

BEYOND CAUSTIC CROSSINGS: PROPERTIES OF BINARY MICROLENSING LIGHT CURVES

CHRISTOPHER NIGHT,¹ ROSANNE DI STEFANO,^{1,2,3} AND MEGAN SCHWAMB⁴

Received 2007 April 24; accepted 2008 May 19

ABSTRACT

Binary microlensing light curves have a variety of morphologies. Many are indistinguishable from point-lens light curves. Of those that deviate from the point-lens form, caustic crossing light curves have tended to dominate identified binary-lens events. Other distinctive signatures of binary lenses include significant asymmetry, multiple peaks, and repeating events. We have quantified, using high-resolution simulations, the theoretically expected relative numbers of each type of binary-lens event, based on its measurable characteristics. We find that a microlensing survey with current levels of photometric uncertainty and sampling should contain at least as many non-caustic crossing binary-lens events as caustic crossing events; in future surveys with more sensitive photometry, the contribution of distinctive non-caustic crossing events will be even greater. To show that this result is robust, we investigate the influence of several physical effects, including blending, sampling rate, and various binary populations.

Subject headings: gravitational lensing

1. INTRODUCTION

1.1. Background

In gravitational microlensing, a lens and a more distant source pass near each other in the sky, causing an apparent brightening (or magnification) of the source that varies with time as the alignment changes. In the simplest case, when the source is a point source and the lens is a point mass, the light curves associated with these events are described by a simple formula (Einstein 1936). Relative to this point-lens form, the light curves produced by a binary lens, such as a binary star, are much more diverse.

General relativity is a nonlinear theory, and binary-lens light curves can therefore differ significantly from the sum of their two point-lens components. Furthermore, although the shape of the light curve associated with any given event can be computed to arbitrary accuracy, there is no simple analytic formula for binary-lens light curves. These light curves can be practically indistinguishable from point-lens light curves, or they can be distinctive, as in the case of caustic crossing events, which exhibit spikes of divergent magnification. Other features not found in point-lens light curves include asymmetry in time, multiple local maxima, and repeated events, which comprise two distinct brightenings separated by a return to baseline.

Because they are easy to identify visually and offer tighter constraints on the binary parameters, caustic crossing events are the most attractive binary-lens events. Therefore, the binary events published to date by surveys such as MACHO and OGLE have tended to be dominated by caustic crossing events. The MACHO data contain 21 known binary-lensing events, of which 14 were caustic crossing; toward the bulge, 12 out of 16 binary events showed caustic crossings (Alcock et al. 2000). The OGLE III 2002–2003 season produced 24 identified binary-lensing events, of which 17 were caustic crossing (Jaroszynski et al. 2004), the 2004 season produced 25 events, of which 22 were caustic crossing (Jaroszynski et al. 2006), and the 2005 season produced

12 events, of which 10 were caustic crossing (Skowron et al. 2007).

1.2. Smoothly Perturbed Events

Our main focus is on binary-lensing events that differ significantly from the point-lens form, but that do not exhibit caustic crossings. We call these events *smoothly perturbed* if they differ significantly from the best-fit point-lens model. The amount by which they must differ from the point-lens form in order to be considered perturbed depends on the photometric uncertainty of a survey. Our main question is this: for typical values of photometric uncertainty, how many such smoothly perturbed events are expected for each detected caustic crossing event? The answer to this and similar questions depends on the characteristics of the binary lens, specifically the mass ratio q and the orbital separation a .

In § 2, we introduce the binary-lensing parameters, and explore the dependence of the rates of different types of light curves on these parameters. In § 3 we describe the simulations we have used to calculate these relative rates numerically. In § 4 we give the results for individual binary lenses over the ranges of our simulation, as characterized by a set of binary parameters. In § 5 we give the results for typical populations of binary lenses and study the implications. We find that current surveys must be capable of identifying significantly more smoothly perturbed binary-lens events than have been found so far.

2. LENSING GEOMETRY

2.1. Lensing Formula

The magnification associated with a binary-lens depends on the binary parameters and the relative locations of the source and of the lens components on the plane of the sky. As the relative positions change, we can compute the magnification at each time (Mao & Paczynski 1991; Schneider et al. 1992; Petters et al. 2001).

It is convenient to express the angular positions and angular distances involved in units of the Einstein angle $\theta_E = \{(4GM/c^2)[(D_S - D_L)/D_L D_S]\}^{1/2}$, where M is the total mass of the lens, D_S is the distance to the source, and D_L is the distance to the lens. For a given M , the lens is described by two parameters, the mass ratio $q = M_2/M_1$ and the instantaneous

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138.

² Kavli Institute for Theoretical Physics, University of California Santa Barbara, Santa Barbara, CA 93106.

³ Department of Physics and Astronomy, Tufts University, Medford, MA 02155.

⁴ California Institute of Technology, Pasadena, CA.




FIG. 1.—Graphical depiction of $\langle s \rangle$ for caustic crossing events for two specific binaries. *Left*: The binary $(q, a) = (0.750, 1.000)$, which has a single connected caustic. The x - y plane shown is the lens plane, with lengths measured in units of the Einstein angle θ_E . For a given value of θ , $s(\theta)$ is the width of the caustic region, as seen from that angle. The minimum bounding polygon (dashed line) is defined by the cusp positions. It can be used to compute $\langle s \rangle$ semianalytically. *Right*: The binary $(q, a) = (0.010, 0.835)$, for which there are three disconnected caustics. In this case, $s(\theta)$ may be the sum of two or three widths. The diagram shows $s(\theta)$ for the same value of θ as the left diagram. Again it is possible to calculate $\langle s \rangle$ from the cusp positions alone.

angular separation a , which is measured in units of θ_E . Also, M_1 and M_2 are the mass of the primary and the secondary, and $M = M_1 + M_2$.

To define the coordinate system, we place the primary at the origin and the secondary on the positive x -axis. A point source is completely described by its position (x_s, y_s) in this coordinate system.

The formula for the magnification, $A(t)$, is simply a function of the binary parameters (q, a) and the angular source position $(x_s(t), y_s(t))$: $A(q, a, x_s(t), y_s(t))$. We use as the units of time the Einstein angle crossing time $t_E = \theta_E/\omega$, where ω is the angular speed of the source relative to the fixed lens.

We consider only the case where the lens is static, so neither a nor the binary's orientation changes. The angular rotation of a binary lens over the course of an event is expected to be small. This is because, although binaries with a wide range of separations may act as lenses, the binary nature is only evident when the projected separation between the components is on the order of θ_E , which is typically larger than an AU for stellar lenses. The orbital periods therefore tend to be longer than the event durations, which are on the order of weeks to months. Dominik (1998) gives more detailed criteria for when binary rotation is likely to be significant.

Thus, we consider the source trajectory to be a straight line in the lens plane, defined by b , the angular distance of closest approach to the lens center of mass, and θ , the angle the trajectory makes with the binary axis (defined to be the x -axis):

$$x_s(t) = -b \sin \theta + t \cos \theta + a_1, \quad (1)$$

$$y_s(t) = b \cos \theta + t \sin \theta, \quad (2)$$

where $a_1 = aM_2/M$ is the x -coordinate of the lens center of mass. Symmetry with respect both to the binary axis and to time reversal allows us to cover the complete parameter space by considering values of θ in the interval $(0, \pi/2)$.

