Exchange Rate Fluctuations, Consumer Demand, and Advertising: The Case of Internet Search

Susan Athey
Harvard and Microsoft Research

Maya Cohen Meidan
Microsoft Research

March 10, 2011

Abstract

This paper addresses the question of how exchange rates affect consumer demand in markets where advertising plays an important role. We identify an effect that has not been emphasized in the existing literature: when foreign exchange rates appreciate, a foreign product becomes more expensive to domestic consumers, but at the same time, advertising becomes cheaper for the foreign advertiser. Thus, demand depends on exchange rates through two channels, the price charged in domestic currency, and the extent of advertising. Our paper attempts to quantify the impact of exchange rates in a market where advertising prominence, its cost to foreign firms, and its impact on consumers can all be directly measured: internet search advertising, a $25 billion international market. We decompose the impact of exchange rate shocks into its impact on the prominence of foreign advertisements and its impact on consumer demand conditional on advertising prominence. We show that in some cases, when foreign currency appreciates, despite foreign products becoming more expensive to domestic consumers and consumers responding to this effect, consumers visit foreign websites more often, due to increases in advertising prominence.

1 Introduction

The question of how exchange rate shocks affect international trade flows is a central issue in international economics. Many authors have found evidence that exchange rate shocks are not fully passed through to consumers in domestic markets.[1]. In his 2002 survey, Engel (2002) identifies three main explanations for "incomplete pass through": (i) the existence of local costs, such as costs for non-traded services, that affect the marginal cost of production, even among goods that are usually considered to be traded; (ii) imperfect competition, which leads manufacturers or retailers to adjust markups, in cases where goods are "priced to market"; and (iii) nominal price rigidities (or "menu costs"), which leave prices fixed in local currencies.

This paper focuses on a channel affecting exchange rate pass through that has not received much attention in the literature to date: advertising. Advertising expenditures are typically priced in the consumer's local currency. Unlike the "local costs" that have been considered in the literature to date, however, advertising is typically a fixed expenditure that has the impact of *shifting* the consumer demand curve. Since exchange rates affect the cost of advertising and thus the optimal choice of advertising, this creates a channel whereby exchange rates, through advertising, shift consumer demand. If the demand curve is endogenous to exchange rates, then in markets where advertising plays an important role, empirical strategies for analyzing the impact of exchange rates may need to be modified, and the interpretation of findings about the impact of exchange rates on prices and output may be qualified.

In the first part of our paper, we develop a simple theoretical model of endogenous advertising. We consider two cases: the "price to market" case and the "world price" case. In the "price to market" case, we show that, for fixed domestic price and domestic advertising

[1]See Goldberg and Knetter (1997) for a survey of earlier literature. More recent contributions emphasize the role of firms-heterogeneity in affecting the aggregate pass-through effect (e.g. Auer and Chaney (2009), Berman, Martin, and Mayer (2010)); Study the impact of exchange rates on prices across industries and countries using disaggregated data (e.g. Bhattacharya, Karayalcin, and Thomakos (2008), Campa and Goldberg (2005)); and examine the role of structural determinants in the degree of exchange rate pass through into prices (e.g. Fitzgerald and Haller (2010), Hellerstein and Goldberg (2010)).

expenditure, an appreciation of the foreign currency decreases the revenue received by the foreign firm and decreases its advertising expenditures proportionally (in terms of foreign currency) while leaving foreign currency marginal production costs unchanged. In turn, this implies that the foreign exchange rate appreciation leads to the same impact on pricing and advertising choices as an increase in marginal production cost. All else equal, an increase in marginal production cost leads to a reduction in the returns to advertising, because the markup goes down. Under fairly general conditions, advertising and price are complements, and a firm's endogenous choice to decrease advertising dampens the incentive of the firm to raise prices, thus decreasing exchange rate pass-through. In this way, endogenous advertising may be part of the explanation for low exchange rate pass through in the "price to market" case.

We contrast this case to one where prices are set on a world market. Many products in the travel industry have this feature: foreign hotels and airfare are typically sold at a single world price, quoted in the foreign producer's currency, and domestic consumers bear the full impact of short-term exchange rate fluctuations.[2] For products that are *not* priced to market, then a foreign exchange rate appreciation leads to proportionally higher domestic prices, without affecting the markup received by the foreign producer. At the same time, the marginal cost of advertising goes down, increasing the returns to advertising. When advertising and price are complements, both effects reinforce one another, and the optimal level of advertising goes up. This in turn mitigates the impact of a foreign exchange rate appreciation on output.

In the absence of price and advertising data, the impact of endogenous advertising might appear to an observer as lower-than-expected pass through: the impact of exchange rate fluctuations on trade flows is dampened. In the presence of price data showing full pass-through, failing to account for the advertising channel would lead one to conclude that demand is less elastic than in reality.

[2] Of course, if a hotel relies *primarily* on international tourism, it may choose to adjust its world price in response to an appreciation in its own country relative to the rest of the world; but it is unlikely to adjust prices in response to an idiosyncratic fluctuation between its currency and the currency of another country that forms a small share of its overall demand.

We then turn to an empirical application where advertising and its impacts can be precisely measured: the internet search advertising market. This is a rapidly growing international market, with over $25 billion in annual advertising revenue worldwide, although this may understate the impact of the advertising; Google recently published estimates suggesting that the value of economic activity mediated by search advertising was up to eight times as large as the advertising fees.[3] Search advertising is dominated by direct marketing, whereby advertisers seek to reach consumers who are ready to make a decision or a purchase. Advertisers recognize that when a consumer enters a search query, she is expressing an intent to find information or take an action at that time, and this may be the last opportunity to reach the consumer prior to a decision.

Search advertising has several interesting features. First, the position of an ad on a search results page has a large and predictable impact on consumers: placing the same ad in a higher position on the page brings it substantially more clicks (that is, consumer visits to the advertiser's website). Second, the ranking of ads and their prices are determined using an auction, where higher bids lead to higher positions, and advertisers pay more to occupy higher positions. Thus, there is a direct, measurable, and immediate impact of changes in advertiser bids on the number of consumers an advertiser can attract. Third, each foreign advertiser places its bids in a single currency, even though the currency in which auctions are held varies with the location of the user; when the advertiser currency differs from the auction currency, the search advertising platform converts their bids using an exchange rate. Thus, an advertiser who does *not* immediately adjust his bid in response to an exchange rate appreciation of his bidding currency, will nonetheless enter a higher *effective* bid into the auction, leading to an improvement in the advertiser's position on the page (and thus an increase in the quantity of advertising purchased) as well as an increase in the number of visitors the advertiser attracts to its website.

All of these features of search advertising make it a rich and interesting environment in which to study the impact of exchange rates on advertising. The immediate and measurable impact of exchange rates on advertising output and consumer responses allow us to precisely estimate exchange rate impacts. We study these impacts using a proprietary dataset from

[3]Reported on Google's web site at http://www .google.com/economicimpact/index.html

Microsoft's search advertising platform, including data on advertising bids, ad placements, and consumer searching and clicking decisions. However, we do *not* have data on the final prices charged to consumers by advertisers.

A complicating feature of search advertising, and the data we have available, is that the relationship between firm prices and consumer choices is less direct than in many markets. Although some search advertisements include product price information, this is relatively unusual. Common search terms entered by users are fairly general, concerning categories of products rather than specific, individual products. In the absence of direct price information, consumers make clicking choices based on their expectations about the value they will get by clicking on a website. These expectations would be based on their knowledge of exchange rates, their familiarity with the web site's country of origin, and their familiarity with the web site's pass through policies. Clearly, we should see a great deal of heterogeneity in the responsiveness of consumer clicking behavior to exchange rate fluctuations. For this reason, in our empirical work, we measure aggregate responses across a random sample of search activity, but we also analyze several specific search advertising "markets" (defined by search terms) in depth. We focus on search-terms that exhibit high international activity, where consumers are more likely to be aware of and respond to exchange rate fluctuations.

Thus motivated, our empirical work proceeds in several steps. For a random sample of world-wide user queries, we study the effects of exchange rate fluctuations on the prominence of advertisements (where we calibrate prominence according to the number of clicks a typical advertiser expects from different positions on the page) as well as on user clicking behavior conditional on advertising prominence. Thus, we decompose the aggregate impact of exchange rates on the volume of user clicks, into the effects due to changes in advertising, and the effects due to exchange rates changing consumer beliefs about the value of clicking on foreign web sites. Due to the great heterogeneity of products, and presumably of exchange rate impacts, encompassed by search advertising, the estimated impacts of exchange rate fluctuations through prices are not statistically significant for the most part.

By contrast, we find that while the advertising prominence effect varies across the main advertiser countries in the data, the elasticity of clicks to a foreign advertiser with respect to a foreign exchange rate appreciation is positive and the effects are mostly statistically

significant. We examine the direct impact of exchange rate fluctuations on the actual bids placed by the advertisers and find that indeed, in accordance with the no pricing to market scenario, when the foreign currency appreciates then advertisers tend to actively increase their bids, and generally the effect is large relative to the price effect, and significant.

Thus, we conclude that in aggregate, the impact of a foreign exchange rate appreciation on user clicks through price, conditional on ad prominence, varies considerably across countries and search phrases, while it leads to an unambiguous *increase* in advertising prominence in most cases. These competing effects lead to substantial heterogeneity in the overall direction and magnitude of the aggregate effect of exchange rate shocks on consumer visits to foreign web sites. Figure 1 illustrates this point: it shows the distribution of the ratio between changes to consumer clicks and changes to user-account country exchange rates in the random sample, by user-account country pairs. The empirical distribution has a (small) negative mean, indicating that on average users tend to click less on foreign ads when the foreign currency appreciates. Yet the distribution of the ratio between click changes and exchange-rate fluctuations is quite dispersed. The dispersion could be attributed in part to the counteracting advertising effect and in part to the heterogeneity of the price effect across countries and markets. Our empirical analysis indicates that, with the exception of the United States and France, the advertising effect is much larger than the price effect.

The next part of our empirical work focuses in on several specific search phrases. In this context, we estimate a discrete choice model of user clicking behavior on advertisements in different positions, studying how the exchange rate between the user and advertiser's country affects clicking behavior, conditional on the advertiser's position on the page. To capture the heterogeneity of user information and beliefs about advertisers, we distinguish between different types of advertisers as well as their country of origin. For the search phrases we select, there are examples of user-country pairs where the effect of exchange rates on consumer choice is more pronounced than in the aggregate data, but in many cases, the impact is still small and not statistically distinguishable from zero.

We then perform a series of counterfactual exercises to decompose the impact of exchange rate shocks into its components. We find that in some cases, the increased prominence of advertising in response to a foreign exchange rate appreciation outweighs the impact of

the appreciation on user clicking, while in other cases, the decreased propensity of users to click conditional on prominence dominates. In particular, it seems that when the sensitivity of users to exchange rate fluctuations is high, advertisers respond less to exchange rate fluctuations, and vice versa.

