Summaries of Published Papers – Summer 2017

Jacob S. Sagi

Because my work spans two distinct areas in economics, decision sciences and financial economics, I divide the description of papers into those two specialties.

Decision theory

A. **"You Need to Recognize Ambiguity to Avoid It"** (2017), Forthcoming in the *Economic Journal* (with S.H. Chew and M. Ratchford).

The Ellsberg two-urn choice problem, first attributable to Keynes (1921), provides a clean and clear test of non-Bayesian behavior and the prevalence of "ambiguity aversion". Most choices featuring ambiguity, in practice, are not as clear-cut as and easy to understand as the Ellsberg two-urn choice problem. This leads one to ask whether individuals who are able to discern the presence of ambiguity in a "complex" situation – one can think of these as sophisticated or high-comprehension decision makers – exhibit ambiguity aversion. Within a large-scale experimental framework, involving a total of 3583 adult subjects with diverse demographics (i.e., not just University students), we address this question. Subjects are presented with screening questions before choosing between two alternatives represented by payoff-matrices which are essentially equivalent to those in Ellsberg's (1961) two-urn problem. When facing this essentially equivalent yet more complex matrix-based choice task, high-comprehension subjects continue to exhibit ambiguity aversion typical of the standard two-urn problem while, as expected, low-comprehension subjects appear to behave randomly. Our design allows us to classify subjects as "probability minded" or "ambiguity minded" based whether they assign probabilities to draws from a deck of cards with unknown composition during the screening phase. High-comprehension subjects who are ambiguity-minded are far more likely to be

ambiguity averse than those who are probability-minded. Significantly, subject "mindedness" appears to explain ambiguity attitudes *an order of magnitude more than all other demographic characteristics, combined*. Contrary to intuition about subjects' sophistication, ambiguity-minded high-comprehension subjects are younger, more educated, more analytic, and more reflective about their choices compared with their probability-minded counterparts. Our findings cast grave doubts on attempts to discount ambiguity aversion (or, more generally, non-Bayesian attitudes toward uncertainty) as mistakes made by less sophisticated decision makers. At the same time, one must conclude that the vast majority of individuals are unlikely to discern ambiguity in complex situations and may therefore exhibit randomization behavior that confounds standard models of ambiguity attitudes.

B. **"An Inequality Measure for Stochastic Allocations"** (2012) *Journal of Economic Theory* **147**, 1517-1544. (with S.H. Chew)

It is self-evident that fairness is something that matters to individuals even when the notion is not directly applicable to themselves or their own welfare. We bemoan the misfortunes of others, generally consider charity a virtue, and devote resources as a society to provide the "have not" segment of society with opportunities for welfare improvement. The seminal works of Harsanyi (1955), Kolm (1969), Atkinson (1970), Rothschild and Stiglitz (1970), and Weymark (1981) contributed to the view that one can measure inequality using measures such as the Gini Coefficient. Diamond (1967) points out that such measures neglect the fact that randomization induces fairness. In other words, neglecting efficiency concerns, if there is to be inequality, it seems "more fair" if inequality is determined randomly rather than via a prejudicial mechanism. This is the notion of "ex-ante" fairness. Ben-Porath, Gilboa, and Schmeidler (1997) point out that correlations also matter when inequality is concerned, and it may be preferable that we all share similar destinies. This is the notion of "ex-post" fairness. Since the

recognition of these various attributes of fairness, it has been seen as "difficult" to obtain a measure of inequality that exhibits an aversion to static inequality, as well as an affinity for ex-post and ex-ante fairness. We provide a set of intuitive axioms that lead to a simple family of inequality measures and which exhibit the desired traits. The measures can be viewed as a one-parameter extension of the Generalized Gini Mean and can be easily used by econometricians. This measure can be used to compare inequality in two countries, where in one the distribution of income is broader and yet the lower income cohorts have better opportunities for advancement. The measure can also be used in other context where one may wish to measure "inequality" but the rankings are not static or deterministic: for instance, competition in an industry.

