Journal of Financial Economics 98 (2010) 415–438

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Differences in beliefs and currency risk premiums ☆

Alessandro Beber [a], Francis Breedon [b], Andrea Buraschi [b,*]

[a] Amsterdam Business School and CEPR, The Netherlands
[b] Imperial College, London, UK

ARTICLE INFO

Article history:
Received 22 June 2009
Received in revised form
9 December 2009
Accepted 5 January 2010
Available online 7 July 2010

JEL classification:
D9
E3
G12

Keywords:
Option pricing
Difference in beliefs
Incomplete markets
Foreign exchange

ABSTRACT

This paper studies the importance of heterogeneous beliefs for the dynamics of asset prices. We focus on currency markets, where the absence of short-selling constraints allows us to perform sharper tests of theoretical predictions. Using a unique data set with detailed information on foreign-exchange forecasts, we construct an empirical proxy for differences in beliefs. We show that this proxy has a strong effect on the implied volatility of currency options beyond the volatility of macroeconomic fundamentals. We document that differences in beliefs impact also on the shape of the implied volatility smile, on the volatility risk-premiums, and on future currency returns.

1. Introduction

Standard asset pricing theories have difficulty explaining episodes that are not simply linked to fundamentals. Notable examples in the dynamics of capital markets are the equity premium puzzle or the excess volatility puzzle. These puzzles have motivated an increasingly large literature over the last couple of decades that explores the general equilibrium implications of uncertainty for asset prices. There are three important directions in this literature. The first has focused on economies with homogeneous investors that are uncertain and learn about the state of the investment opportunity set (e.g., David, 1997; Veronesi, 1999; Brennan, 1998; Brennan and Xia, 2001a; Brennan and Xia, 2001b; David, 2008). A second strand of the literature has investigated the effect of knightian uncertainty in economies in which a single agent makes decisions that are robust to model misspecification doubts about the real economy (e.g., Hansen and Sargent, 2005; Maenhout, 2004; Anderson, Hansen, and Sargent, 2003; Cagetti, Hansen, Sargent, and Williams, 2002; Hansen and Sargent, 2007; Leippold, Trojani, and Vanini, 2008).[1] The third stream explores the implications of multiple agents with different beliefs about the growth rate of the economy. In these models, the interaction and

☆ We thank an anonymous referee, Karim Abadir, Philippe Bacchetta, Gurdip Bakshi, Bruno Biais, Michael Brandt, Mikhail Chernov, Joost Driessen, Darrell Duffie, Bernard Dumas, Tim Johnson, Chris Jones, Antonio Mele, and conference participants at the American Finance Association 2008 meeting in New Orleans, CEPR meeting in Gerzensee, the 2008 Duke-UNC Asset Pricing Conference, the European Finance Association 2007 meeting in Llubjana, the Swiss Finance Meetings, HEC Paris, Imperial College, London Business School, London School of Economics, Queen's University, University of Amsterdam, and University of Oxford for comments. We would like to thank both Citigroup and Lehman Brothers for their help collecting the market data used in this study. We are also grateful to Reuters for their forecast database. The usual disclaimer applies.

* Corresponding author. Tel.: +44 7932 728656; fax: +44 020 7594 9210.
E-mail address: a.buraschi@imperial.ac.uk (A. Buraschi).

0304-405X/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jfineco.2010.07.001

[1] An alternative tool for handling model misspecification doubts is based on multiple recursive priors. Early contributions of that work are Gilboa and Schmeidler (1989), Epstein and Wang (1994), Epstein and Schneider (2003), Anderson, Hansen, and Sargent (2003), Chen and Epstein (2002), Epstein and Miao (2003).

intertemporal risk sharing of agents with different beliefs affect asset prices and equilibrium risk premia (e.g., Scheinkman and Xiong, 2003; Buraschi and Jiltsov, 2006). In Buraschi and Jiltsov (2006), this interaction impacts the stochastic discount factor thus affecting both expected returns and also the volatility smile and volatility risk premium. Differences in beliefs are motivated as the result of heterogeneous learning about some unobservable and uncertain fundamentals of economy, such as the growth rate of dividend. Depending on the learning behavior of agents, the difference in beliefs is *stochastic* and may directly affect asset prices.

Within the class of heterogeneous belief models, one can identify at least two different approaches with somewhat different empirical implications. The first (neoclassical) approach considers rational, risk-averse agents, with identical preferences and endowments, no trading frictions, but incomplete and heterogenous information. In this case, even if dividends have constant volatility, agents have different optimal portfolio demands and in equilibrium the stochastic discount factor varies with different levels of difference in beliefs.[2] As a result, differences in beliefs directly increase expected returns (e.g., Buraschi, Trojani, and Vedolin, 2008), generate excess realized volatility (e.g., Dumas, Kurshev, and Uppal, 2009), and impact option-implied volatility (e.g., Buraschi and Jiltsov, 2006). The second approach builds on the interaction between behavioral biases and trading frictions, the Miller (1977) conjecture. For example, Scheinkman and Xiong (2003) study a model with overconfident and risk-neutral agents. They show that, in this context, short-selling constraints can support rational asset price bubbles in equilibrium. Empirically, Diether, Malloy, and Scherbina (2002) use equity returns to show that negative opinions are not fully revealed and thus difference in beliefs has a negative impact on expected stock returns. Understanding the link between uncertainty, differences in beliefs and asset prices is very important. However, the empirical evidence is mixed. Diether, Malloy, and Scherbina (2002) focus on equity expected returns and find supporting evidence for Miller's (1977) conjecture. On the other hand, Anderson, Ghysels, and Juergens (2005) and Buraschi and Jiltsov (2006) use a different time period and empirical methodology and find supporting evidence for a neoclassical (i.e., risk-based) interpretation of the impact of differences in beliefs on asset prices.

Knowing whether differences in beliefs matter even in the absence of short-selling constraints is important to improve our understanding of the link between uncertainty and asset prices. The goal of this paper is to study empirically this link and to disentangle the different implications of heterogeneous agent models. The analysis entails three key elements. The first is the ability to measure empirically the dispersion of beliefs. Uncertainty is a precondition for differences in beliefs to matter in theoretical models, but it is considerably harder to measure economic uncertainty than dispersion of beliefs. The second key aspect is to focus on implications for both the first moment and the second moment, given that these are markedly different for different models. Finally, we select a setting that gives the best empirical opportunity to neoclassical models by focusing on the foreign-exchange (FX) market. We choose this market for a number of reasons. First, the marginal investor in the FX market is unaffected by short-selling constraints. Second, there is a large literature that documents the relative unimportance of macro fundamentals in explaining the dynamics of exchange rates (e.g., Meese and Rogoff, 1983), thus leaving open the question of what might possibly help explain its dynamics. Third, exchange rates are unlikely to be driven by private information and are thus an ideal laboratory to investigate a different source of information-driven trades, such as heterogeneous beliefs in the absence of private information. In this context, a rejection of the neoclassical risk-based null hypothesis that differences in beliefs explain asset price dynamics would be very important as one may argue that, a fortiori, this hypothesis would be rejected even more strongly in less favorable contexts. Obviously, failure to reject the null hypothesis in FX markets would not necessarily rule out the explanatory power of the Miller's (1977) argument in other markets with binding short-selling constraints or indeed, the possibility of some other, as yet unspecified, behaviorial explanation.

The description of a noteworthy event in the foreign exchange (FX) markets is useful to illustrate intuitively the potential link between exchange rates and differences in beliefs. At the beginning of 2004, the FX market was approaching an 'overwhelming consensus' that the yen would continue to appreciate against the dollar.[3] The combination of a historically small interest rate differential, large and opposing current account imbalances, and strong economic growth in Japan generated the belief that only the Bank of Japan's persistent currency intervention was stopping further declines in the dollar/yen exchange rate from becoming a 'one-way bet.'[4] Consistent with this scenario, at the beginning of 2004 about 75% of forecasters in our sample were predicting yen appreciation over the coming three months. Fig. 1, Panel 1, shows an histogram of dollar/yen forecasts at the beginning of January 2004 that illustrates the low disagreement of professional investors. At the same time, dollar/yen implied volatility fell to a multi-year low. Over the subsequent months, however, the Bank of Japan intervened massively spending more than $100 billion buying dollars and selling yen. These record levels of intervention, coupled with veiled U.S. criticism of Japan's actions, created active disagreement in the market over the future of the dollar/yen exchange rate.[5] With

[2] The implications of heterogeneous beliefs on the equilibrium risk premium and interest rates have originally been studied by Detemple and Murthy (1994), Zapatero (1998), and Basak (2000). See Basak (2005) for a survey of this literature.

[3] See, for example, *Institutional Investor* February 12, 2004. For further details on the FX market in the first quarter of 2004, see the report of the Federal Reserve Bank of New York at http://www.newyorkfed.org/newsevents/news/markets/2004/fxq104.pdf.

[4] See, for example, *Financial Times* January 9, 2004.

[5] Shortly after Japan's most significant intervention, U.S. Treasury Secretary John W. Snow warned of the dangers of propping up currencies artificially. For an account of these events, see *Business Week*, March 22, 2004, "Don't Let Japan's Mr. Dollar Get Away With It", http://www.businessweek.com/magazine/content/04_12/b3875047.htm.

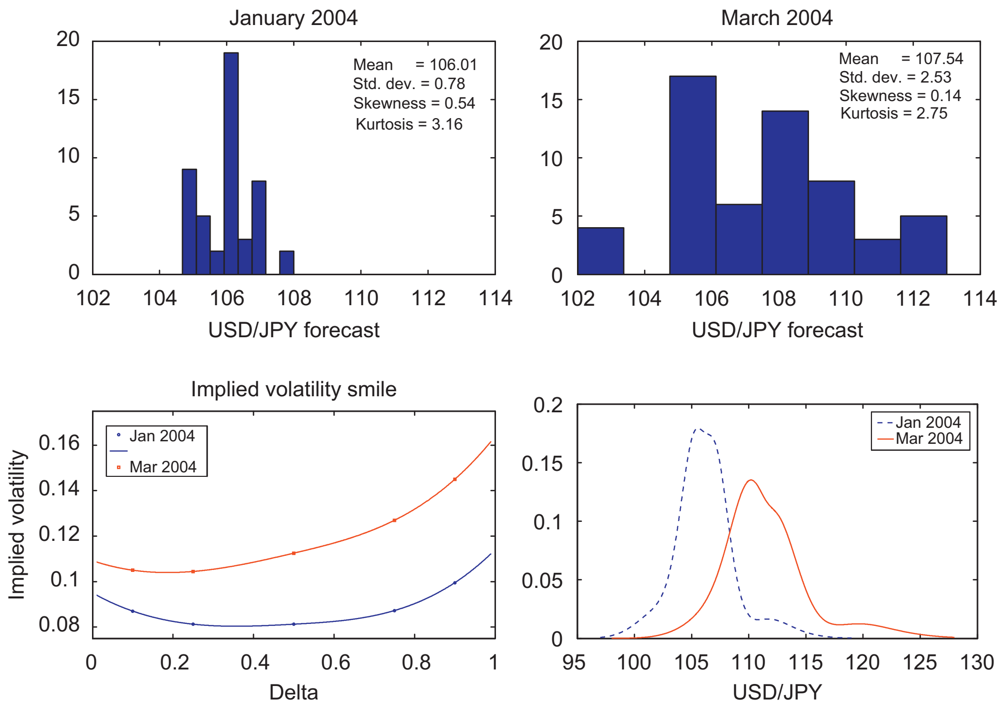

Fig. 1. Beliefs, implied volatility smiles, and risk-neutral densities. The first and second plot show histograms of the one-month foreign-exchange forecasts for the USD/JPY at the beginning of January 2004 and March 2004. The plots also report the mean, standard deviation, skewness, and kurtosis of the forecasts. The third panel is the cubic-spline interpolated implied volatility of one-month currency options as a function of the delta for January and March 2004. The fourth panel shows the option-implied state price density of the USD/JPY in January and March 2004.

unchanged economic fundamentals, the impact of these events on the level of the exchange rate itself was not dramatic (the yen weakened about 3% against the dollar during February having appreciated by almost 2% in January). Its impact on market expectations, however, was dramatic, as the yen appreciation consensus was broken and in March the dispersion of FX forecasts increased dramatically (see Fig. 1, Panel 2). At the same time, the implied volatility on one-month options increased by 50% (about 4%) and the slope of the implied volatility smile increased substantially, with important implications on the option-implied state price densities (see Fig. 1, Panels 3 and 4).[6] At least in this example, disagreement among market participants was strongly positively correlated with the implied volatility. Is this an isolated event, or is this an example of a more generalized and deeper link between uncertainty and asset prices? Studying such a link is important from a theoretical perspective since the implied volatility depends not only on the characteristics of the distribution under the physical probability measure but also on the properties and dynamics of the price of risk.

The empirical analysis of this paper is thus based on two unique data sets. The first one is based on over 10 years of daily options data for three currency pairs (USD/EUR, USD/JPY, and GBP/USD), at four different constant maturities, and different moneyness levels. Option markets are an ideal testing ground, because the estimation of the relevant variables is simple and the neoclassical approach with heterogeneous agents provides specific predictions on volatility under the risk-neutral measure. The second data set is a detailed record of currency forecasts made by a large cross-section of professional market participants that we use to obtain a proxy for the differences in beliefs. More specifically, at each survey date, from the distribution of FX forecasts we compute measures of cross-sectional dispersion, skewness, and kurtosis. These measures capture different characteristics of the differences in beliefs in the FX market. It is important to highlight that another key advantage of our FX focus is that our data allow us to measure dispersion of beliefs directly on the underlying asset. In contrast, other papers in the literature construct proxies for the differences in beliefs about factors that are only indirectly related to the underlying asset prices. For example, Anderson, Ghysels, and Juergens (2005) construct proxies of differences in beliefs about stock prices using forecasts about earnings, which require some additional modeling assumptions linking earnings to asset prices.

The design of our empirical analysis is the following. We first study the effects of difference in beliefs on currency option markets and then turn to an analysis of the effects

[6] We obtain the option-implied state price density as the second derivative of the option pricing formula with respect to the strike price (Breeden and Litzenberger, 1978). The option pricing formula is the Black-Scholes formula, where we replace the constant volatility parameter with a cubic-spline interpolated implied volatility smile (see, among others, Shimko, 1993).

on FX underlying returns. We find that differences in beliefs have statistically significant and economically important effects on the implied volatility of currency options. This finding is valid at four different horizons and on three different underlying exchange rates. For instance, for the USD/JPY exchange rate, our proxy for the difference in beliefs explains up to 64% of the contemporaneous variation in at-the-money implied volatilities. More specifically, a high dispersion of beliefs (one-standard deviation above the mean) increases implied volatility by about 2%. Furthermore, we show that the impact of dispersion of beliefs on currency implied volatility goes beyond the effect of the current uncertainty about macroeconomic fundamentals, proxied by the volatility of 12 macroeconomic indicators. Lastly, we also investigate the previous link in the context of predictive regressions and find that, even after controlling for the current level of volatility, the differences in beliefs exhibit predictive power in explaining the dynamics of future implied volatility.

We also find that differences in beliefs explain the shape of the implied volatility smile. For instance, our empirical measures of the dispersion of beliefs for the USD/EUR exchange rate explain, depending on the horizon, between 11% and 21%, of the contemporaneous variation in the slope of the implied volatility smile. Furthermore, in periods with high dispersion of beliefs, the slope of the implied volatility smile changes by about one-third. We obtain similar results for the other two currency pairs. This result is interesting since it implies that agents with different subjective valuations about market fundamentals engage in risk-sharing trading, thus affecting the relative prices of options with different moneyness. This finding also provides an economic explanation for the evidence of strong and time-varying asymmetry of the conditional risk-neutral distribution of currency returns shown in Carr and Wu (2007).