Taken together, our empirical results suggest that in the context of internet search, advertising levels (the prominence of foreign advertisements) do change substantially in response to exchange rate shocks, and these changes have a large impact on the volume of international transactions. In addition, despite the lack of precise consumer information about how exchange rates will impact prices, we find that consumers do respond to exchange rate shocks in their choices among competing advertisements, but the consumer response is quite heterogeneous. From the perspective of the broader literature on exchange rate pass through, our results can be taken as evidence that advertising is both responsive to exchange rates and important in terms of causing changes to consumer demand.

In future work, we plan to assess the importance of "sticky advertising bids" on outcomes. We further plan to study measures of user activity on advertiser web sites after the click, to control for the fact that exchange rates affect not just traffic to the web site, but also their propensity to purchase once they arrive at a web site.

2 A Simple Model of Advertising by Foreign Firms, Exchange Rates and Domestic Demand

Consider a foreign producer who sells his product to domestic consumers at a price p (given in domestic-currency units). The price of the product in the foreign producer's home market will be denoted by p^f. We shall assume that the foreign producer can advertise his product in the domestic market and thus increase the demand. Let $K(a)$ denote the cost (in domestic currency) of a units of advertising, and assume that $K_a > 0$.

The demand for the foreign product is denoted by $q(p, a)$, where q is twice-differentiable with $q_p < 0$, $q_a > 0$. We will assume that $q_{pa} \geq 0$: higher advertising levels mitigate the impact of raising price on demand. For simplicity, assume that all other costs are incurred in

the foreign market and hence c^f, the marginal cost, is given in units of the foreign currency.

In this setup, exchange rate fluctuations affect the domestic demand through two channels: price and advertising by the foreign advertiser. Price changes generate a movement *along* the demand curve, while advertising-quantity changes lead to a *shift* of the domestic demand function.

Let e denote the exchange rate in units of domestic currency per foreign currency. The foreign producer's objective is to maximize his profits from domestic consumers. We shall differentiate between cases where the foreign producer sets a price that is unique to the domestic market (pricing-to-market), and those where the price of the product is determined in the foreign country and domestic consumers face a price that is either pegged-to or quoted in foreign currency: $p = p^f e$ (no pricing-to-market). An example of the first case would be a U.S. consumer buying imported beer or cheese in his local supermarket. An example of the second case would be a U.S. consumer purchasing a vacation package in Europe.

2.1 pricing-to-market

In the pricing-to-market case, consumers face a domestic price p, which is set by the foreign producer specifically for their market and is quoted in units of the domestic currency. A foreign producer's profits can be written:

$$\Pi(p, a; e, c^f, K(\cdot)) = q(p, a)(\frac{p}{e} - c^f) - \frac{K(a)}{e} \tag{1}$$

Maximizing $\Pi(p, a; e, c^f, K(\cdot))$ is equivalent to maximizing

$$\tilde{\Pi}(p, a; e, c^f, K(\cdot)) = e \cdot \Pi(p, a; c^f e, K(\cdot)) \tag{2}$$

so we will focus on this problem instead, as it clarifies that a change in the exchange rate affects the optimization problem in a way that is equivalent to a shift in marginal cost. In particular, even though a domestic exchange rate depreciation (increase in e) lowers the marginal cost of advertising for the foreign producer, it lowers the effective price received by the producer proportionally. This implies that the returns to advertising go down when the domestic exchange rate depreciates, since the profit the firm receives from each sale goes

down (due to the effective increase in marginal cost). To see this formally, the first-order necessary conditions for (a^*, p^*) to be optimal choices are:

$$\frac{\partial \tilde{\Pi}}{\partial p} = q_p(p^*, a^*)(p^* - c^f e) + q(p^*, a^*) = 0 \tag{3}$$

$$\frac{\partial \tilde{\Pi}}{\partial a} = q_a(p^*, a^*)(p^* - c^f e) - K_a(a^*) = 0 \tag{4}$$

To understand the impact of exchange rate changes on the optimal choices of price and advertising, observe that:

$$\tilde{\Pi}_{pa} = q_{pa}(p, a)(p - c^f e) + q_a(p, a) > 0 \tag{5}$$

$$\tilde{\Pi}_{ae} = -q_a(p, a)c^f < 0 \tag{6}$$

$$\tilde{\Pi}_{pe} = -q_p(p, a)c^f > 0 \tag{7}$$

Thus, an increase in e, acting as an increase in marginal cost, has the direct effect of decreasing the returns to advertising and increasing the returns to a higher price. However, because price and advertising are complements, additional structure is required to derive monotone comparative statics on price and advertising. Despite this, so long as q_{ap} is not too large, we can conclude that the optimal level of advertising decreases in e, which in turn implies that prices rise more in response to an increase in e when advertising is held fixed than when the producer selects it optimally. Thus, endogenous advertising is a potential explanation for incomplete pass-through of exchange rate shocks.

Lemma 2.1 *In markets where advertising increases demand and there is pricing-to-the domestic market by the foreign firms, then if* $0 \leq q_{pa}$ *and* q_{pa} *is not too large and the profit function is concave in* (a, p)*:*

1. $a^*(e)$ *is decreasing, that is, the optimal domestic advertising by the foreign producer decreases when the domestic currency depreciates (e rises).*

2. *If* $\hat{p}(e; a) = argmax_p \tilde{\Pi}(p, a; e, c^f, K(\cdot))$*, then* $\hat{p}(e; a)$ *is increasing in e and a: when advertising is exogenous, the optimal price increases in the level of advertising and the exchange rate.*

3. *For any* $e_H > e_L$, $p^*(e_L) = \hat{p}(e_L; a^*(e_L)) < \hat{p}(e_H; a^*(e_L))$, *and further,*

$$\hat{p}(e_H; a^*(e_L)) > \hat{p}(e_H; a^*(e_H)) = p^*(e_H). \tag{8}$$

That is, if advertising is held fixed when e *rises, then the optimal price goes up, but when advertising is permitted to adjust as well, price falls relative to the fixed-advertising optimal level of price: endogenous advertising mitigates the exchange rate pass-through.*

Proof Part (i): The implicit function theorem and concavity of the objective imply that

$$\begin{aligned} sign\frac{\partial a^*(e)}{\partial(e)} &= sign(-\tilde{\Pi}_{ae}\tilde{\Pi}_{pp} + \tilde{\Pi}_{pe}\tilde{\Pi}_{ap}) &(9)\\ &= sign(q_a(p,a)c^f(q_{pp}(p,a)(p-c^f e) + q_p(p,a)) - q_p(p,a)q_{pa}(p,a)c^f(p-c^f e)) &(10)\end{aligned}$$

Our assumptions imply that this is negative if q_{pa} is sufficiently small.

Part (ii) follows because, as derived above, $\tilde{\Pi}_{pa} > 0$ and $\tilde{\Pi}_{pe} > 0$. Part (iii) is a direct consequence of parts (i) and (ii).

When the foreign currency depreciates, the firm's optimal response is the same as if it experienced a marginal cost increase. A marginal cost increase decreases the optimal advertising level so long as the complementarity between advertising and price in the demand function is not so strong to outweigh the direct effects of the cost increase on advertising and price. Since exogenous increases in advertising and the exchange rate e both increase price, the effect of e on price with advertising fixed is positive, while the indirect effect of e on price through a reduction in advertising is negative. This formalizes the idea that advertising mitigates exchange-rate pass through.

2.2 No pricing-to-market

In the absence of pricing-to-market, the price of the product is determined in the foreign market, and is given exogenousely to domestic consumers as p^f in units of the foreign currency. Hence, consumers face a domestic price $p = p^f e$: there is full pass through of exchange-rate fluctuations to the price that domestic consumers are facing. Since price is is determined outside of the domestic market by assumption, the foreign producer maximizes his profits

from domestic consumers as a function of a only, as follows:

$$\max_a \Pi(a; p^f, e, c^f, K(\cdot)) = \max_a q(p^f e, a)(p^f - c^f) - \frac{K(a)}{e} \tag{11}$$

The optimal choice of a, the level of advertising in the domestic market, is given by the F.O.C.:

$$\frac{\partial \Pi}{\partial a} = q_a(p^f - c^f) - \frac{K_a}{e} \quad = \quad 0 \tag{12}$$

In contrast to the previous case, when there is no pricing-to-market, the price faced by domestic consumers fully reflects foreign exchange rate fluctuations. An increase in e decreases the marginal cost of advertising, which favors an increase in advertising; and, unlike the pricing-to-market case, the price-cost margin (which determines the incentive to shift demand out at a given consumer price) is assumed to stay fixed. Further, since the price faced by consumers in domestic currency goes up, and since advertising and price are assumed to be complements at increasing demand, there is a reinforcing effect in favor of increasing advertising.

Formally, since there is only one choice variable in this case, $sign\frac{\partial a*(e)}{\partial(e)} = sign\frac{\partial \Pi}{\partial e \partial a}$, and

$$\frac{\partial \Pi}{\partial e \partial a} = q_{ap} p^f (p^f - c^f) + \frac{K_a}{e^2} > 0. \tag{13}$$

This implies:

Lemma 2.2 *In markets where advertising is important and there is no pricing-to-the domestic market by foreign firms ($p = p^f e$):*

1. *$a^*(e)$ is increasing, that is, the optimal domestic advertising by the foreign producer increases when the domestic currency depreciates.*

2. *$p'(e) = p^f$, by assumption, there is full pass-through of exchange-rates fluctuations to the price faced by domestic consumers*

3. *For $e_H > e_L$, if e increases from e_L to e_H, then $q(p^f e_H, a^*(e_L)) < q(p^f e_H, a^*(e_H))$: when advertising adjusts optimally, output is higher relative to the case where advertising is fixed.*

3 Online-Search Advertising with International Users and Advertisers

The online search market is a $\$25B$ international market. The importance of international (cross-border) transactions in this market has been constantly increasing. Google's international revenues surpassed its domestic revenues for the first time in 2008.[4]. International ecommerce is generally on the rise: for example, eBay (one of the largest internet search advertisers) recently reported that its cross-border transaction are "an increasingly important source of both revenue and profits." [5]

The online-search market provides a natural environment for quantifying the impact of exchange rate fluctuations on domestic demand in markets where advertising is important. Online-search engines create a virtual marketplace where users and advertisers from all over the world can find one another relatively easily. In fact, most online-advertising platforms enable an advertiser to target his ad at users from multiple foreign countries at no additional up-front cost (showing the ad to users from anywhere in the world is the default when advertising on Microsoft's internet search platform). Because of this, many advertisers on online-search platforms advertise world-wide, but do not tailor their landing pages or product prices specifically to users from different countries.

In this paper, we analyze data from Microsoft's search advertising platform in the period January, 2009 to June, 2010. During that period, Microsoft's share of the U.S. online search market was around $10 - 12\%$ and its worldwide share was around 5%; Bing's yearly revenues were around $\$1B$. Advertisers from anywhere in the world could open an account with Microsoft's online-advertising services, and users from anywhere in the world could access Microsoft's online-search websites. However, only in certain countries did Microsoft's ad platform serve a high volume of users. During our sample period, Microsoft operated online-search and advertising platforms in the following countries (also the countries where Microsoft properties had a relatively large market share of users overall): United States, Canada, France, the United Kingdom and Singapore. In other countries, Microsoft typically

[4]Reported on Google's web-site at http://investor.google.com/financial/tables.html.