C. **"Small Worlds: Modeling Attitudes Towards Sources of Uncertainty"** (2008) *Journal of Economic Theory* **139**, 1-24. (with S.H. Chew)[1]

Nearly 50 years after Savage's Foundations of Statistics (1954), empirical and experimental evidence on individual decision making leaves little doubt that most of Savage's assumptions are violated in practice. Arguably, violations of probabilistic sophistication (e.g., the Ellsberg paradox), presents the most difficulties – as evidenced by the large body of theoretical literature on the subject. Our earlier paper on probabilistic sophistication finds the weakest conditions currently known that ensure probabilistic sophistication. The advantage to this parsimonious approach is that it leaves no doubt as to which behavioral assumption (or axiom) one ought to relax to depart from probabilistic sophistication so as to accommodate 'Ellsbergian' behavior. In this paper, we relax this one assumption and derive a model and representation of behavior in which distinct sources of uncertainty (e.g., an urn with 50 red and 50 black balls versus an urn with an unknown mixture of red and black balls) correspond to probabilistic sophistication within the distinct domains (e.g., red and black from the same urn are equally likely) but

[1] This was the lead article in the *Journal of Economic Theory*, in the first issue of 2008.

not across domains (e.g., I'd rather bet on red from the known versus red from the unknown urn). We show that this leads to a very simple yet intuitive and easily applied model of Ellsbergian behavior (preferring a bet on the known urn to the same bet on the unknown urn). The notion of "source dependent" preferences is now coming into vogue and our paper will hopefully play a key part in the theoretical development of this literature.

D. **"Event Exchangeability: Probabilistic Sophistication without Continuity or Monotonicity"** (2006), ***Econometrica*, 74**, 771-786 (co-authored with S.H. Chew).

Savage's 'Foundations of Statistics' (1954) is recognized as one of the most influential and impressive treatises on the theory of decision making. In it, Savage shows that any individual who adheres to six behavioral principles (or axioms) will make decisions *as if* she assigns subjective probabilities to uncertain events. The latter is termed 'probabilistic sophistication' and is at the heart of much of modern Bayesian statistics and economic theory. One problem with Savage's approach is that his axioms are too strong. In particular, they rule out probabilistic behavior even in cases where a decision maker is *obviously* probabilistically sophisticated. One simple example is an individual with mean-variance preferences. Such an individual violates *three* of Savage's axioms even though her preferences are explicitly probabilistic. This paper provides the heretofore simplest behavioral basis for either determining or requiring probabilistic sophistication. In comparison with other theories, our approach subsumes the largest set of non-pathological preferences consistent with a subjective Bayesian approach to uncertainty. As such, it is a fundamental contribution to social science and statistics and extends the work of Machina and Schmeidler (1992) and Grant (1995).

Because it offers the most parsimonious rationalization of probabilistic attitudes towards uncertainty to-date, the paper also paves the way to understanding what assumptions ought to be relaxed in order to

accommodate the growing experimental evidence against probabilistic sophistication (e.g., the Ellsberg paradox).

E. **"Anchored Preference Relations"** (2006) *Journal of Economic Theory*, vol. 130, pp. 283-295.

Perhaps the most cited and descriptively successful theory of static individual choice is Cumulative Prospect Theory – a variant on Kahneman and Tversky's Nobel winning (1979) paper. One of its central tenets is the notion that losses loom larger than gains (loss aversion) and to model this one assumes that the utility for a prospect is measured relative to some status quo benchmark, such as current wealth. Even aside from this application, the idea of a status quo bias in decision-making plays an important role elsewhere. An example is the well documented endowment effect in which decision makers opt for a mug rather than $5 if they have the mug, or prefer to keep the $5 if they do not have the mug. Other examples in the form of an often large disparity between the willingness to pay for something versus the willingness to accept cash for something are widely documented in the consumer behavior literature.

All of these phenomena are attributed to reference dependence in the way decision makers approach choice problems. A potential problem with a generic model of reference dependent preferences, however, is that it may inadvertently allow for the possibility of systematic manipulation. My paper poses the following question: What model restrictions are implied if one is to rule out situations in which a decision maker has a choice between two prospects, selects one, subsequently changes her mind and selects the other – even if the change is costly. Although quite appealing and simple, my paper shows that such a condition places very stringent constraints on any theory that models reference dependence. In particular, it rules out cumulative prospect theory and virtually every other model commonly used. This is a strong result and one that may not be welcomed by some behavioral economists. However, my purpose is not to prove that reference dependence

is unreasonable. Quite to the contrary, I believe that reference dependence plays an important role in decision making even while decision makers do not systematically expose themselves to manipulation. To demonstrate that it is indeed possible to 'have your cake and eat it too,' I explicitly show the existence of a class of reference dependent models that do not exhibit the cycling described.

F. **"Inter-temporal Preference for Flexibility and Risky Choice"** (2006), *Journal of Mathematical Economics* 42, pp. 698-709 (with Alan Kraus).