Another important result obtained in the currency option markets is the positive impact of dispersion of beliefs on the volatility premium (i.e., the difference between implied volatility and realized volatility). For example, for the USD/JPY exchange rate, the dispersion of foreign-exchange forecasts alone explains between 37% and 60% of the variation of the volatility spread. This finding is important since it suggests that differences in beliefs directly affect the pricing kernel and thus, the difference between the physical and the risk-neutral probability measure is correlated with differences in beliefs. Our results here are consistent with the findings of Driessen and Perotti (2003), who use currency option prices to study the evolution of investor confidence over the probability that individual currencies would eventually converge to the euro. The convergence risk, which may reflect uncertainty over policy commitment as well as exogenous fundamentals, induces a level of implied volatility in excess of actual volatility.

As a capstone to our analysis, we test whether our three main results above hold also in the cross-section of currencies. We find that the currency pairs characterized by larger disagreement are also the ones with larger implied volatility, steeper implied volatility smiles, and larger volatility premiums. These findings are valid for the full sample, for different subsamples, and within a difference-in-difference regression specification.

Finally, we turn to the effects of differences in beliefs on FX underlying returns. The empirical analysis proceeds along two related lines. We first look at carry trades, a simple investment strategy based on borrowing in low interest rate currencies and investing in high interest rate currencies. Generally, the Sharpe ratios of such trades have been higher than that of an investment in the Standard and Poor (S&P) 500 index and a few papers provide a risk-based explanation (e.g., Lustig, Roussanov, and Verdelhan, 2008). Our empirical analysis shows that large differences in beliefs are strongly correlated with carry trade expected returns, consistent with the interpretation of the difference in beliefs being a priced risk factor. This result is important since it suggests a simple structural explanation that links contemporaneous changes in FX volatility and carry trade expected returns, which has been one of the most thought-provoking empirical regularities in the FX literature. We also investigate this question within the extensive literature on the foreign-exchange expectation hypothesis puzzle. In particular, we show that the dispersion in beliefs explains future currency returns beyond what is embedded in the forward premium. More specifically, the excess returns obtained from the violation of the expectation hypothesis at times of low uncertainty, low differences in beliefs, and normal funding liquidity compensate the appreciation of the low-yield currencies at times of high uncertainty, large differences in beliefs, and funding liquidity problems. In states of the world in which the differences in beliefs are the largest, low interest rate currencies appreciate the most and carry trades suffer large negative returns. These results support our findings in the currency option market and are consistent with the interpretation of the differences in beliefs being a component of a time-varying foreign-exchange risk premium.

The remainder of the paper proceeds as follows. In Section 2, we analyze the theoretical link between heterogeneous beliefs and asset prices in the literature to derive testable implications. Section 3 provides further details about the data and conducts some preliminary analysis. In Section 4, we show our empirical results. Section 5 concludes.

2. Theoretical framework and literature review

Neoclassical models with uncertainty (i.e., incomplete information about fundamentals) and heterogeneous agents study the link between uncertainty about dividend growth rates and differences in beliefs. For example, Proposition 1 in Scheinkman and Xiong (2003), Section I.C in Buraschi and Jiltsov (2006), and Lemma 2 in Dumas, Kurshev, and Uppal (2009), characterize the link between the volatility of the expected growth of aggregate dividends (i.e., economic uncertainty) and the stochastic process for difference in beliefs. Buraschi, Trojani, and Vedolin (2008) study a heterogeneous economy with learning and distinguish between average subjective uncertainty and the difference in subjective uncertainty.

They show that "when agents perceive different levels of uncertainty, the higher the average subjective uncertainty across agents, the higher is the average disagreement," thus suggesting an empirically viable way to build a proxy for average uncertainty. They show that this property is consistent with the empirical observation that during periods of high economic uncertainty, the average level of the dispersion in beliefs increases. At the same time, these models show that the interaction of agents with different beliefs has important implications for equity returns, implied volatility (e.g., Buraschi and Jiltsov, 2006), and credit spreads (Buraschi, Trojani, and Vedolin, 2008).[7]

In this class of models, Buraschi and Jiltsov (2006) provide a simple generalization of the standard general equilibrium Lucas economy with rational agents, identical preferences and endowments, but incomplete and heterogeneous information. Agents select their optimal consumption plans based on two different sets of beliefs (on the probability distribution dQ), which creates a wedge in equilibrium between their marginal utilities. They argue that when two agents have different beliefs $dQ_1 \neq dQ_2$, the socially optimal ex ante allocation between the two agents must satisfy $0 = \int [u'(c_1) + u'(c_2)\, dQ_2/dQ_1]\, dQ_1$. In the special case of homogeneous beliefs, that is when $dQ_1 = dQ_2$, one obtains the standard result that the distribution of consumption depends exclusively on preference characteristics, such as discount rates, risk aversion, the elasticities of intertemporal substitution, and differences in the distribution of initial wealth. However, in the more general case in which $dQ_1 \neq dQ_2$, changes in the difference of beliefs dQ_2/dQ_1 affect the relative weights of the two agents in the representative agent utility function. This has important asset pricing implications. When agents disagree, the equilibrium stochastic discount factor is affected by two components. The first component is related to the aggregate endowment shocks, as in standard Lucas endowment economies, since these affect the total amount of resources $c_1 + c_2$ that can be shared by the two agents. The second component is the difference in beliefs dQ_2/dQ_1: since this affects the relative discount factors of the two agents, given the total amount of resources available, it induces agents to engage in additional state-contingent risk-sharing contracts with a number of important implications. The agent who is more optimistic about the endowment process provides, in equilibrium, insurance to the pessimistic one and thus, the changes in difference in beliefs affect the observed distribution of consumption in equilibrium. Moreover, the aggregate stochastic discount factor can be volatile even when the endowment process is smooth, which makes differences in beliefs a priced source of risk.

This basic theoretical framework can be extended in two main directions. First, the model can be extended to economies with multiple goods. This extension has been studied in Buraschi, Trojani, and Vedolin (2009) in the context of a Lucas orchard economy to investigate the occurrence of endogenous co-movement and time-varying correlations in periods of uncertainty and large differences in beliefs. When preferences in the two consumption goods are separable, the main implications of the original model can be extended to a multiple goods economy. The second extension is to incorporate money in the utility functions of the two agents (e.g., Basak, 2005; Croitoru and Lu, 2008). In this setting, heterogeneity of beliefs concerning monetary policy impacts both the price of money and stock price volatility. Since an exchange rate is the relative price of a domestic and foreign good, each expressed in terms of their numeraire, the appropriate theoretical framework to analyze the effects of differences in beliefs on exchange rates should incorporate both extensions, multiple goods (domestic and foreign), and money in the utility function. Although it would be interesting to build on the current literature and study such a model with both features combined, developing a structural general equilibrium FX model is beyond the scope of the current paper. When preferences are separable in both money and consumption goods, it is realistic to expect that the theoretical relations between differences in beliefs and asset prices borne out in the simpler models are similar.

The understanding of the impact of differences in beliefs is also an important theme in the behavioral finance literature. In this strand of the literature, however, the theoretical link between differences in beliefs and expected returns is reversed. Disagreement affects asset prices not because of the existence of a risk-based channel, but because of a limits-to-arbitrage argument based on Millers (1977) conjecture: when frictions prevent the revelation of negative opinions (e.g., short-sale constraints), an increase in the dispersion in beliefs decreases expected returns. Although this theory does not have explicit predictions for the volatility of returns, Diether, Malloy, and Scherbina (2002) obtain empirical results that are consistent with Miller's model for U.S. stock market expected returns. In particular, they find that in the period 1983–1991, the dispersion of analysts short-term earnings forecasts (obtained from the Institutional Brokers Estimate System (IBES)) is negatively correlated with future returns. However, other studies (e.g., Qu, Starks, and Yan, 2003; Anderson, Ghysels, and Juergens, 2005) find opposite results when using different time periods and a different empirical approach.[8]

We contribute to this literature in several respects. First, since equity analysts most often provide forecasts of future earnings as opposed to future returns, these studies need to make assumptions on a pricing model to map earnings forecasts to stock price forecasts.[9] In our study,

[7] These approaches are part of a theoretical literature where heterogeneous beliefs of multiple agents enter the stochastic discount factor (e.g., Detemple and Murthy, 1994; Zapatero, 1998; Basak, 2000); for a survey, see Basak (2005)

[8] Dumas, Kurshev, and Uppal (2009) extend this literature by studying an economy in which some agents overreact to a public signal, introducing an additional risk factor in the economy, which increases both the equity premium and the volatility of stock and bond returns. In contrast to Miller-type behavioral theories, however, this paper does not rely on short-selling constraints.

[9] Anderson, Ghysels, and Juergens (2005), for instance, consider a constant dividend growth model and assume that all earnings are distributed as dividends.

we work directly with price (exchange rate) forecasts. Thus, our results are not sensitive to these types of modeling assumptions. Second, there is considerable evidence that analysts' forecasts of earnings are affected by agency issues between firms and investment banks (Rajan and Servaes, 1997; Daniel, Hirshleifer, and Teoh, 2002), which has recently led to the introduction of several regulatory measures. Given the nature of the FX market, agency biases are much less of a concern in our analysis. Moreover, the limited role of private information in FX markets also reduces these agency concerns. Third, we investigate the link between dispersion in beliefs and the volatility risk premium (the spread between the option-implied risk-adjusted volatility and the realized physical volatility). This allows us to study more directly the extent to which the dispersion in beliefs is a priced risk factor. Fourth, the FX market is very liquid and not affected by short-selling constraints, so that we can investigate the role of differences in beliefs in the absence of key market frictions and provide a clean environment in which to examine neoclassical models with heterogeneous agents. Finally, our FX survey features an average of about 50 market participants' forecasts, which is more than three times larger than the number of analyst forecasts included in the First Call/Thomson equity survey. The computation of our dispersion measure is thus likely to be more reliable.

In this paper our empirical focus is on currency options markets though we also look at underlying spot and forward markets. Options are a natural testing ground not only because their state-contingent payoffs are well-suited to study agents with heterogeneous beliefs, but also because they offer a simple way to measure volatility. More specifically, we can investigate four empirical questions.

First, is the level of the option-implied volatility affected by the difference in beliefs? This link is natural in neoclassical models with differences in beliefs, given that heterogenous agents engage in state-dependent risk-sharing contracts and implied volatility represents the price of these contracts. More specifically, the pricing kernel used in option pricing depends on both the cash flow and difference in beliefs processes.[10] However, in Miller-type behavioral models where dispersion of beliefs is not an additional risk factor, there are no predictions about implied volatility. Documenting a significant relation would thus lend support to the risk-based explanation of difference in beliefs outlined in neoclassical models. As for the sign of this relation, all the extant neoclassical models cited above derive a positive relation between differences in beliefs and the volatility of stock returns. While one may expect to obtain the same positive relation in foreign-exchange markets, we leave the determination of the sign and magnitude of the relation as an empirical question, since one could potentially assume different preferences for the two agents on the domestic and foreign good that would generate offsetting effects. In this sense, our empirical results will provide some guidance for any theoretical model studying differences in beliefs in an international context.

Second, is the cross-section of option prices (e.g., the slope of the volatility smile) correlated with the difference in beliefs? Since marginal rates of substitution across states depend on agents' beliefs, the subjective valuation of out-of-the-money options differs across agents with different beliefs. In neoclassical models with differences in beliefs, in equilibrium agents engage in state-contingent risk sharing with the pessimists purchasing protection from the optimists. The larger the difference in beliefs, the higher the equilibrium cost of protection required by the optimists to provide insurance, with immediate implications for equilibrium option prices in the cross-section of moneyness. In FX markets, the definition of a bad scenario that agents want to hedge against depends on several factors and it could be country-specific. Thus, it becomes an empirical question to determine the sign of the relation between differences in beliefs and the relative pricing of out-of-the-money puts and calls and its variation over different business cycles.

Third, is the size of the implied-realized volatility spread (the "volatility risk premium") affected by the level of the dispersion in beliefs? The difference between the physical and risk-neutral probability measure depends on the characteristics of the stochastic discount factor of the representative agent, which in a neoclassical economy with heterogeneous beliefs depends on the relative weights of the two agents in the representative agent utility function. Given that the relative weights of the two agents depend on the difference in beliefs, these models imply a link between disagreement and the volatility spread. This evidence is consistent with the findings in Buraschi and Jackwerth (2001) who show that the smile is priced and the dynamics of the option implied volatility is not spanned by traditional single factor models.

Last, is there a relation between differences in beliefs and first moments in FX markets? The literature on the forward premium puzzle has found that forward exchange rates are biased predictors of future spot rates possibly because of a time-varying risk premium (Fama, 1984). In neoclassical models with heterogeneous agents, dispersion of beliefs enters the stochastic discount factor and is thus a natural candidate to explain future currency returns.[11] For example, the impact of changes in differences in beliefs on the stochastic discount factor could explain currency returns through the liquidity spiral model of Brunnermeier and Pedersen (2009). In this setting, shocks to uncertainty/differences in beliefs affect the risk tolerance of traders, so that when funding constraints are hit, the lower risk tolerance induced by higher differences in beliefs triggers the unwinding of carry positions with a consequent impact on FX returns.

[10] In Dumas, Kurshev, and Uppal (2009), the overreaction of overconfident investors to a public signal generates excess volatility under the physical measure. This model predicts a positive relation between dispersion of beliefs and implied volatility only to the extent that implied volatility is correlated with physical volatility.

[11] This is also the same logic that guided the earlier heterogeneous beliefs models (e.g., Detemple and Murthy, 1994) in explaining the equity premium puzzle.

Thus, we investigate the existence of a link between periods of appreciation of low interest rate currency and large uncertainty/differences in beliefs.

3. Data and preliminaries

In this section, we describe our data sets in detail, we show the approach we use to obtain an empirical proxy for differences in beliefs, and perform some preliminary analysis.

Currency option data: Our data set contains the daily currency option prices traded in the over-the-counter (OTC) market over the period April 1992 through December 2006, kindly supplied by Citigroup. The options are European-style and they are written on the Japanese yen price of the U.S. dollar (USD/JPY), the dollar price of the British pound (GBP/USD), the German mark price of the dollar (USD/DEM) before 1999, and the euro price of the dollar (USD/EUR) thereafter.[12] In the remainder of the paper, we will also denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999.

We keep track of the USD/DEM and USD/EUR as different currencies for at least two reasons. First, the German mark represented only about 32% of the European currency unit before the euro was introduced. Second, the European Monetary System during our sample period (1993–1999) was a limited-flexible exchange rate system, with defined bands in which the bilateral exchange rates of the member countries could fluctuate. When a market exchange rate reached either of these intervention points, the central banks were compelled to support these rates indefinitely through open market operations. In the traditional heterogeneous belief economy, the cost for the optimist to insure the pessimist is increasing in the level of disagreement. However, if currencies cannot effectively fluctuate beyond a prespecified band, it is a third agent, the central bank, that will take on the risk of the most extreme changes. For this reason, some of our empirical results, in particular the relation between the slope of the implied volatility smile and difference in beliefs, might be weaker in the case of the German mark.

Options on each currency pair have four constant times to maturity, at one month, three months, six months, and one year. The constant time-to-maturity feature of the OTC currency option market is useful as it allows us to sidestep the issue of interpolating between adjacent maturities that typically occurs with the preset maturity dates of exchange-traded options.