[5]eBay 10-K, filed 2/17/2010, p. 23.

had agreements with other companies (such as Yahoo!) to provide advertising for users of its country-specific sites, and so users from those countries saw Microsoft ads only in special circumstances (such as navigating to specific Microsoft sites). Given that, the share of advertisers coming from countries where Microsoft operates a platform was much greater than the share of advertisers from other countries.

During late 2010, Microsoft began to take over both algorithmic and advertising search services for Yahoo!, beginning in the U.S. and expanding worldwide. In preparation for its substantial expansion, many aspects of Microsoft's domestic and international advertising platforms were updated and changed, and our description here of the historical system should not be taken as indicative of current practices.

In order to participate in search advertising on AdCenter (Microsoft's search advertising platform), the advertiser placed bids on particular keywords indicating which user search terms were appropriate for the advertisement. The advertiser could choose the exact user search terms that match its keywords, but it could also select "broad match" and other types of matching, whereby the ad platform used algorithms to match the advertisement to search terms that were similar or that contain the advertiser's keywords. The bids were placed on a *per-click* basis: the advertiser bid the amount that they would be willing to pay if a user were to click on their ad. Whenever a user entered a search phrase into Microsoft's search engine, the platform ranked the ads of all the advertisers that placed a bid on a relevant keyword. Moving an ad to a higher position increased the number of clicks it received, and the rate of decay of clicks as the ad shifted to lower positions was fairly substantial. Thus, advertisers were willing to bid more to attain higher positions.

The ads were ranked based on their "quality-weighted" bids. *Quality scores* were used by all three major search advertising platforms (Google, Yahoo! and Bing). While the exact formula for calculating quality scores was proprietary, a major factor in an ad's quality score was its historical "clickability" (roughly, how many times the ad has been clicked out of the total number of times it had been shown, adjusted for the position it was shown in, so that clickability gives the expected number of clicks if the ad were to be shown in the top position).[6] In that sense, there was a feedback loop from past clicking behavior to quality

[6]See Wikipedia's article on quality scores, http://en.wikipedia.org/wiki/Quality_Score.

scores and ad rankings. If quality scores are *equal to* the ad's clickability, we can think of the search engine as ranking ads according to the expected revenue it gets from placing the ad in a given position, since the revenue is equal to the number of clicks times the price per click.

Given the rankings of advertisements for the particular auction, the platform placed the winning ads on the page in decreasing order (from highest to lowest quality-weighted bid), first on the top of the page and then on the side-bar (the space to the right of algorithmic results). The payments were determined by a *generalized second-price auction*. In particular, whenever a user clicked on an ad in position i, advertiser i paid the minimum price that would put his ad in its current position: the quality-weighted bid of the advertiser in position $i+1$, $s_{i+1}b_{i+1}$ divided by advertiser i's quality score, s_i. For example, if the user clicks on the ad in position 1, advertiser 1 is charged a "price per click" (PPC) equal to $\frac{s_2 b_2}{s_1}$. If the quality score is equal to the "clickability," then the revenue collected from the first advertiser is expected to be s_1 times the PPC, or $s_2 b_2$, the revenue that would have been collected from the second advertiser if that advertiser had been placed in the first position instead. The real-world auction included other features as well, including factors such as reserve prices and minimum thresholds on quality scores for appearing on a page.

Equilibria in generalized second price auctions have been studied by Varian (2007) and Edelman, Ostrovsky, and Schwartz (2007). These authors modeled the game as a one-shot auction with complete information, and showed that while truthful bidding is not optimal for bidders, the outcomes in a particular equilibrium selection are equivalent to Vickrey outcomes. In particular, advertisers are ranked in order of their score-weighted values per click. In addition, Athey and Nekipelov (2010) developed a richer model that incorporates the realistic feature that advertisers hold their bids fixed for periods of time, during which their bids are applied to many user queries, quality scores are updated by the ad platform algorithms, and competitors pause and resume campaigns. This uncertainty leads to a unique equilibrium under fairly general conditions, and Athey and Nekipelov (2010) develop a structural econometric model to estimate bidder valuations from bidding data.

The advertisers' bids were placed in units of their chosen *billing currency*.[7] The adver-

[7]Billing currency is chosen by the advertiser when they first set up an AdCenter account. Once the

tisers could change their bid as often as they liked, although in practice the majority rarely did so. One important aspect of auctions with international bidders was how the platform treated bids in foreign currency. For example, auctions for ads on Microsoft's North America online-search domains were run in U.S. dollars, while auctions for the French market were held in Euros. Bids in foreign currencies were converted to the market currency using Microsoft's billing exchange rate, thus generating the *effective bids* for the auction. Importantly, during our historic sample period Microsoft's billing exchange rates were updated roughly *once every thirty days*. Advertisers were not informed as to the relevant billing exchange rate or the dates of the billing-cycle period.

The billing-cycle generated mechanical changes to foreign bids. Although idiosyncratic, this environment creates a convenient setup for estimating the impact of exchange rates on domestic demand through advertising: the exogenous shocks to advertising enable us to obtain a clean estimate of the impact of advertising on demand by assessing the impact of billing exchange-rate changes on advertisers' effective bids, position and the quantity of ad-clicks by users. The data also enable us to contrast the mechanical bid changes with actual bid changes by foreign advertisers in response to these shocks, in terms of the magnitude and the direction of mechanical versus voluntary bid changes.

When considering how the search advertising setting fits into our theoretical model, note first that the presence of the "mechanical" bid changes introduces a new feature not present in the general theoretical model: given "sticky" advertising bids (presumably due to transactions costs faced by bidders, such as the time and effort to monitor campaigns and adjust bids), a domestic exchange rate depreciation leads to an increase in foreign advertising (by promoting the foreign ads to higher positions) that are not necessarily optimal for the bidder. In some cases, these changes can be quite pronounced, and they can potentially dominate other factors. We call this effect the *billing cycle effect.*

Once a bidder does adjust its bid, generating what we call the *bid change effect,* the bid determines its position (actually, a distribution over positions, in an environment with uncertainty), which in turn determines the quantity of clicks. So holding fixed the user's beliefs about the value of clicking on a particular advertiser, we can think of an advertiser's

account is set, the advertiser cannot change their billing currency.

bid as selecting a quantity of expected clicks, which would look like an increase in a in the formal model, one that comes with an advertising cost and that shifts out demand. From the perspective of an individual foreign bidder bidding against a set of domestic bidders, a depreciation of domestic currency leads to change in billing exchange rates that results in a proportional decrease in the cost of every advertising position, and thus a proportional decrease in the cost of any level of advertising a, just as in the theoretical model. However, in practice, if advertisers from different countries compete and the billing exchange rate from each country changes, the shape of the advertising cost curve faced by the advertisers may change. In addition, as we will discuss in more detail below, a change in exchange rates affects the user's desire to click on the ad, and if the advertiser's price in domestic currency changes, the user may be less likely to purchase after a click. In this paper, we will not attempt to directly measure post-click user behavior, but we plan to address this in the future.

Finally, there is a third component of the advertising channel. As discussed above, an ad's "clickability" enters directly into both the ad's rank and the ad's price. If consumers click less frequently on an ad, then the estimated "clickability" will fall, and with it the advertiser's quality score. This effect will typically counteract the billing cycle and bid change effects, since a domestic exchange rate depreciation will improve an ad's ranking through the billing-cycle and bid effects, but if the users click less, then this effect, which we call the *quality score effect,* will worsen the ad's ranking. Table 1 summarizes the expected impact of a foreign-currency appreciation on user-clicks in the online-search market through the price and advertising channels, breaking the latter into its three components. Table 2 does the same for the case of pricing-to-market.

4 Microsoft Online-Search Auction Data and Descriptive Statistics

In order to examine the impact of exchange-rate fluctuations on the behavior of online-search users, we have first collected data on daily exchange-rates between January 2009 and

June 2010 (our data period) from Global Financial Data. Figures 2, 3 and 4 plot the daily exchange rates for three major user country-currencies in our data: the United States and Europe, the United States and Canada and the United States and the United Kingdom. The plots demonstrate that during our sample period, there have been considerable exchange-rate fluctuations for all three country-pairs. In Table 4, we present the standard deviations, minimum and maximum values of mean-normalized exchange-rates for the major user- and account-currencies in our data. The table reiterates the fact that exchange-rate fluctuations were by no means negligible during this period.

To assess the impact of exchange rate fluctuations on demand, we use two rich, proprietary datasets on U.S. and non-U.S. user-searches on Microsoft's online-search websites. The first is a random sample of search activity, which has been collected over the period January 2009 through June 2010, every Monday, Wednesday and Saturday. This is a representative sample of searches conducted through Microsoft. The representative sample is used to assess the overall impact of exchange-rates on the online-search market. The second dataset focuses on particular search phrases with high revenues and high international activity (in terms of foreign ads and clicks by domestic users on foreign ads). The specific search phrase sample is used to assess how clicks respond to exchange rate fluctuations and quantify the price versus advertising channel effects for each search phrase and user-country separately.

The variables we collected include: user country, local date and time, the user search term, the position and identity of the ads shown on the page, the text of the ads and their landing pages, the billing country of each advertiser, and the billing exchange-rates. In addition, we have rich auction-level data, including bids, quality scores, prices paid and rankings on the page, as well as ad clicks. There are also some variables available that pertain to user engagement with the advertiser's site after the click, but we do not utilize that data in this version of the paper.

Table 3 provides a summary of ad exposures and ad clicks on foreign and domestic ads for the United States, major non-U.S. markets and other non-U.S. markets in the representative sample of Microsoft's search activity. International (non-U.S.) users represented about 40% of search activity on Microsoft's websites. Notably, the United States had the highest share of domestic (U.S.) ad exposures, clicks on domestic ads and revenues from domestic

advertisers. For major non-U.S. markets (where AdCenter operates a platform), about 30% of ad clicks were on foreign (non user-country) advertiser ads. For smaller non-U.S. markets (no AdCenter websites), almost all of the clicks and revenues came from ads of advertisers foreign to the user's market. Overall, AdCenter is biggest in the United States, where its user base is large enough to attract a large number of local advertisers. In foreign markets, AdCenter does not have sufficient users to attract as large of a local advertiser base, leaving more room for international advertisers. In addition, many of the foreign markets rely more on international trade for the kinds of products sold in search than the U.S. market.

5 Empirical Models and Estimation Results

Our goal in this section is to examine empirically how exchange rate fluctuations affect domestic demand in our online-search dataset. Unfortunately, we do not at present have data on the prices charged by online advertisers to users after they click on ads. In what follows, we focus on clicks on ads by users as a measure of the volume of transactions for products advertised through search.

In the theoretical model, we discussed two primary channels through which exchange rate fluctuations may affect domestic transaction volume. The first was the pricing channel. In the context of our online-search datasets, this would be the immediate effect of a change in the user-per-advertiser country exchange rate on the quantity of clicks, all else being equal (advertising level in particular). We observe that the pricing channel operates slightly differently in the search advertising context than in a typical product market, since users typically cannot directly observe the prices offered by the advertiser when they select an ad. The impact of an exchange rate shock on user clicking behavior depends on the user's awareness of exchange rates, of the advertiser's country of origin, and of the advertiser's prices. We expect these factors to vary widely across search terms and advertisers. Users may be more aware of exchange rates and their impact on prices when planning vacation travel than when buying inexpensive consumer goods.

