator. For a given new contract, we first use the aforementioned regression model to determine which range the GP Delta is most likely to be in (say $[0,5]$ with 85% probability).

Next, we predict the GP Delta of the new opportunity by taking a weighted average of the GP Deltas of the similar historical contracts, whose GP Deltas fall into that particular (say $[0,5]$) bucket, where the weights refer to contract similarity $[0$ through $1]$.

We compare the accuracy of these two approaches in our pilot on Slide 12.

Risk Prediction & Quantification

- Knowing that a given opportunity is likely to become unprofitable is often not enough
 - Service Providers need to know:
 - what the potential risks are
 - how to quantify these potential risks
- Risk prediction and quantification can also benefit from contract similarity

Given target opportunity i

For each similar contract j

For each risk r_k in contract j

$$1. \text{Similarity}(i,j) = 1 - \text{Dist}(i,j)$$

$$2. r_impact_k = \sum_{j=1 \text{ to } N} ((\text{GP Delta}_j / \text{numberOfRisks}_j) * \text{Similarity}(i,j)) / \text{TotalSimilarity}(i,N)$$

$$3. r_probability_k = \sum_{j=1 \text{ to } N} (\text{Similarity}(i,j) / \text{TotalSimilarity}(i,N))$$

© 2013 S. Güven, S. Tao and S. Goma

For a service provider, knowing that a given opportunity is likely to become unprofitable is often not enough. They also need to know what the potential risks are as well as how to quantify these potential risks to be able to mitigate them before they materialize.

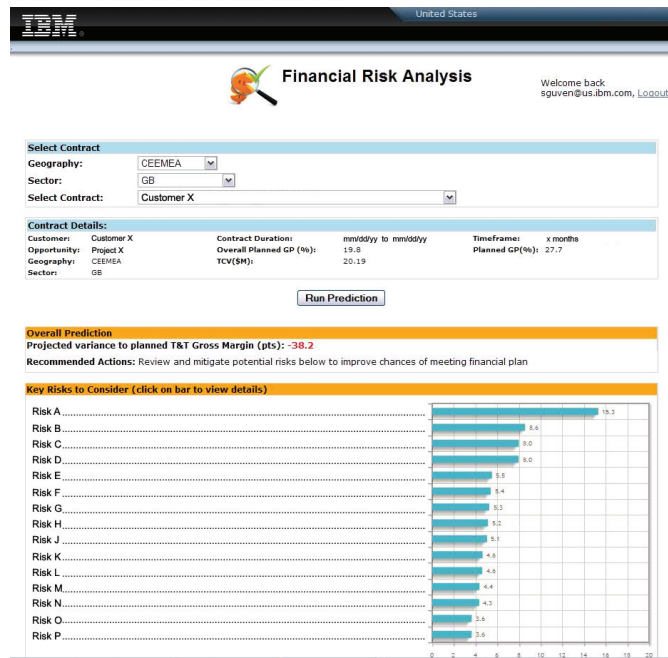
Risk prediction and quantification can also benefit from contract similarity, as previously introduced on Slides 5 and 6. For each target opportunity, we first determine a set of similar contracts (see Slide 6) along with a degree of similarity [0 through 1], as shown in step 1 above.

For each reported risk (or root cause) of a historical similar contract, we calculate the potential impact by dividing the GP Delta of this similar contract by the number of risks observed for this similar contract. (Note that this is an approximation due to the lack of more accurate impact assignment data, which can be improved if the risk management process requires certain impact values to be assigned to each reported risk). We then take a weighted average of all calculated impacts for this particular risk observed across all similar contracts such that the weight is determined by the degree of contract similarity (step 2).

Step 3 shows how to calculate the probability of a given risk (or root cause) happening for the target opportunity: essentially by taking a weighted average of its number of occurrences across all similar contracts such that the weight is, again, determined by the degree of contract similarity.

FRA Tool

- Provide predictive analytics through:
 - Regression
 - Similarity
- Need the following data from historic contracts to train the model:
 - Client Risk Assessment
 - Technical Risk Assessment
 - Contract Risk Assessment
 - Reported Risks (Root Causes)
 - Financials
- For a potential new client, FRA predicts:
 - Contract profitability
 - Potential key risks that are likely to materialize



© 2013 S. Güven, S. Tao and S. Goma

The above Figure shows a screenshot of our FRA Tool implementation. First, the user selects a Geography and a Sector to narrow down the set of available opportunities to analyze ahead of contract signing. Once the opportunity of interest is selected (say Customer X), the user presses the *Run Prediction* button to display the results of the FRA analysis.

FRA predicts the contract profitability (GP Delta) as well as the top 15 potential risks for the target opportunity. For example, for Customer X, FRA predicts a GP Delta of -38.2 points, indicating a 38.2% less profitability than the plan.

For each predicted risk, FRA also displays a probability, as shown in the above Figure. For example, risk D has 8% probability of happening for the selected target opportunity (Customer X).

Customer name, dates and particular risk details have been anonymized for confidentiality reasons.

