An Interactive Marketing Primer

Adam Marsh
EconoMeta.com

1 Introduction

Advertising is an important part of making the Web work well: it helps developers and writers get paid, gives users more choice in how they support sites, and reinforces the shift in power to participatory users and small creators. It also certainly has its problems, but the first step towards solving these problems is to understand how the whole thing works.

Overview

If you have a web site that you'd either like to put advertising on or buy advertising for, or if you just have an interest in how this process works, you're faced with an entire industry complete with specialized jargon and complex relationships. This primer is a set of notes collected by a relative newcomer to the industry trying to make sense of it all.

Assumptions

This primer is written for people who are familiar with the Web and its basic technology components, but unfamiliar with advertising and how it works on the Web. Concepts such as link, HTML, JavaScript, redirect, cookie, and search engine will be assumed to be familiar.

Contents

1 Introduction	**1**
Overview	1
Assumptions	1
Contents	2
Terminology	3
Scope	3
2 Advertisers And Publishers	**4**
Ad Inventory	4
Types Of Ads	4
Ad Targeting	5
Compensation	6
Ad Placement Overview	6
3 Outsourcing Ad Buying And Selling	**7**
The Advertiser Viewpoint	7
The Publisher Viewpoint	8
Ad Flow Overview	8
4 Search Engines And Intermediaries	**9**
Search	9
Search Engine Marketing	9
Intermediaries	9
Intermediary Overview	10
5 Google, Auctions, And Arbitrage	**11**
Google	11
Keyword Auctions	12
Keyword Arbitrage	12
Tracking And Reporting	13
Google Advertising Overview	14

Terminology

To keep things consistent, we'll use the following terminology:

- **User** (AKA consumer): a person surfing the web
- **Browser**: the user's web browser
- **Publisher**: a web site that wants to sell advertising space
- **Advertiser** (AKA marketer): a web site that wants to buy ads
- **Ad network**: a company who buys advertising space from publishers and sells it to advertisers

A "publisher" refers to any web site selling ad space, and could be an application, service, or site that might not actually "publish content" in the traditional sense. An advertiser is assumed to have a web site, since most Web ads link to a site where the user can find more information, purchase a product, etc. Finally, note that a single web site might be both a publisher and an advertiser.

Scope

Many terms are used to refer to the general activity of marketing and selling on the Internet, including:

- Interactive marketing
- Online marketing
- Internet advertising

There's a lot more to the Internet than just the Web, and there's a lot more to marketing than just advertising; but here we'll focus on the placement of ads on web pages (which should probably be called "Web advertising," but for some reason nobody seems to use this term). Interactive marketing in general has a more general scope, some aspects of which include:

- **Email marketing**: ads are emailed to the user
- **Adware**: ads are shown by a software application downloaded by the user
- **Permission marketing**: the user volunteers to receive ads, by email or otherwise

Email marketing sometimes takes place without obtaining user permission and/or providing the ability to stop receiving the ads, in which case it's called **spam**. Adware sometimes collects data or displays ads without user consent, knowledge and/or the ability to uninstall, in which case it's called **spyware**.

2 Advertisers and publishers

The basics of how advertising is defined and paid for.

Ad inventory

A publisher has a certain amount of space on pages within the site that are set aside for ads; this is called the **ad space**. Each page with ad space is served to a user some number of times during a given time period; this is called the **page views**.

The publisher's **ad inventory** is then the ad space multiplied by the expected page views in a given time period. So a publisher's monthly ad inventory would be the set of ads expected to be viewed by users during that month.

This seems like a pretty simple situation, but it can be complicated by various technicalities. One complication is in counting page views. For example, by page views we really mean views by a user (not for example by a bot or a search engine crawler). Some concepts that are used in calculating an accurate ad inventory include:

- **Hit**: a request for a file on the site
- **Page view**: the serving of a web page, which usually includes many files, and therefore many hits
- **Session**: a unique destination of page views within a certain time period; usually a single IP address that has received some number of page views within 24 hours

Besides distinguishing robots from humans, counting users and ad inventory can be made more accurate by knowing if hits are actually delivered to the browser and viewed, avoiding counting spurious hits such as page refreshes, and mapping users to sessions in light of dynamic IP addresses, etc.