The second channel highlighted in the theoretical model is the advertising channel. Above, we noted that there are actually three components to the advertising channel in

search advertising: the mechanical billing cycle effect, the bid change effect, and the quality score effect. We will attempt to decompose the two main channels, as well as the three components of the advertising channel, in the empirical work.

5.1 Aggregate Effects of Exchange Rates on a Representative Sample

As a first step, we examine our representative sample of internet search activity to see whether actual normalized exchange rates (user-country per advertiser-country currency, normalized to have mean one) have the expected effect (negative) on user clicks, holding advertising fixed. Similarly, we use the sample to examine whether normalized billing exchange rates (user-market currency per billing-currency, normalized to have mean one) have the expected effect on clicks (positive).

We start by looking at the total effect of exchange rate fluctuations on user clicks. Figure 1 depicts, for all user-advertiser country pairs in the data, the ratio between the percentage change in clicks and the percentage change in exchange rates from one billing cycle to the next. While on average it seems that users tend to click less on foreign ads when their currency depreciates (i.e. the exchange rate increases), there is a significant amount of dispersion in the data.

Next, we use regression analysis to measure the total effect of exchange rate shocks on clicks in the random data. Note that since online-auction on Microsoft's Bing are run only in Euros or in US Dollars (depending on the user's market), in many cases there is a disconnect between the currency conversion faced by the user and the conversion faced by the advertiser. In particular, when a user ponders clicking on a (foreign) ad, they consider the exchange rate between their own currency and the advertiser country currency. By contrast, when an advertiser bids for an ad space in the user's market, his bid is converted from his billing currency of choice to *the platform's user-market currency* (either US Dollars or Euros). Consequently, with the exception of U.S. and French users, the billing exchange rate faced by the advertiser is generally different from the user currency per advertiser currency exchange rate. For this reason, one cannot capture the total effect of exchange rates on

clicks using the user-advertiser exchange rates alone: the billing exchange rates must also be included whenever the two differ.

In Table 5, we examine the overall impact of mean normalized user-advertiser and billing exchange rates (in units of user currency per advertiser country currency and user market currency per billing currency, respectively) on clicks using three different regression specifications. The dependent variable is the log of user clicks on a given advertiser's ads over a billing cycle period. We control for the user-advertiser currency and the billing exchange rates as well as other advertiser characteristics that are not affected by exchange rates, but might affect the ad's position or the user's tendency to click on it (or both). These include the average of a measure of the ad's relevance to the search phrase; the average of an indicator variable that equals one when there is an exact match between the phrase and the keyword; and the average of an indicator variable that equals one for "premium" AdCenter accounts (a measure of the client's advertising activity and complexity on AdCenter, where premium is the category for the highest-spending advertisers). In addition, we control for ads that belong to the largest specialized search engines (e.g. Shopzilla and Calibex for consumer products; Kayak.com and Cheapflights.com for hotels and flights).[8].

In addition, the specifications include controls for the average change in rivals' billing exchange rates, the share of "domestic" rivals (i.e. advertisers whose billing currency is identical to the auction currency and hence are unaffected by billing exchange rate fluctuations) and the interaction of the two. In all three specifications, the sign of the coefficient on rivals' exchange rate is negative, indicating that when rival's currencies appreciate (average rival exchange rate increases), they advertise more and this leads to lower clicks on own-ads, all else being equal. Similarly, the coefficient on the share of "domestic" rivals is negative as well, indicating that users are more likely to click on advertisers from their own country. Once we control for individual-advertiser effects, the coefficient on the interaction between rival exchange rates and share of domestic rivals is negative, indicating that the higher the

[8]The rationale behind creating this separate category is that for these ads, there is no clear impact of exchange-rates on clicking behavior. When a user navigates to a specialized search-engine website, he is faced with multiple products, offered by multiple sellers, and often more sponsored links on the page. For this reason, we do not interact this category with exchange rates at all.

share of domestic advertisers on the page, the higher the effect of rivals' exchange rates on own-ad clicks (increased competition between foreign advertisers).

in specifications 1 and 2 we force the impacts of the mean-normalized exchange rates to be identical across all user countries; the coefficients on the exchange rates are mostly small and insignificant. In specification 3 we allow the mean-normalized exchange rates to vary by the major user countries in the data. For the United States and France, the user-market currency and the user currency are one and the same, and so the corresponding user-advertiser currency exchange rates capture the *total* effect of exchange rates on user clicks. By contrast, for the United Kingdom, Canada and all other countries combined, both the billing and the user-advertiser exchange rates must be included in order to capture the total effect of exchange-rate fluctuations on user clicks. Note that for the latter group of countries, the user-advertiser exchange rate coefficient captures the combined impact of the price effect *and* the quality-score effect (which is a direct outcome of the price effect), while the billing exchange rate coefficient captures the combined impact of the billing-cycle effect and the bid-change effect. The total effect in this case can be thought of as the sum of the country's coefficients on the billing and the user-advertiser currency exchange rates: it is the total impact on clicks of an identical percentage change in both the billing and the user-advertiser currency exchange rates.[9] Thus, the table provides estimates for the total effect of exchange rates on clicks in the random sample.

In particular, according to table 5, the (total) impact of exchange rate fluctuations on clicks by French users is not significantly different than zero. For U.S. users, the price effect dominates the advertising effect since the exchange rate coefficient is negative and significant at the 10% confidence level. For the other three user countries, the coefficients on the billing exchange rates are positive, significant (at the 10% confidence level for Canada) and generally much larger than the corresponding coefficients on the user-advertiser exchange rates. Thus,

[9]During our sample period, the average percentage change in user-advertiser exchange rates were slightly smaller but very similar to the average percentage change in billing exchange rate for each of the major user countries in the data: for Canada, the former was 0.1% while the latter was 0.4%; for France, −0.2% versus −0.3%; for the United Kingdom, −0.1% versus 0.6%; for the United States, the majority of advertisers are domestic, so the overall average change in exchange rates is very close to zero; for all other countries combined, 0% versus 0.1%

the total effect of exchange rates on clicks seems to be dominated by the (positive) advertiser effect in these cases.

Table 6 focuses on the impact of exchange rates on advertising and through it on user-clicks. All three components of the advertising channel have the same ultimate impact: changing the advertiser's average position on the page. Since the ad position is an ordinal rather than cardinal concept, and the impact of changing positions is much greater at the top of the page than at the bottom, we assign to each position a "position discount," which is our estimate of the ratio of the clicks a given advertiser would expect in a given position, relative to what they would get in the top position. (The literature on search advertising typically maintains the assumption that the clicks an advertiser expects in a given position can be decomposed into the product of an "advertiser effect" and a "position effect." We rely on a similar assumption to estimate the position effects: we have estimated the probability of a click on an ad in a given position after controlling for individual advertiser effects.)

In particular, the dependent variable in Table 6 is the logarithm of an advertiser's expected position discount over a billing cycle period. The variable is constructed by taking each user query where the advertiser appeared, calculating the position discount for the position where the advertiser's ad was displayed, and finally summing over queries in the billing cycle and weighing the sum by the advertiser's mean quality score.

The explanatory variables include the mean-normalized billing currency exchange rate (in units of user market currency per billing currency), as well as the mean-normalized user-advertiser currency exchange rate (in units of user currency per advertiser country currency). The former feeds into the advertising effect through both the billing cycle effect and the bid change effect. The latter affects advertising through its impact on user clicks which in turn feed into the ad's quality score (the quality score effect). As previously noted, there is a disconnect between billing exchange rates and user-advertiser currency exchange rates in most user countries in the data (with the exception of the United States and France). For this reason, we use a specification similar to the one in table 5: the total advertising effect can be measured by the user-advertiser exchange rate only for the United States and France. For all other user countries, the advertising effect is separated into the quality-score effect, which is captured by the user-advertiser exchange rate; and the billing cycle and bid change

effects, which are captured by the billing exchange rate.

Other control variables are identical to the ones used in the total effect specifications above. In the first two specifications, the exchange rates are restricted to have the same impact on expected clicks across all advertisers, while in the last specification the impact of exchange rates on clicks is allowed to vary by user country. In all three specifications, the sign of the coefficient on rivals' billing exchange rate is negative, indicating that when rival's currencies appreciate (average rival billing exchange rate increases), they advertise more and this leads to lower positions of own-ads, all else being equal. Similarly, the coefficient on the share of "domestic" rivals is negative and significant (with the exception of specification 2). Also, once we control for individual-advertiser effects, the coefficient on the interaction between rival exchange rates and share of domestic rivals is negative, indicating that the higher the share of domestic advertisers on the page, the higher the effect of rivals' exchange rates on own-ad position (increased competition among foreign advertisers).

According to our theoretical model, the coefficient on the user-advertiser currency exchange rate, which captures the quality-score effect, should be negative. The coefficient on the billing exchange rate, which captures both the billing cycle and the bid change effects, should be positive for the case of no pricing to market and negative when there is pricing to market, as summarized in tables 1 and 2.

In the first two specifications, the average impacts of exchange rates on advertising prominence are not significant. The last specification separately estimates the quality-score effect (the coefficients on the user-advertiser exchange rate) and the combined billing-cycle and base-bid change effects for Canada, the United Kingdom and all other countries. The quality-score effect is negative and significant for Canada, and quite small for the other two. By contrast, all three coefficients on the billing exchange rates are much larger and highly significant. This is consistent with the no pricing to market scenario. For France and the United States, only the total advertising effect can be estimated. For both the total advertising effect is negative but not significant, and quite small for the United States. a negative coefficient is consistent with the case of pricing to market, where the total advertising effect is unambiguously negative. It is also consistent with no pricing to market, if the quality-score effect is large enough to dominate the billing cycle and bid change effects combined. This

ambiguity can be resolved by examining the link between bid changes and billing exchange rates: according to our model, bids should unambiguously decrease [increase] in response to a foreign currency appreciation when there is [no] pricing to market.

Table 7 examines the link between advertiser *base bids* and fluctuations of the billing exchange rate (in units of the user market currency per billing currency), controlling for ad-relevancy, match type and advertiser-size (Premium account dummy) as well as individual-advertiser effects. The base bid is the original bid placed by the advertiser, in units of the billing currency. Unlike the effective bid, it is unaffected by mechanical billing exchange rate changes (the billing cycle effect) or by *targeting*. Targeting enables advertisers to change their effective bid selectively for particular auctions, based on parameters such as user location, time of day or day of the week. This results in automatic incremental increases of the effective bid whenever a user query meets any of the targeting conditions. By contrast, the *base bid* would only change if the advertiser were to actively modify it (no mechanical effects).