2.2. The Relative Rate Measure $\langle s \rangle$

We consider a light curve to contain an event if its maximum magnification is greater than a certain value, A_{cut} . The rate for lensing events is given by $\sigma s \omega$, where σ is the source density, and s is the width of the lensing region, that is, the region bounded

by the isomagnification contour in the lens plane corresponding to the cutoff magnification. Then for a point lens, with $A_{\text{cut}} = 1.34$, the lensing region is a circle of radius θ_E , and so $s = 2\theta_E$. For a given binary lens, the lensing region is not circular, so s depends on the angle of approach θ . The value $s(\theta)$ is the linear size of the lensing region as seen from the angle θ . The average value of $s(\theta)$, $\langle s \rangle$, for a given binary represents the lensing region's average angular width. If a large number of sources were to pass behind this binary, the event rate would be $\sigma \langle s \rangle \omega$. Since the only dependence on the intrinsic properties of the lens is through the factor $\langle s \rangle$, it is this quantity we will compute for different lenses. Taking σ and ω to be the same for all potential lenses, the ratio of rates does not depend on them. We can therefore use the ratios of the values of $\langle s \rangle$ to compare relative average rates for the total numbers of events generated by different lenses.

For a given binary lens, different source paths produce light curves with different characteristics. For example, paths that cross the caustics produce light curves with distinctive wall-like structures, the so-called caustic crossing light curves. The rate of such events is determined by the dimensions (in angular units in the lens plane) of the caustics (see Fig. 1). The value $s_{\text{cc}}(\theta)$ measures the extent of the caustics perpendicular to the direction of approach specified by θ .

Even event types that do not correspond to crossing an easily defined region, such as smoothly perturbed events, may be thought of as having an effective width, which $s(\theta)$ measures. For a given q, a , and θ , there is a range of values of b that correspond to light curves with any property we choose to investigate, such as being smoothly perturbed. This range may be connected or disconnected, but generally, the bigger the size of the range (or the sum of the sizes of the parts), the more likely that a randomly chosen b will correspond to a light curve with the given property. The size of this range, then, is $s(\theta)$. [In set-theoretic terms, $s(\theta)$ is the 1D Lebesgue measure of the set of all values of b that correspond to the property in question, for fixed q, a , and θ .]

We can therefore compute $s(\theta)$ for any given property by generating many different source paths with incidence angle θ , and counting the fraction of all events in which the light curve exhibits that property. Specifically, we generate light curves for many values of b for each value of θ . For each such trajectory, we sample the magnification at a sequence of times t_i , compute the




FIG. 2.—Peak count and asymmetry parameter in the light curve corresponding to a given b and θ , with $(q, a) = (1.0, 1.0)$ held fixed. The graph is *not* shown in x - y coordinates like in Fig. 1; b and θ are similar to the polar coordinates in that plane. *Left*: Number of peaks color-coded. Colors gray, red, yellow, orange, green, and blue correspond to 1, 2, 3, 4, 5, and 6 peaks, respectively. *Right*: Asymmetry parameter k (see § 3.3.2), color-coded logarithmically, with black the most symmetric ($k \approx 0$) and white the least symmetric ($k \approx 1$). In both panels, black sinusoidal curves show the b -value of the caustic cusps as a function of θ ; any trajectory with values of b and θ between these black curves will produce a caustic crossing light curve. The black curves are easy to describe in terms of the cusp locations in the x - y plane (see Fig. 1), and so $\langle s \rangle$ for caustic crossings may be computed from them. For other light curve characteristics, however, $\langle s \rangle$ is more complicated. For instance, $\langle s \rangle$ for single-peaked events is given by $1/\pi$ times the area of the gray region in the left panel. That is, $\langle s \rangle = (1/\pi) \int_0^\pi \int_{-\infty}^{+\infty} N(b, \theta) db d\theta$, where $N(b, \theta)$ is 1 at any point in the gray region and 0 elsewhere. In the right panel, $\langle s \rangle$ for light curves with a certain amount of asymmetry corresponds to the size of a contour for a certain brightness. Any other $\langle s \rangle$ is similarly defined in terms of a region on the b - θ graph.

corresponding $A_i = A(t_i)$, and examine the resulting light curve. The sampling rate Δt is the difference between consecutive values of t . As a default we choose $0.01\theta_E$, but in § 4.4 we explore how the results vary for different sampling rates.

Once we have determined $s(\theta)$, the next step is to average over the possible values of θ : $\langle s \rangle = (1/2\pi) \int_0^{2\pi} s(\theta) d\theta$. For some event types, there is a simple expression for $s(\theta)$, so we may compute $\langle s \rangle$ analytically. Generally, however, no such simple expression exists, so we must perform the integration numerically with simulations.

From the values of $\langle s \rangle$ for various properties, we may compute the relative rates for light curves with those properties. For instance, suppose that for a fixed q and a , the value for caustic crossing events is $\langle s_{cc} \rangle$, and the value for all events (i.e., light curves that exceed our minimum cutoff magnification A_{cut}) is $\langle s_e \rangle$. Then, for this binary alone, the probability of an event being a caustic crossing event is $\langle s_{cc} \rangle / \langle s_e \rangle$, and this also equals the relative rates for caustic crossing events to all binary events. Realistically, however, there is a population of binary lenses described by a probability distribution function $P(q, a)$. In that case, the relative rate we seek is given by $\int \langle s_{cc}(q, a) \rangle P(q, a) dq da / \int \langle s_e(q, a) \rangle P(q, a) dq da$. Our results for individual binaries are given in § 4.1, and our results for populations of binaries are given in § 4.3.

2.3. Caustic Crossing $\langle s \rangle$

For caustic crossing light curves, there is a simple analytic expression for $s(\theta)$ and thus for $\langle s \rangle$. This is because caustic crossing

light curves are defined by whether the trajectory crosses the caustic region. We may therefore determine whether a trajectory corresponds to a caustic crossing light curve without generating the light curve itself. So for caustic crossings, $s(\theta)$ is the width of the caustic region as seen from the angle θ (see Fig. 1), and $\langle s \rangle$ is the average width of this region.

It can be shown that the average width of any convex region (and thus $\langle s \rangle$ for any property that corresponds to crossing the region) is equal to $1/\pi$ times the region's perimeter. Thus, we can efficiently compute $\langle s \rangle$ for caustic crossing light curves in terms of the cusp locations, which are given by the simultaneous roots of two polynomials of degree 11 and 10 (see Schneider et al. 1992, their eq. [6.23]). Practically speaking, however, simulation can give almost as good as results in a reasonable time, and so we use this method simply as a check.

2.4. The Value of $\langle s \rangle$ for Other Light Curve Features

The rate for any type of event that corresponds to crossing a specific region, such as events above a certain maximum magnification, or repeating events (Di Stefano & Mao 1996), can also be computed semianalytically, without generating any light curves. In general, however, there is no easier way to compute $\langle s \rangle$ than numerically. For example, Figure 2 shows that $s(\theta)$ for multi-peaked events and asymmetric events are more complicated than for caustic crossing events. Therefore, we compute the event rates we seek using simulations.

3. SIMULATION SPECIFICATIONS

Our simulation generates a set of light curves for a given q and a by selecting values of b and θ , one at a time. The light curve is defined by a discrete set of times t_i and the corresponding amplifications $A_i = A(t_i)$. If the light curve contains an event [$\max(A_i) > A_{\text{cut}}$], certain light curve parameters are measured. Our default value of A_{cut} is 1.1, that is, a 10% magnification above baseline. (For comparison with such surveys as OGLE, $A_{\text{cut}} = 1.34$ may be more appropriate; see § 4.4.) Among light curves with events, running totals of all light curves with certain parameter values and characteristics are kept.