The OTC currency option market has very specific trading conventions.[13] Unlike bond or equity options that are typically traded in terms of option premiums at different strike prices, currency options trade in terms of implied volatilities. More specifically, implied volatility quotes are most commonly available in the form of three types of option combinations: the delta-neutral straddle, the risk reversal, and the strangle. A straddle is a portfolio of a call option and a put option with the same strike price and maturity. For the straddle to be delta-neutral, the strike price needs to be sufficiently close to the forward price to make the implied volatility quote of the straddle an at-the-money implied volatility (ATM). The risk reversal (RR) measures the difference in implied volatility between an out-of-the-money call option and an out-of-the-money put option. The moneyness level is defined in terms of the Black-Scholes delta of the option and is conventionally set at 25-delta.[14] The strangle (STR) corresponds more precisely to a butterfly spread and measures the difference between the average volatility of the two 25-delta options and the delta-neutral straddle implied volatility. In summary

$$\sigma_{ATM} = \sigma_{50\Delta Call}, \tag{1}$$

$$\sigma_{RR} = \sigma_{25\Delta Call} - \sigma_{25\Delta Put}, \tag{2}$$

$$\sigma_{STR} = \frac{\sigma_{25\Delta Call} + \sigma_{25\Delta Put}}{2} - \sigma_{ATM}. \tag{3}$$

Eqs. (1)–(3) show that the straddle is a measure of the level of the implied volatility, the risk reversal is a measure of the slope of the implied volatility smile, and the strangle is a measure of the curvature of the implied volatility smile. From the three quotes, we can derive the level of implied volatilities at the three levels of delta:

$$\sigma_{25\Delta call} = \sigma_{ATM} + \sigma_{STR} + \tfrac{1}{2}\sigma_{RR}, \tag{4}$$

$$\sigma_{25\Delta put} = \sigma_{ATM} + \sigma_{STR} - \tfrac{1}{2}\sigma_{RR}. \tag{5}$$

For our empirical analysis, we use implied volatilities from market quotes and obtained through Eqs. (4) and (5). As a result, we do not need to match the currency option data with underlying FX rates and interest rates, thus avoiding potential asynchronicity biases.

Table 1 contains summary statistics for our sample of currency options related to the three option combinations defined by Eqs. (1)–(3). The sample averages for at-the-money implied volatilities show that USD/JPY is the most volatile currency pair in our sample, followed by USD/EUR$^+$ and GBP/USD. The average term structure of at-the-money implied volatility is upward-sloping for all the underlyings. The analysis of the risk reversal combination reveals that the sample average for the USD/JPY and the USD/EUR$^+$ is negative at all maturities, implying that the out-of-the-money put options are, on average, more expensive than the corresponding out-of-the-money call options during our sample period. In contrast, for the GBP/USD, the sample average of the risk reversal is close to zero. In all cases, the standard deviation is large, suggesting that risk reversals vary greatly over time, as shown also in Carr and Wu (2007). The strangle combination is significantly positive for all currency pairs and maturities, peaking for the USD/JPY

[12] The currency pairs are all expressed with the low-yielding currency as the numeraire.

[13] We describe the trading conventions that are key to understanding the properties of our data set. For further details on the currency option market practices see, among others, Malz (1997).

[14] A 25-delta call corresponds to a call option with Black-Scholes delta of 0.25, and a 25-delta put corresponds to a put option with Black-Scholes delta of −0.25. Another level of moneyness that is typically traded in the currency option market is the 10-delta.

Table 1
Currency options summary statistics.

This table presents summary statistics for our sample of currency options written on USD/JPY, GBP/USD, USD/EUR$^+$ (USD/DEM before 1999, USD/EUR thereafter). We report the mean, standard deviation, and number of observations on the straddle (σ_{ATM}), on the 25-delta risk reversal (σ_{RR}), and on the 25-delta strangle (σ_{STR}), for four option maturities. Data are daily from May 1993 to December 2006. We denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999.

		USD/JPY	GBP/USD	USD/EUR$^+$	USD/JPY	GBP/USD	USD/EUR$^+$
		One-month maturity			Three-month maturity		
σ_{ATM}	Mean	10.79	8.41	10.13	10.94	8.75	10.35
	Std.dev.	2.90	1.68	2.05	2.62	1.52	1.78
	Obs.	3597	3597	3597	3597	3597	3597
σ_{RR}	Mean	−0.75	−0.02	−0.24	−0.75	0.01	−0.28
	Std.dev.	0.90	0.33	0.45	0.82	0.28	0.38
	Obs.	3122	3122	3122	2515	2515	2515
σ_{STR}	Mean	0.35	0.21	0.24	0.36	0.21	0.24
	Std.dev.	0.13	0.06	0.08	0.13	0.05	0.06
	Obs.	3122	3122	3122	2515	2117	2515
		Six-month maturity			One-year horizon		
σ_{ATM}	Mean	11.07	8.95	10.47	11.18	9.15	10.56
	Std.dev.	2.56	1.45	1.66	2.50	1.44	1.57
	Obs.	3597	3597	3597	3597	3597	3597
σ_{RR}	Mean	−0.72	0.02	−0.29	−0.74	0.03	−0.30
	Std.dev.	0.86	0.27	0.36	0.92	0.26	0.34
	Obs.	2515	2515	2515	2515	2515	2515
σ_{STR}	Mean	0.37	0.22	0.24	0.37	0.22	0.25
	Std.dev.	0.12	0.04	0.06	0.12	0.04	0.06
	Obs.	2515	2515	2028	2515	2421	2515

Table 2
Foreign-exchange forecasts summary statistics.

This table presents summary statistics for our sample of foreign-exchange forecasts on USD/JPY, GBP/USD, USD/EUR$^+$ (USD/DEM before 1999, USD/EUR thereafter). For each forecasting horizon, we show the number of individual forecasts, the number of monthly surveys, and the number of forecasters for each survey. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999.

		USD/JPY	GBP/USD	USD/EUR$^+$
One-month horizon				
Individual forecasts	No.	7566	6987	7679
Monthly surveys	No.	156	142	156
	Max no.	66	65	66
Forecasters	Ave no.	47.8	48.6	49.2
	Min no.	27	30	27
Three-month horizon				
Individual forecasts	No.	7640	7006	7729
Monthly surveys	No.	156	142	156
	Max no.	67	66	67
Forecasters	Ave no.	48.2	48.6	49.5
	Min no.	30	29	27
Six-month horizon				
Individual forecasts	No.	7562	6948	7643
Monthly surveys	No.	156	142	156
	Max no.	66	65	66
Forecasters	Ave no.	47.7	48.2	49.0
	Min no.	27	29	27
One-year horizon				
Individual forecasts	No.	7521	6896	7581
Monthly surveys	No.	156	142	156
	|Max no.	65	64	65
Forecasters	Ave no.	47.4	47.8	48.6
	Min no.	27	29	27

currency pair. This is an indication that the risk-neutral distribution of all three currencies exhibits fat tails. The strangle spreads are much more stable than the risk reversals, as the lower standard deviations demonstrate.

The FX forecasts: Our data set on differences in beliefs is based on the full set of forecasts that make up the Reuters survey of FX forecasts. The data are sampled at a monthly frequency from May 1993 to December 2006. At the beginning of each month, about 50 market participants provide their forecasts of future exchange rates. The Reuters survey is generally conducted on the first Tuesday of the month and released on Wednesday (we have information on all the survey collection dates). In our empirical analysis, we use the timing convention of the day in which the survey is conducted. We also check the effect of using as alternative convention—the time the survey is released: all results hold with no significant differences. The forecast horizons are set to be one month, three months, six months, and 12 months, respectively. Table 2 contains summary statistics for the FX forecasts. For all forecasting horizons and currency pairs, we have more than 6,000 individual forecasts over 140 monthly surveys. The average number of forecasters in each survey is almost 50 in all cases and we always have at least 27 major market participants in the survey. This fares well against surveys on company earnings that typically feature less than one-third of this number of respondents.

Besides offering a meticulous archive of individual forecasts (the longest uninterrupted sample available), the Reuters survey has a number of advantages over other FX forecast surveys such as those undertaken by Consensus Economics, Wall Street Journal (WSJ), Centre for European Economic Research (ZEW), Blue Chip, and Forecasts Unlimited (formerly the Financial Times currency forecasts and the Currency Forecast Digest). First, since it is conducted by the key FX news provider, it is very much focussed on FX market participants whereas other surveys often include many other forecasters such as professional forecast firms, corporations, and academic institutions. We estimate that around 95% of contributors to the Reuters survey are active market participants, compared to 85% for Consensus Economics and even less for the other major surveys. This is important since, as Ito (1990) finds, these other forecasters are not comparable with those actively trading in foreign exchange. Second, the pool of forecasters is relatively constant. Other surveys have both gaps in coverage (missing individual months and in some cases years) and a relatively rapid turnover of contributors. Third, it is the only survey that collects one-, three-, six-, and 12-month-ahead forecasts, thus offering the most complete short-term coverage.

FX forecast data have already been used to help understand market expectations. Frankel and Froot (1987) use FX forecasts and find that expectational errors play a significant role in explaining the forward premium puzzle. More recently, Bacchetta, Mertens, and Van Wincoop (2006) find a strong relationship between forecast errors and excess returns, confirming the link between FX forecasts and underlying market behavior. On the other hand, studies that focus on the predictive power of the forecasts (e.g., Bofinger and Schmidt, 2004) generally find no informational content in FX forecasts that helps in predicting future exchange rates.

A maintained assumption in our surveyed theoretical interpretation of the link between beliefs heterogeneity and currency markets is that the average agent has no private information on the future dynamics of the exchange rate. We briefly investigate the robustness of this assumption by estimating the following regression:

$$\hat{S}_T - S_t = \alpha + \beta(\hat{S}_{t,T} - S_t) + e_t, \tag{6}$$

where S_t represents the foreign-exchange rate at time t and $\hat{S}_{t,T}$ represents the average foreign-exchange forecasts on day t with horizon T. We also compute the root mean squared error (RMSE) of the average forecast and compare it with the RMSE of a random walk (RMSErw).

We find that professional forecasters do not predict future exchange rates better than a random walk (Table 3). The explanatory power of the forecasting regression is generally very low, with the difference between the average forecast and the spot rate never explaining more than 6% of the change in future exchange rates. The slope coefficient is significantly different from zero for only one currency pair, GBP/USD, and is always significantly different from one. The RMSE is always higher than the RMSE of a random walk forecasting model.[15] Not surprisingly for such a liquid market, these results lend support to the conjecture that currency prices are mainly determined by the flow of public information about cash flows and discount rates and by the mechanism of price discovery, i.e., by the aggregation of heterogeneous private information (or heterogeneous interpretation of public information) through trading. Clearly, even if exchange rates are not predictable by the average trader, the observed heterogeneity in beliefs may still affect, as the theoretical literature shows, the risk-neutral measure and derivative asset prices.

Table 3

Do foreign-exchange forecasts have predictive power?

This table shows the results of estimating the following regression:

$$S_T - S_t = \alpha + \beta(M_{t,T} - S_t) + e_t,$$

where S_t represents the foreign-exchange rate at time t and $M_{t,T}$ represents the average foreign exchange forecasts on day t with horizon T. We analyze three currency pairs: USD/JPY, GBP/USD, USD/EUR$^+$ (USD/DEM before 1999, USD/EUR thereafter). We also compute the root mean squared error (RMSE) of the average forecast and compare it with the RMSE of a random walk (RMSErw). The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.

	USD/JPY	GBP/USD	USD/EUR$^+$
One-month horizon			
Constant	0.2730	0.0002	0.0006
β	0.0655	0.3075	0.1651
R^2	0.0015	0.0120	0.0155
RMSE/RMSErw	1.3045	1.0494	1.3816
Three-month horizon			
Constant	0.1985	0.0055	0.0079
β	−0.2675	0.4667**	−0.0618
R^2	0.0139	0.0272	0.0008
RMSE/RMSErw	1.3000	1.0139	1.2315
Six-month horizon			
Constant	0.1644	0.0107	0.0213
β	−0.0538	0.3597	−0.2350
R^2	0.0005	0.0151	0.0086
RMSE/RMSErw	1.2089	1.0467	1.2067
One-year horizon			
Constant	0.6375	0.0287**	0.0434**
β	−0.2013	0.7904**	−0.3935
R^2	0.0090	0.0629	0.0186
RMSE/RMSErw	1.3156	0.9678	1.1704

An empirical proxy for differences in beliefs: We construct an empirical measure for the difference in beliefs implicit in the cross-sectional distribution of FX forecasts. Heterogeneous belief models (e.g., Basak, 2005; Buraschi and Jiltsov, 2006) derive a sufficient statistic for the difference in beliefs that affects risk premiums. Specifically, they define the stochastic process ψ_t as the disagreement between the two agents about the growth rates of the observable processes scaled by their respective volatilities:

$$\psi_t = \Omega^{-1}(m_t^1 - m_t^2), \tag{7}$$

where m_t^n is the perceived drift of agent n of the observable processes and Ω is their volatility. The natural empirical proxy for Eq. (7) is the mean absolute deviation (MAD), i.e., the average absolute deviation of all forecasts from the mean forecast, scaled by its local volatility. In our empirical analysis, we measure the perceived drift of different agents m_t^n using the forecasts in terms of log returns with respect to the spot exchange rate. Expressing the forecast in log returns rather than levels has the desirable property of removing level effects from the time-series of the FX forecasts.

While Eq. (7) is an appropriate measure in the context of the models mentioned above, alternative proxies could be more suitable if these models are misspecified or in different theoretical settings. In our empirical analysis, we thus also consider higher-order moments and additional statistics of difference in beliefs, such as a simple standard deviation measure and an entropy measure.[16]

Table 4 shows summary statistics. Since the average number of forecasters in each survey is about 50

[15] The results of estimating Eq. (6) do not change when we use the median foreign-exchange rate forecast instead of the average foreign-exchange rate forecast.

[16] In the information theory literature, entropy is widely accepted as a measure of uncertainty in information systems. Entropy has been related in economics to the accuracy of predictions (Sandroni, 2000) and model uncertainty (e.g., Anderson, Hansen, and Sargent, 2003). In the discrete case, entropy has its maximum value when all probabilities are equal (uniform distribution, in the case of bounded supports), and the resulting value for entropy is the logarithm of the number of states. The intuition is that when any random event can occur with equal probability, one faces a situation with the highest uncertainty.

Table 4
Difference in beliefs in foreign-exchange forecasts.

This table presents summary statistics for the mean absolute deviation (MAD), the skewness, and the kurtosis of the cross-sectional distribution of forecasts. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR+ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999.