Focusing on the estimated impact of exchange rates on base bids in table 7, the coefficient on the billing exchange rate in specification 1 is positive but not significant. Once we add individual-advertiser effects, the coefficient becomes much larger and significant, indicating that on average advertisers increase their base bids in response to an appreciation in their billing currency (i.e. an increase in the billing exchange-rate). Specification 3 reveals considerable variation in the effects of billing exchange rate changes across the major advertiser countries in the data. The coefficients on billing exchange rates are positive and significant for all user countries but France, indicating that a foreign currency appreciation would lead the advertiser to increase their billing currency base bid in almost all user markets. These findings are consistent with a no pricing to market scenario for all user countries but France. Note that for the latter, the coefficient on the billing exchange rate has large standard errors, possibly indicating variations in advertiser pricing-to-market policies within that market. Overall, the table provides some suggestive evidence of a positive impact of billing rates on advertising quantity.

Table 8 estimates the impact of exchange rates on actual clicks by users, controlling for the position of the advertiser (that is, controlling for the average position effects, the dependent variable from table 6.). For each advertiser in the data, we sum the total number

of user clicks on his ads during a billing cycle, and use this as the dependent variable. By controlling explicitly for position-induced expected clicks, we are in effect holding advertising fixed and examining how clicking behavior is affected by exchange-rate fluctuations, all else being equal (other control variables are identical to the one used in the specifications above). Specifications 1 and 2 in the table estimate the overall effect of exchange rates on clicks, pooling all user and advertiser countries in the data together, with user country, advertiser-country and advertiser fixed effects, respectively. The next specification allows the impact of exchange rates to vary by user country (for major user countries in the data, and all other countries grouped together).

The results indicate the existence of considerable variation in exchange-rate pass through across user countries. In the first two specifications, the coefficient estimates are quite small and not significantly different than zero. When the impact of exchange rates on users is allowed to vary for the major user countries in the data (specification 3), the coefficient is negative for France, the United Kingdom and the United States, but significant only for the latter; For France, the coefficient has a large confidence interval; Thus, we conclude that there seems to be a lot of variation in the price-effects across user countries in the data. We also note that the elasticity of advertising prominence on clicks (as measured by the coefficient on the log of position-induced clicks) is positive and significant for all five user-country categories, and ranges from less than 0.1 for all other countries combined to 0.3 for France and the United States.

Note that there are many questions that are not answered by these aggregate regressions; for example, the estimates may reflect composition effects, since the set of search phrases where an advertiser appears may change from month to month. Our discrete choice analysis of user clicking behavior on specific search phrases below attempts a more structured and granular approach that circumvents these problems.

5.2 Effects of Exchange Rates on Selected Search Phrases

Next, we turn to assess and quantify the impact of exchange-rate fluctuations for a selected sample of search phrases with high revenue and high international activity. By focusing in on individual search phrases, we are able to control for some of the confounding factors that

arise in aggregate regressions, and further we can provide some more specific interpretations of our findings. In addition, we are able to estimate a structural model of user choice that enables us to do counterfactual experiments decomposing the various effects of exchange rates. By selecting search phrases with high proportions of foreign ads and foreign-ad clicks, it is likely that we are focusing on markets where user awareness regarding foreign prices and exchange rates is high. Consequently, we may expect user clicks to be more sensitive to exchange-rate fluctuations relative to the random sample (the random sample includes many phrases with low international activity).

Since the vast majority (over 97%) of users in the selected sample click on no more than one ad per session, we use a discrete-choice model. For each of the major non-U.S. user countries in the data (France, United Kingdom and Canada), we estimate a nested logit model of the consumer's choice among ads, where the nests are: mainline (top of the page) ads; sidebar ads; and no clicks on ads. The choice variable is a click (or no click), and the explanatory variables include: position fixed effects, the advertiser's average quality score (throughout the data period), a measure of ad-relevancy, a dummy for exact-match between the advertiser's keyword and the user's search phrase and a dummy for premuim AdCenter accounts, and major advertiser-country fixed effects interacted with a dummy for "big" advertisers (advertisers whose ads account for more than 5% of the total ads from the advertiser country in the user's country for the given search phrase). The base category in each specification is big accounts from the user's country.

Of interest to us are the interactions of mean-normalized daily exchange-rates (in units of user country per advertiser country currency) with the interaction of country fixed effects and the big advertiser dummy. The coefficients on these variables measure the impact of exchange rates on clicks through price changes. Table 9 displays the estimates of exchange-rate interactions for selected search phrases and country pairs. We have only included estimation results for pairs that have sufficient clicks (in general and on foreign ads in particular), and where there are sufficient observations so as to identify the variables of interest. For pairs without sufficient data, the relevant exchange-rate interactions are omitted from the model.

The table demonstrates that there is significant variation in the impact of exchange-rate

fluctuations across advertiser countries and user pairs. For those advertiser countries where the exchange-rate interaction is significant, it is (almost) always negative (the exceptions are for smaller countries, which are included under the "rest of the world" category, which may be imprecisely classified). To interpret the magnitude of these variables, note that a coefficient of -0.022 on small U.S. advertisers for Canadian users and "tire" searches translates to an odds-ratio of $e^{-0.022} = 0.98$. That is, the probability of clicking on a small U.S. advertiser-ads relative to a big Canadian advertiser-ads increases by 2% when the Canadian-per-U.S. dollar exchange rate falls by 1%.

Using our estimates from the nested logits, we can compare the relative magnitudes of the pricing effect versus the advertising channel (decomposed into the billing exchange rate effect, the quality score effect, and the bid change effect) when the relevant exchange rates change. The nested logit provides us with an estimate of how each of these channels (exchange rates, advertising) affects the probability of a click. Given our rich auctions data, we can isolate and quantify the contribution of each channel to the total effect of exchange rate fluctuations on clickling behavior. The idea is to simulate the impact of a given exchange rate change on each channel while holding the other channels fixed at their pre-change level.

Since the billing exchange rates are changed once every 30 days, we start by calculating the corresponding average actual user-per-advertiser country exchange rate for the same period and predict the *benchmark* number of clicks for our counterfactuals by replacing the actual user-per-advertiser daily exchange rates with the average. Then, for each user country and search phrase specification, we fix a given user-country u, advertiser-country a, and billing currency m triplet in a particular billing-cycle period t. We then examine the impact of a change in this user-advertiser country's average exchange rate on clicks on ads from that advertiser country.

To isolate the *pricing-channel* effect, we hold bids, positions and all other control variables fixed. Then, we change the user-per-advertiser currency exchange rate to its average level in period $t + 1$. We predict the number of clicks with the new exchange rate and calculate the percentage difference relative to benchmark case.

Next, we allow the *quality scores* to change and reflect the impact of the predicted change in clicks (due to the pricing-channel effect). This change affects the score-weighted bids of

ads from the particular advertiser country. We simulate all the auctions with the new quality scores but without changing the bids, calculate the new positions of all the ads and use our nested logit model to predict the number clicks using the new positions.

To isolate the *billing currency* effect, we hold both the user-per-advertiser currency exchange-rate and the quality score at their benchmark levels. We fix the bids (in billing currency units) at their benchmark levels, and allow the billing-per-market currency exchange-rate to change. We then calculate the new effective bids (in units of the market currency) and simulate the results of the auctions, all else fixed at its benchmark levels. These new positions are then fed into the nested-logit model to predict the number of clicks in this scenario.

As previously noted, we may also expect advertisers to actively change their bids in response to billing exchange rate fluctuations. In order to assess how *bid changes* affect advertising prominence and through it user clicks, we hold the user-per-advertiser currency exchange-rate, the quality score and the billing exchange rate fixed at their benchmark levels, and allow the bids of affected advertisers to change from their current billing-cycle levels to the next billing cycle's levels. Since advertiser bids may change more than once during a billing-cycle period, in response to pre-set user characteristics (targeting) or otherwise, we use the average billing-cycle bid for the benchmark, and replace it with the next billing cycle's average bid for the bid change counterfactual. Using the next billing cycle's average bids, we calculate the effective bids using the benchmark billing exchange-rate levels and simulate the results of the auctions, all else fixed at its benchmark levels. The new positions are again fed into the nested-logit model to predict the number of clicks in this case.

Lastly, we estimate the compounded effect of all four together, by changing the user-advertiser exchange rates, the billing exchange rates and the quality scores all at once. Table 10 summarizes which variables are changed and which are held fixed in each of the four experiments detailed above.

In what follows, we run two sets of counterfactual experiments. In the first set, our goal is to compare between the mechanical billing-exchange rate and pricing-channel impacts, abstracting from any bid change effects. This comparison is particularly interesting since in our experience, many advertisers on Microsoft's online-advertising platform rarely change

their bids. To facilitate this comparison, we change the billing exchange rate according to the percentage change we used in the price and quality score counterfactuals. In the second set of counterfactual experiments, we simulate the full effect of all four channels (price, quality score, billing exchange rate and bid change) for a subset of search phrases, user countries and major advertisers whose billing currency is different than the bidding currency. The second set of counterfactuals allows us to compare between the impacts of the full advertising and price effects on user clicks across several different search phrases, user and advertiser countries.

Table 11 summarizes the results of the first set of counterfactual experiments for search phrases containing the word "wig" (English speaking countries only). Table 12 does the same for search phrases that contain "tire" ("pneu" for France). Table 13 does the same for search phrases that contain "Cyprus" ("Chypre" for France). Since we are interested in comparing the mechanical billing exchange-rate effect with the price-channel effect, we have simulated an advertising effect for a couple of cases where in reality there is no exchange-rate effect since the billing currency is identical to the market currency. However, since all our model estimates is the impact of a position change on clicks, and since the billing-exchange rate change is really translated into a change in position, these numbers give the hypothetical magnitude of a corresponding billing exchange-rate change.

From the counterfactual tables it is clear that the quality score effect reinforces the pricing channel effect, but the billing effect channel mitigates the pricing channel effect. However the magnitude of all effects relative to the initial exchange-rate change varies widely across user countries, advertiser countries and search phrases. For example, for U.K. users and small U.S. advertisers on wigs, a 3.4% increase in the GBP-USD rate decreases the clicks through the price-channel by 1.4%. The quality score effect is much lower and reduces clicks only by an additional 0.3% more. The corresponding hypothetical billing exchange rate effect increases advertising and consequently clicks increase by 4.7%. Overall, the billing exchange rate effect dominates in this case. For Canadian users and big U.S. wig advertisers, a 3.3% decrease in the CAD-USD exchange rate increases clicks by 4.5%. The quality-score effect adds 4.63% more clicks, while the decrease in advertising leads clicks to decrease by 4.55%: in the absence of the quality-score effect, the price and advertising effect would have canceled

one another, but the quality-score effect (an additional increase of 4.6% in clicks) eventually leads to an overall increase in clicks of 6.25%. . It should be noted that often the effect of an increase in an ad's rank on the page would be larger than the impact of an identical decrease in the rank (and vice versa), since there may be large differences in the click-through-rates of mainline ads versus sidebar ads.

A similar variation across user-countries and advertisers is evident for the search phrases containing "tire." In particular, French users exhibit the smallest price-effect (less than 100%), while Canadian users increase their clicks on U.S. ads by 5.2% in response to a 3.3% decrease in the CAD-USD exchange rate. Both the quality score effect and the advertising effect are much smaller than the price effect for Canadian and French users, but for U.K. users and big German advertisers, the price and quality score effects are of similar magnitudes (2.3% and 2.6%, respectively). For all three user countries, the price effect dominates the advertising effect.