3.1. Trajectory Selection

We wish to select values of b and θ in such a way as to mimic the way they are randomly “selected” in nature. This is done simply enough by realizing that the probability distribution function for the trajectories is uniform in both b and θ . We select 1600 equally spaced values for θ between 0 and $\pi/2$, and for each of these determine the set of values for b that will completely cover the set of events. This set is all trajectories that cross the isomagnification contours of $A(x_s, y_s) = A_{\text{cut}}$ in the lens plane, as shown in Figure 3.

Although it is essential to completely cover these isomagnification contours, it is desirable to limit the number of light curves generated that do not cross the contours. Starting with a given binary, defined by the values (q, a) , and considering an individual value of θ , there is a limited range of values of b that yield events. For a given θ , we first identify a value of b for which we know there is an event. We accomplish this by starting with the fact that, for any value of A_{cut} , there are isomagnification contours that enclose caustics, as in Figure 3 with two separate contours, each containing one caustic. We therefore begin sampling b by choosing a value that hits a caustic cusp. We then continue by choosing a sequence of values of b , separated by Δb , from the original one. We sample in both perpendicular directions (“up” and “down”) from the original trajectory, keeping θ fixed. Therefore, by starting at each caustic and moving in both directions by an amount Δb (equal to 0.002 for most binaries in the simulation) for as long as events are generated, we can be sure to cover the entire contour, while limiting nonevent light curves. (This corresponds to stopping when the white area is reached in the left panel of Fig. 2.) The caustics, as mentioned in § 2.3, are bounded by their cusps, which can be located to high precision before the simulation is run. Double sampling is avoided by maintaining a running list of the range of b sampled so far.

For each event, as the source gets far enough away from the lens, the magnification drops toward 1, so for any trajectory there is a certain range of t for which $A(t) > A_{\text{cut}}$. The limits of t should be chosen to at least cover this range, but because the tails with $A(t) < A_{\text{cut}}$ are ignored, it is a good idea to reduce these limits as much as possible while still ensuring that all regions above A_{cut} are covered. Therefore, as with b , we chose to stop light curve generation when $A(t) < A_{\text{cut}}$ and all the caustics have been passed (notice the rectangular shape of the dashed lines in Fig. 3). Starting from $t = 0$ and moving in both directions, this allows us to avoid computing many unneeded magnification values.

3.2. The Magnification Function

The time-consuming part of light curve generation is determining the value of the magnification function $A(q, a, x_s, y_s)$. To compute it, one determines the locations in the lens plane of the lensing images, from which A can be computed analytically. The




FIG. 3.— Coverage of sampling. The x - y plane shown is the lens plane, measured in units of θ_E , as in Fig. 1. For $(q, a, \theta) = (0.562, 4.217, 0.400)$ (chosen as a typical example), each dashed line is a trajectory for a light curve that was generated, each with a different value of b . Solid lines correspond to the parts of the trajectory that were identified as events. Isomagnification contours corresponding to $A_{\text{cut}} = 1.1$ and locations of the caustics are also shown. The sampling in b used in the simulations is 50–100 times denser than that shown in this figure.

positions (x_i, y_i) of the images, however, cannot be computed analytically. They are expressed in terms of the simultaneous roots of two polynomials in x_i and y_i of degree 5 and 4, or alternately in terms of a single complex polynomial of degree 5 (Erdl & Schneider 1993).

The method we choose to determine these roots numerically is Newton’s method with polynomial deflation (Press et al. 1992) on the real fifth-degree polynomial given in Asada et al. (2004).⁵ Because it uses real (rather than complex) arithmetic, this method, combined with some other minor optimizations, allows for a great improvement in speed over root finding for the complex polynomial.

Finally, after up to several hundred calls to the function, the light curve is complete, and stored as an array of t_i and A_i values. We may then analyze it, and record its characteristics and statistical parameters.

3.3. Defined Parameters

We seek to quantify the amount by which a binary-lens light curve differs from a point-lens light curve. To this end we define three parameters. For a point-lens light curve, the value of each parameter is exactly zero, and light curves with more pronounced deviations from the point-lens form have higher parameter values. Our three parameters relate to the least-squares fit of a point lens, the asymmetry of the curve, and whether the light curve exhibits more than one local maximum, i.e., whether it is multi-peaked.

3.3.1. Best-fit σ_p

The most straightforward way to evaluate whether a light curve is consistent with a certain model is to fit it to that model using a least-squares fit. Therefore, we begin by fitting each binary-lens light curve with a standard three-parameter point-lens form, known as the Paczyński model. (We assume that in observations,

⁵ The relevant equations in Asada et al. are eqs. (2.25)–(2.34). Note that what are here called (m, a, x_s, y_s) are there called (ν, l, a, b) . Note the following two typos: eq. (2.32) is missing a factor of b outside the parentheses, and eq. (2.33) is missing an overall factor of b^2 .

the baseline can be determined arbitrarily well, so it is not considered a fit parameter.) As usual, we define the closest point-lens model, $A_{\text{PL}}(t)$, to be the one which minimizes the sum of the squared differences between the values of the simulated binary-lens light curve, A_i , and the model:

$$\sum_{i=1}^n [A_i - A_{\text{PL}}(t_i)]^2. \quad (3)$$

The binary-lens light curve offers a certain amount of intrinsic deviation from the point-lens model, characterized by this sum-of-squares quantity. In real observations, there will be additional deviation due to the photometric uncertainty. Roughly speaking, when the intrinsic deviation is larger than the photometric deviation, the binary nature of the light curve will be evident through least-squares fitting. When the intrinsic deviation is smaller, A_{PL} will provide an acceptable fit to the binary-lens light curve.

Thus we define our first parameter, σ_P , as a measure of how large the photometric uncertainty would need to be (in flux space, as a fraction of the baseline flux) in order for the observed light curve to be well fit by a point-lens model. This parameter is given by the rms difference between the simulation light curve and the best-fit light curve:

$$\sigma_P = \left\{ \frac{1}{n-3} \sum_{i=1}^n [A_i - A_{\text{PL}}(t_i)]^2 \right\}^{1/2}. \quad (4)$$

Note that σ_P does not actually quantify the goodness of fit of a point-lens model to a simulated binary-lens light curve; since there are no uncertainties associated with the simulated light curve, such a statistic is meaningless. For an observed light curve with associated uncertainties, the goodness of fit is quantified by the reduced χ^2 test statistic. Its formula is very similar to our formula for σ_P :

$$\chi^2 = \frac{1}{n-3} \sum_{i=1}^n \frac{[A_i - A_{\text{PL}}(t_i)]^2}{\sigma_i^2}. \quad (5)$$

The t_i , A_i , and σ_i are the times, magnifications, and uncertainties of the observed light curve, of which there are n , and $A_{\text{PL}}(t)$ is the three-parameter point-lens fit chosen such that χ^2 is minimized. For our purposes, we will assume that the uncertainties σ_i are equal to a constant fraction of baseline in flux space, denoted by σ_{phot} , an observational parameter that quantifies the photometric uncertainty of the survey.