Types of ads

Ad inventory is characterized by various general attributes, including:

- **Premium**: the ad will appear on a page which is considered to be valuable beyond the number of users viewing it, e.g. the home page of a portal
- **ROS** (run of site): the ad will appear on any page within the site; for an ad network this is called **RON** (run of network)
- **ROC** (run of category): the ad will appear on any page within a specified category, e.g. business or music
- **Targeted**: the ad will appear in page views that are targeted by either content or user, e.g. an ad for golf clubs is targeted to pages that have several instances of the word "golf" or users who have expressed an interest in golf

In addition, ads come in various formats, many of which have been standardized. These include:

- **Text**: a short textual description and a link
- **Banner**: an image placed horizontally on the page
- **Vertical banner**: an image placed vertically on the page
- **Skyscraper**: a vertical banner that is taller than normal
- **Interstitial**: an ad that loads between two content pages
- **Pop-up**: an ad that creates a new browser window
- **Pop-under**: a pop-up that is created behind the active window
- **Rich media**: an ad using animation or audio/video, such as Flash

Ad targeting

Ads can be targeted in two main ways: by the content near which the ad appears (**content targeting**), and by the user who is viewing it (**user targeting**). These two techniques are part of what is sometimes called **relevance marketing** (as opposed to mass marketing).

Content targeting has several variants, including:

- **Site targeting**: targets ads to "vertical" sites that focus on a specific topic or user segment related to the ad
- **Category targeting**: targets ads to a category within a site (see ROC above)
- **Contextual targeting**: the text on a given page is analyzed and an algorithm decides which ads best match this text

Unlike content targeting, user targeting depends upon knowing something about the user's interests. Various methods exist for obtaining user data, thereby gaining the ability to offer user targeted ad space to advertisers:

- **Registration**: based on data provided by the user, usually in return for premium or personalized content
- **Technical targeting**: based on data that can be obtained from the user's HTTP page request, such as domain, ISP, connection speed, operating system, and browser type
- **Demographic targeting**: based on the user's demographics such as zip code, age, gender, etc., usually obtained via registration or surveys
- **Geo targeting**: based on the user's location, usually obtained via registration or from the IP address
- **Behavioral targeting** (BT): based on the user's behavior, e.g. which kinds of pages have been viewed in the recent past
- **Search keyword targeting**: based on keywords the user provides in a specific search, either on a publisher site or on a search engine

In order to allow advertisers to match ads to user data, users with similar interests are often grouped into **user segments**.

Compensation

The advertiser can pay the publisher according to several commonly used measures:

- **CPM** (cost per thousand): the advertiser pays a fixed amount per thousand impressions, an **impression** being a user seeing the ad
- **CPC** (cost per click): the advertiser pays a fixed amount per click, a click being a user clicking on the ad to visit the advertiser's site
- **CPA** (cost per action): the advertiser pays a fixed amount per user action; the action can be an inquiry or lead (**CPI**, **CPL**), a sale (**CPS**), or any other kind of specified transaction (**CPT**)

The latter non-CPM measures are called **performance-based**, since the advertiser only pays based upon the ad performing, that is, generating an interaction with the user. Publishers often convert performance-based measures into an **effective CPM** (**eCPM**) by taking into account the percentage likelihood an impression will yield a click or an action; for clicks, this is called the **click through rate** (**CTR**). For example, for a CPC ad, we have:

eCPM = CPC * CTR * 1000.

For a CPA ad, the advertiser must measure the **conversion rate**, or the percentage likelihood a click will yield an action. If this is denoted CR, we have:

eCPM = CPA * CR * CTR * 1000.

Ad placement overview

In summary, the following figure depicts how publishers sell ad space to advertisers.