For French Users and search phrases containing the word "Chypre" (Cyprus), the advertising effect is much bigger than the price effect (4.4% decrease versus 0.4% increase). For U.K. users and big Cyprus advertisers, the advertiser effect clearly dominates the price and quality-score effects combined. In particular, the elasticity of U.S. user-clicks with respect to exchange-rates is very low: for a 1.80% increase in the exchange rate, user-clicks decrease by 0.9% only. With such a low price effect, the additional impact on quality scores is extremely low as well, allowing the advertising impact to dominate and the overall effect of the increase in the exchange rate to be positive rather than negative.

Table 14 presents the results of the second set of counterfactual experiments (including the bid-change effect). As in the previous counterfactual experiments, we focus on a particular advertiser country and billing currency combination for a given user country. Previously, we have compared the impact of an *identical* change to the user-advertiser and billing exchange rates on clicks. However, since our goal in table 14 is to decompose the advertising effect into its actual components (quality score effect, billing cycle effect and bid change effect), we use the actual billing exchange rates for our billing exchange rate and bid change counterfactuals. Since bids may oscillate mechanically over time based on particular user, location or time characteristics (targeting), we calculate the benchmark by replacing the actual bid of affected

advertisers with their corresponding average bids over the billing cycle. The price effect, quality score and the billing exchange rate counterfactuals are then calculated relative to the average-bid benchmark. the bid change effect is calculated by replacing the current billing cycle's average bid with the average bid from the next billing cycle. Lastly, we calculate the effect of all four counterfactuals combined.

We focus on cases where the billing and auction currencies are different, since in these cases it is more likely that bids will change in response to billing exchange-rate changes. In particular, we examine Canadian users and big Canadian advertisers for "wig" (billing currency: CAD, auction currency: USD); United Kingdom users and big German advertisers for "tyre" (billing currency: EUR, auction currency: USD); French users and big U.K. advertisers (GBP-EUR) and U.K. users and big Cyprus advertisers (EUR-USD) for "Cyprus". The results of these four experiments again display large variations across countries and search phrases in the response of advertisers to exchange-rate changes.

Interestingly, it seems that when the price effect is large (that is, when the price elasticity of user-demand is high), the overall advertising effect tends to be low. In particular, the largest price effect is a 4.4% decrease in clicks in response to a 2.4% increase in the EUR-GBP exchange rate for French users and big British "pneu" (tire) advertisers. Note that the "net" billing exchange rate effect (billing cycle and bid change effects, combined) on clicks in this case is less than 2% ($0.00\% + 1.62\% = 1.62\%$). Similarly, for U.K. users and big German "tyre" advertisers, a 0.4% decrease in the exchange rate causes clicks to rise by 1.1%. The net billing exchange rate effect in this case is only about $0.3 + 0.2 = 0.5\%$. By contrast, for U.K. users and big Cyprus advertisers, a 5.9% increase in the exchange rates causes clicks to drop only by 2.8%; Yet the corresponding advertiser-side effect (billing exchange rate decrease of 4.8% and the average bid change) leads to a 5.59% increase in clicks. Similarly, for Canadian users and big Canadian "wig" advertisers, there is no price effect, but the net billing exchange rate effect, in response to a 3.2% decrease in the CAD-USD exchange rate increase, is an overall 6.5% increase in clicks. These findings are reminiscent of the Dorfman and Steiner (1954) rule, which states that, all else equal, when the price elasticity of demand is low, a monopolist would choose a high advertising intensity.

6 Conclusion

In this paper, we have highlighted the fact that exchange rate fluctuations can lead to changes in advertising by foreign firms, and this can lead to a shift in user demand that may work against the more standard impact of changes in domestic prices. We believe that this effect may be important for many product markets. However, in general both advertising expenditures and the impact of those advertising expenditures are notoriously difficult to measure.

By studying the case of internet search advertising, we are able to directly measure the impact of exchange rates on advertising, as well as the impact of changes in advertising on changes in visits by domestic consumers to foreign web sites. We have established that the impacts of exchange rates on advertising and user visits conditional on advertising are extremely heterogeneous, but the effects can be large, important, and operate in opposite directions.

There are several important next steps for this research. First, we need to further explore advertiser bid changes, or lack thereof, in response to billing exchange rate shocks. To do so, we plan to utilize measures of post-click user activity in order to assess the impact of exchange rates on the value of a consumer website visit. Then, we will estimate the returns to the advertiser from changing their bid. We can back out the implied transaction costs for advertisers, and assess whether those costs are plausible. In addition, we can look at the inefficiencies created by "sticky bids."

Another interesting avenue to explore would be to gather end-user price data for selected search phrases. This would allow us to identify advertisers who price to market, versus those who maintain a world price; and in addition, we can relate consumer responses to exchange rates, to advertiser pricing policies and changes in final good prices.

More broadly, internet commerce is growing in importance, and international trade in e-commerce is growing even more rapidly. Online advertising plays a central role in determining which web sites consumers visit, and some of the largest international e-commerce web sites (such as eBay) have the features that prominence of their listings matter, and firms can pay for more prominent placement. Thus, we believe that there is much fertile ground for future

research in studying the interaction between advertising and trade flows in e-commerce.

References

Athey, S., and D. Nekipelov (2010): "A Structural Model of Sponsored Search Advertising Auctions," .

Auer, R., and T. Chaney (2009): "Exchange Rate Pass-Through in a Competitive Model of Pricing-to-Market," *Journal of Money, Credit and Banking*, 41(1), 151–175.

Berman, N., P. Martin, and T. Mayer (2010): "How do different exporters react to exchange rate changes? Theory, empirics and aggregate implications1," .

Bhattacharya, P. S., C. Karayalcin, and D. Thomakos (2008): "Exchange rate pass-through and relative prices: An industry-level empirical investigation," *Journal of International Money and Finance*, 27(7), 1135–1160.

Campa, J. M., and L. S. Goldberg (2005): "Exchange-Rate Pass Through Into Import Prices," *Review of Economics and Statistics*, 87(4), 679–690.

Dorfman, R., and P. O. Steiner (1954): "Optimal Advertising and Optimal Quality," *American Economic Review*, 44, 826–836.

Edelman, B., M. Ostrovsky, and M. Schwartz (2007): "Internet Advertising and the Generalized Second Price Auction," *AER*, 97(1), 242–259.

Engel, C. (2002): "Expenditure Switching and Exchange-Rate Policy," *NBER Macroeconomic Annual*, 17, 231–300.

Fitzgerald, D., and S. Haller (2010): "Pricing to Market: Evidence from Producer Prices," .

Goldberg, P. K., and M. M. Knetter (1997): "Goods Prices and Exchange Rates: What Have We Learned," *JEL*, 35(3), 1243–1272.

Hellerstein, R., and P. K. Goldberg (2010): "A Structural Approach to Identifying the Sources of Local-Currency Price Stability," .

Varian, H. (2007): "Position Auctions," *International Journal of Industrial Organization*, 25(7), 1163–1178.

Table 1: Pricing To Market: Impact of Foreign Currency Appreciation[a]

Name	Impact	Demand (clicks on ad)
Price Effect	dom.price+	−
Advertising Effect		
Quality Score	score+	−
	ad rank−	
Billing cycle	eff.bid+	+
	ad rank+	
Bid effect	bid+	+
Total Effect	price+	fall by less
		relative to no-advertising case

[a]The table summarizes the predicted effect of a foreign currency appreciation on user clicks, through the price and advertising channel, when there is pricing-to-market. The advertising channel is broken into its three components: the quality score effect (feedback created by user clicks); the billing-cycle effect (the mechanical change to the billing exchange-rate) and the bid effect (changes to the billing-currency bids by advertisers).

Table 2: No Pricing To Market: Impact of Foreign Currency Appreciation[a]

Name	Impact	Demand (clicks on ad)
Price Effect	dom.price+	−
Advertising Effect		
Quality Score	score−	−
	ad rank−	
Billing cycle	eff.bid+	+
	ad rank+	
Bid effect	bid−	−
Total Effect	price+	−

[a]The table summarizes the predicted effect of a foreign currency appreciation on user clicks, through the price and advertising channel, when there is no pricing-to-market. The advertising channel is broken into its three components: the quality score effect (feedback created by user clicks); the billing-cycle effect (the mechanical change to the billing exchange-rate) and the bid effect (changes to the billing-currency bids by advertisers).

Table 3: Characteristics of online-search markets[a]

User-Country Name	% of AdCenter user queries	% of ad-exposures from advertisers foreign to the users in these markets	% of clicks on ads from advertisers foreign to the users in these markets
United States	62.3%	6.7%	5.5%
Major non-U.S. markets (United Kingdom, Singapore, France, Canada, Australia)	31.0%	39.2%	29.0%
Other non-U.S. markets	6.8%	99.7%	99.7%

[a]Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010.

Table 4: Mean-normalized exchange rate statistics for the major currencies in the data[a]

	User Country Name			
Advertiser Country	Canada	France	United Kingdom	United States
Australia	0.13	0.11	0.10	0.08
	0.84	0.84	0.85	0.91
	1.42	1.31	1.33	1.22
Norway	0.04	0.0	0.05	0.05
	0.88	0.92	0.89	0.89
	1.07	1.09	1.08	1.10
Canada	0.00	0.07	0.07	0.05
	1.00	0.90	0.89	0.85
	1.00	1.20	1.14	1.08
France		0.00	0.03	0.05
		1.00	0.94	0.86
		1.00	1.06	1.09
United Kingdom			0.00	0.05
			1.00	0.91
			1.00	1.08
United States				0.00
				1.00
				1.00

[a]Each cell contains: standard deviation, minimum and maximum value over the period January 2009-June 2010.

Table 5: The impact of exchange-rate changes on clicks: total effect[a]

	Specification		
log (Actual clicks per advertiser, billing cycle)	I	II	III
log normalized user-adv ex.rate	0.009	0.011*	
	(0.005)	(0.005)	
Canada X log normalized ex.rate			-0.061
			(0.044)
France X log normalized ex.rate			-0.067
			(0.071)
ROW X log normalized ex.rate			0.014*
			(0.006)
United Kingdom X log normalized ex.rate			-0.035
			(0.041)
United States X log normalized ex.rate			-0.05†
			(0.026)
log normalized billing ex.rate	0.053	0.019	
	0.034	(0.039)	
Canada X log billing ex.rate			0.122†
			(0.064)
ROW X log billing ex.rate			0.358***
			(0.053)
United Kingdom X log billing ex.rate			0.303***
			(0.053)
prevalence of search term on adv. web site	0.003	0.019***	0.019***
	0.002	(0.003)	(0.003)
log mean quality score	0.067***	–	–
	0.001	–	–
Dummy {=1 for exact match}	0.035***	0.031***	0.031***
	0.001	0.002	0.002
Dummy {=1 for premium account}	0.134***	0.072***	0.072***
	0.001	0.005	0.005
Same user and advertiser country	0.04***	0.151***	0.15***
	0.002	0.002	0.002
Same billing and auction currency	0.033***	0.024***	0.026***
	0.002	0.003	0.003
log average rival normalized ex.rate	-0.197***	-0.005	-0.233***
	0.049	0.054	0.055
Share of "domestic rivals"	-0.033	-0.012***	-0.012***
	0.002	0.002	0.002
log avg. rival norm. rate X share "domestic rivals"	0.263***	-0.426*	-0.492**
	0.155	0.17	0.17
Intercept	0.257	-0.018**	0.015**
	0.004	0.006	0.004
Fixed effects	Adv country	Advertiser	Advertiser
	User ctry	Adv country	Adv country
		User ctry	User ctry
No. Obs	519697	519678	519678
F	2137.01	2135.46	1625.91

[a] † $p < 0.1$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$;
Base category: United States users and advertisers.
Dependent variable: Total number of clicks in a billing-cycle per advertiser
Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010.
normalized exchange rate in units of user currency per advertiser-country currency
The average effects of actual and billing exchange rates on actual clicks at the monthly level is positive but mostly not significant at the 5% confidence level in specifications 1 and 2. Once we break the impact of exchange rates by user countries, the coefficient on the billing exchange rate is positive and (mostly) highly significant for those countries where the billing exchange rate effect can be separately identified (Canada, the United Kingdom and all other countries combined), while the user per advertiser country exchange rate effect is mostly negative but not significant, with the exception of all other countries combined. For the two countries where the billing and user-advertiser country exchange rates cannot be separately identified (France and the United States), the total effect of the exchange rate on actual clicks is negative (indicating that the user effect dominates the advertising effect), although not significant at the 5% confidence level.