Thus, roughly speaking, for a given binary-lens light curve, σ_P is the critical value for the photometric uncertainty: if $\sigma_{\text{phot}} > \sigma_P$, the light curve will be well fit by a point-lens model, and if $\sigma_{\text{phot}} < \sigma_P$, it will not. Strictly speaking, if a binary-lens light curve with intrinsic deviation σ_P is observed with uniform, uncorrelated, Gaussian random photometric errors with standard deviation σ_{phot} , and then fit with a point-lens model, the reduced χ^2 statistic will have an expected value (for $n \gg 1$) of $\chi^2 = 1 + (\sigma_P/\sigma_{\text{phot}})^2$. Therefore, for $\sigma_P > \sigma_{\text{phot}}$, the expected value of χ^2 exceeds 2, so the point-lens model would probably be identified as a bad fit. This is the basis of our rule of thumb that a survey with an average photometric uncertainty of σ_{phot} should discriminate as a non-point lens a binary-lens light curve with $\sigma_P > \sigma_{\text{phot}}$.

3.3.2. Asymmetry Parameter k

Model fitting reveals how much a binary-lensing light curve differs from the point-lens model, but not in what way. Because




FIG. 4.—Comparison of residuals for an asymmetric binary-lens light curve and a point-lens light curve. The binary-lens light curve in the left panel (dots, corresponding to simulated observations) has a best-fit point-lens light curve as shown by the solid curve. It has a best-fit parameter of $\sigma_P = 0.09$, the rms deviation from this fit, but an asymmetry parameter of $k = 0.21$. The center panel shows the same observations with 10% Gaussian error added; if this light curve were fit assuming 10% photometric uncertainty, the best-fit test statistic would be $\chi^2 < 2$. The right panel shows the same best-fit point lens, with enough Gaussian error added to give it a comparable χ^2 . Below each light curve are shown the residuals after subtracting out the best-fit curve. Although the average size of the residuals is comparable in the two right panels, the residuals of the center panel show much more apparent structure, as expected.

point-lens light curves are always symmetric in time, i.e., $A(t)$ is an even function, one obvious way in which we expect some binary-lens light curves to differ is by exhibiting some asymmetry with respect to time reversal. Therefore, we define an asymmetry parameter k that is 0 for a symmetric light curve, and larger for more asymmetric ones. The parameter k is based on the Chebyshev coefficients T_n , which are the convolution of the light curve function $A(t)$ with Chebyshev polynomials, as in Di Stefano & Perna (1997). For even n the polynomial is an even function, and for odd n the polynomial is an odd function. Therefore, a function with even symmetry has $T_n = 0$ for all odd n , and a function with some asymmetry must have T_n significantly nonzero for some odd n . We then define the asymmetry parameter to be the ratio of the rms average of the odd coefficients to the even coefficients:

$$k = \left(\frac{\sum_{n=1}^{\infty} T_{2n+1}^2}{\sum_{n=1}^{\infty} T_{2n}^2} \right)^{1/2}. \quad (6)$$

One use of the asymmetry parameter is to characterize the way in which a particular light curve deviates from the point-lens form. We note, however, that when we apply model fits to light curves with pronounced asymmetries, the residuals to the best point-lens fit are not randomly distributed in time, but instead exhibit some structure. That is, the deviations from the point-lens light curve are correlated, even when σ_P is small. Thus, the parameter k can identify light curves with small but correlated residuals, whereas the parameter σ_P cannot. An example of two light curves with identical χ^2 but different k is shown in Figure 4.

In this paper we identify light curves that are smoothly perturbed from a point-lens form by their values of σ_P . We note, however, that in future work it may be possible to use the asymmetry parameter to identify binary-lens light curves, even if the best-fit parameter fails to distinguish them from point-lens events. Figure 5 shows typical light curves with various intermediate values of σ_P and k . Extremely point-lens-like or extremely perturbed events will easily stand out with any choice of statistics.

3.3.3. Multipeak Parameter p

This parameter is identically 0 for any light curve with a single peak. For light curves with two peaks, it is the difference between the first peak and the intervening minimum, or the difference between the second peak and the intervening minimum,




FIG. 5.—Variety of binary-lens light curves along with their best-fit point-lens models. In the upper right of each panel is shown the asymmetry parameter k followed by the best-fit parameter σ_P . The top, middle, and bottom rows show light curves with asymmetry parameter k equal to approximately 5%, 10%, and 20%. The error bar in the upper left has a value of $2\sigma_P$; thus, it is the size of the error bar necessary to make the best-fit χ^2 equal unity. Generally speaking, the greater the parameter values, the further the light curves are from their point-lens models.

whichever is smaller. The idea is that if differences of less than p were undetectable, the multi-peaked nature of the light curve would go unnoticed. For light curves with three or more peaks, it is the largest such value for any of the minima in the light curve:

$$p = \max_i [\min(A_{\max,i}, A_{\max,i+1}) - A_{\min,i}], \quad (7)$$

where $A_{\min,i}$ is the magnification at the i th minimum and $A_{\max,i}$ is the magnification at the i th maximum.

4. RESULTS

4.1. Results for Individual Binaries

Figures 6 and 7 depict the relative event rates $\langle s \rangle$ as a function of q and a . Figure 6 plots these rates for several fixed values of q , and Figure 7 shows contour plots for the entire range of q and a covered by our simulation. The graphs in Figure 6 are vertical cross sections of the contour plots in Figure 7.

The total rate of events (light curves for which $A_{\max} > A_{\text{cut}}$) is shown as the top, solid black curve (“all”) in each panel of Figure 6, and in the top left panel of Figure 7. The values we have computed can be checked analytically both for low values of a , where the lens essentially acts as a single point lens, and for high values of a , where the lens acts approximately as two independent point lenses. In the limit of small a , the event rate is determined simply by our cutoff magnification. For $A_{\text{cut}} = 1.34$, corresponding to $b \leq \theta_E$ (or $b \leq 1$ in our units) for a point lens, the rate would be $\langle s \rangle = 2$. For our default value of $A_{\text{cut}} = 1.1$, this rate is $\langle s \rangle = 3.35$. On the other hand, for a very large, the masses act as two separate point lenses, each with its own Einstein radius proportional to the square root of its mass (Griest et al. 1991).

Thus, the event rate increases by a factor of $(1 + q^{1/2})/(1 + q)^{1/2}$. This factor is maximized with a value of 1.414 for $q = 1$, producing an overall event rate of $\langle s \rangle = 4.73$. This is the value approached, for large a , by the solid black curves in Figure 6.

Also plotted are rates for the two main types of binary perturbations: non-caustic crossing but with $\sigma_P \geq 10\%$ (smoothly perturbed events), and caustic crossing. These are shown in Figure 6 with a dotted orange curve and a solid blue curve, respectively, and in the middle two panels of Figure 7. All events that are not perturbed in either of these two ways are shown as “point-lens-like” events, with a red dashed curve in Figure 6, and in the top right panel of Figure 7. The rate for “all” events, that is, the total rate, is the sum of the rates for smoothly perturbed, caustic crossing, and point-lens-like events.

For all values of q , there are values of a for which the smoothly perturbed events outnumber the caustic crossing events. In Figure 6, this is where the blue solid curve is above the orange dotted curve. In Figure 7, this can be seen by comparing the overlaid numbers at the top and bottom of the middle two panels. An increase or decrease in a from $a = 1$ results in the smoothly perturbed event rate exceeding the caustic crossing rate. The fact that it occurs on both sides of $a = 1$ suggests that a binary distribution function weighted toward either high or low a will not tend to favor caustic crossing events over smoothly perturbed events.

Binary perturbations are most frequent close to $a = 1$ for all values of q , and there tends to be more variation of event rates over the range of a covered by our simulation than over the range of q . This can be seen by the largely horizontal contours in Figure 7.