			USD/JPY	GBP/USD	USD/EUR+	USD/DEM	USD/EUR
		max	0.0311	0.0263	0.0248	0.0248	0.0224
	MAD	ave	0.0142	0.0101	0.0134	0.0122	0.0139
		min	0.0058	0.0059	0.0070	0.0076	0.0070
		max	1.5594	6.7005	1.2899	1.2899	1.2784
One-month horizon	Skewness	ave	−0.2151	−0.0478	0.1495	0.1490	0.1513
		min	−2.7184	−2.2878	−1.8880	−1.8880	−1.2530
		max	14.0450	46.6212	7.7532	7.7532	7.5144
	Kurtosis	ave	3.7382	4.4633	3.2631	3.4293	3.1949
		min	1.9067	1.8743	1.8088	1.8088	2.1372
		max	0.0530	0.0322	0.0373	0.0373	0.0372
	MAD	ave	0.0256	0.0175	0.0236	0.0222	0.0242
		min	0.0147	0.0102	0.0124	0.0128	0.0124
		max	1.1200	6.2487	2.1533	2.1533	1.3044
Three-month horizon	Skewness	ave	−0.1059	−0.1131	0.2482	0.2444	0.2847
		min	−2.8092	−5.5905	−2.0457	−2.0457	−1.3025
		max	13.1263	42.0335	10.4946	10.4946	9.4833
	Kurtosis	ave	3.3926	4.1744	3.2151	3.5426	3.0937
		min	1.7834	2.0986	1.7866	1.7866	1.9998
		max	0.0654	0.0382	0.0499	0.0448	0.0499
	MAD	ave	0.0375	0.0241	0.0333	0.0319	0.0341
		min	0.0228	0.0165	0.0181	0.0245	0.0181
		max	2.8892	5.6869	1.6646	1.6151	1.6646
Six-month horizon	Skewness	ave	0.1300	−0.1263	0.2189	0.0060	0.3689
		min	−1.7103	−3.7290	−1.6808	−1.6808	−1.2503
		max	15.8295	37.3594	8.2769	6.0783	8.2769
	Kurtosis	ave	3.3414	3.8324	3.2045	3.2129	3.1777
		min	1.8746	2.0347	1.4880	1.4880	1.6805
		max	0.0836	0.0470	0.0769	0.0549	0.0769
	MAD	ave	0.0518	0.0324	0.0448	0.0405	0.0471
		min	0.0264	0.0221	0.0246	0.0261	0.0246
		max	1.8026	5.0023	1.7149	1.2321	1.7149
One-year horizon	Skewness	ave	0.4185	−0.2039	0.2398	−0.2753	0.5386
		min	−1.3730	−2.1288	−2.1723	−2.1723	−0.3676
		max	8.6445	31.8274	11.5625	11.5625	9.1081
	Kurtosis	ave	3.5670	3.9591	3.5135	3.3906	3.5436
		min	1.6948	2.2310	1.8886	2.0596	1.8886

(see Table 2), the sample mean absolute deviation is a reliable estimate of the dispersion of FX forecasts. The maximum and minimum values for the mean absolute deviation show a relatively wide range around its mean. This indicates a significant time-variation in the dispersion of beliefs and that the dispersion is larger at longer forecast horizons. The average sample skewness alternates sign for different currencies/horizons, but the large maximum and minimum values indicate that it is positive for some periods and negative for some other periods in our sample. Summary statistics for kurtosis suggest that extreme forecasts are more likely at shorter horizons and that they occur more often for the GBP/USD currency pair.

Fig. 2 illustrates the time-series dynamics of the mean absolute deviation, skewness, and kurtosis of beliefs at different forecast horizons. The time-variation of the mean absolute deviation of beliefs is significantly different across our sample currency pairs, suggesting that they are not driven by a single factor and that a panel of data for different exchange rates adds important information. The skewness of beliefs switches sign over our sample period, and is less correlated at different forecast horizons. The kurtosis of beliefs is subject to frequent extreme peaks.

4. Empirical results

In our empirical analysis, we investigate the implications of the theoretical literature surveyed in Section 2. First, we analyze the relation between at-the-money implied volatilities and the differences in beliefs of foreign-exchange forecasters. Second, we extend this analysis to the entire implied volatility smile. Third, we provide evidence on the links between the currency volatility spread and FX differences in beliefs. After investigating the characteristics of the second moments of FX returns under both the risk-neutral and the physical measure, we finally show the extent to which this priced risk factor can also help explain the time-variation in

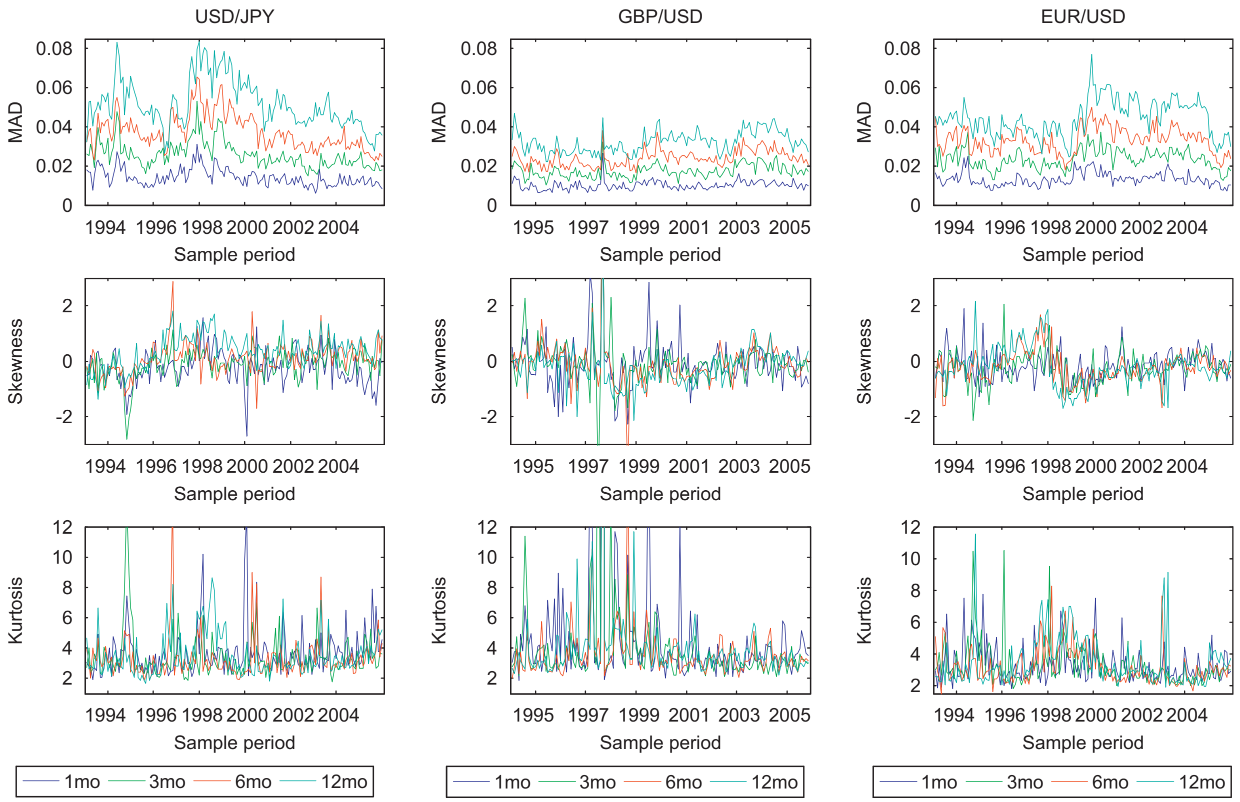

Fig. 2. Measures of foreign-exchange differences in beliefs. These plots show the dynamics of the mean absolute deviation (MAD), skewness, and kurtosis of foreign-exchange forecasts for the three currencies (USD/JPY; GBP/USD; USD/DEM before 1999, USD/EUR thereafter) and four forecast horizons in our sample.

conditional first moments with an analysis of carry trade returns and the FX forward premium puzzle.

4.1. At-the-money implied volatility and differences in beliefs

One of the main implications of the neoclassical models described in Section 2 is the link between the implied volatility of currency options and differences in beliefs. We measure empirically the difference in beliefs $\psi_{t,T}$ defined in Eq. (7) by computing the mean absolute deviation (MAD) of foreign-exchange forecast returns, i.e., the average absolute deviation of all forecasts from the mean forecast. We then estimate the following linear regression:

$$IV_{t,T} = \alpha + \beta_{\psi}\psi_{t,T} + \beta_{sk}SK_{t,T} + \beta_{ku}KU_{t,T} + \beta_{iv}IV_{t-30,T-30} + e_{t,T}, \quad (8)$$

where $IV_{t,T}$ represents the implied volatility of at-the-money (ATM) options on day t with maturity T.[17] Since currency options have fixed times to maturity, rather than fixed maturity dates, we observe every day implied volatilities for options with maturities of one month, three months, six months and one year. As a result, the horizon of the foreign-exchange forecasts matches precisely the maturity of the option and our empirical analysis is not affected by the data synchronicity issues shown in Cavaglia, Verschoor, and Wolff (1994). We complement the specification suggested by the theory with three additional regressors. To account for a comprehensive description of market participants' beliefs, we also include the skewness and the kurtosis of the cross-sectional distribution of the foreign-exchange forecasts on day t with horizons T, $SK_{t,T}$, and $KU_{t,T}$, respectively. These two variables also control for potential uninformed extreme forecasts that could introduce noise into the estimate of dispersion. Moreover, since there is evidence that stock market implied volatility has predictive power for the future level of volatility, beyond past realized volatility (e.g., Christensen and Prabhala, 1998), we also include as a control variable the lagged at-the-money implied volatility. This addresses the possibility that $\psi_{t,T}$ could simply pick up the effect of persistent implied volatility shocks.

Table 5 shows the results of estimating Eq. (8). In the first five columns, we only use the mean absolute difference of foreign-exchange forecast returns as a regressor. The coefficient is positive and highly statistically significant in all cases. The results for USD/JPY are striking. The MAD of foreign-exchange forecasts in isolation explains between 54% and 64% of the variation of at-the-money implied volatilities. These results are also economically important. A one-standard deviation increase in differences in beliefs

[17] The implied volatility is the daily implied volatility prevailing precisely on the day in which the FX survey is conducted. The results do not change if implied volatility is measured on the day in which the FX survey results are released.

Table 5
At-the-money implied volatility and difference in beliefs.

This table shows the results of estimating the following regression:

$$IV_{t,T} = \alpha + \beta_{\psi} MAD_{t,T} + \beta_{sk} SK_{t,T} + \beta_{ku} KU_{t,T} + \beta_{iv} IV_{t-30,T-30} + e_{t,T},$$

where $IV_{t,T}$ represents the implied volatility of at-the-money (ATM) options on day t with maturity T, and $MAD_{t,T}$, and $SK_{t,T}$, and $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign-exchange forecasts on day t with horizon T, respectively. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR^{+} the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.

	USD/ JPY	GBP/ USD	USD/ EUR^{+}	USD/ DEM	USD/ EUR	USD/ JPY	GBP/ USD	USD/ EUR^{+}	USD/ DEM	USD/ EUR
					One-month horizon					
α	4.32***	5.44***	4.91***	3.88***	5.06***	1.25**	1.51***	1.07*	0.71	1.73**
β_{ψ}	466.73***	283.45***	407.25***	532.07***	374.18***	241.83***	209.32***	246.78***	398.34***	167.30***
β_{sk}						0.07	−0.13	0.35*	0.24	0.62**
β_{ku}						0.07	−0.02	0.24***	0.28**	0.05
β_{iv}						0.55***	0.57***	0.50***	0.37***	0.58***
R^2	0.54	0.23	0.46	0.60	0.44	0.72	0.59	0.68	0.72	0.72
					Three-month horizon					
α	3.45***	6.39***	5.95***	5.82***	5.55***	0.09	1.34***	0.60	0.55	0.53
β_{ψ}	294.79***	123.71***	192.02***	215.25***	198.49***	118.83***	63.38***	93.53***	129.22***	90.09***
β_{sk}						0.04	−0.01	0.24*	0.18	0.44***
β_{ku}						0.13*	−0.01	0.16***	0.09*	0.25***
β_{iv}						0.68***	0.72***	0.67***	0.64***	0.65***
R^2	0.58	0.14	0.30	0.32	0.34	0.85	0.66	0.76	0.73	0.81
					Six-month horizon					
α	2.74***	6.81***	6.76***	6.29***	6.52***	0.37	1.11**	0.87*	−0.22	1.04*
β_{ψ}	225.08***	78.87***	113.77***	138.10***	115.33***	57.97***	29.17***	29.83***	77.03***	24.41**
β_{sk}						−0.09	−0.04	0.18	−0.08	0.65***
β_{ku}						0.01	0.01	0.06	−0.03	0.05
β_{iv}						0.77***	0.79***	0.80***	0.80***	0.78***
R^2	0.64	0.10	0.21	0.16	0.28	0.88	0.70	0.77	0.74	0.84
					One-year horizon					
α	3.36***	7.13***	7.25***	3.74***	6.85***	0.23	1.03**	0.49	0.22	0.41
β_{ψ}	152.92***	53.83***	75.61***	173.58***	78.09***	21.40***	11.45*	13.90**	52.85**	17.43***
β_{sk}						−0.06	0.04	0.08	0.16	0.30
β_{ku}						0.02	−0.01	0.07**	0.08	0.06
β_{iv}						0.88***	0.84***	0.87***	0.76***	0.85***
R^2	0.57	0.09	0.22	0.37	0.29	0.91	0.77	0.84	0.81	0.87
Obs.	156	142	156	60	96	155	141	155	59	95

implies an increase in implied volatilities between 1.9% and 2.1%. The explanatory power for GBP/USD is lower and decreasing for longer horizons. However, all the coefficients are statistically significant at the 1% level. The economic significance is also weaker, ranging between 0.7% and 0.3% for a one-standard deviation shock in the differences in beliefs, although the underlying currency is unconditionally less volatile. For the USD/EUR^{+} exchange rate, the results are again highly significant at all horizons. The increase in implied volatility for a one-standard deviation increase in the standard deviation of foreign-exchange forecasts is between 1.4% at short horizons and 0.7% at long horizons. These economic effects are very similar for the USD/DEM and the USD/EUR considered as separate currency pairs.

The last five columns of Table 5 show the results of estimating the full specification of Eq. (8). This includes controlling for the skewness and the kurtosis of the distribution of foreign-exchange forecasts, as well as the at-the-money implied volatility of the previous month. In fact, it could be argued that part of the relevant characteristics of the differences in beliefs is given by the asymmetry of market forecasts or by the frequency of extreme forecasts, as reflected in the thickness of the tails of the distribution. The statistical significance of the slope coefficient of $\psi_{t,T}$ is always preserved. The role of the asymmetry is particularly relevant for USD/EUR. The skewness coefficient is statistically significant at least at the 5% level at all but the one-year forecast horizon. In this case, the asymmetry of foreign-exchange views reflects some features of the difference in beliefs that impact implied volatility. The kurtosis of the foreign-exchange forecasts explains implied volatility more sparsely and with somewhat lower statistical significance levels. Finally, as expected, due to the persistence of implied volatility, the coefficient of lagged implied volatility is highly significant and the level of significance is increasing in the horizon. In summary, the difference in beliefs proxied empirically by $\psi_{t,T}$ explains the current level of at-the-money implied volatility beyond the information in lagged at-the-money implied volatility.