Table 6: Impact of ex-rate changes on position-induced clicks: advertising-channel[a]

	Specification		
log(expected total billing-cycle clicks by position per advertiser)	I	II	III
log normalized user-adv ex.rate	0.005	0.006	
	(0.004)	(0.004)	
Canada X log normalized ex.rate			-0.13***
			(0.032)
France X log normalized ex.rate			-0.069
			(0.051)
ROW X log normalized ex.rate			0.008*
			(0.004)
United Kingdom X log normalized ex.rate			0.032
			(0.030)
United States X log normalized ex.rate			-0.019
			(0.019)
log normalized billing ex.rate	0.031	0.03	
	(0.026)	(0.028)	
Canada X log billing ex.rate			0.346***
			(0.046)
ROW X log billing ex.rate			0.241***
			(0.038)
United Kingdom X log billing ex.rate			0.256***
			(0.038)
prevalence of search term on adv. web site	-0.007***	0.002	0.002
	(0.002)	(0.002)	(0.002)
log mean quality score	0.083***	–	0.008***
	(0.000)	–	(0.001)
Dummy {=1 for exact match}	0.017***	0.074***	0.074***
	(0.001)	(0.003)	(0.003)
Dummy {=1 for premium account}	0.135***	0.142***	0.14***
	(0.001)	(0.002)	(0.002)
Same user and advertiser country	0.04***	0.021***	0.022***
	(0.001)	(0.002)	(0.002)
Same billing and auction currency	0.025***	-0.035***	0.018***
	(0.002)	(0.039)	(0.006)
log average rival normalized ex.rate	-0.158***	-0.007	-0.235***
	(0.037)	(0.001)	(0.039)
Share of "domestic rivals"	-0.038***	0.044***	-0.008***
	(0.001)	(0.122)	(0.001)
log avg. rival norm. rate X share "domestic rivals"	0.502***	-0.056	-0.032
	(0.117)	(0.002)	(0.122)
Intercept	0.339***	0.037***	0.038***
	(0.003)	(0.003)	(0.003)
Fixed effects	Adv country	Adv country	Adv country
	User ctry	User ctry	User ctry
		Advertiser	Advertiser
No. Obs	519696	519677	519677
F	4578.82	4219.59	3212.11

[a] * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$;

Base category: United States users and advertisers.

Dependent variable: log estimated expected clicks by position (no advertiser effect), weighted by advertiser average quality score

Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010. normalized billing exchange rate in units of user market currency per billing currency

0ggregate impacts of the exchange-rate on advertising prominence are not significant in specifications 1 and 2. Once we break the impact of exchange rates on advertising prominence by user countries, the coefficient on the billing exchange rate is positive and highly significant for those countries where the combined billing-cycle and bid-change effects can be separately identified (Canada, the United Kingdom and all other countries), while the quality-score effect (captured by the user-advertiser exchange rate) is mostly negative but not significant (except for Canada), with the exception of all other countries combined. For the two countries where the three effects that comprise the advertising channel cannot be separately identified (France and the United States), the total effect of the exchange rate on actual clicks is not significant.

Table 7: The impact of billing exchange-rates on advertiser base bids[a]

	Specification		
log (Average billing-cycle base bid in units of the billing currency)	I	II	III
log norm. avg. billing cycle ex.rate	0.023	0.329***	
	(0.104)	0.069	
Canada X log normalized ex.rate			0.549***
			(0.121)
France X log normalized ex.rate			-0.087
			(0.209)
ROW X log normalized ex.rate			0.221*
			(0.092)
United Kingdom X log normalized ex.rate			0.442***
			(0.097)
United States X log normalized ex.rate			0.326**
			(0.124)
log average rival normalized billing ex.rate	0.433**	-0.491***	-0.542***
	(0.149)	(0.096)	(0.100)
Share of "domestic rivals"	-0.01	0.063***	0.063***
	(0.005)	(0.003)	(0.003)
log avg. rival norm. rate X share "domestic rivals"	-4.427***	-0.031	0.012
	(0.469)	(0.300)	(0.300)
prevalence of search term on adv. web site	0.519***	0.269***	0.269***
	(0.007)	(0.005)	(0.005)
Dummy {=1 for exact match}	0.563***	0.403***	***
	(0.003)	(0.008)	(0.008)
Dummy {=1 for premium account}	0.042***	-0.019***	-0.019***
	(0.004)	(0.003)	(0.003)
log mean quality score	-0.7***	–	–
	(0.002)	–	–
Same billing and auction currency	0.228***	0.37***	0.37***
	(0.006)	(0.006)	(0.006)
Intercept	1.561***	3.716***	3.717***
	(0.010)	(0.007)	(0.007)
Fixed effects		Advertiser	Advertiser
	Adv.ctry	Adv. Ctry	Adv. Ctry
	User ctry	User ctry	User ctry
No. Obs	472197	472185	472185
F	15114.02	1644.07	1331.53

[a] † $p < 0.1$ * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$;

Base category: United States advertisers and users.

Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010. normalized exchange rate in units of market-currency per billing-currency

The table indicates that when the billing exchange rate increases (the foreign currency appreciates), the base bid (given in units of the billing currency) generally increases, and this increase is significant once individual-advertiser effects are controlled for (with the exception of France in specifications 3).There is also some indication that an appreciation of rivals' exchange rates, which leads them to increase their advertising, decreases own advertising.

Table 8: Impact of ex-rate changes on user clicks conditional on position: price-channel[a]

	Specification		
log (Actual clicks per advertiser, billing cycle)	I	II	III
log normalized ex.rate	0.004	0.007	
	(0.004)	(0.004)	
Canada X log normalized ex.rate			0.057
			(0.036)
France X log normalized ex.rate			-0.008
			(0.058)
ROW X log normalized ex.rate			0.008
			(0.005)
United Kingdom X log normalized ex.rate			-0.041
			(0.034)
United States X log normalized ex.rate			-0.056***
			(0.021)
prevalence of search term on adv. web site	0.004*	0.018***	0.019***
	(0.001)	(0.002)	(0.003)
log mean quality score	0.041***	–	–
	(0.000)	–	–
Dummy {=1 for exact match}	0.026***	0.019***	0.018***
	(0.001)	(0.001)	(0.001)
Dummy {=1 for premium account}	0.007***	-0.074**	0.022***
	(0.001)	(0.023)	(0.004)
Same user and account country	0.002	-0.004**	-0.039***
	(0.001)	(0.001)	(0.002)
Position-induced expected clicks	0.204***	0.222***	
	(0.000)	(0.000)	
Canada X log expected clicks			0.2***
			(0.002)
France X log expected clicks			0.303***
			(0.002)
ROW X log expected clicks			0.09***
			(0.002)
United Kingdom X log expected clicks			0.273***
			(0.001)
United States X log normalized ex.rate			0.301***
			(0.001)
Intercept	0.05***	-0.056***	-0.127***
	(0.002)	(0.009)	(0.003)
Fixed effects	Adv. Ctry User ctry	Advertiser Adv. Ctry User ctry	Advertiser Adv. Ctry User ctry
No. Obs	827632	827632	521674
F	23909.31	18991.82	11215.55

[a†] $p < 0.1$ * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$;
Base category: United States users and advertisers.
Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010.
Normalized exchange rate in units of user-currency per advertiser-country currency
The table indicates that the price effect of exchange rate fluctuations on user-clicks is mostly not significant.
The impact of exchange rates on user clicks is negative and significant for U.S. users in specification 3.

Table 9: Summary of coefficients of interest for selected nested-logit specifications[a]

	Search phrase	Cyprus		Wigs		Tires		
Advertiser Country	User Country	United Kingdom	France	United Kingdom	Canada	United Kingdom	France	Canada
	Acct size							
UK	Small	–	-0.001	–		–		
		–	0.044	–		–		
	Big	–	-0.02*	–		–		
		–	0.01	–		–		
US	Small			-0.009	–			-0.022***
				0.01	–			0.005
	Big			0.013	-0.014			-0.014***
				0.011	0.01			0.002
Cyprus	Small	0.008						
		0.008						
	Big	-0.006						
		0.007						
Germany	Small					–	–	
						–	–	
	Big					-0.025*	–	
						0.011	–	
Switzer	Small						–	
-land							–	
	Big						-0.013***	
							0.003	
ROW		0.037*	–	-0.012	-0.054	-0.119***	0.007**	-0.004
		0.015	–	0.135	0.048	0.023	0.003	0.003
No. Obs		1600348	109704	348552	105959	2133335	1803000	1383652

[a]* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$;

All specific user-country search-phrase specifications include the following controls, in addition to the exchange rate interactions shown here: advertiser country interacted with a small advertiser dummy and a big advertiser dummy (big advertisers are defined as having at least 5% of all ad exposures of ads from their advertiser country in the user-country, for the particular search phrase); a measure of the prevalence of the search term on the advertiser's web site; log of the advertiser's mean quality score; a dummy that equals one for exact match; and a dummy that equals one for premium AdCenter accounts.

The base category is big advertisers from the user's country.