Small q corresponds to a low mass secondary, approaching the planet in the limit $q \ll 1$. For the most part, for small q , both




FIG. 6.—Relative rate measure $\langle s \rangle$ for various types of binary-lensing events for certain fixed values of q . For each of the four values of q as indicated in the left panels, $\langle s \rangle$ is plotted against a . The types of events identified are caustic crossing events (orange dotted curve), smoothly perturbed events ($\sigma_P > 0.1$, blue solid curve), point-lens-like events ($\sigma_P < 0.1$, red dashed curve), and the total of these three categories (black solid curve). Note particularly that there are many values of q and a where the smoothly perturbed events outnumber the caustic crossing events. The right panels show the same data as the left, except on a logarithmic plot. The right panels also show the relative rates for two more types of events, repeating events ($A_{\min} < A_{\text{cut}}$, solid yellow curve) and asymmetric events ($k > 0.1$, dashed green curve). There is some overlap between these two categories, and between these categories and the others.

event characteristics and event rates approach those of point lenses. Nevertheless, there are two regions of parameter space where planetary systems can provide distinctive lensing signatures. The first is for values of a close to unity. In this “resonant” regime (Mao & Paczynski 1991; Gould & Loeb 1992), the most obvious evidence of the existence of the planet is provided by caustic crossings, which punctuate an otherwise point-lens-like light curve. The second is for larger values of a , where asymmetry in otherwise point-lens-like events, and also repeating events provide evidence of the low mass companion (Di Stefano & Mao 1996; Di Stefano & Scalzo 1999a, 1999b). In this paper we focus on the regime of binary stars. As q increases from near 0, binary signatures become more frequent and more pronounced, with relatively little effect for varying q greater than ~ 0.2 (see Fig. 7). Because values of q near or larger than 0.2 are favored by binary distributions, e.g., the distribution studied by Duquennoy & Mayor (1991), we expect the distribution in a to have a more pronounced effects than the choice of realistic distribution in q .

Although the strongest binary signatures occur within an order of magnitude of $a = 1$, there is a dip in binary signatures in the center of this range, very near to $a = 1$. This can be seen in Figure 6 as a dip in total event rate and a smaller dip in caustic crossing event rate. This can be explained in terms of the geometric considerations of the lens plane from § 2. When a is very close to 1, the caustics (and thus the region of binary perturbations) are roughly equal in size in both dimensions. As a increases or decreases, the caustics stretch in one dimension or the other, and so they have a larger extent in the lens plane. The caustic crossing event rate peaks just as the caustics transition from connected to disconnected, and the perturbed event count peaks farther from $a = 1$, with the value of a at the peak determined by the value of A_{cut} .

Although our coverage in a is not infinite, spanning two orders of magnitude from $a = 0.1$ to $10 \theta_E$, it includes the range with significant binary signatures. For any value of q , all binary effects will go to 0 as a becomes very large or very small. In Figure 6, this is seen in all panels, as all curves but the top two




FIG. 7.—Contour plots of $\langle s \rangle$ vs. q and a for the entire range covered by our simulations. Each panel corresponds to a different type of event: all events ($A_{\max} > A_{\text{cut}}$), point-lens-like events ($\sigma_p < 10\%$), smoothly perturbed events ($\sigma_p > 0.1$), multi-peaked events at 10% (multi-peaked $k > 0.1$), and finally repeating events ($A_{\min} < A_{\text{cut}}$). Consecutive contours are separated by 0.1 with dark contours every 0.5. Overwritten numbers show the same data in tabular form. Notice that generally there are more variations with varying a than with varying q , and that the curves exhibiting binary signatures exist most in the “critical zone” around $a = 1$. For a point-lens light curve, $\langle s \rangle$ for the total number of curves would be 3.35.

go to 0 in both directions. In Figure 7, this is seen in the lower four panels, as all tabulated values go to 0 at the top and bottom of the panel.

Repeating events (multi-peaked events with $A_{\min} < A_{\text{cut}}$), and asymmetric events ($k > 0.10$) are shown in the right-hand (log

plot) panels of Figure 6, as gold solid curves and dashed green curves, respectively. Repeating events are one type of binary behavior that does not fall off with the size of the caustics for large a . Rather, two separated point lenses can still produce a significant number of repeating events. This can be computed analytically as

in § 2. If R is the rate for a point lens of the same mass ($R = 3.35$ for $A_{\text{cut}} = 1.1$), then to leading order in R/a , the rate for repeating events is approximately $[1/(2\pi)]\text{sech}[\frac{1}{2}\ln(q)](R^2/a)$. This expression goes as a^{-1} , whereas the caustic crossing event rate goes as a^{-2} .

Although asymmetric and repeating events are significantly less frequent than caustic crossing events and smoothly perturbed events, they are expected to constitute a few percent of binary-lens events. We should therefore find them in any large sample of binary-lensing events.

4.2. Optical Depth

The optical depth τ is commonly used to express the contribution of various lens populations to the microlensing signal. It is defined as the angular area on the sky covered by the lensing region of the lenses in question (i.e., the region in the lens plane where $A > A_{\text{cut}}$).

All of our main results are in terms of relative rates, so they cannot be compared to any values of the optical depth, just as a perimeter of one shape cannot be compared to the area of another. However, as all of the events we study are due to a single population, namely binaries, we can ask how the optical depth due to this population compares to that of an equivalent point-lens population.

To compute τ for the total of all binary-lensing events together, including caustic crossing, smoothly perturbed, and point-lens-like, we can measure the area of the lensing region as a function of q and a . It will be identical to that of an equivalent point lens for large and small values of a , but it can differ for intermediate values of a . If the difference is significant enough, then the binary fraction might be an important consideration in interpreting the overall microlensing optical depth.

We find, however, that this is not the case. In order to efficiently calculate the area of the region, we run a modified quad-tree algorithm. First, the lens plane is divided into squares of side length h . Each of these is painted either white (if the magnification at the square's center exceeds A_{cut}) or black (if it does not). Next, every white square within $3h$ of a black square (using the L_∞ metric) and every black square within $3h$ of a white square is identified. Each of these squares is subdivided into four squares of side length $h/2$, which are individually painted according to the magnification function at their centers, and the process is repeated. The total area of the lens plane painted black is our optical depth (in units of square θ_E). We divide this value by the corresponding value for an equivalent point lens (for $A_{\text{cut}} = 1.1$, this is approximately $8.799\theta_E^2$).

The ratio we determine is shown in Figure 8. At most, it differs from 1.00 by a few percent around the critical zone. Furthermore, the regions where the ratio exceeds 1.00 are approximately equal in magnitude to the regions where the ratio is less than 1.00, so for any reasonable distribution of binary parameters, the effect will largely cancel out. For a log-uniform distribution in both q and a , for example, the average value on the grid in Figure 8 is 0.996. Therefore we conclude that, while a significant fraction of light curves are due to binaries, the binary fraction has no more than a marginal effect on the total microlensing optical depth.

4.3. Results for Populations of Binaries

In order to use these results to make predictions for surveys, we must consider a population of binary lenses whose properties (mass ratios and projected orbital separations) are drawn from a realistic distribution given by a probability distribution function




FIG. 8.—Contour plot of the optical depth τ due to microlensing, shown as the ratio between the optical depth for a binary lens and for an equivalent point lens. Contours correspond to values of 0.94, 0.98, 1.02, and 1.06. Overwritten numbers show the same data in tabular form.

$P(q, a)$. Then, for instance, the overall relative rate for a given type of event is

$$\langle s_x \rangle = \iint P(q, a) \langle s_x(q, a) \rangle dq da. \quad (8)$$

Note that previously we have used $\langle s \rangle$ to denote an average only over b and θ , but here and in subsequent sections, we use it to denote an average over all four binary parameters. As a default, we assume a distribution uniform in q and log-uniform in a , but in § 4.4 we explore how the results vary for different distributions.