We extend the inter-temporal choice model of Kreps and Porteus (1978) to allow for the possibility that individuals are not sure of their own futures tastes. Such models have widely been seen as offering a potential approach to bounded rationality, where the decision maker realizes that (s)he does not know everything about the world, and that something could happen to alter his/her views – which for all intents and purposes means that their choice behavior might not be completely predictable ex-ante. We derive, using a parsimonious set of axioms, an inter-temporal utility model. The model departs from the traditional recursive utility paradigm in that at each decision node the decision maker must take into account several possible continuation (or indirect, or "Bellman") utility functions rather than just one.

G. **"What is an 'Endogenous State Space'?"** (2006), *Economic Theory* **27**, 305-320.

Kreps (1979) and Dekel, Lipman, and Rustichini (2001), among others, derive models in which the decision maker is not sure about his/her futures preferences because of "unforeseen contingencies". Their axiomatic models derive utility representations in which the decision maker behaves as if there are a multitude of future "utility states" and each state corresponds to possible future utility function. The set of states thus derived is interpreted as a subjective state space even though subsequent rankings need not conform to any one of the aggregated utilities. This paper proposes a definition for a subjective state space under unforeseen contingencies that is topologically

unique, derives its existence from preference primitives as opposed to the representation of preferences, and does not commit to an interpretation in which states correspond to future realized rankings. Roughly, a unique endogenous state space, is defined in the paper to exist if the maximal Pareto Frontier of every choice problem can be made to look topologically the same by making infinitesimal changes in the choice set. The maximal Pareto Frontier can then be identified with the state space. Such an approach allows one to give the endogenous state space a more objective sense, because it can be, in principle, observed through experiment.

Financial Economics

A. **"Estimating Oil Risk Factors Using Information from Equity and Derivatives Markets"** (2015), *Journal of Finance*, Vol. 70, 769-804 (with I.E. Chiang and W.K. Hughen)

   Economists have been fascinated with energy prices and their impact on the macro economy since at least the oil crisis of 1973. In their influential paper testing Ross' (1976) Arbitrage Pricing Theory or "APT", Chen, Roll and Ross (1986) were the first to posit that oil price risk cannot be diversified away and exposure to it should therefore command a risk premium. The empirical evidenced amassed since then has been decidedly unconvincing. We take the view that there is no "monolithic" notion of oil price risk to which other assets might or might not be exposed. Rather, investors can be legitimately concerned about various types of *distinct* oil price *risks*: temporary (or short-term) shocks, persistent shocks, long-term shocks (those only affecting supply and demand considerations in the distant future), and volatility shocks. We demonstrate that one can decompose a four-factor affine term-structure model with unspanned stochastic volatility into exactly such a hierarchy of shocks. We then estimate this model (via MCMC) using data from oil futures, oil options, and oil stock portfolios. Using oil stocks is important because they contain information that may not be contained in

those derivatives that are frequently traded. We demonstrate that by employing oil stocks, the resulting estimated oil factors have significant explanatory power for non-oil stocks. Non-oil stocks are typically exposed to between 14% and 20% of the oil risk to which the oil industry itself is exposed. This amounts to an average oil risk premium of 0.70% for the typical non-oil stock. It is important to note that the standard methodology for picking up risk premia (via Fama-MacBeth regressions or using noisy proxy portfolios) cannot typically pick up a risk premium of this magnitude. We succeed in identifying the risk premia because we combine information from various asset markets. The APT provides the intuition that a handful of systematic factors should suffice to account for all priced risk. While we only focus on one fundamental source of systemic risk (i.e., energy), our methodology can be applied more broadly and paves the way to a more structured approach to the pricing of securities in an APT framework.

B. "Managed Distribution Policies in Closed-End Funds and Shareholder Activism" (2014), *Journal of Financial and Quantitative Analysis*, Vol. 49, 1311-1337 (with M. Cherkes and J.Z. Wang)

The tradeoff theory of the closed-end fund discount, as developed in my 2009 paper with Cherkes and Stanton and in Berk and Stanton (2007), posits that a closed-end fund's price premium to NAV corresponds to the present value of the benefits provided by the fund structure and managerial skill less the liability posed by the cost imposed by management through fees and misaligned interests. When a fund is trading at a discount, free-ridership among shareholders entrenches management, thereby exacerbating the discount in equilibrium. The presence (or potential presence) of activist shareholders, on the other hand, reduces the equilibrium discount because of the threat of fund liquidation. Managers, on the other hand, are not passive onlookers and can take action to manage the size of the discount and, therefore, the likelihood that an activist might get involved. One way to do this is by committing to an increase in dividends (effectively, a slow

liquidation of the fund) – i.e., a managed distribution policy. In this paper we provide empirical evidence for a model-deduced equilibrium interaction between managerial choices, activist shareholder choices, and the closed-end fund discount. Relative to the total market size of all traded closed-end funds in the U.S., there is an astonishing amount of academic focus on this investment vehicle. Our paper provides further evidence for the tradeoff-based, rather than behavioral, approach to understanding this market.