The literature of heterogeneous agent models surveyed in Section 2 shows that differences in beliefs can generate endogenous volatility of asset returns independently from the volatility of fundamentals. However, one could argue that the volatility of fundamentals is empirically correlated with the uncertainty on the growth rate of the aggregate dividend process, which defines the degree of information incompleteness in economies with uncertainty. In that case, the results of Table 5 could be spurious and just reflect that correlation. To investigate this possibility, we collect a comprehensive data set of economic indicators for the U.S., Japanese, European, and British economy from the Organisation for Economic Cooperation and Development (OECD). The single source for this data is important to guarantee homogeneity of measurement and accuracy across countries. More specifically, we obtain indicators for the gross domestic product (GDP), industrial production (IP), consumer price index (CPI), unemployment rate (UR), broad money (BM), and balance of payments (BP). To estimate the volatility of the macroeconomic time-series, we use the methodology of Schwert (1989). We first estimate a 12th-order autoregression for R_t, the first differences of the macroeconomic series in logarithms, including dummy variables D_{jt} to allow for different monthly returns:

$$R_t = \sum_{j=1}^{12} \alpha_j D_{jt} + \sum_{i=1}^{12} \beta_i R_{t-i} + \varepsilon_t. \tag{9}$$

We then estimate a 12th-order autoregression for the absolute values of the errors from (9),

$$|\hat{\varepsilon}_t| = \sum_{j=1}^{12} \gamma_j D_{jt} + \sum_{i=1}^{12} \rho_i |\hat{\varepsilon}_{t-i}| + u_t. \tag{10}$$

The fitted values from (10) estimate the conditional standard deviation of R_t, given information available before month t. We can now complement the differences in beliefs proxies used in the empirical analysis so far with a measure of volatility of current macroeconomic fundamentals:

$$\begin{aligned} IV_{t,T} = {} & \alpha + \beta_{\psi} MAD_{t,T} + \beta_{sk} SK_{t,T} + \beta_{ku} KU_{t,T} + \beta_{iv} IV_{t-30,T-30} \\ & + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{GDP,h}| + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{IP,h}| + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{CPI,h}| \\ & + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{UR,h}| + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{BM,h}| + \sum_{h=1}^{2} \beta_h |\tilde{\varepsilon}_t^{BP,h}| + e_t, \end{aligned} \tag{11}$$

where we include for each currency and each horizon estimates of the volatility of economic indicators for the home economy ($h=1$) and for the foreign economy ($h=2$). For example, we complement differences in beliefs at a one-month horizon for the USD/JPY currency pair with the volatility of GDP, IP, CPI, UR, BM, and BP of both the U.S. and Japan.[18]

Table 6 presents the key results. We observe in general a small increase in the explanatory power of the regressions, which is an indication that currency implied volatility reflects to some extent volatility of current fundamentals. Typically, two of the 12 regressors that describe macroeconomic uncertainty are statistically significant. However, the proxy for differences in beliefs remains highly statistically significant for all currencies and all horizons, suggesting that it contains additional information with respect to what is currently observed in the economy.

The level, rather than the volatility of macroeconomic fundamentals, could be an alternative channel through which economic conditions can impact FX implied volatilities. For example, in the model of Verdelhan (2010), the pricing kernel is more volatile during bad times in the home country. We thus complement the regressors with the level of the same macroeconomic variables used to compute the volatility of fundamentals.[19] The results of Table 6 do not change. The levels of fundamentals are generally not statistically significant at any horizon or for any underlying currency. We find only sporadic cases of statistical significance, such as the European industrial production in forecasting six-month and 12-month implied volatility of the USD/EUR$^+$, or the U.S. unemployment in forecasting 6-month USD/JPY implied volatility.

In summary, differences in beliefs contain additional information besides the level of volatility of current fundamentals. It might still capture volatility of future fundamentals or might just better aggregate and measure information with respect to low-frequency macroeconomic indicators. If this is indeed the case, our results suggest the importance of explicitly taking into account differences in beliefs, in addition to current changes in fundamentals.

4.1.1. *Specification and robustness*

The documented relation between difference in beliefs and implied volatility can potentially be nonlinear. More specifically, disagreement could play a different role in periods of market stress characterized by large implied volatility and this could affect our previous findings within a linear specification. We thus estimate Eq. (8) using a quantile regression approach. In Table 7, Panel A, we report the results for the median implied volatility. The findings are very similar to the standard regression results, suggesting that they are robust to nonlinearities or large outliers. Furthermore, Table 7, Panel B, shows the results for the top quartile. We notice that the relation between difference in beliefs and large implied volatility is still positive and strongly significant. In all but one case, the relation is even stronger in periods of high market stress.

We test further whether our regression model is appropriate using two additional specification tests. The first is the Ramsey regression equation specification error test (RESET) to understand whether nonlinear combinations of the estimated values help explain the endogenous

[18] In the case of the USD/EUR$^+$ currency pair, we use the volatility of the fundamentals of the U.S. and of the euro area. Replacing the volatility of the fundamentals of the euro area with the fundamentals of Germany before 1999 does not change the results.

[19] We also account for potential non-stationarities in the macroeconomic series and estimate the model using first differences rather than levels. Furthermore, we introduce recession dummies when the GDP variables for each country are in the bottom quartile in our sample period. Both alternatives do not change our findings on the effect of differences in beliefs on implied volatility.

Table 6
Implied volatility, difference in beliefs, and economic uncertainty.
This table shows selected results of estimating the following regression:

$$\begin{aligned} IV_{t,T} &= \alpha+\beta_{\psi}MAD_{t,T}+\beta_{sk}SK_{t,T}+\beta_{ku}KU_{t,T}+\beta_{iv}IV_{t-30,T-30} \\ &+ \sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{GDP,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{IP,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{CPI,h}| \\ &+ \sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{UR,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{BM,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{BP,h}|+e_t, \end{aligned}$$

where $IV_{t,T}$ represents the implied volatility of at-the-money (ATM) options on day t with maturity T, $MAD_{t,T}$, $SK_{t,T}$, $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign-exchange forecasts on day t with horizon T, respectively, and $|\tilde{\varepsilon}_t^{ECO,h}|$ represents estimates of the volatility of economic indicators ECO={GDP,IP,CPI, UR,BM,BP} for the home (h=1) and the foreign (h=2) economy. *key ECO* represent the proportion of economic variables that is significant at least at the 10% level. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.

	USD/ JPY	GBP/ USD	USD/ EUR$^+$	USD/ DEM	USD/ EUR	USD/ JPY	GBP/ USD	USD/ EUR$^+$	USD/ DEM	USD/ EUR
			One-month horizon					Three-month horizon		
α	0.23	1.42	1.06	−1.70	2.24	−0.60	1.25	0.93	0.43	1.55
β_{ψ}	237.04^{***}	214.21^{***}	268.90^{***}	410.33^{***}	170.21^{***}	115.81^{***}	68.10^{***}	92.08^{***}	130.09^{***}	71.66^{***}
β_{sk}	0.05	−0.15	0.37^{*}	0.25	0.64^{***}	0.01	0.01	0.27^{**}	0.41^{**}	0.43^{**}
β_{ku}	0.06	−0.02	0.22^{***}	0.34^{***}	0.04	0.10	−0.02	0.11^{**}	0.04	0.20^{**}
β_{iv}	0.53^{***}	0.60^{***}	0.47^{***}	0.42^{***}	0.57^{***}	0.67^{***}	0.72^{***}	0.67^{***}	0.63^{***}	0.69^{***}
key ECO	2/12	3/12	2/12	3/12	1/12	2/12	1/12	1/12	1/12	1/12
R^2	0.76	0.65	0.73	0.79	0.76	0.86	0.70	0.78	0.78	0.82
			Six-month horizon					One-year horizon		
α	0.26	0.87	1.20	−0.56	1.43	0.20	0.90	1.05	0.05	1.11
β_{ψ}	50.76^{***}	38.02^{***}	32.44^{***}	72.99^{***}	16.66^{**}	15.32^{**}	19.91^{**}	18.12^{***}	55.46^{***}	12.53^{**}
β_{sk}	−0.09	−0.06	0.28^{**}	0.11	0.59^{***}	0.03	0.05	0.16^{**}	0.17	0.37
β_{ku}	−0.01	0.01	0.01	−0.09	0.01	−0.02	−0.01	0.05^{*}	0.05	0.01
β_{iv}	0.78^{***}	0.79^{***}	0.78^{***}	0.78^{***}	0.82^{***}	0.89^{***}	0.82^{***}	0.84^{***}	0.76^{***}	0.87^{***}
key ECO	2/12	3/12	2/12	1/12	1/12	3/12	1/12	2/12	2/12	1/12
R^2	0.89	0.75	0.80	0.80	0.86	0.92	0.80	0.86	0.86	0.88
Obs.	155	141	155	59	95	155	141	155	59	95

variable. The results are reported in Table 7, Panel C. We find that this test never rejects the null hypothesis that the model is correctly specified.[20]

As a capstone to our specification analysis, we also compute an asymptotic Generalized method of moments (GMM) test, in which the moment conditions are obtained by conditioning on the three different terciles of the difference in beliefs distribution, low, medium, and high. We use as instruments the lagged values of difference in beliefs and the current volatility of three series of macroeconomic fundamentals. The choice of the instruments is based on their level of significance determined by the regression results of Table 6. This setting gives rise to a Chi-square test with nine degrees of freedom. Table 7, Panel D, summarizes the results. We find that the impact of difference in beliefs is strongly statistically significant. Furthermore, the *J*-specification tests for overidentifying restrictions never reject the model for any currency pair, suggesting that the relation between differences in beliefs and implied volatility is stable at all levels of disagreement.

[20] We also experiment with a similar specification test whereby implied volatility is regressed on its estimated value from Eq. (8) and on its estimated value squared. We find that the coefficient on the squared term is never significant for any currency pair at any horizon, suggesting that the original model is correctly specified.

As a series of additional robustness checks, we consider specifications of Eq. (11) with different lags for the regressors. For example, we lag by one week the MAD, the higher-order moments of the distribution of the foreign-exchange forecasts, and implied volatility. The results are virtually unchanged.

Another interesting check is to re-estimate Eq. (11) using the beliefs characteristics lagged one month. This alternative specification recasts the previous empirical question in the context of predictive regressions for future implied volatility. This empirical exercise addresses potential concerns of endogeneity, since differences in beliefs, implied volatility, and all the other regressors are measured at the same point in time. Again, the horizon of the forecast and the maturity of the option perfectly match, so that the analysis is unaffected by the synchronicity issues described in Cavaglia, Verschoor, and Wolff (1994). Table 8 shows the results for the one-month horizon. For all currencies, differences in beliefs predict future implied volatility beyond current implied volatility. These results hold also when the regressors include uncertainty about current fundamentals, which marginally increase predictability. As a final robustness check, given the documented high persistence of implied volatility, we estimate the specification (11) using first differences in implied volatility as the dependent variable.

Table 7
Specification tests.
This table shows the results of four tests of the following specification:

$$IV_{t,T} = \alpha + \beta_{\psi} MAD_{t,T} + \beta_{sk} SK_{t,T} + \beta_{ku} KU_{t,T} + \beta_{iv} IV_{t-30,T-30} + e_{t,T},$$

where $IV_{t,T}$ represents the implied volatility of at-the-money (ATM) options on day t with maturity T, and $MAD_{t,T}$, $SK_{t,T}$, and $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign-exchange forecasts on day t with horizon T, respectively. Panels A and B show quantile regression estimates for the median and the top quartile, respectively. Panel C reports the results of a Ramsey RESET test. Panel D shows parameter estimates, t-statistics, and a J-specification test for overidentifying restrictions in a standard GMM framework where different moment conditions correspond to three different percentiles of MAD and instruments are lagged differences in beliefs and macroeconomic volatility. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR^{+} the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st,1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively, obtained by bootstrap replications in Panels A and B and using the Newey-West method in Panel D.

	USD/ JPY	GBP/ USD	USD/ EUR^{+}	USD/ DEM	USD/ EUR
Panel A: Quantile regression: median					
α	1.55**	1.68***	1.22***	0.71	1.92***
β_{ψ}	252.70***	162.96***	222.18***	347.91***	246.84***
β_{sk}	0.29	0.03	0.36**	0.09	0.63***
β_{ku}	0.04	−0.04	0.13**	0.15	0.04
β_{iv}	0.52***	0.61***	0.53***	0.46***	0.45***
Pseudo-R^2	0.48	0.36	0.47	0.46	0.51
Panel B: Quantile regression: top quartile					
α	1.24	2.02***	1.21*	0.01	1.64***
β_{ψ}	318.89***	221.82***	242.22***	498.81***	142.10***
β_{sk}	0.30	−0.19	0.43*	0.20	0.78***
β_{ku}	0.11	−0.01	0.10	0.09	0.10
β_{iv}	0.53***	0.55***	0.58***	0.46***	0.65***
Pseudo-R^2	0.54	0.35	0.44	0.49	0.53
Panel C: RESET specification test					
F-test	1.18	0.48	1.54	0.88	0.52
(*P-value*)	(0.86)	(0.92)	(0.22)	(0.57)	(0.90)
Panel D: GMM test, 3 parameters, 12 moment conditions					
β_{ψ}^{low}	460.63***	269.64***	423.18***	546.77***	366.58***
β_{ψ}^{med}	453.62***	266.30***	379.58***	530.75***	377.38***
β_{ψ}^{high}	474.16***	310.04***	413.66***	517.52***	408.88***
J-stat	7.45	5.76	9.38	5.34	6.44
(*P-value*, 9 d.f.)	(0.59)	(0.76)	(0.40)	(0.80)	(0.70)

The results (not reported) show again that the differences in beliefs significantly explain the change in implied volatility for all currencies and at all horizons.

In summary, we show that a simple empirical proxy of the difference in beliefs of foreign-exchange market participants, $\psi_{t,T}$, is an economically and statistically important driver of implied volatility in FX markets. This finding holds for different currency pairs, at different forecasting horizons, controlling for the volatility of current fundamentals, and with different empirical specifications.

4.2. The option-implied volatility smile

The second set of empirical implications of the literature surveyed in Section 2 is related to the effects of the structure of beliefs on the relative prices of derivative contracts. More specifically, in the case of options, this should occur in terms of the cross-section of implied volatilities for different moneyness levels, i.e., the implied volatility smile. This theoretical prediction is especially relevant, given that the characteristics of the option-implied volatility smile are closely related to the characteristics of the foreign-exchange state price density (e.g., Bates, 2001; Bakshi, Kapadia, and Madan, 2003).

We obtain the implied volatility of at-the-money options directly from market quotes and compute implied volatilities of options with other moneyness levels from risk reversal and strangle combinations, as described in Section 3.[21] Then, we investigate the link between differences in beliefs and the slope of the implied volatility smile by estimating the following regression:

$$\left(\frac{IV_{t,T}^{25-\text{delta}}}{IV_{t,T}^{75-\text{delta}}}\right) - 1 = \alpha + \beta_{\psi}\psi_{t,T} + \beta_{sk} SK_{t,T} + \beta_{ku} KU_{t,T} + e_{t,T}, y, \tag{12}$$

where $IV_{t,T}^{x-delta}$ represents the implied volatilities at $x-delta$ moneyness for options on day t with maturity T and the regressors are the same as those of the previous specifications.

Table 9 (first five columns) shows these results in the univariate regression case. We observe that higher dispersion of beliefs influences the slope of the volatility smile significantly at three-, six-, and 12-month horizons for USD/JPY, and at all horizons for USD/EUR. This evidence confirms the predictions of the theoretical literature. It is remarkable that in this setting the results for the USD/DEM are much weaker than for the USD/EUR currency pair. This different finding could be related to the role of the German central bank in the limited-flexible exchange rate system with defined bands prevailing before the introduction of the euro. In the traditional heterogeneous belief economy, the cost for the optimist to insure the pessimist is increasing in the level of disagreement. However, if currencies cannot effectively fluctuate beyond a prespecified ban, it is a third agent, the central bank, that will take on the risk of the most extreme changes.

The results of Table 9 are also economically important. For example, the unconditional slope of the implied volatility smile for USD/JPY is about −7% for longer horizons. However, the slope drops to about −4% when the differences in beliefs are one-standard deviation above the mean. The economic significance is similar for the USD/EUR exchange rate.

The positive sign implies that higher dispersion of beliefs increases the risk-neutral probability of a Japanese yen or euro appreciation against the dollar. Carr and Wu (2007) show a strong correlation between the slope of the smile and the returns on the underlying currency. Our result would thus occur naturally, for example, if there is a relation between dispersion of beliefs and underlying returns. In Section 4.4 we will specifically investigate this hypothesis.

In the remainder of Table 9, we estimate the full specification of Eq. (12). The results are very similar to the

[21] Data on the risk reversal and strangle combinations for the USD/DEM currency pair were not available for the six-month maturity.