For every caustic crossing light curve, there are a certain number of smoothly perturbed light curves, as shown in Figure 9. We want to find this ratio. Our default definition for smoothly perturbed events is $\sigma_P \geq 0.10$, but we vary this cutoff parameter value over a wide range, from 0.01 to 1.0, and show the ratio for all cutoff values.

Figure 10 shows the major results. For various cutoff σ_P values, it shows the ratio of smoothly perturbed to caustic crossing events. For a cutoff value of σ_P of 0.10, this ratio is slightly greater than unity. Roughly speaking, this means that a survey with 10% photometry has the potential to detect more smoothly perturbed events than caustic crossing events. Naturally, the better the photometry, the larger this ratio should be. For $\sigma_P = 1\%$, the theoretical ratio is as high as 3. Because this ratio is much lower than the ratios of known events for both MACHO and OGLE (as long as σ_P is less than 35), it may be smoothly perturbed events that have not yet been identified.

This conclusion, that undiscovered events must exist in real data, depends on the fact that the discrepancy between calculated and observed event rate ratios cannot be explained in terms of any of the assumptions made in our simulations. While we did not assume any particular photometric precision, we did assume default values of the sampling rate and event detection threshold, values that are more ideal than could be expected of a current survey. It appears, however, that the most obvious such possible explanations cannot be responsible for the discrepancy. We may see this by determining the effect that modifying these assumptions has on the results.

4.4. Robustness of Results

Ultimately, we are interested in the relative rates of smoothly perturbed and caustic crossing events, as they would appear in




FIG. 9.—Random sampling of non-caustic crossing, nonrepeating light curves perturbed at the 10% level. Each light curve is shown with its best-fit point-lens model. Residual plots below show the patterns of the differences between the light curve and the model. The three light curve parameters are shown in the upper right of each panel; these light curves were selected from the population that has at least one of these three parameters greater than 0.1. Horizontal (time) axis tick marks for each light curve are in units of 0.5 times the Einstein angle crossing time θ_E/ω . Vertical axis tick marks in the residual plots are in units of σ_p , and the axis ranges from $-3\sigma_p$ to $+3\sigma_p$. Our simulations show that, typically, for every 13 caustic crossing events, there are 15 events that are smoothly perturbed with $\sigma_p > 10\%$, and a total of 24 events for which at least one of the three parameters exceeds 0.10.

real data. It is therefore important to note that, although we do not include the effects of binary-lens rotation, parallax, binary source rotation, or finite source size, they all occur in nature. Each of these effects can significantly change the shape of the light curve, and they would preferentially perturb light curves away from rather than toward the point-lens form. Only rarely, however,

will they change the number of caustic crossing light curves. By ignoring these effects, we potentially compute values of the relative rates of smoothly perturbed events to caustic crossing events that are lower than the values that should be observed. Thus, if anything, our results are a lower limit on the relative rates of smoothly perturbed to caustic crossing events.




FIG. 10.—Curve showing the expected rate ratio of smoothly perturbed to caustic crossing events. The y -axis shows the rate $\langle s \rangle$ of events with σ_p greater than the corresponding x -axis value, normalized to $\langle s \rangle$ for caustic crossing events. The horizontal solid line shows $\langle s \rangle$ for caustic crossing events. The horizontal solid line corresponds to equal rates of smoothly perturbed and caustic crossing events. Under our default assumptions, a survey with perfect detection efficiency of events perturbed at the $\sigma_p = 15\%$ level would expect to find as many smoothly perturbed events as caustic crossing events. Shown as dashed horizontal lines are the ratios of known binary events from MACHO (5.7 yr data; Alcock et al. 2000; 4 perturbed events out of 16 binaries), OGLE III 2002–2003 season (Jaroszynski et al. 2004; 7 perturbed candidates out of 24 binary candidates), and OGLE III 2004 season (Jaroszynski et al. 2006; 3 perturbed events out of 25 binary candidates). Implicit in this graph are the default assumptions for our simulation, which are the selection criterion $A_{\text{cut}} = 1.10$, the sampling rate $\Delta t = 0.01\theta_E$, the distribution function $P(q, a)$ uniform in q and $\log(a)$, and the assumption of no blending. Figs. 11–14 show how this graph is affected when these are changed.

A binary source or an extended source would have a more complicated form, but we assume these complications to be negligible. Only rarely is it necessary to invoke a binary source model to fit a binary-lens event, and finite source sizes only significantly affect the shape of caustic crossings, without changing whether the event is caustic crossing or not.

Furthermore, we considered only the simplest possible model of uncertainty, with uncertainty in flux space equal to a constant fraction of the baseline flux. Real uncertainties are much more complicated, but our results should not depend greatly on the model of uncertainty. This is because, for events of low magnification, the flux of the source stays fairly constant, so the uncertainty will remain about the same whether constant in flux space, constant in magnitude space, or something more complicated. Perturbed high magnification events, meanwhile, tend to be significantly perturbed, so they too will not depend strongly on the uncertainty. In particular, since we assume that caustic crossing events are detected with perfect efficiency, their detection rate depends in no way on the uncertainty model. Therefore, as long as a representative uncertainty value can be used to describe the photometric precision of a survey, our model, although simple, should make a reasonable comparison.

Beyond those simplifications, our simulation involved several assumptions, in the form of default parameter values, that were better than a realistic modern survey. Could one of these assumptions cause our simulations to produce unrealistic rate ratios? Here we consider variations of the four most significant assumptions by changing certain parameter values. Figure 10, showing the rate ratio versus the cutoff parameter value, as modified by these variations, is followed by a separate figure for each one. These variations are event detection cutoff A_{cut} (Fig. 11), sampling rate




FIG. 11.—Event rate ratio for smoothly perturbed to caustic crossing events, but considering a variety of event detection thresholds. Starting with the top curve on the left and proceeding down are values of A_{cut} of 1.1, 1.2, 1.34, 1.5, 2, 3, 5, and 10. The line for $A_{\text{cut}} = 1.1$ shows the same data as Fig. 10.

(Fig. 12), binary population distribution function (Fig. 13), and blending parameter (Fig. 14).

Decreasing the sampling rate (Fig. 12) and considering a different distribution of binary lenses from log-uniform (Fig. 13) will affect both the number of caustic crossing light curves detected as well as the number of perturbed non-caustic crossing light curves. Since we can only consider ratios of event types detected rather than absolute numbers, to evaluate the effects of these changes we must divide by the new rate of perturbed events by the new rate of caustic crossing events.

Increasing the event detection threshold A_{min} (Fig. 11) and decreasing the blending parameter (Fig. 14), on the other hand, selectively affect non-caustic crossing light curves. These both generally tend to weaken perturbations and decrease σ_p for a given light curve. There are, however, some exceptions, such as a light curve that is more similar to a point lens on its wings than at its peak. Such light curves will actually produce worse least-squares fits when the wings are cut off by blending or increasing




FIG. 12.—Event rate ratio for smoothly perturbed to caustic crossing events, but considering a variety of light curve sampling rates. The curve for the default rate of $\Delta t = 0.01\theta_E$ (showing the same data as Fig. 10) is on bottom on the left of the plot, and each curve going up increases Δt by a factor of 2.