C. "Trading Relative Performance with Alpha Indexes" (2011), *Financial Analysts Journal*, November/December Issue, Vol. 67, No. 6, 77-93 (with Bob Whaley)

On October 11, 2010, the NASDAQ began disseminating calculations of NASDAQ OMX Alpha Indexes. These are proprietary relative performance indices each of which tracks the relative performance of a target traded security (e.g., Apple Computers) against that of a benchmark (say, the S&P 500 ETF). On April 18, 2011 trading began in Alpha Index Options. To date about 16,000 have been traded, corresponding to $800M in notional terms. Bob and I designed these indices, hoping that they will appeal both to investors who wanted to focus on firm specific performance without having to worry about timing the market, and to quants who wanted to trade correlation. As Marty Gruber said in a recent NYU conference dedicated to Alpha Indexes, "…if you're on an actively managed mutual fund board, the first thing you hear about is relative performance. That's all directors talk about. That's all portfolio managers talk about." In this paper, we provide the motivation for trading in such indices, the valuation analysis for futures and options, and the hedging ratios for risk management. One of the most important aspects of these products is that one can back out implied correlations from index option price. In particular, one can calculate forward-looking measures of betas. This promises a vast improvement on the risk-return analysis relative to current approaches (which primarily rely on

regressions using historical data). In time, we hope that this work will impact a profound effect on both practitioners and academics.

In many ways, this paper represents the meeting of my two seemingly unrelated research areas. The indexes are designed to appeal to individual traders, investors, and market makers. Yet, their overall role in informing us about how markets work has tremendous potential.

D. "A Liquidity-Based Theory of Closed-End Funds" (2009), Review of Financial Studies 22, 257-297 (with Martin Cherkes and Richard Stanton)[2]

A closed-end fund (CEF) is a publicly traded firm that invests in securities. While investors can, in principle, trade either in the CEF's shares or directly in the underlying securities, a CEF rarely trades at a price equal to the value of the securities it holds (its Net Asset Value, or NAV). CEFs usually trade at a discount to NAV, though it is not uncommon for them to trade at a premium. The existence and behavior of this discount, usually referred to collectively as the "closed-end fund puzzle'", poses one of the longest standing anomalies in finance: Why do CEFs generally trade at a discount, and why are investors willing to buy a fund at a premium at its IPO, knowing that it will shortly thereafter fall to a discount?

Our paper develops a rational, liquidity-based model of CEFs that provides an economic motivation for the existence of this organizational form: CEFs offer a means for investors to buy illiquid securities, without facing the potential costs associated with direct trading and without the externalities imposed by an open-end fund structure. In the paper, we first establish that a liquidity rationale for the existence of CEFs is indeed present in the data. We then develop a model based on this insight in which there is a tradeoff between the liquidity benefits of investing in the CEF and the fees charged by the fund's managers. In particular, the model predicts that IPOs will occur in waves in certain sectors at a time, that funds will be issued at a

[2] This paper won "First Prize" at the 2006 Utah Winter Finance Conference.

premium to net asset value (NAV), and that they will later usually trade at a discount.

We also collect data from a rich variety of sources in order to investigate the model both qualitatively and quantitatively. Overall, we find support for both the liquidity motive for issuing CEFs and the predicted patterns in the premium/discount. Moreover, we find little or no support for an alternative explanation based on investor sentiment. However, we do document one feature of the data that our model cannot explain: the return of newly issued funds under-performs that of seasoned CEFs, though only in CEFs managing fixed-income securities. This overpricng at the IPO suggests that a full explanation of the discount may also require behavioral considerations.

Ours is one of only <u>two</u> theoretical papers examining the CEF puzzle. Given the huge literature on this subject, we expect this paper to have a significant impact. Moreover, recent papers have started to corroborate our findings in related areas (e.g., U.K. traded hedge funds).

E. "Predicting Risk from Financial Reports with Regression" (2009), in *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 272–280, Boulder, Colorado, Association for Computational Linguistics. (with S. Kogan, D. Levin, B. R. Routledge, and N. A. Smith).