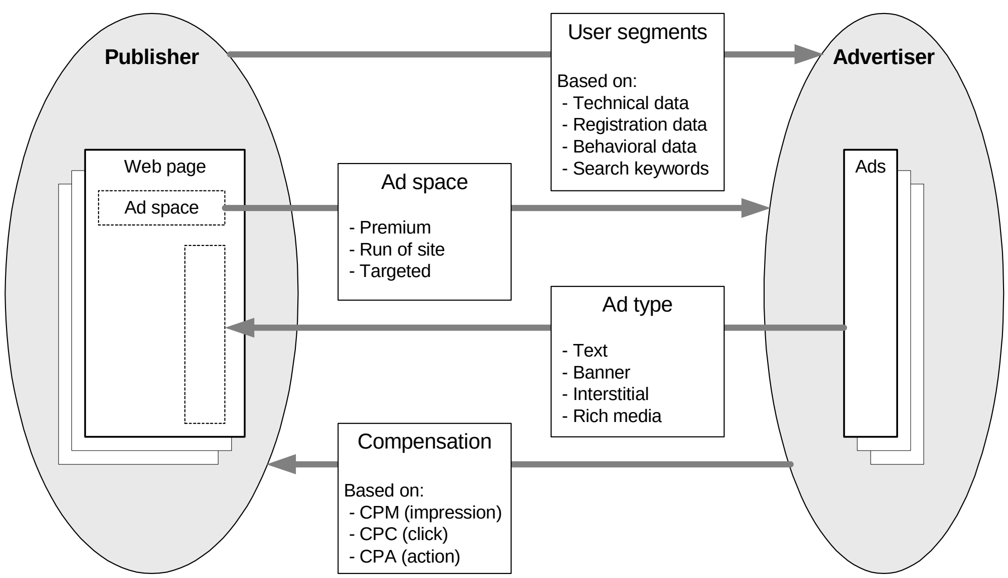

Note that this is a conceptual diagram, not a technical one, e.g. the publisher does not send user segment data to the advertiser, instead segments to which ads may be delivered are part of the sale of ad inventory.

3 Outsourcing ad buying and selling

Here we dive more deeply into the particular viewpoints of advertisers and publishers, and how they translate into motivations in the ad economy.

The advertiser viewpoint

From the advertiser viewpoint, the key question is:

"How do I increase sales at the lowest cost?"

Advertisers have traditionally viewed their activities as falling into two categories: **branding** and **direct response**. Branding builds general awareness and future likelihood to buy, while direct response attempts to directly lead to a sale. Although this line is blurring, branding is generally associated with graphical CPM ads, while direct response is associated with CPC or CPA ads.

Ads are sometimes part of **campaigns** that attempt to optimize effectiveness by controlling the order, duration, frequency, and context in which users see the various ads in the campaign. Campaigns can extend across various media including the Web, and managing them is often outsourced to an **ad agency**. This outsourcing includes coming up with actual advertisements (**creative**) and/or placing these ads and tracking their performance or **return on investment** (**ROI**).

ROI, which is usually expressed as the ratio of gain over investment, has traditionally been difficult to measure, since there was no way to know what ad(s) influenced a given purchase. In the case of advertising investments, ROI is sometimes called **return on ad spend** (**ROAS**), and is calculated as follows:

ROI = (Value gained from ads – Cost of ads) / Cost of ads.

One of the advantages of the Web is that ROI is immediately apparent in the case of direct response ads that lead to a sale. The average profit per impression is

CTR * CR * Average profit per conversion

so the ROI is

ROI = (CTR * CR * Average sale per conversion * 1000 – CPM) / CPM.

This calculation is even more straightforward if the ad space is paid for on a CPC or CPA basis; if the action is a sale in the latter case, the ROI is simply

ROI = (Average sale – CPA)/CPA.

Thus CPC and CPA compensation shifts the risk of mis-estimating CTR and/or CR to the publisher, even though these measures are mostly affected by the creative and conversion efforts of the advertiser.

The publisher viewpoint

From the publisher viewpoint, the key question is:

"How do I maximize the compensation I get for my ad inventory?"

The first answer is to simply maximize the quantity of ad inventory. Ad inventory depends upon page views and ad space, so publishers can either try to attract more viewers or increase the number of ads per page (which tends to adversely affect eCPM).