FIG. 13.—Event rate ratio of smoothly perturbed to caustic crossing events, considering eight different binary-lens distribution functions $P(q, a)$. There are two distributions of q considered, uniform and log-uniform, and four different distributions of a considered, log-uniform, increasing power law ($a^{0.4}$), decreasing power law ($a^{-0.4}$), and Gaussian in log space [centered at $\log(a) = 0$ with a standard deviation of 1 dex]. Each of the eight $P(q, a)$ is a product of two of these. The curves are largely indistinguishable, as shown in the bottom panel, which is the residuals of the top panel divided by the default curve of uniform in q and Gaussian in $\log(a)$. The lowermost curve in this panel represents the distribution uniform in q and Gaussian in $\log(a)$.

the detection threshold, thus potentially changing their classification from point-lens-like to smoothly perturbed.

Figure 11 shows various event detection thresholds. The value $A_{\text{cut}} = 1.34$ is often used because for point-lens light curves it corresponds to $b = 1$. Other values to consider are $A_{\text{cut}} = 1.46$, corresponding to $b = 0.75$, and $A_{\text{cut}} = 2.03$, corresponding to




FIG. 14.—Event rate ratio of smoothly perturbed to caustic crossing events, considering a variety of blending parameters. Each curve corresponds to a different value of f as shown in the plot, varying from $f = 1.0$ to 0.125 , with consecutive values separated by a factor of 0.707 . Two curves, corresponding to $f = 1.0$ and 0.354 , are shown with thicker lines for contrast. Although there is some overlap at small values of σ_p , for $\sigma_p > 0.1$, decreasing f (increasing the blending) systematically decreases the event rate ratio, as expected.






FIG. 15.—Histograms showing the cumulative distribution function for the three parameters for nonrepeating, non-caustic crossing light curves. For any particular value of a parameter, the graph shows the value $\langle s \rangle$ for all light curves with at least the given value. The data in the top panel are identical to Fig. 10.

$b = 0.3$. Only at very high thresholds, well above these values, does the event rate ratio become significantly different for $\sigma_p = 10\%$.

Figure 12 shows what happens with a varying sampling rate. Again, at $\sigma_p = 10\%$, the ratio changes very little for several factors of 2 away from our default assumption.




FIG. 16.—Scatter plot showing the relative distribution of the best-fit parameter σ_P and the asymmetry parameter k . Horizontal and vertical lines denote parameter values of 0.10, and along the diagonal line the two parameters are equal. Note that there are a significant number of light curves for which one parameter is greater than 0.10 but not the other. Sample light curves are shown to represent three different places in the parameter space: top are four curves with significant asymmetry but good point-lens fits; right are four symmetric curves without good point-lens fits; top and right are four curves with high values of both parameters. Beneath each light curve is shown the residual plot for the best point-lens fit. For clarity, this scatter plot shows only 1 in every 10 points compared with the next two plots.

Figure 13 shows variations due to various binary-lensing populations. Since being a binary has no effect for values of a very far from unity, only the distribution in a near unity matters for our rate ratio. So for instance, although a log-uniform distribution in a could not technically extend to infinity, for our simulation it does not matter where the lower and upper cutoffs are, as long as they include the orders of magnitude near unity. In addition to the log-uniform distribution in a , we test two power-law distributions, with power-law exponents of 0.4 and -0.4 . Generally, these have little to no effect on the rate ratios we predict. We also test a distribution designed to maximize caustic crossings, namely a distribution Gaussian in $\log(a)$, centered at $\log(a) = 0$, with a standard deviation of 1. Even this unrealistic distribution (shown in the bottom curve of Fig. 13) only lowers the rate ratio by a factor of about 0.8.

Finally, Figure 14 shows blending considerations. There is a significant difference in the rates when blendings is taken into account. Even for a particularly low blending parameter of $f = 0.250$ (where f is the fraction of baseline light from the lensed

star), however, there are still expected to be as many non-caustic crossing light curves with $\sigma_P > 10\%$ as caustic crossing curves.

4.5. Additional Ways to Detect Perturbations

In the work presented above we have used the value of σ_P to determine whether a light curve should be viewed as smoothly perturbed from the point-lens form. Other parameters, such as the asymmetry parameter, k , and the multipole parameter, p , have been defined and calculated to provide ways of characterizing deviations from the point-lens form. It is important to note, however, that by using σ_P , we are employing a measure of the deviation of the light curve as a whole. Light curves with short-lived deviations from the point-lens form may not be characterized as smoothly perturbed for values of $\sigma_P \sim 0.1$, even if current observational setups can detect the short-lived deviations at a high level of significance.

In other words, by using σ_P to identify smoothly perturbed light curves, we are taking a conservative approach, since even this approach leaves many smoothly perturbed light curves to be




FIG. 17.—Scatter plot showing the relative distribution of the best-fit parameter σ_p and multipeak parameter p . Horizontal and vertical lines denote parameter values of 0.10, and along the diagonal line the two parameters are equal. Again, there are a significant number of light curves for which one parameter is greater than 0.10 but not the other. Sample light curves are shown to represent three different places in the parameter space: top are four curves with significant secondary peaks but good point-lens fits; right are four curves with only very small secondary peaks, but without good point-lens fits; top and right are four curves with high values of both parameters. Beneath each light curve is shown the residual plot for the best point-lens fit. Single-peaked events have $p = 0$ and so are not represented in this scatter plot.

discovered. Nevertheless, in future work it may be appropriate to extend the definition of smoothly perturbed light curves to include some with small values of σ_p , but large values of k or p .

We therefore consider the distribution of the three parameters σ_p , k , and p that is produced by the default binary distribution in § 4.3. We consider the parameters independently (Fig. 15) and also in relation to one another (Figs. 16, 17, and 18). These figures suggest that it would be worthwhile to augment least-squares fitting by considering other avenues of identifying binary-lens events. If this is done, some events that would otherwise be classified as point-lens-like will instead be classified as smoothly perturbed. The predicted values of $\langle s_{sp} \rangle / \langle s_{cc} \rangle$ would then be even larger.

5. CONCLUSIONS

We have explored and categorized the full range of microlensing light curves produced by binary lenses. The three mutually exclusive categories of binary-lens light curves we have studied are point-lens-like, smoothly perturbed, and caustic crossing. Smoothly perturbed light curves are defined to be those that

exhibit continuous deviations from the point-lens form; this naturally excludes caustic crossings. In this paper, smoothly perturbed light curves have been identified by their failure to be fit by a point-lens model with a least-squares metric. We have also determined whether each light curve we have computed exhibits asymmetry with respect to time reversal, and/or multiple peaks. For the purposes of this paper, we have used an asymmetry parameter and a multipeak parameter only to characterize the deviations from the point-lens form, not to identify smoothly perturbed light curves.

Because we have sampled a wide range of binary separations, we have also covered repeating events (Di Stefano & Mao 1996). Our simulation results are consistent with the analytically calculated rates (Di Stefano & Scalzo 1999b) and predict that repeating events will form a significant part of data sets that are sensitive to deviations at the few percent level. The rate of repeating events increases with the size of the lensing region, and so it increases significantly with improved photometric precision. The sensitivity of existing data sets should be enough for repeating events to constitute a few percent of binary-lens events, so as the number




FIG. 18.—Scatter plot showing the relative distribution of the multipole parameter p and the asymmetry parameter k . Horizontal and vertical lines denote parameter values of 0.10, and along the diagonal line the two parameters are equal. Again, there are a significant number of light curves for which one parameter is greater than 0.10 but not the other. Sample light curves are shown to represent three different places in the parameter space: top are four curves with very small secondary peaks but high asymmetry; right are four symmetric curves with significant secondary peaks; top and right are four curves with high values of both parameters. Beneath each light curve is shown the residual plot for the best point-lens fit. Single-peaked events have $p = 0$ and so are not represented in this scatter plot.

of binary events approaches 100 and sensitivity continues to improve, repeating events become inevitable.