Table 8
Forecasting implied volatility with difference in beliefs.
This table shows selected results of estimating the following regression:

$$\begin{aligned} IV_{t+30,t+60} &= \alpha+\beta_{\psi}MAD_{t,T}+\beta_{sk}SK_{t,T}+\beta_{ku}KU_{t,T}+\beta_{iv}IV_{t,t+30} \\ &+ \sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{GDP,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{IP,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{CPI,h}| \\ &+ \sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{UR,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{BM,h}|+\sum_{h=1}^{2}\beta_h|\tilde{\varepsilon}_t^{BP,h}|+e_t, \end{aligned}$$

where $IV_{t,T}$ represents the implied volatility of at-the-money (ATM) options on day t with maturity T, $MAD_{t,T}$, $SK_{t,T}$, and $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign-exchange forecasts on day t with horizon T, respectively, and $|\tilde{\varepsilon}_t^{ECO,h}|$ represents estimates of the volatility of economic indicators $ECO=\{GDP,IP,CPI,UR,BM,BP\}$ for the home (h=1) and the foreign (h=2) economy. *KeyECO* represent the proportion of economic variables that is significant at least at the 10% level. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR$^+$ the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively, with robust standard errors.

	USD/ JPY	GBP/ USD	USD/ EUR$^+$	USD/ DEM	USD/ EUR	USD/ JPY	GBP/ USD	USD/ EUR$^+$	USD/ DEM	USD/ EUR
	No economic fundamentals					With economic fundamentals				
α	2.15***	2.64***	2.51***	3.97***	1.58**	4.65***	2.68***	4.55***	4.74*	3.44**
β_{ψ}	86.01**	73.02**	137.65***	100.29*	196.62***	74.70**	93.28***	156.87***	136.76*	189.77***
β_{sk}	0.19	0.06	0.55***	0.03	0.79***	0.29	0.07	0.46**	0.04	0.80***
β_{ku}	−0.01	−0.03	−0.06	−0.35	0.17	0.05	−0.04	−0.05	−0.29	0.21*
β_{iv}	0.70***	0.60***	0.59***	0.61***	0.51***	0.68***	0.59***	0.57***	0.61***	0.48***
key ECO						3/12	3/12	1/12	0/12	0/12
R^2	0.65	0.50	0.61	0.52	0.74	0.67	0.56	0.63	0.58	0.76
Obs.	155	141	155	59	95	155	141	155	59	95

univariate specification for USD/JPY and GBP/USD, except that the higher-order moments of the distribution of beliefs become statistically significant when the mean absolute deviation is not significant (e.g., the one-month horizon for USD/JPY). For USD/EUR, the higher-order moments improve explanatory power significantly. The R^2 almost doubles on average, with the largest increase at the one-month horizon. Both skewness and kurtosis of beliefs affect the slope of the implied volatility smile.

In summary, the results show that the differences in beliefs are an important determinant of the shape of the option-implied volatility smile, and thus, of the shape of the risk-neutral foreign-exchange state price density. This is important evidence that the difference in beliefs drives differential pricing of state-contingent claims, precisely as the neoclassical literature with heterogeneous agents predicts and it supports the idea that dispersion of beliefs enters the stochastic discount factor. This last aspect will become apparent in the empirical tests of the volatility risk premium in the next subsection.

4.3. Volatility risk premium and difference in beliefs

In the heterogeneous agents neoclassical literature, beliefs enter the stochastic discount factor and thus, have a direct effect on the difference between the risk-neutral and physical probability measure. A natural test for this theoretical prediction is thus to investigate whether measures of the differences in beliefs impact the volatility spread, i.e., the difference between implied and realized volatility.

Britten-Jones and Neuberger (2000), among others, show that the *risk-neutral* expected value of the quadratic variation of returns is expressed in a model-free fashion by a particular portfolio of options. In practice, this measure must be constructed from a finite number of strike prices and, unfortunately, the three moneyness levels that are available in our data set for the FX market do not allow for an accurate estimation. We thus approximate the *risk-neutral* expected value of the return quadratic variation from time t to time T with the time t at-the-money implied volatility of currency options maturing at time T.

The quadratic variation of returns under the *physical* measure is typically estimated using squared returns. This approach has strong theoretical underpinnings and good empirical performance (see, among others, Andersen, Diebold, and Labys, 2003). Realized squared returns are used to produce the expectation at time t of the realized volatility between time t and time T (e.g., Bollerslev, Tauchen, and Zhou, 2008; Drechsler and Yaron, 2008). More specifically, as an empirical proxy for the *physical* expected value of the return quadratic variation from time t to time T, we use an exponential moving average of the quadratic returns on the underlying according to the following specification:

$$E_t[RV_{t,T}]=\sqrt{(1-\alpha_{T-t})(r_{t-1}^2+\alpha_{T-t}r_{t-2}^2+\alpha_{T-t}^2r_{t-3}^2+\cdots)}, \quad (13)$$

where r_t is the log return on the underlying asset on day t and α_{T-t} is the smoothing factor that depends on the horizon.[22] The use of exponentially decaying weights is a useful device in our setting to increase the number of observations, since we use daily rather than intraday returns,

[22] The exponential moving average can also be calculated in a recursive fashion as $E_t[RV_{t,T}]=\sqrt{\alpha_{T-t}r_{t-1}^2+(1-\alpha_{T-t})E_{t-1}[RV_{t-1,T-1}]^2}$. The smoothing parameter decreases with the horizon and is equal to 0.03, 0.01, 0.005, and 0.003 for a one-, three-, six-, and 12-month horizon, respectively.

Table 9
The implied volatility smile and difference in beliefs.
This table shows the results of estimating the following regression:

$$\left(\frac{IV_{t,T}^{25-delta}}{IV_{t,T}^{75-delta}}\right)-1=\alpha+\beta_{\psi}MAD_{t,T}+\beta_{sk}SK_{t,T}+\beta_{ku}KU_{t,T}+e_{t,T},$$

where $IV_{t,T}^{x\text{-}delta}$ represents the implied volatilities at *x-delta* moneyness for options on day *t* with maturity *T* and $MAD_{t,T}$, $SK_{t,T}$, $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign exchange forecasts on day *t* with horizon *T*, respectively. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR^{+} the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively, with robust standard errors.

	USD/JPY	GBP/USD	USD/EUR^{+}	USD/DEM	USD/EUR	USD/JPY	GBP/USD	USD/EUR^{+}	USD/DEM	USD/EUR
					One-month horizon					
α	−0.06***	−0.01	0.05***	−0.01	0.08***	−0.06***	−0.02	0.06***	0.01	0.08***
β_{ψ}	−0.52	0.92	1.60*	−2.01	3.90***	0.38	1.29	2.20**	−1.80	4.08***
β_{sk}						−0.03***	−0.01**	−0.02***	−0.01	−0.03***
β_{ku}						−0.01	0.01*	−0.00	−0.01	−0.01
R^2	0.01	0.01	0.02	0.03	0.09	0.06	0.05	0.09	0.03	0.21
Obs.	138	137	138	44	94	138	137	138	44	94
					Three-month horizon					
α	−0.11***	0.02	0.09***	0.12***	0.09***	−0.13***	0.05**	0.05**	−0.02	0.07***
β_{ψ}	1.49*	−0.86	2.38***	4.69**	2.40***	1.63*	−1.55	2.04***	3.34**	2.19***
β_{sk}						−0.02	0.00	−0.00	0.01	−0.01
β_{ku}						0.01	−0.01**	0.01**	0.04**	0.01
R^2	0.02	0.01	0.09	0.17	0.10	0.04	0.07	0.14	0.65	0.12
Obs.	112	94	112	18	94	112	94	112	18	94
					Six-month horizon					
α	−0.18***	0.01	0.08***		0.08***	−0.19***	0.02	0.11***		0.11***
β_{ψ}	3.01***	−0.36	1.46***		1.46***	2.94***	−0.73	1.72***		1.72***
β_{sk}						−0.01	0.01	−0.02*		−0.02*
β_{ku}						0.01	0.00	−0.01*		−0.01*
R^2	0.15	0.00	0.09		0.09	0.16	0.02	0.14		0.14
Obs.	112	112	91	none	91	112	112	91	none	91
					One-year horizon					
α	−0.19***	0.01	0.05***	0.10**	0.07***	−0.17***	0.06***	0.012	−0.05	0.05**
β_{ψ}	2.23***	−0.26	0.52*	2.26*	0.81***	2.21***	−1.18***	0.21	0.12	0.63**
β_{sk}						0.03	0.02***	0.03***	−0.02	0.03***
β_{ku}						−0.01	−0.00***	0.01***	0.01	0.01***
R^2	0.12	0.00	0.02	0.12	0.06	0.13	0.18	0.10	0.62	0.11
Obs.	112	107	112	18	94	112	107	112	18	94

and corresponds to a widely used approach among practitioners (see Morgan's RiskMetrics, 1996).[23]

We then regress the volatility spread on the characteristics of the distribution of foreign-exchange forecasts:

$$IV_{t,T}-E[RV_{t,T}]=\alpha+\beta_{\psi}\psi_{t,T}+\beta_{sk}SK_{t,T}+\beta_{ku}KU_{t,T}+e_{t,T}, \qquad (14)$$

where $IV_{t,T}$ represents the at-the-money implied volatilities for options on day *t* with maturity *T*, $E[RV_{t,T}]$ is an exponential moving average of realized volatility obtained with the formula in Eq. (13) for a horizon matching the maturity of the option, and $\psi_{t,T}$, $SK_{t,T}$, and $KU_{t,T}$ are defined as before.[24]

Table 10 shows the results of estimating Eq. (14). As before, the first five columns show our findings for the univariate specification with the mean absolute deviation of the foreign-exchange forecasts on the right-hand side. For USD/JPY and USD/EUR^{+}, $\psi_{t,T}$ has a significantly positive effect on the volatility risk premium. In particular, for USD/JPY, the dispersion of beliefs explains a striking average 50% of the variation in the volatility risk premium across forecasting horizons. Unconditionally, the difference between implied and realized volatility is about 3.3%, but it increases to about 5.3% when the dispersion of beliefs is one-standard deviation above average. For USD/EUR^{+}, the volatility risk premium is unconditionally lower, about 1.4%, but it increases to about 1.8% with high dispersion of beliefs. For the USD/EUR considered as a single currency, the economic effects are even stronger: large disagreement leads to a 75% increase in the volatility risk premium.

The remainder of Table 10 shows the results of estimating the full specification of Eq. (14). The findings of the previous analysis are confirmed. There is only some evidence that the higher-order moments of

[23] Our results are robust to the choice of model to proxy for expected realized volatility. As alternative estimators, we used the exponential moving average of Eq. (13) with different smoothing parameters, a centered moving average with squared differences between returns and mean returns, and a GARCH(1,1) model.

[24] We also estimate a version of Eq. (14) where the volatility risk premium is the ratio of implied to realized volatility instead of the difference between them. We obtain similar results, suggesting that our findings are not driven by the positive correlation between the general level of volatility and the volatility risk premium expressed as a difference.

Table 10
The volatility risk premium and difference in beliefs.

This table shows the results of estimating the following regression:

$$IV_{t,T} - E[RV_{t,T}] = \alpha + \beta_{\psi} MAD_{t,T} + \beta_{sk} SK_{t,T} + \beta_{ku} KU_{t,T} + e_{t,T},$$

where $IV_{t,T}$ represents the at-the-money implied volatilities for options on day t with maturity T, $E[RV_{t,T}]$ is an exponential moving average of realized volatility over the previous month, and $MAD_{t,T}$, $SK_{t,T}$, and $KU_{t,T}$, represent the mean absolute deviation, the skewness, and the kurtosis of foreign-exchange forecasts on day t with horizon T, respectively. The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR^{+} the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively, with robust standard errors.

	USD/JPY	GBP/USD	USD/EUR^{+}	USD/DEM	USD/EUR	USD/JPY	GBP/USD	USD/EUR^{+}	USD/DEM	USD/EUR
					One-month horizon					
α	-2.90^{***}	0.54	0.76^{*}	1.18	-0.56	-2.98^{***}	0.65	-0.57	-0.20	-1.17^{*}
β_{ψ}	443.05^{***}	23.18	50.08^{*}	89.02	105.58^{***}	431.71^{***}	-6.74	59.01^{*}	88.35^{*}	106.97^{***}
β_{sk}						0.37	-0.02	0.38^{*}	0.30	0.57^{**}
β_{ku}						0.09	0.04^{*}	0.35^{***}	0.39^{***}	0.16
R^2	0.37	0.002	0.01	0.04	0.08	0.38	0.02	0.09	0.16	0.15
					Three-month horizon					
α	-4.54^{***}	1.26^{***}	0.67	0.89	-0.61	-5.85^{***}	1.31^{***}	-0.59	0.48	-2.51^{***}
β_{ψ}	305.32^{***}	-18.77	31.41^{*}	63.67^{**}	60.17^{***}	306.86^{***}	-25.63	45.15^{***}	74.51^{**}	78.95^{***}
β_{sk}						-0.13	0.07	0.43^{**}	0.47^{**}	0.63^{***}
β_{ku}						0.37	0.02	0.25^{***}	0.02	0.41^{***}
R^2	0.47	0.004	0.01	0.07	0.08	0.51	0.02	0.13	0.15	0.34
					Six-month horizon					
α	-6.13^{***}	1.31^{***}	0.62	-0.66	0.04	-5.88^{***}	1.25^{***}	-0.21	0.67	-2.03^{***}
β_{ψ}	251.47^{***}	-11.57	22.00^{*}	85.78^{**}	25.27^{*}	252.04^{***}	-13.94	29.24^{**}	72.19^{**}	44.99^{***}
β_{sk}						-0.20	-0.07	0.18	0.11	0.96^{***}
β_{ku}						-0.08	0.03	0.17^{**}	0.28	0.33^{***}
R^2	0.60	0.003	0.01	0.16	0.03	0.60	0.01	0.05	0.24	0.40
					One-year horizon					
α	-5.53^{***}	0.85^{*}	0.37	-1.04	-0.39	-5.52^{***}	0.47	-0.82	-1.52	-2.19^{***}
β_{ψ}	168.63^{***}	6.17	19.05^{**}	64.76^{***}	29.37^{**}	161.47^{***}	12.20^{*}	29.62^{***}	76.05^{***}	47.84^{***}
β_{sk}						0.80^{***}	-0.22^{**}	0.03	0.26	1.28^{***}
β_{ku}						0.02	0.03	0.21^{***}	0.01	0.06
R^2	0.55	0.01	0.02	0.13	0.06	0.57	0.03	0.10	0.16	0.40
obs.	156	142	156	60	96	156	142	156	60	96

foreign-exchange forecasts distribution explain the volatility risk premium beyond $\psi_{t,T}$. More specifically, for USD/EUR, the kurtosis of the distribution of beliefs has a systematically significant positive effect on the difference between implied and realized volatility.

Our empirical measure of the volatility risk premium uses the expectation of realized volatility constructed with information available at time t. However, some papers in the literature have measured realized volatility on the interval from time t to time T spanned by the maturity of the option (e.g., Carr and Wu, 2009). In this case, the volatility risk premium coincides with the payoff of a volatility swap. We repeat our empirical analysis above using this alternative definition, obtaining very similar results.

An additional robustness check is motivated by the frequent assumption in the literature that the volatility risk premium is linear in volatility (e.g., Heston, 1993). Bakshi and Kapadia (2001), among others, find empirical support for this hypothesis for the volatility risk premium in the stock market. We thus estimate a more general version of Eq. (14), where we also control for the lagged level of implied volatility on the right-hand-side. The findings of Table 10 on the effect of dispersion of beliefs on the volatility risk premium are unchanged (results not reported).

