1086-90-883 **Daniel Bienstock*** (dano@columbia.edu), Dept. of IEOR, Columbia University, 500 W 120th St, New York, NY 10027, and **Alexander Michalka** (admichalka@gmail.com), Dept. of IEOR, Columbia University, 500 W 120th St, New York, NY 10027. *Optimizing convex functions over non-convex sets.*

A recurring theme in optimization theory is that convexity (in the objective function and in the constraints) usually leads to provably good, i.e., polynomial-time, algorithms. However, many important problems in science and engineering inconveniently do not adhere to this prototype. Whereas the objective function is frequently convex (even quadratic), corresponding for example to least-squares estimation, or energy minimization, the constraints of the problem, which reflect fundamental physical or engineering relationships, are not convex. In such cases we would still prefer to leverage the techniques of convex optimization. Such a situation will ensue if, in the appropriate reformulation of the problem, we can solve the classical "separation problem" efficiently – that is to say, given an arbitrary vector which is not feasible, separate it from the feasible set by means of a linear inequality. In the non-convex case this cannot be done unless the problem is suitably reformulated. The emphasis on linear separating inequalities is justified in terms of numerical stability of the overall optimization algorithm, and solution speed. Thus we can leverage fast, robust, linear programming solvers. In this talk we describe ongoing work that develops this paradigm. (Received September 14, 2012)