The second answer is to maximize eCPM. When buying ad space, advertisers are willing to pay more to reach the right users at the most receptive moment. Thus the key to increasing eCPM is user targeting and content targeting, as previously detailed. Content targeting is straightforward enough, but user targeting runs two risks:

- Asking for user data via a registration barrier is a hassle for users, and drives many of them to leave or enter false information
- Inferring user data by tracking or correlating to an offline database can raise privacy concerns

Clearly then, a key challenge for online marketing is to approach the end goal of individualized or one-to-one marketing while preserving the user's privacy and sense of control.

Just as advertisers outsource to agencies, publishers can outsource ad sales to ad networks. Ad networks aggregate ad inventory across many publishers and provide a single place where advertisers (or agencies) can buy ad space in volume. This helps lower the cost of selling for publishers with ROS or "remnant" inventory and smaller publishers with less inventory, while advertisers and agencies can more easily buy in volume, lowering their cost of buying.

Ad flow overview

In summary, the following figure depicts how ads flow from advertisers to publishers, including the roles of ad agencies and ad networks.

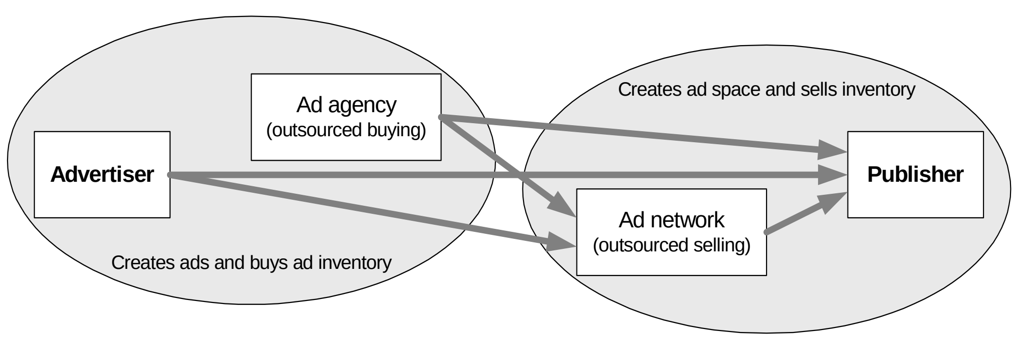

4 Search engines and intermediaries

Search

According to our terminology, search engines are publishers: they are a web site that sells advertising space. So are sites that incorporate search in their use, such as shopping sites.

A huge advantage of search is that in performing it, a user automatically provides the highest quality user targeting data known: search keywords. Keywords tell the publisher what the user is interested in at that very moment, without causing either user inconvenience or undue privacy concerns.

In addition, search has become an important de facto interface to the Web for many users, so that ranking highly in search results and advertising on search engines has become an important factor in marketing success for many companies. The result is that search has quickly become a huge component of online ad sales, spawning an entire sub-industry in the process.

Search engine marketing

Search engine marketing (**SEM**) has two aspects, both of which can become complex and so are often outsourced to SEM specialist companies. The **search engine results page** (**SERP**) has two sections:

- **Organic results** (AKA natural results): this is the list of links that the search engine found most relevant to the search keywords according to its presumably objective search algorithm
- **Sponsored links** (AKA paid placements): this is the list of ads shown along with the organic results; they are usually text ads, which may either avoid distraction from or encourage confusion with organic results

Attempting to increase a web site's rank in the organic results is called **search engine optimization** (**SEO**). SEO can range from modifying a site to be more accurately indexed by search engines to manipulating the search engine algorithm by creating false keywords, links, or other factors used to determine ranking in organic results.

The other aspect of SEM is managing **paid search**, that is, buying of keywords on search engines and calculating the resulting ROI. **Paid inclusion** is a program some search engines have that allows a site to pay in order to be "guaranteed" inclusion in organic results.

Intermediaries

If you use a search engine to search for a common product, for example digital cameras, you'll find several different types of sites in the organic and paid results:

- Manufacturer sites (e.g. Kodak)
- Distributor sites (e.g. Circuit City)
- Shopping comparison sites (e.g. PriceGrabber)

- Directories and research sites (e.g. CNET)
- Vertical specialist sites (e.g. digitalcamera-hq.com)

Distributors and **value added resellers** (**VAR**s) are a normal part of the offline world of sales and fulfillment: they are **vendors** who buy product from the manufacturer and sell it directly to the customer. But **intermediaries** such as comparison, research, and specialist sites make their money in a completely different way. Instead of selling product themselves, they link the user to a vendor for purchase, and for this they are compensated.