In summary, the evidence in this subsection is an important empirical validation of the main theoretical result of neoclassical economies with heterogeneous beliefs: the stochastic discount factor depends on the relative weights of different agents in the representative agent utility function, which is a function of differences in beliefs. If there was solely a trading friction channel for dispersion of beliefs to impact asset prices, we would not find any significant relation between the volatility spread and disagreement. This finding has broad implications. For instance, since differences in beliefs are time-varying, this can generate time-varying risk-premiums and expected returns, which could help to explain several stylized facts in the FX markets.

4.4. *Differences in beliefs and currency options in the cross-section*

The relation between differences in beliefs and currency options has been analyzed so far in the time-series and in isolation for each currency pair. However, the theoretical predictions could also be empirically tested in the cross-section, that is, currency pairs characterized by different disagreement should also exhibit different implied volatilities, different slopes of the implied volatility smile, and different volatility spreads.

The most serious limitation for the identification of the effects of disagreement in the cross-section is the availability of only three currency pairs in our setting, with the USD/EUR$^+$ viewed as a combination of USD/DEM and USD/EUR. Nonetheless, we observe these three pairs at different points in time over more than 10 years and we can thus exploit the identification power of a repeated cross-section.

The starting point of the cross-sectional analysis is to look at the average disagreement and currency option characteristics over the full sample period. As the summary statistics in Table 4 show, the ranking of average differences in beliefs, from higher to lower, is USD/JPY, USD/EUR$^+$, GBP/USD for all forecasting horizons. We observe the same ranking for the average implied volatility (see Table 1), for the absolute average magnitude of the risk reversal (see Table 1), and for the volatility spread as defined in Section 4.3 (results not reported).[25] At least at this very aggregate level, the relations between disagreement and currency option markets that hold in the time-series, also hold in the cross-section.

We take this analysis further by identifying two subsamples where the ranking of disagreement changes. Between 1995 and 1999, we observe that the one-month MAD for the USD/JPY currency pair is 40% higher than the corresponding disagreement for USD/EUR$^+$, which is in turn 30% higher than the MAD for GBP/USD. Consistent with theory and time-series evidence, the implied volatility of USD/JPY is 27% higher than the implied volatility of USD/EUR$^+$, which is itself 26% higher than the implied volatility of GBP/USD. Similar relations hold for the slope of the smile, where the risk-neutral probability of appreciation of the low-yield currency is higher for high-disagreement currency pairs. Finally, the average volatility spread in this subsample is 5.8%, 2.1%, and 1.2%. Between 2000 and 2004, however, USD/EUR$^+$ becomes the currency pair with the largest differences in beliefs, with a one-month MAD 15% higher than USD/JPY (and 48% higher than GBP/USD). USD/EUR$^+$ implied volatility is now 7% higher than USD/JPY and 29% higher than GBP/USD implied volatility. Again, we observe the same ranking across currency pairs for the slope of the implied volatility smile and the volatility spread.

Motivated by the evidence in the subsample qualitative analysis, we estimate a difference-in-difference specification, where the explanatory variable is the difference in disagreement between two currency pairs and the dependent variable is the difference in either the level of implied volatility of the two currencies, or the difference in the slope of the two volatility smiles, or the difference in the volatility spreads. In all three cases, we find that the difference in disagreement is statistically significant at the 1% level with robust standard errors.[26]

In summary, the cross-sectional analysis fully confirms our findings in the time-series, lending further support to the notion that differences in beliefs are an important factor that help to explain currency option dynamics.

4.5. *Carry trades, exchange rates predictability, and dispersion of beliefs*

The empirical evidence presented so far is consistent with the predictions of the neoclassical heterogeneous agent literature reviewed in Section 2 that the dispersion of beliefs is a priced source of risk in FX markets. The first papers in the literature to investigate risk factors in currency markets are related to the failure of the uncovered interest parity (e.g., Hansen and Hodrick, 1980; Fama, 1984). More recent research has focused on carry trades: borrowing in a low interest rate currency and investing in a high interest rate one. These strategies are widely used by market practitioners to exploit the predictable excess returns implied by the forward premium puzzle. The Sharpe ratios of such trades tend to be higher than the Sharpe ratio of an investment in the S&P500 index. For example, a simple developed country carry trade strategy proposed in Lustig, Roussanov, and Verdelhan (2008) generates a Sharpe ratio of 1.12 as compared with 0.57 for the S&P500 between 1982 and 2007. The academic literature has tried to reconcile these puzzling returns with risk-based explanations. Lustig and Verdelhan (2007), for example, show that U.S. consumption growth risk can explain predictable returns in currency markets. Given our empirical evidence that dispersion of beliefs is a priced risk factor, it becomes natural to ask how heterogeneity of beliefs affects carry trade returns and, more generally, the forward premium puzzle.

Given the sample of currencies for which we can compute measures of dispersion of beliefs, we focus on carry trade returns for the USD/JPY currency pair, which is the most popular (and discussed) carry trade strategy among practitioners.[27] We proceed as follows. At each monthly survey date, we borrow JPY and invest in USD and close out the position on the day before the following survey release. This approach allows us to relate precisely carry trade returns to dispersion of beliefs and implied volatility. We decompose total carry returns into returns generated by the interest rate differentials and the returns generated by exchange-rate appreciation/depreciation. Fig. 3 plots average returns in high/low dispersion of beliefs states, high/low implied volatility states, and high/low volatility

[25] Within the USD/EUR$^+$ currency pair, it is more difficult to compare USD/DEM to USD/EUR in this cross-sectional exercise, because they refer to a different sample period.

[26] We obtain similar evidence for a standard panel regression with time fixed-effects. The results of both estimation techniques are not reported, but are available upon request.

[27] The interest rate differential is lower and switches sign for the USD/EUR$^+$ and GBP/USD pairs. Furthermore, these are not typical carry trade pairs among practitioners.

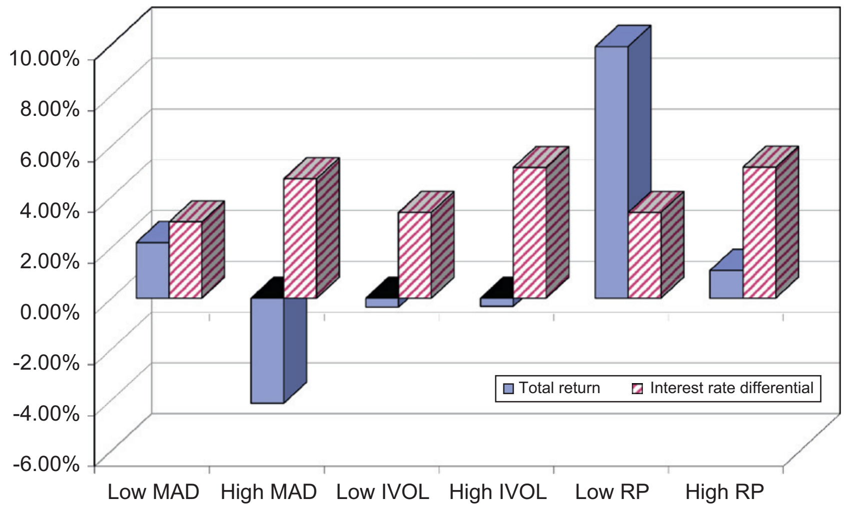

Fig. 3. USD/JPY carry trade returns. This plot shows average total carry trade returns (annualized) and interest rate differential when the MAD for USD/JPY one-month forecast is low or high, USD/JPY one-month implied volatility (IVOL) is low or high, the volatility risk premium (implied-realized one-month volatility, RP) is low or high. For all measures, low (high) is bottom (top) quartile.

risk-premium states. We can readily see that the interest rate differential is larger when dispersion of beliefs, implied volatility, or the volatility risk premium are high. However, when we look at total carry trade returns, they seem to be virtually unaffected by low/high implied volatility states. In contrast, we notice that total returns are much lower when the volatility risk premium is high and they even turn negative when dispersion of beliefs is high. Whenever market participants disagree about the future value of currencies, the carry compensation seems to be insufficient to make up for adverse spot currency movements.

We test more formally for the relation between USD/JPY carry trade returns and the dispersion of beliefs by regressing total carry trade returns on the dispersion of beliefs, implied volatility, and the volatility risk premium. Table 11, Panel A, shows the results. The graphical intuition from Fig. 3 is confirmed. Carry trade returns are virtually unaffected by the level of implied volatility. The coefficients are not statistically different from zero. In contrast, carry returns are significantly lower when dispersion of beliefs is high. The economic significance is striking: a one-standard deviation shock to dispersion of beliefs leads to a negative annualized returns of 7.7% on the carry trade. The magnitude of these negative total returns is remarkable considering the positive average 3.7% interest rate differential between USD and JPY in our sample. The effect of the volatility risk premium is not statistically significant at conventional levels, but economically a one-standard deviation shock to the volatility risk premium generates a negative 5% annualized return. These findings emphasize again that the extent of the disagreement about future exchange rates is a key variable that affects asset returns, volatilities, and risk premiums.

These results are in line with the evidence on the implied volatility smile presented in Section 4.2. When disagreement about the USD/JPY increases, the risk-neutral probability of the Japanese yen appreciation also increases and carry trade returns turn out to be severely penalized.

A straightforward statistical interpretation of the negative relation between dispersion of beliefs and carry trade returns would be that the distribution of beliefs has some forecasting power for the large negative returns that periodically hit the carry trade strategy. Obviously, a potential shortcoming of this empirical exercise is the availability of a sample of currency returns and FX forecasts that is long enough to contain a sufficient number of FX jumps for statistical inference. In our sample, we identify only 18 instances in which FX returns are larger than two-standard deviation bands. The dispersion of beliefs at the one-month horizon shows some predictive power for these negative large returns, albeit not at the conventional statistical levels of significance.

An economic interpretation of our findings could be cast within the "liquidity spirals" arising in the model of Brunnermeier and Pedersen (2009), where an exchange-rate correction only occurs when sufficiently many traders unwind their carry trade position. Because of a negative shock or of a change in risk tolerance by traders, funding constraints are hit and this triggers unwinding of carry positions with a consequent price impact, trading losses, a further increase in the funding problems, volatility, and margins. In our setting, it is precisely the effect of differences in beliefs on the pricing kernel that has an impact on risk tolerance of the representative agent.

We perform an additional empirical exercise to better understand whether our findings could be interpreted within this theoretical framework. More specifically, we try to isolate empirically two cases in which a change in risk tolerance would have the largest effects. First, when funding liquidity is low, a decrease in risk tolerance translates directly into unwinding of carry positions.

Table 11

Carry trade returns, currency returns, and dispersion of beliefs.

Panel A shows the results of estimating the following regression:

$$carry_{t,t+1} = \alpha + \beta_{\psi} MAD_{t,t+1} + \beta_{iv} IV_{t,t+1} + \beta_{rp} RP_{t,t+1} + e_{t,t+n},$$

where $carry_{t,t+1}$ represents USD/JPY carry trade returns between time t and t+1, $MAD_{t,t+1}$ is the mean absolute deviation of USD/JPY one-month forecasts at time t, $IV_{t,t+1}$ is the implied volatility of USD/JPY one-month options, and $RP_{t,t+1}$ is the volatility risk premium with one-month horizon.

Panel B shows estimates of the same specification of Panel A, conditioning on TED spread in the top quartile (high TED), in the bottom three quartiles (normal TED), and on carry trade returns of the previous period being positive or negative.

Panel C shows the results of estimating the following regression:

$$s_{t+n} - s_t = \alpha + \beta_{fwd}(f_{t,t+n} - s_t) + \beta_{\psi} MAD_{t,t+n} + e_{t,t+n},$$

where s_t represents the logarithm of the spot rate at time t, $f_{t,t+n}$ represents the logarithm of the forward rate at time t with maturity t+n, $MAD_{t,t+n}$ is the mean absolute deviation of foreign-exchange forecasts at time t for the horizon t+n, and n is one-month, three-months, six-months, and one-year.

Panel D shows estimates of the same specification of Panel C for the one-month horizon, conditioning on TED spread in the top quartile (high TED), in the bottom three quartiles (normal TED). The data are sampled at a monthly frequency from May 1993 to December 2006. We denote with USD/EUR^{+} the joint exchange rate represented by the USD/DEM and the USD/EUR before and after January 1st, 1999. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively, with robust standard errors.

Panel A

α	0.0169*	0.0083	0.0069*	0.0046
β_{ψ}	−0.9519*			−1.3905*
β_{iv}		−0.0004		0.0021
β_{rp}			−0.0009	−0.0013
R^2	0.0151	0.0016	0.0088	0.0267

Panel B

	High TED	Normal TED	$carry_{t-1,t} > 0$	$carry_{t-1,t} < 0$
α	0.0305	−0.0185	−0.0249	0.0177
β_{ψ}	−2.8228**	−0.7214	−1.7779**	−0.5658
β_{iv}	0.0015	0.0026	0.0037	−0.0005
β_{rp}	−0.0007	−0.0002	−0.0024	−0.0006
R^2	0.0900	0.0227	0.0607	0.0339

Panel C

	USD/JPY		GBP/USD		USD/EUR^{+}	
One-month horizon						
α	−0.0049	0.0055	−0.0006	−0.0074*	0.0038*	0.0074
β_{fwd}	−1.9377	−3.0161**	−1.3137	−1.4692	−4.5970***	−4.0266**
β_{ψ}		−1.0099**		−0.6869*		−0.8052*
R^2	0.0058	0.0185	0.0027	0.0080	0.0388	0.0487
Three-month horizon						
α	−0.0168*	0.0180	−0.0025	−0.0209	0.0081	0.0119
β_{fwd}	−2.3323*	−3.6999***	−2.0814	−2.3989	−3.4795**	−3.5397**
β_{ψ}		−1.8992**		−1.4011**		0.1532
R^2	0.0271	0.0607	0.0245	0.0438	0.0758	0.0760
Six-month horizon						
α	−0.0549***	−0.0054	−0.0026	0.0496	0.0235***	−0.0353
β_{fwd}	−6.3486***	−8.2463***	−4.1094	−4.7260	−10.5510***	−9.3788***
β_{ψ}		−1.7592**		−2.2483*		−1.6443*
R^2	0.0970	0.1215	0.0513	0.0918	0.2953	0.3180
One-year horizon						
α	−0.0850***	−0.0331	−0.0173	−0.0165**	0.0414*	−0.0845*
β_{fwd}	−2.4804**	−3.4377***	−2.7106*	−2.7128**	−4.6206***	−3.7687***
β_{ψ}		−2.9331**		−0.0243		−2.5645**
R^2	0.1451	0.2341	0.1687	0.1687	0.4737	0.5060

Panel D

	USD/JPY		GBP/USD		USD/EUR^{+}	
	High TED	Normal TED	High TED	Normal TED	High TED	Normal TED
α	0.0759	−0.0145	−0.0209	0.0107	−0.0278	−0.0067
β_{fwd}	10.1584	−4.2448**	2.5469	−4.5870	6.8976	−5.7916**
β_{ψ}	−2.2147**	0.4750	−0.5311*	−1.4506	−1.0010*	−0.7112
R^2	0.0561	0.0507	0.0115	0.0212	0.0223	0.0603

We thus regress carry trade returns on disagreement, conditional on a large (upper quartile) TED spread.[28] Second, the potential for delivering in case of a decrease in risk tolerance is larger when carry trade speculating positions are more substantial. Brunnermeier, Nagel, and Pedersen (2008) show that the lagged performance of carry trades describes speculator positions better than measures obtained from currency futures exchanges, because speculator positions tend to grow with a positive performance of the strategy. We thus use a conditioning dummy variable equal to one when the performance of the carry trade in the previous period was positive and the potential for delivering is thus likely to be substantial.

