Ryan Cumings-Menon
Research Statement

My research focuses on the formulation of nonparametric quantitative methods and their applications. My job market paper uses optimal transport to estimate a density function that is capable of leveraging the information provided by the shape of the density during estimation. I also have written a paper on a nonparametric estimator of a mode conditional on covariates. Both of these papers can be found on my website at `https://sites.google.com/view/ryancumingsmenon`. My final paper, which I am working on with Professor Elizabeth Powers, estimates the mental health impacts of having a third child. These papers are described below in more detail.

Shape Constrained Density Estimation Via Optimal Transport

My job market paper proposes a method to estimate a density function subject to a constraint on the shape of the estimator. Some examples of these estimators in the literature include those that constrain the density to be concave or log-concave. When one has reason to believe that a density estimator satisfies a particular shape constraint, these estimators provide a method to leverage this information during estimation for a gain in accuracy. In this sense, shape constrained density estimators can be viewed as a middle ground between the parametric estimators and the fully nonparametric estimators, such as the kernel density estimator. However, the classic shape constrained density estimator, which is defined by maximizing a likelihood function, requires the shape constraint to be sufficiently strong, and attempting to relax this constraint results in either computational difficulties or the nonexistence of the estimator. For example, when constraining the density to be $\rho-$concave, finding the estimator requires the solution of a nonconvex optimization problem when the constraint is weaker than log-concavity, which corresponds to $\rho = 0$, and the estimator does not exist when $\rho < -1$.

My job market paper proposes an estimator that allows for estimation subject to arbitrary shape constraints. There are four primary advantages of this method. First, the computational efficiency of the proposed estimator compares favorably to that of the maximum likelihood estimator, so it is well suited for analysis of large datasets. Second, many random variables of interest in economics have thick tails, such as income, firm size, changes in stock market prices, and many more. Thus, these densities do not satisfy log-concavity, but the proposed estimator can still be used with a weaker shape constraint, like quasiconcavity or constraining the density estimator to have a single mode. Third, when estimating a density as part of an economic model that assumes the density satisfies a regularity condition, the method allows for the estimation of the density subject to this condition, without the requirement of additional parametric assumptions. Lastly, the paper provides a method to test whether a population density satisfies a given condition, including the commonly used regularity conditions in economic models.

The paper provides an application to demonstrate these last two advantages in particular. Specifically, the paper uses data from the California Department of Transportation to explore if their decision to award construction contracts with the use of a first price auction is cost minimizing. Myerson (1981) provides a model that requires the density of firms' valuation for these contracts to satisfy a regularity condition for this to be the case. Given a density, $f(x)$, with a distribution denoted by $F(x)$, Myerson's (1981) regularity condition requires the function

$x - (1 - F(x))/f(x)$ be strictly increasing. Note that log-concavity is a sufficient condition to guarantee that this condition holds, but since Myerson's (1981) regularity condition does not rule out thick tail behavior, log-concavity is stronger than required. The paper shows that a kernel density estimate of the firms' valuations does not satisfy this condition. However, the proposed estimator, subject to the constraint that Myerson's (1981) regularity condition is satisfied, appears to match this kernel density estimator fairly closely, and the proposed test also fails to reject the null hypothesis that the regularity condition is satisfied.

The proposed estimator is defined by an optimization problem, and the method allows for the use of weak shape constraints by using an objective function that was recently introduced in the optimal transportation literature, known as the entropic regularized Wasserstein distance. This objective function converges to a commonly used metric between density functions, the Wasserstein distance, as the sample size diverges. In finite sample sizes, the regularization term ensures that the objective function has a global unconstrained minimum, so problems related to the nonexistence of the optimum, and thus also the estimator, are avoided. This unconstrained minimum also corresponds to a well founded density estimator; specifically, it is a kernel density estimator.

There are four primary contributions of the paper. First, the paper provides results on consistency, limiting distribution, and asymptotic rate of convergence under two different sets of assumptions. Second, a proof that the proposed test is consistent as well as the limiting distribution of the test statistic is formulated. Third, the paper provides sufficient conditions for convexity of the optimization problem defining the estimator. When using the estimator to increase the accuracy of an estimator, rather than to ensure that a regularity condition holds, $\rho-$concavity constraints provide a rich class of shape constraints, and this optimization problem is always convex in this case. A corollary shows that the optimization problem is also convex when the shape constraint is sufficiently weak. Lastly, the paper proposes a computationally efficient method to calculate the estimator.

A Nonparametric Modal Regression

The inherently local nature of the mode makes estimators of the mode one of the more robust notions of central tendency. For example, estimators of the univariate mode are robust to outliers moving further away from the mode, switching from one side of the mode to the other, or being removed from the dataset entirely. Since estimators of the mode rely on a proportion of the data that converges to zero, this property also holds asymptotically for any datapoint that is not located in the neighborhood of the mode. Partly for this reason, these estimators have also been generalized to estimators of the mode of a dependent variable, conditional on covariates, which is known as a modal regression. These regressions can be viewed as estimators of the most common or typical value of the dependent variable given the covariates.

This paper proposes an alternative starting point for nonparametric modal regressions. Specifically, a graph is constructed using the data, which is used to define paths that balance a tradeoff between the distance between consecutive points and curvature. Afterward, these paths are combined to define the final regression.

This approach lends three primary advantages to the method. First, it allows for the formulation of an algorithm to compute the proposed regression with a computational efficiency that improves on that of the classic estimator, which is based on the kernel density estimator.

Second, while the method allows for the specification of a smoothing parameter for increased accuracy, this is not a requirement. Third, even without specifying this smoothing parameter, the method's accuracy appears to compare favorably to methods based on kernel density estimation when the conditional mode and the value of the joint density at the conditional mode are both nonlinear.

After providing a consistency result and simulations, the paper moves onto an application using data on high school GPA, University of Illinois at Urbana-Champaign (UIUC) GPA, and ACT score of all undergraduates that attended UIUC between 2000 and 2010. The paper estimates the most typical value of UIUC GPA conditional on high school GPA and ACT score, and finds a relatively weak association between ACT score and the most typical UIUC GPA.

The Mental Health Effects of Having a Third Child

Professor Elizabeth Powers and I are also working on a paper in which we estimate the causal effect of having a third child on the mental health of parents. We use the identification strategy that was first described by Angrist and Evans (1998). Specifically, since parents in the United States are more likely to have a third child after having two children of the same gender, the gender of the first two children can be used to construct instrumental variables for having a third child. We find evidence that having a third child negatively impacts the happiness of mothers but not fathers.

References

Angrist, J. D., & Evans, W. N. (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *The American Economic Review, 88*(3), 450-477.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research, 6*(1), 58-73.