The relationship between intermediary and vendor depends upon whether any user data is included:

- **Affiliate marketing**: If the intermediary simply causes the user to follow a link to the vendor site, the link is called a **referral**, and the intermediary is called an **affiliate**.
- **Lead generation**: If the intermediary collects and passes along data about the user such as contact information and details concerning what the user is looking to purchase, the user following a link to the vendor site is called a **lead**, and the intermediary is called a **lead generator**.

The line between affiliates and lead generators is not always distinct, but in general affiliates are smaller and more numerous, and are usually provided with a standardized way to make referrals (such as Amazon affiliate product links). Both are paid on a CPC or more often a CPA basis.

Since intermediaries are compensated by generating interest in a vendor, they can be viewed as outsourced marketing, in the same way that distributors and VARs can be viewed as outsourced sales. And just as distributors often supply VARs, affiliates can join an **affiliate network**, which gives them a central place to manage and track affiliate relationships with many vendors.

Intermediary overview

In summary, the following figure depicts how vendors use intermediaries as outsourced marketing in the same way that manufacturers use distributors and resellers as outsourced sales.

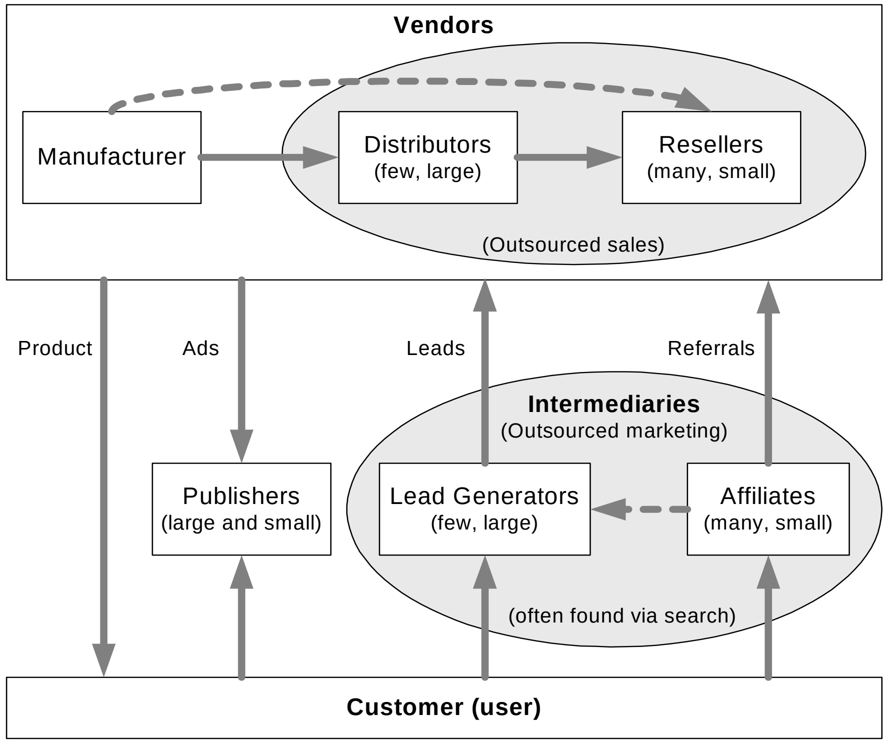

5 Google, auctions, and arbitrage

Google

Since Google has been so successful in advertising, and since AdSense is currently the main player in placing ads on smaller publisher sites, it is worth going into some detail on how Google advertising solutions work. Many other companies use elements of these solutions as well, so a lot can be learned about the market in general by studying Google.

Google has two advertising programs, one for advertisers and one for publishers:

- **AdWords**: This program allows advertisers to buy ad inventory based on keywords. The ads may appear on Google SERPs, within Google applications, on other search engines such as AOL, and on publisher sites subscribing to the AdSense program.
- **AdSense**: This program allows publishers to sell ad inventory to advertisers participating in the AdWords program. Ads may be targeted either contextually or based upon search keywords entered in a site-specific Google searchbox.