FRA Tool

- Provide risk details:
 - description
 - probability
 - impact
- Suggest mitigation plans for identified risks
- Enable review of similar contracts

Risk Details

Risk Num: Risk D
Risk: Title of Risk D
Description: Detailed description of Risk D

Probability (%): 8
Potential impact based on historical observations: -3.6 pts

Suggested Mitigation Plan

1. Mitigation step 1
2. Mitigation step 2
3. Mitigation step 3

Similar Historical Contracts which Observed Selected Root Cause

Contract (ordered by similarity)	Root Causes Observed / Reported
1. Contract Name	Risk D Risk P Risk J
2. Contract Name	Risk B Risk K Risk D Risk H Risk C
3. Contract Name	Risk P Risk A Risk C Risk G Risk D Risk J
...	

@ 2013 S. Güven, S. Tao and S. Goma

Selecting a particular risk from our top 15 list reveals more information about that risk. For example, the above screenshot reveals additional information for Risk D after the user has selected it through the interface shown on Slide 10. In addition to more detailed description of Risk D, FRA also shows the probability (8%) and the potential impact (-3.6 points) of that risk for the target opportunity.

Another important step in risk management is risk mitigation. For each predicted risk, FRA shows a set of mitigation steps the user can take to proactively manage that risk before it materializes.

Finally, the user is presented with a set of similar historical contracts along with their observed risks to enable a more detailed investigation of potential risks, if needed.

Detailed risk definitions, associated mitigation steps and similar contract names have been anonymized for confidentiality reasons.

Pilot I: Quantitative – GP Delta Prediction

- Conducted a retrospective pilot on ~60 contracts to test the effectiveness of the GP Delta and risk prediction models.
- Regression model test results:

Test Category	Metric	Result
Gross Profit Margin Direction	Predict whether contract will meet GP plan or not	95%
GP Delta Precision (Predicted vs. Actual GP)	Prediction within +/- 5 pts of Actual GP	48%
	Prediction within +/- 10 pts of Actual GP	75%

- The added (direction) guidance from the regression model:
 - provides a more accurate result compared to the regression model when the regression prediction is in the right direction

Contract ID	Actual GP Delta	Original GP Delta Prediction	New GP Delta Prediction with Guidance
Contract X	6.105	9.985380085	6.139090199

- performs worse than the regression model when the regression prediction is off

Contract ID	Actual GP Delta	Original GP Delta Prediction	New GP Delta Prediction with Guidance
Contract Y	-6.873	8.729026151	11.88203175

© 2013 S. Güven, S. Tao and S. Goma

We trained our predictive model with about 500 historical contracts (whose gross profitability and observed risks are known) and conducted a retrospective pilot on about 60 historical contracts to test the effectiveness of the GP Delta and risk prediction approaches defined on Slides 7 through 9.

The results show that Approach I (the regression model) is able to:

- correctly predict 95% of the time whether a given historical contract will meet the GP plan
- predict the correct GP Delta within +/- 10 points for 75% of the historical contracts
- predict the correct GP Delta within +/- 5 points for 48% of the historical contracts

The results also show that, with Approach II (the regression model with incorporated contract similarity), the predictions are much more accurate when the regression model is able to determine the GP Delta direction correctly (see above). This result is intuitive as the weighted average would only be applied to the GP Deltas of the similar contracts falling within the correctly determined bucket range.

When the regression prediction is off (which is only 5% of the time), however, the results are worse than using the regression model alone (see above). This result is also intuitive given that the weighted average would then only be applied to the GP Deltas of the similar contracts falling within the incorrectly determined bucket range.

Pilot I: Quantitative – Risk Prediction

- Risk prediction model test results:

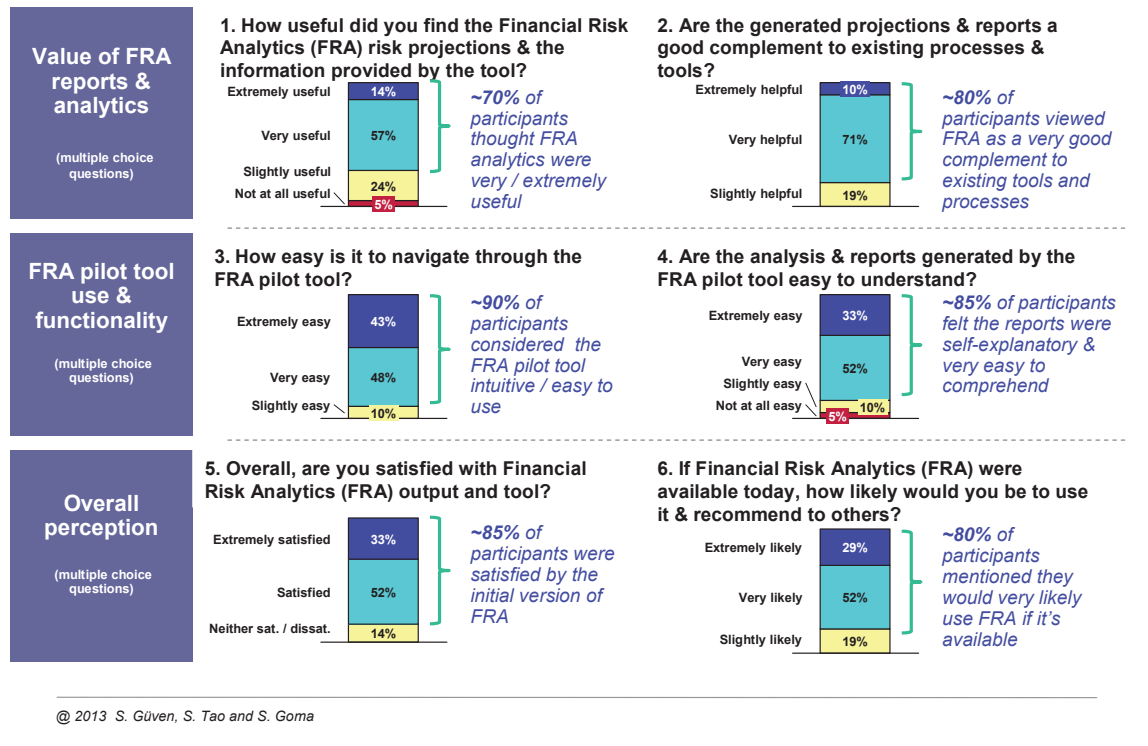
Test Category	Metric	Result
	Percentage of contracts with at least 1 actual risks predicted	100%
Individual Risk Prediction	Percentage of contracts with at least 1 actual risks predicted on top-15 list	77%
	Average percentage of actual risks predicted on top-15 list	43%

@ 2013 S. Güven, S. Tao and S. Goma

As for our contract similarity model, our test results reveal the following performance metrics:

- for all test contracts, 1 or more actual risks are predicted
- for 77% of the test contracts, 1 or more risks are predicted in our top 15 list
- for all test contracts, 43% of actual risks are predicted in our top 15 list

Pilot II: Qualitative



We also conducted field pilots across 13 target SO opportunities prior to contract signing & solicited feedback from 40+ pilot participants, whose roles ranged from Quality Analysts, to Risk Managers and Project Executives.

The results of the filed pilots are shown in the above diagram. A number of positive observations were also highlighted by the pilot participants:

“ It provides another lens through which risk information is viewed in the form of \$ impact. “

“I believe that it adds another set of data to the arsenal of information a Risk Manager would use to attain stakeholder attention and drive remediation action.”

“ Really good start on meeting a difficult objective ”

Conclusions and Future Work

- The Financial Risk Analytics (FRA Tool is designed to assist Risk Managers and Quality Analysts to:
 - identify key characteristics in services contracts that often led to delivery issues
 - predict GP performance of a target opportunity based on similar contracts executed in the past
 - determine and quantify the impact of engagement/delivery risks
 - recommend and prioritize risk mitigation actions
- FRA also helps influence the pricing of contracts with better analytical insights derived from historical data
- Pilot results confirm that FRA predictions are crucial to the contract risk management and pricing processes
- Future work: Include enhancements based on pilot feedback
 - **Include additional contract features** to improve similarity measure to historical contracts
 - **Add a clip level for historical contracts** used to train the predictive model
 - **Portfolio level predictions:** highlight risks that span multiple contracts which will generate / drive mitigation actions across contracts

© 2013 S. Güven, S. Tao and S. Goma

We have presented the Financial Risk Analytics (FRA) tool, which is designed to assist Risk Managers and Quality Analysts to adequately identify and manage contract risks prior to contract signing. FRA also helps quantify the identified risks and thereby influences the pricing of contracts with better analytical insights derived from historical data.

We have piloted FRA both *quantitatively* across 60 historical contracts as well as *qualitatively* with 13 SO engagements prior to contract signing, and the pilot results are promising.

As the next step, we would like to incorporate some new functionality based on pilot feedback. Some immediate action items are to:

-include additional contract features, such as Growth Market vs. Major Market, GP target, what kind of requirements are stated in the solution or the contract, etc , to improve our similarity measure to historical contracts.

-add a clip level for historical contracts used to train the predictive model. For example, only consider contracts from the past 2 years.

-provide portfolio level predictions. Namely, highlight risks that span multiple contracts which will generate / drive mitigation actions across contracts.

References

- [1] M. Harris, D. Herron and S. Iwanicki, *"The Business Value of IT: Managing Risks, Optimizing Performance and Measuring Results,"* Auerbach Publications, 2008.
- [2] K. Goolsby, *"Building Your Business Case for Best-practice IT Services Delivery,"* Oracle White Paper, 2009.
- [3] A. Field, *"Discovering Statistics Using SPSS,"* Sage Publications, 2009.
- [4] E. Deza and M. Deza, *"Encyclopedia of Distances,"* Springer, 2009.