SimQL research questions

- How little of the model needs to be exposed?
- How can defaults be set rationally?
- How should expected execution cost be reported?
- How should uncertainty be reported?
- Are there differences among application areas that require different language structures?
- Are there differences among application areas that require different language features?
- How will the language interface support effective partitioning and distribution?