So Google is both a publisher and an ad network. One reason for Google's success as an ad network is that every advertiser customer as a publisher automatically becomes an advertiser customer as an ad network. This neatly solves the problem all "two-sided markets" have, which is how to prime one side of the market before the other exists.

Keyword auctions

Another factor in Google's success is its pioneering use of keyword auctions in selling ad inventory. As we noted previously, a huge advantage of search is that it automatically generates very high quality targeting data, the search keywords themselves. A keyword auction increases market efficiency by linking pricing to targeting data, and also lowers Google's cost of sales by setting prices and fulfilling purchases automatically.

The auction itself operates in a straightforward way. For each set of keywords, advertisers choose:

- The ad that will appear
- Additional targeting criteria (geo and/or specific sites)
- The maximum CPC that will be bid
- The maximum daily dollar amount of ad inventory that will be bought

For a given ad space, Google then collects all active bids on the keywords associated with the search or web page, and places ads taking into account the maximum bid amounts. To encourage advertisers to enter high bids, Google reduces the CPC of the top bidder to be just over the next highest bid.

However, a key innovation in Google's ad placements is that the CTR for a given ad is taken into account along with the maximum bid in determining placement position. This means that an ad with a high CTR can be placed more prominently than another ad with a higher bid, but a lower CTR. This has two consequences:

- Users see more ads that other users have previously clicked on, and which are therefore presumably more relevant and helpful.
- The overall CTR of displayed ads is higher, increasing transaction volume and thus Google's income, which is a percentage of CPC.

Keyword arbitrage

It is neither difficult nor capital-intensive to set up an intermediary site that is paid based upon leads or referrals, and then to drive users to the site by buying keyword ads. This has resulted in arbitrageurs trying to take advantage of market inefficiencies to "buy users" for less than they "sell" them.

There are several ways to approach this arbitrage, including the following:

- By acquiring better knowledge of the auction system, one can buy users that vendors have missed for a CPC that is less than vendors are willing to pay for a referral.
- By acquiring better knowledge of user behavior, one can buy users based upon non-obvious low-CPC keywords and sell them for a higher CPC or CPA to vendors.
- By adding value to the user experience, one can buy users at a CPC and sell them to vendors at a CPA that generates a profit.

It's probably worth noting that while they do act to decrease inefficiencies in the market, the first two techniques also can tend to pollute search results with spurious links that deteriorate the user experience. However, they also increase the number of bidders in the keyword auction, thus keeping prices high and benefiting the search engine, so it's not clear where motivations lie on this topic.

Tracking and reporting

Along with the auction mechanism itself, Google and others also provide tools for both advertisers and publishers to track and calculate various relevant quantities, including CPC, CTR, ROI, and eCPM. In order for advertisers to take advantage of these tools, they must consider two more pages in the marketing process:

- **Landing page**: the page that users see after clicking on an ad.
- **Conversion page**: the "thank you" page that users see after completing the action desired by the advertiser, such as entering information or purchasing a product.

Since the JavaScript that a publisher places on a web page in order to insert ads comes from Google, Google can do several things:

- Identify the publisher page on which the ads will appear
- Count the impression against the selected ads
- Determine technical targeting data for the user
- Place a cookie in the user's browser

When a user clicks on an ad, the link is actually to Google, who then redirects the browser to the advertiser's landing page. This lets Google:

- Count the click against the ad and the publisher page the ad appears on
- Identify the ad and publisher page in the user's cookie

Finally, to track ROI advertisers must place a Google beacon on the conversion page. This lets Google read the user's cookie and count the action for use in ROI calculations.

It's probably worth noting that during this process, Google can collect information on ads clicked and conversions achieved per user. This comprises additional valuable targeting data on the user, and if the user is logged into a Google user account, can potentially be aggregated across browsers and computers. Google requires advertisers to make the conversion page beacon visible as a link marked "Google site stats," which leads to a page explaining why the beacon is present and for what purposes the data will be used.

Google advertising overview

In summary, the following figure depicts how the Google advertising system works.