This paper introduces the methodology of "text regressions". Simply described, the idea is similar to how spam filters operate, but instead of forecasting whether a document is spam or not (a binary forecast), our algorithm forecasts return volatility (a continuous forecast). We download 10-k reports from the SEC and isolate the Managers' Discussion and Analysis section from each. As long recognized by the accounting literature, these sections contain forward-looking statements and are therefore candidates for testing whether mandatory reports contain information that is useful to investors. We find that the algorithm succeeds in forecasting firms'

return volatilities out of sample as well, and sometimes better, than forecasts based on realized volatility history. We follow this up with a paper that analyzes the data in more detail and concludes that in the period subsequent to the passage of the Sarbanes-Oxley Act of 2002 MD&A sections in 10-k reports have become more informative about firm risk, as measured by future realized volatility.

F. "Firm-Specific Attributes and the Cross-Section of Momentum" (2007) *Journal of Financial Economics* 84, 389-434 (with Mark S. Seasholes).

Among the 450 or so papers on return momentum, there are a few theoretical papers that explain firm-level return-momentum; they do not, however, tie the effect to firm-specific observables. There are also a few papers that show the possibility of momentum portfolio strategy profits; they do not, however, manage to recreate the correct magnitude or term-structure of the phenomenon. Our paper does several things that no other theory paper to-date has managed to do: (i) tie the momentum effect to firm specific attributes, (ii) use that to produce realistic momentum strategy profits, (iii) predict new momentum effects, and (iv) demonstrate that the predictions hold in the data.

The paper begins by documenting intuitive examples in which an increase in the firm's value (good news) is followed by an increase in expected returns. This can be thought of as 'return-momentum' at the single firm level. We investigate this by modeling a firm with production, operating costs, expansion and abandonment options. Numerical analysis suggests that the options component and costs dominate the effects mentioned. Real options contribute to return-momentum while firm leverage (financial or operational) can reverse return-momentum and even lead to reversals.

We follow this analysis with a simulation of an economy populated by many firms. This allows us to construct *momentum portfolios*: these are portfolios that are long on recent winners and short

on recent losers. Historically, portfolios constructed using this strategy have proven quite profitable and their profitability has hitherto been seen as unexplainable by standard efficient market arguments. We show that we can calibrate our model economy to exhibit (i) realistic momentum portfolio profits, and (ii) momentum profits that are higher when the portfolio consists only of high revenue-volatility, low book-to-market, or low-cost firms. The additional prediction made by our model are, amazingly enough, borne out in data.

G. "Asset Pricing Under Unforeseen Contingencies" (2006) *Journal of Financial Economics*, vol. 82, pp. 417-453 (with Alan Kraus).

The Breeden-Lucas-Rubinstein model of a representative agent with time-separable expected utility does not appear to match observed asset return dynamics (Mehra and Prescott, 1985). Given the historically smooth aggregate consumption time-series, a preference-based explanation of asset returns requires more volatile state-prices than is afforded by time-separable expected utility. It is now recognized that one needs to introduce either path dependence at the representative agent level, or non-stationarity in order to capture the magnitudes and dynamics of asset-pricing moments (see the papers of Campbell and Cochrane (1999), Barberis, Huang, and Santos (2001)) Gordon and St-Amour (2000), Melino and Yang (2003), and Routledge and Zin (2004)).

Our paper asks whether the extra source of volatility in state prices can arise from "taste shocks." Moreover, we tackle this issue in a more sophisticated way than similar works in the literature that simply assume a representative agent with time-varying preference parameters (Mehra and Sah (2000), Gordon and St-Amour (2000) and Melino and Yang (2003)). Our paper examines an economy of heterogeneous agents facing individual taste shocks. Agents' taste shocks have an idiosyncratic as well as (possibly) a common component. Because taste shocks are private, only their common components across agents can impact prices.

To avoid moral hazard issues (or unrealistic degree of individual monitoring) we only allow trade on claims over aggregate demand and price-contingent events (i.e., 'public' events). Thus agents in our economy generally cannot fully insure against private taste shocks and markets are incomplete. Moreover, since available securities affect portfolio choice, which in turn affects prices over which agents can then contract, the equilibrium in the type of economy we consider must *simultaneously* determine both the equilibrium mix of tradeable assets and their prices.

In an economy of heterogeneous agents, we characterize the appropriate equilibrium concept when contingent claims can only trade on price- or aggregate demand-contingent events. To my knowledge, this is quite an original equilibrium concept – the only other macro asset pricing model that endogenizes the asset mix in an incomplete market setting is an, unfortunately, uncited paper by Kerry Back in *Economic Theory*. Surprisingly, one can aggregate across our agents' demand for securities if their preferences are sufficiently homothetic. In this case, we establish the existence of an equilibrium for a class of parametric models. We then calibrate the model to fit a litany of stylized facts. The overall conclusion of the paper is that both taste shocks and non-standard preference over taste-shock risk are needed to yield a model consistent with asset returns data.