Table 10: Summary of counterfactual experiments

Counterfactual name	Change in user ex.rate	Change in billing ex.rate	Change in quality score	Effective Bid	Ad Rank	Predicted effects from nested logit
Price effect	Yes	No	Fixed	Fixed	Fixed	Price-effect
Quality score	Yes	No	Predicted effect from price effect	Fixed	Simulate auctions with new quality score	Quality-score effect
Billing Ex. Rate	No	Yes	Fixed	Adjusted by billing ex. rate	Simulate auctions with new effective bid	Advertising-effect
Bid effect	No	No	Fixed	New bid	Simulate auctions with new bids	Advertising-effect
Aggregate	Yes	Yes	Predicted effect from price effect	New bid, adjusted by billing ex. rate	Simulate auctions with new quality score, eff. bid	Price, quality score and advertising effects

Table 11: First set of counterfactual experiments: "wig" [a]

circa September-October 2009: U.K. users, small U.S. advertisers				
Counterfactual name	Change in ex.rate (GBP/USD)	Change in billing ex.rate (USD/USD)	percent change in clicks	Change in rank (no. positions)
Price effect	3.38%	–	-1.4%	–
Quality score	3.38%	–	-0.3%	0.03
Billing ex. rate	–	-3.27%	4.7%	-0.15
All together	3.38%	-3.27%	0.7%	-0.08
circa July-August 2009: Canadian users, big U.S. advertisers				
Counterfactual name	Change in ex.rate (CAD/USD)	Change in billing ex.rate (USD/USD)	percent change in clicks	Change in rank (no. positions)
Price effect	-3.3%	–	4.50%	–
Quality score	-3.3%	–	4.63%	-0.07
Billing ex. rate	–	3.4%	-4.55%	0.10
All together	-3.3%	3.4%	6.25%	-0.01

[a]All 3-word search-phrases containing "wig". The counterfactual experiments are based on the nested-logit coefficient estimates.

Figure 1: Changes in clicks over changes in user-advertiser normalized exchange rates, by user-advertiser country pairs[b]

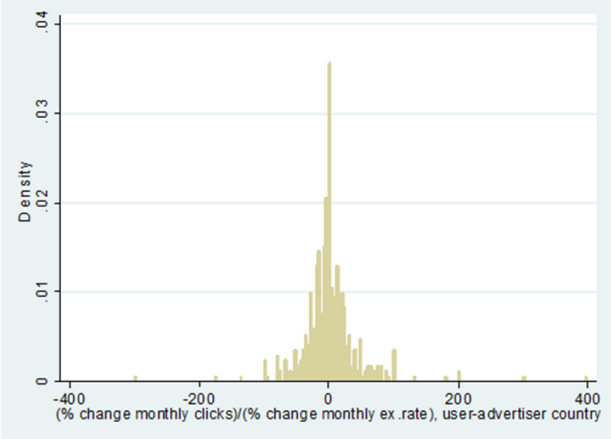

[a]The ratio is calculated as the percentage difference in total clicks, per user-advertiser country pair, over the percentage change in average exchange rates from one cycle to the next.
Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010.
normalized billing exchange rate in units of billing currency per market currency
The histogram shows that on average, foreign currency appreciation results in fewer user clicks, but there is substantial dispersion.

[b]The ratio is calculated as the percentage difference in total clicks, per user-advertiser country pair, over the percentage change in average exchange rates from one cycle to the next.
Source: random sample of all search activity on Microsoft's websites, 3 days a week January 2009-June 2010.
normalized billing exchange rate in units of billing currency per market currency
The histogram shows that on average, foreign currency appreciation results in fewer user clicks, but there is substantial dispersion.

Table 12: First set of counterfactual experiments: "tire" [a]

circa July-August 2009: Canadian users, big U.S. advertisers				
Counterfactual name	Change in ex.rate (CAD/USD)	Change in billing ex.rate (USD/USD)	percent change in clicks	Change in rank (no. positions)
Price effect	-3.3%	–	5.17%	–
Quality score	-3.3%	–	0.66%	-0.04
Billing ex. Rate	–	3.4%	-0.15%	0.02
All together	-3.3%	3.4%	5.51%	-0.02
circa September-October 2009: U.K. users, big German advertisers				
Counterfactual name	Change in ex.rate (GBP/EUR)	Change in billing ex.rate (EUR/USD)	percent change in clicks	Change in rank (no. positions)
User effect	0.8%	–	2.27%	–
Quality score	0.8%	–	2.58%	-0.05
Billing ex. Rate	–	-0.8%	-0.02%	0.00
All together	0.8%	-0.8%	4.43%	-0.04
circa January-February 2010: French users, big Swiss advertisers				
Counterfactual name	Change in ex.rate (EUR/CHF)	Change in billing ex.rate (EUR/EUR)	percent change in clicks	Change in rank (no. positions)
Price effect	8.1%	–	-6.88%	–
Quality score	8.1%	–	-0.91%	0.06
Billing ex. Rate	–	-7.5%	1.32%	-0.07
All together	8.1%	-7.5%	-6.86%	-0.01

[a]All 3-word search-phrases containing "pneu" (in France) or "tire". The counterfactual experiments are based on the nested-logit coefficient estimates.

Figure 2: Daily Euro-per-USD exchange rates,January 2009-June 2010[a]

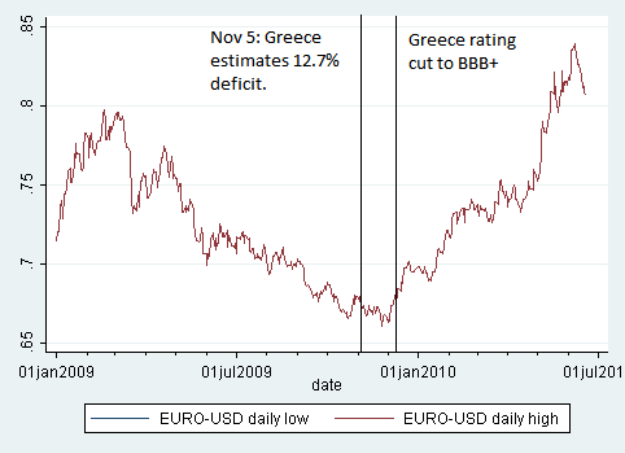

[a]Source: Global Financial Data. The vertical lines indicate the onset of the recent Greece-crisis which has had a dramatic effect on the Euro.

Table 13: First set of counterfactual experiments: "Cyprus" [a]

circa July-August 2009: U.K. users, big Cyprus advertisers				
Counterfactual name	Change in ex.rate (GBP/EUR)	Change in billing ex.rate (EUR/USD)	percent change in clicks	Change in rank (no. positions)
Price effect	1.8%	–	-0.88%	–
Quality score	1.8%	–	-0.09%	0.02
Billing ex. Rate	–	-1.8%	1.61%	-0.06
All together	1.8%	-1.8%	-0.01%	-0.04
circa November-December 2009: French users, big U.K. advertisers				
Counterfactual name	Change in ex.rate (EUR/GBP)	Change in billing ex.rate (GBP/EUR)	percent change in clicks	Change in rank (no. positions)
Price effect	2.4%	–	-4.40%	–
Quality score	-5.6%	–	-0.04%	0.04
Billing ex. Rate	–	-2.3%	0.04%	-0.01
All together	2.4%	-2.3%	-4.44%	0.02

[a]All 3-word search-phrases containing "Chypre" (in France) or "Cyprus". The counterfactual experiments are based on the nested-logit coefficient estimates.

Figure 3: Daily CAD-per-USD exchange rates,January 2009-June 2010[a]

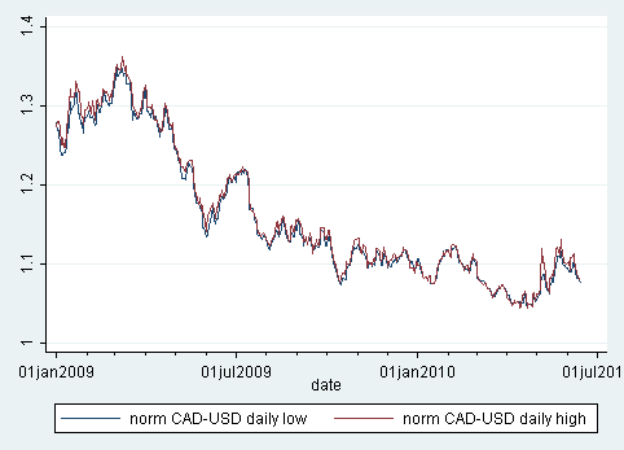

[a]Source: Global Financial Data.

Table 14: Second set of counterfactual experiments[a]

Counterfactual name	Change in ex.rate (CAD/CAD)	Change in billing ex.rate (CAD/USD)	percent change in clicks	Change in rank (no. positions)
circa July-August 2009: "wig", Canadian users, big Canadian advertisers				
Price effect	0.00%	–	–	–
Quality score	0.00%	–	–	–
Billing ex. rate	–	-3.2%	4.13%	-0.07
Bid effect	–	-3.2%	5.77%	-0.10
All together	0.00%	-3.2	6.45%	-0.11
circa May-June 2009: "Cyprus", U.K. users, big Cyprus advertisers				
Counterfactual name	Change in ex.rate (GBP/EUR)	Change in billing ex.rate (EUR/USD)	percent change in clicks	Change in rank (no. positions)
Price effect	5.9%	–	-2.81%	–
Quality score	5.9%	–	-0.38%	0.04
Billing ex. rate	–	-4.8%	4.76%	-0.21
Bid effect	–	-4.8%	1.22%	-0.06
All together	5.9%	-4.8%	0.77%	-0.14
circa November-December 2009: "Cyprus", French users, big U.K. advertisers				
Counterfactual name	Change in ex.rate (EUR/GBP)	Change in billing ex.rate (GBP/EUR)	percent change in clicks	Change in rank (no. positions)
Price effect	2.4%	–	-4.39%	–
Quality score	2.4%	–	0.00%	0.03
Billing ex. rate	–	-1.7%	0.00%	-0.00
Bid effect	–	-1.7%	1.62%	-0.02
All together	2.4%	-1.7%	-2.84%	-0.01
circa January-February 2010: "tyre", U.K. users, big German advertisers				
Counterfactual name	Change in ex.rate (GBP/EUR)	Change in billing ex.rate (EUR/USD)	percent change in clicks	Change in rank (no. positions)
User effect	-0.4%	–	1.08%	–
Quality score	-0.4%	–	0.04%	-0.00
Billing ex. Rate	–	5.2%	-0.28%	0.01
Bid effect	–	5.2%	-0.20%	-0.01
All together	-0.4%	5.2%	0.31%	0.04

[a]The counterfactual experiments are based on the nested-logit coefficient estimates, using actual exchange rates for the price effect and billing exchange rate for the billing ex.rate effect. Actual bids are replaced by average bid per billing cycle to calculate the benchmark. To calculate the bid effect, the average bid per billing cycle is replaced with the next billing cycle's average bid.

Figure 4: Daily GBP-per-USD exchange rates,January 2009-June 2010[a]

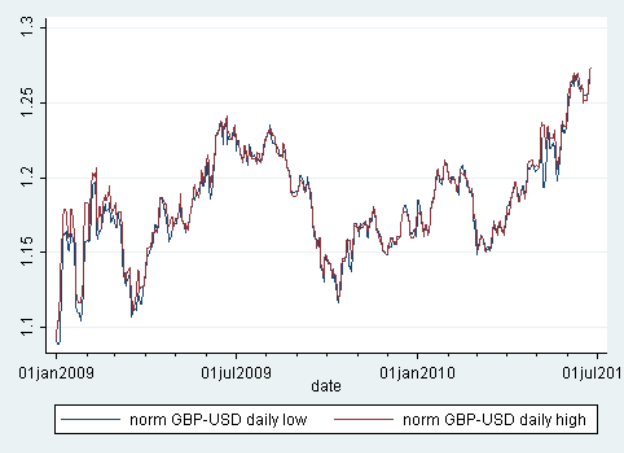

[a]Source: Global Financial Data.