Table 11, Panel B, presents the carry trade conditional results. When funding liquidity problems are more serious (high TED spread) or speculator positions are large, the effect of disagreement on carry trade returns is very strong. In particular, when we use a large TED spread as a conditioning dummy variable, the economic significance of disagreement doubles and the explanatory power of the regression triplicates with respect to the unconditional case.[29] In contrast, when funding liquidity is normal or when there is a smaller potential for delivering, differences in beliefs are not a significant determinant of carry trade returns.[30] This evidence suggests that the interaction between funding liquidity and shocks to risk tolerance induced by changes of differences in beliefs has a large impact on carry trade returns.

We extend this analysis and study the link between currency returns and beliefs beyond the USD/JPY exchange rate and the one-month horizon. To do this, we cast our investigation in the context of expectation hypothesis regressions and the extensive literature on the forward premium puzzle. We thus estimate the following specification:

$$s_{t+n}-s_t=\alpha+\beta_{fwd}(f_{t,t+n}-s_t)+\beta_{\psi}\psi_{t,t+n}+e_{t,t+n}, \qquad (15)$$

where s_t represents the logarithm of the spot rate at time t, $f_{t,t+n}$ represents the logarithm of the forward rate at time t with maturity $t+n$, $\psi_{t,t+n}$ is the mean absolute deviation of foreign-exchange forecasts at time t for the horizon $t+n$, and n is either one-month, three-months, six-months, or one-year. This analysis is clearly related to the carry trade strategy. The dependent variable is one of the two components of the carry trade returns. Furthermore, the interest rate differential (the other component of carry trade returns) is controlled for on the right-handside through the forward premium.

Table 11, Panel C, shows the results of estimating Eq. (15) using Newey-West standard errors with 12 lags.[31] The first column for each currency and horizon just replicates the findings in previous literature and shows that the *forward premium puzzle* is generally present in our sample period, for our sample currency pairs, and at our sampling frequency.[32] In the second column, we add our empirical proxy for dispersion of beliefs. We observe that for all currencies and at almost all horizons, the dispersion of beliefs has significant predictive power for subsequent currency returns beyond the forward premium. More specifically, when we also include the MAD among the regressors, we obtain a substantial increase in the R^2. The currency pair for which MAD matters the most in terms of explanatory power is the USD/JPY, where we register an average increase across horizons of more than 100%. The forecast horizon where the MAD plays the strongest role across currencies is at one-month, where the R^2 increases, on average, by about 1.5 times. The additional explanatory power decreases monotonically for longer horizons.

The relation between dispersion of beliefs and currency returns is negative for all currencies. Higher dispersion of beliefs has a negative effect on FX returns, consistent with the results presented in Panel A for the USD/JPY exchange rate. The economic magnitude of the coefficients is remarkable. For example, a one-standard deviation shock to dispersion of beliefs causes a negative 3.3% annualized return on the USD/EUR$^+$ exchange rate at the one-month horizon. The economic significance monotonically decreases for longer horizons, much in the same way the explanatory power did.

To better understand the forward premium puzzle interaction with differences in beliefs and link it to the previous evidence on carry trades, we re-estimate the one-month regressions conditioning on a normal or large TED spread. Table 11, Panel D, illustrates intriguing results. When funding liquidity is at normal levels, the forward premium puzzle shown in the literature is present for all the three currency pairs we consider and differences in beliefs do not play a significant role. However, when the TED spread is in the upper quartile, the puzzle disappears and differences in beliefs become a strong predictor of an appreciation of the low-yield currency. This evidence suggests that the failure of the expectation hypothesis and the resulting puzzling excess returns to investing in high-yield currencies can be seen as a compensation for the potential risk of a sudden appreciation of the low-yield financing currency at times of an increase of aggregate implied risk aversion.

[28] The TED spread is the difference between the three-month London Interbank Offered Rate (LIBOR) Eurodollar rate and the three-month T-bill rate. The LIBOR rate reflects uncollateralized lending in the interbank market, which is subject to default risk, while the T-bill rate is riskless since it is guaranteed by the U.S. government. The TED spread has been traditionally used as an empirical proxy for funding liquidity problems, when it typically increases (e.g., recently in Brunnermeier, Nagel, and Pedersen, 2008).

[29] The correlation between the TED spread and differences in beliefs is not statistically significant (it is -0.08 in our sample) and therefore, it is an appropriate conditioning variable.

[30] A number of papers use the Chicago Board Options Exchange (CBOE) Volatility Index (VIX) option—implied volatility index as an empirical proxy for the amount of risk capital devoted to carry trade strategies. Conditioning on large VIX delivers similar but weaker results than conditioning on a large TED spread. VIX itself as an additional regressor is not statistically significant when MAD is included.

[31] We also check the robustness of the results using Newey-West standard error corrections with lags matching (a) the forecast horizons, and (b) the forecast horizon minus one, similar to Cavaglia, Verschoor, and Wolff (1994). The results are virtually unchanged.

[32] In this case, we do not present the results for USD/DEM and USD/EUR separately to save space. These results are qualitatively similar and available on request.

In summary, we show that the dispersion of beliefs appears to be a risk factor in currency markets with important asset pricing effects that goes beyond the strong effects shown for currency options. More specifically, the degree of dispersion of beliefs explains carry trade returns for the USD/JPY currency pair and, more generally, explains spot currency returns beyond the forward premium. This supports the predictions of the heterogeneous agent neoclassical literature, where disagreement impacts the first moment of asset returns, much in the same logic used to rationalize the equity premium puzzle by the earlier models (e.g., Detemple and Murthy, 1994). Given that short-selling constraints are absent in the FX market, this result is not consistent with the limits to arbitrage explanations of the behavioral literature.

5. Conclusions

Financial prices reflect uncertainty and reveal information about its resolution. Many causes of uncertainty are exogenous and are naturally resolved over time. Some other sources are instead of a strategic nature, as the timing of their resolution is a strategic variable for some agents with incentives to hide their information. Independent of the nature of uncertainty, however, when agents have different expectations of future market fundamentals, they can engage in dynamic trading either for speculative reasons or for dynamic hedging. In this paper, we test the link between differences in beliefs and asset prices using a measure of dispersion in beliefs that is based directly on the distribution of forecasts of future exchange rates provided by market participants. Combining this with an extensive data set on OTC currency options, we investigate a number of empirical questions that the theoretical literature has put forth.

We obtain four important results. First, we find that the differences of beliefs derived from currency forecasts help explain the level of implied volatility of currency options. This finding is statistically significant and robust, economically important, valid at different horizons, and goes beyond the effect of volatility of current fundamentals. Second, we find that disagreement affects the relative pricing of currency options in the cross-section, i.e., the shape of the implied volatility smile. Third, differences in beliefs impact the volatility risk premium. Last, our empirical proxy for differences in beliefs is also related to FX underlying returns. During periods of high differences in beliefs, carry trades are unwound and low interest rate currencies appreciate. This is consistent with the hypothesis that differences in beliefs affects the stochastic discount factor and the price of risk. Our empirical findings are valid both in the time-series and in the cross-section. All our results support the predictions of the heterogeneous agent neoclassical literature, where differences in beliefs represent an additional risk factor and highlight the role of uncertainty in asset pricing. The results are not consistent with some of the behavioral finance models that suggests a non-risk-based role for heterogeneous beliefs founded on the interaction between short-selling constraints and Miller's (1977) conjecture.

References

Andersen, T., Diebold, F., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71, 579–625.
Anderson, E., Ghysels, E., Juergens, J., 2005. Do heterogenous beliefs matter for asset pricing? Review of Financial Studies 18, 875–924.
Anderson, E., Hansen, L.P., Sargent, T.J., 2003. A quartet of semi-groups for model specification, robustness, prices of risk, and model detection. Journal of the European Economic Association 1, 68–123.
Bacchetta, P., Mertens, E., Van Wincoop, E., 2006. Predictability in financial markets: what do survey expectations tell us. Unpublished working paper. Study Center Gerzensee 06.04.
Bakshi, G., Kapadia, N., 2001. Delta-hedged gains and the negative market volatility risk premium. Review of Financial Studies 14, 527–566.
Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and differential pricing of individual equity options. Review of Financial Studies 16, 101–143.
Basak, S., 2000. A model of dynamic equilibrium asset pricing with heterogeneous beliefs and extraneous risk. Journal of Economic Dynamics and Control 24, 63–95.
Basak, S., 2005. Asset pricing with heterogenous beliefs. Journal of Banking and Finance 29, 2849–2881.
Bates, D., 2001. The market for crash risk. Unpublished working paper, NBER 8557.
Bofinger, P., Schmidt, R., 2004. Should one rely on professional exchange rate forecasts? CEPR Discussion Paper 4235.
Bollerslev, T., Tauchen, G., Zhou, H., 2008. Expected stock returns and variance risk premia. Review of Financial Studies 22, 4463–4492.
Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. Journal of Business 51, 621–651.
Brennan, M., 1998. The role of learning in dynamic portfolio decisions. European Finance Review 1, 295–306.
Brennan, M., Xia, Y., 2001a. Stock return volatility and equity premium. Journal of Monetary Economics 47, 249–283.
Brennan, M., Xia, Y., 2001b. Assessing asset pricing anomalies. Review of Financial Studies 14, 905–942.
Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes and stochastic volatility. Journal of Finance 55, 839–866.
Brunnermeier, M., Pedersen, L., 2009. Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.
Brunnermeier, M., Nagel, S., Pedersen, L., 2008. Carry Trades and Currency Crashes. NBER Macroeconomics Annual 23.
Buraschi, A., Jackwerth, J., 2001. The price of a smile: hedging and spanning in option markets. Review of Financial Studies 14, 495–527.
Buraschi, A., Jiltsov, A., 2006. Model uncertainty and option markets with heterogenous agents. Journal of Finance 61, 2841–2897.
Buraschi, A., Trojani, F., Vedolin, A., 2008. Economic uncertainty, disagreement and credit markets. Unpublished working paper, Imperial College, London.
Buraschi, A., Trojani, F., Vedolin, A., 2009. When uncertainty blows in the orchard: differences in beliefs and equilibrium volatility risk premia. Unpublished working paper, Imperial College, London.
Cagetti, M., Hansen, L., Sargent, T., Williams, N., 2002. Robustness and Pricing with Uncertain Growth. Review of Financial Studies 15, 363–404.
Carr, P., Wu, L., 2007. Stochastic skew in currency options. Journal of Financial Economics 86, 213–247.
Carr, P., Wu, L., 2009. Variance risk premiums. Review of Financial Studies 22, 1311–1341.
Cavaglia, S., Verschoor, W., Wolff, C., 1994. On the biasedness of forward foreign exchange rates: irrationality or risk premia? Journal of Business 67, 321–343.
Christensen, B., Prabhala, N., 1998. The relation between implied and realized volatility. Journal of Financial Economics 3, 125–150.
Croitoru, B., Lu, L., 2008. Asset pricing in a monetary economy with heterogeneous beliefs. Unpublished working paper, McGill University.
Daniel, K., Hirshleifer, D., Teoh, S., 2002. Investor psychology in capital markets: evidence and policy implications. Journal of Monetary Economics 49, 139–209.
David, A., 1997. Fluctuating confidence in stock markets: Implications for return and volatility. Journal of Financial and Quantitative Analysis 32, 427–462.

David, A., 2008. Inflation uncertainty, asset valuations, and the credit spreads puzzle. Review of Financial Studies 21, 2487–2534.
Detemple, J., Murthy, S., 1994. Intertemporal asset pricing with heterogeneous beliefs. Journal of Economic Theory 62, 294–320.
Diether, K., Malloy, C., Scherbina, A., 2002. Differences of opinion and the cross-section of stock returns. Journal of Finance 57, 2113–2141.
Drechsler, I., Yaron, A., 2008. What's vol got to do with it? Unpublished working paper, University of Pennsylvania.
Driessen, J., Perotti, E., 2003. Confidence building on euro convergence: theory and evidence from currency options. Unpublished working paper, University of Amsterdam.
Dumas, B., Kurshev, A., Uppal, R., 2009. Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. Journal of Finance 64, 579–629.
Epstein, L., Miao, J., 2003. A two-person dynamic equilibrium under ambiguity. Journal of Economic Dynamics and Control 27, 1253–1288.
Epstein, L., Schneider, M., 2003. Recursive multiple-priors. Journal of Economic Theory 113, 1–31.
Epstein, L., Wang, T., 1994. Intertemporal asset pricing under Knightian uncertainty. Econometrica 62, 283–322.
Fama, E., 1984. Forward and spot exchange rates. Journal of Monetary Economics 14, 319–338.
Frankel, J., Froot, K., 1987. Using survey data to test some standard propositions regarding exchange rate expectations. Unpublished working paper, NBER 1672.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. Journal of Mathematical Economics 18, 141–153.
Hansen, L., Hodrick, R., 1980. Forward exchange rates as optimal predictors of future spot rates: an econometric analysis. Journal of Political Economy 88, 829–853.
Hansen, L., Sargent, T., 2001. Robust control and model uncertainty. American Economic Review 91, 60–66.
Hansen, L., Sargent, T., 2005. Robust estimation and control under commitment. Journal of Economic Theory 124, 258–301.
Hansen, L., Sargent, T., 2007. Robustness. Princeton University Press.
Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Review of Financial Studies 6, 327–343.
Ito, T., 1990. Foreign exchange expectations: micro survey data. American Economic Review 80, 434–439.
Leippold, M., Trojani, F., Vanini, P., 2008. Learning and asset prices under ambiguous information. Review of Financial Studies 21, 2565–2597.
Lustig, H., Roussanov, N., Verdelhan, A., 2008. Common risk factors in currency markets. Unpublished working paper, UCLA.
Lustig, H., Verdelhan, A., 2007. The cross-section of currency risk premia and U.S. consumption growth risk. American Economic Review 97, 89–117.
Maenhout, P., 2004. Robust portfolio rules and asset pricing. Review of Financial Studies 17, 951–983.
Malz, A., 1997. Estimating the probability distribution of the future exchange rate from option prices. Journal of Derivatives, 18–36.
Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies. Journal of International Economics 14, 3–24.
Miller, E., 1977. Risk, uncertainty, and divergence of opinion. Journal of Finance 32, 1151–1168.
Morgan, J.P., 1996. RiskMetrics Technical Document. Unpublished working paper, J.P. Morgan.
Qu, S., Starks, L., Yan, H., 2003. Risk, dispersion of analyst forecasts and stock returns. Unpublished working paper, University of Texas Austin.
Rajan, R., Servaes, H., 1997. Analyst following of initial public offerings. Journal of Finance 52, 507–529.
Sandroni, A., 2000. Do markets favor agents able to make accurate predictions? Econometrica 68, 1303–1341.
Scheinkman, J., Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of Political Economy 111, 1183–1219.
Schwert, W., 1989. Why does stock market volatility change over time? Journal of Finance 44, 1115–1153.
Shimko, D., 1993. Bounds of probability. RISK 6, 33–37.
Verdelhan, A., 2010. A habit-based explanation of the exchange rate risk premium. Journal of Finance 65, 123–146.
Veronesi, P., 1999. Stock market overreaction to bad news in good times: a rational expectations equilibrium model. Review of Financial Studies 12, 975–1007.
Zapatero, F., 1998. Effects of financial innovations on market volatility when beliefs are heterogeneous. Journal of Economic Dynamics and Control 22, 597–626.