As with repeating events, other types of exotic events that are expected to constitute just a few percent of binary-lens events will also become inevitable with growing data sets.

5.1. The Missing Smoothly Perturbed Light Curves

Our main result relates to the relative number of caustic crossing events. We find that under most assumptions about the binary-lens population and about the observational sampling, there are more smoothly perturbed events than caustic crossing events. Identified binary-lens events are dominated by caustic crossing events, only because smoothly perturbed events have not been exhaustively searched for.

To understand the dependencies of this discrepancy, we have conducted a range of simulations. We find that the ratio is not very sensitive to changes in the binary population, specifically the distribution of binary parameters. The reason for the discrepancy must therefore be related to the observational setup or to the analysis. We also find, however, that the ratio is not very

sensitive to changes in the cutoff magnification or sampling frequency. It is of course sensitive to changes in the photometric uncertainty, but it would have to be much larger than 10% to account for the discrepancy. Note that we have considered only the case in which the fractional photometric uncertainties are constant. On the other hand, the observational uncertainties near peaks in magnification may tend to be smaller. If so, for a given value of the uncertainty at baseline, more deviations from the point-lens form could be observed. This would further increase the fraction of events that are recognized as being smoothly perturbed.

Is it possible that some of the physical effects we have neglected in our simulations could produce more caustic crossing events and fewer smoothly perturbed events? The effects we have neglected include binary rotation, finite source size effects, binary sources, and parallax. All of these effects can produce perturbations from the point-lens form. Only rotation is likely to lead to more caustic crossings, and this is expected to happen only rarely. Therefore, calculations that include all of the expected physical effects will likely produce even larger rate ratios.

Only one physical effect seems to have a significant influence, and that is blending. Blending cannot obscure the wall-like features that mark caustic crossings, but it can smooth out the distinctive features of a smoothly perturbed event, making it more likely that a point-lens fit will be successful. Blending is therefore expected to decrease the ratio $\langle s_{sp} \rangle / \langle s_{cc} \rangle$. Nevertheless, blending would have to be severe, with much less than 12% of the baseline light coming from the lens star, in order for blending to be responsible for the discrepancy between the predictions and the results derived so far.

5.2. Searching for the Missing Events

The discussion above indicates that the lensing programs are certain to have detected smoothly perturbed binary-lens light curves, which were not identified as such. Instead, some smoothly perturbed events may have been identified and published as point-lens light curves. If this were the case for all smoothly perturbed light curves, then the experimentally derived lensing event rate would be correct, but it would be difficult to assess the contributions of binary lenses.

On the other hand, the perturbations of some smooth binary-lens light curves may be so pronounced that they have prevented the events from being identified as lensing candidates at all (Di Stefano & Perna 1997). This was especially likely during the early phase of microlensing studies, when it was important for selection criteria to be conservative, to be certain that the selected events were truly associated with microlensing.

There are several important reasons to attempt to identify all binary-lens events. Beyond the direct considerations of determining the rate of events, the interpretation of the events can help to determine the locations of significant populations of lenses. Already several binary lenses in the direction of the Magellanic Clouds have been located by the caustic crossing events they caused. As these events with measurable lens distances are only a fraction of all events caused by the lens population, a large number of all lenses can be located indirectly. Di Stefano (2000) used this argument to demonstrate the possibility that most of the Magellanic Cloud events were caused by self-lensing. To confirm or refute this argument, we must be able to identify the smoothly perturbed binary-lens light curves.

Furthermore, it is important to identify all binary-lens events to learn about the characteristics of the binary-lens and planet-lens populations. One important strength of microlensing planet searches is that, of all methods of planet and binary detection, microlensing alone can identify the locations and measure some of the properties of planets in distant stellar systems such as external galaxies. To conduct a genuine population study, however, we must be able to identify a wide range of events, and to understand the detection efficiencies for each type of event.

5.3. Future Monitoring Programs

The ability of monitoring programs to detect lensing events has increased dramatically since the first published events. The OGLE team now routinely identifies roughly 500 lensing events per year, compared with a total of 9 events identified in its first two years (Udalski et al. 1994). New monitoring projects will be focused on microlensing, while wide field monitoring conducted by Pan-STARRS (Kaiser et al. 2002) and LSST (Tyson 2002) will identify thousands of events per year caused by lenses in the source galaxies, nearby lenses, and also by MACHOs, should they exist (Di Stefano 2007). The photometric sensitivity of these projects will approach the level of 1%. As illustrated in Figure 10, this means that these programs should be able to identify about 3 times as many smoothly perturbed binary-lens light curves as caustic crossing light curves, using least-squares fitting alone. If, in addition, pronounced asymmetries or correlated residuals can be used to quantify the probability that some light curves with acceptable point-lens fits were caused by binaries, smoothly perturbed light curves may play an even larger role in the data sets.

In order for future monitoring projects to place constraints on the form of dark matter in MACHOs and to study the underlying characteristics of each lens population, including planet lenses, methods to detect the full range of binary-lens and planet light curves must be developed, and the relevant detection efficiencies must be quantified.

R. D. would like to thank Rosalba Perna and Nada Petrovic for conversations. Funded in part by NASA NAG5-10705, PHY05-51164, and a grant for SAO Internal Research and Development.

REFERENCES

- Alcock, C., et al. 2000, *ApJ*, 541, 270
 Asada, H., Kasai, T., & Kasai, M. 2004, *Prog. Theor. Phys.*, 112, 241
 Di Stefano, R. 2000, *ApJ*, 541, 587
 ———. 2007, preprint (arXiv: 0712.3558)
 Di Stefano, R., & Mao, S. 1996, *ApJ*, 457, 93
 Di Stefano, R., & Perna, R. 1997, *ApJ*, 488, 55
 Di Stefano, R., & Scalzo, R. A. 1999a, *ApJ*, 512, 564
 ———. 1999b, *ApJ*, 512, 579
 Dominik, M. 1998, *A&A*, 329, 361
 Duquenois, A., & Mayor, M. 1991, *A&A*, 248, 485
 Einstein, A. 1936, *Science*, 84, 506
 Erdl, H., & Schneider, P. 1993, *A&A*, 268, 453
 Gould, A., & Loeb, A. 1992, *ApJ*, 396, 104
 Griest, K., et al. 1991, *ApJ*, 372, L79
 Jaroszynski, M., et al. 2004, *Acta Astron.*, 54, 103
 ———. 2006, *Acta Astron.*, 56, 307
 Kaiser, N., et al. 2002, *Proc. SPIE*, 4836, 154
 Mao, S., & Paczynski, B. 1991, *ApJ*, 374, L37
 Petters, A. O., Levine, H., & Wambsganss, J. 2001, *Singularity Theory and Gravitational Lensing* (Boston: Birkhauser)
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in FORTRAN* (Cambridge: Cambridge Univ. Press)
 Schneider, P., Ehlers, J., & Falco, E. E. 1992, *Gravitational Lenses* (Heidelberg: Springer)
 Skowron, J., et al. 2007, *Acta Astron.*, 57, 281
 Tyson, J. A., et al. 2002, *Proc. SPIE*, 4836, 10
 Udalski, A., et al. 1994, *Acta Astron.*, 44, 165