

## TODO List

Page 4 **erights**: Consider Bruno's suggestion on simplifying this.

Page 22 **all**: Write this section

Page 22 **all**: Write this section

Page 22 **all**: Write this section

Page 22 **all**: Write this section

# Caja

## Safe active content in sanitized JavaScript

Mark S. Miller      Mike Samuel      Ben Laurie      Ihab Awad      Mike Stay

June 7, 2008

### Abstract

Using Caja, web apps can safely allow scripts in third party content.

The computer industry has only one significant success enabling documents to carry active content safely: scripts in web pages. Normal users regularly browse untrusted sites with JavaScript turned on. Modulo browser bugs and phishing, they mostly remain safe. But even though web apps build on this success, they fail to provide its power. Web apps generally remove scripts from third party content, reducing content to passive data. Examples include webmail, groups, blogs, chat, docs and spreadsheets, wikis, and more; whether from Google, Yahoo, Microsoft, HP, Wikipedia, or others.

Were scripts in an object-capability language, web apps could provide active content safely, simply, and flexibly. Surprisingly, this is possible within existing web standards. Caja represents our discovery that a subset of JavaScript is an object-capability language.

## 1 Introduction

An *object-capability language* is essentially a memory-safe object language with encapsulation, with additional restrictions that *protect the outside world from the objects*.<sup>1</sup> In a memory-safe object language such as JavaScript, object A can only invoke object B if A has a reference to B. If A already has references to

B and C, A can invoke B passing C as an argument, giving B access to C. Memory-safe object languages with encapsulation, such as Java, *protect objects from the outside world*. The clients of an encapsulated object can make requests using its public interface, but how an object reacts to a request is up to the object. An encapsulated object can ensure that the only way to invoke its code or change its state is through its public interface.

In an object-capability language, an object can only cause effects outside itself by using the references it holds to other objects. Objects have no powerful references by default, and are granted new references only by normal message passing rules. Object references thereby become the sole representation of rights to affect the world, and normal message passing (method invocation) is the only rights transfer mechanism. An object can be denied authority simply by not giving it those references which would provide that authority.

The browser sandbox already mostly protects the world outside the browser from scripts running on web pages. A great virtue of JavaScript is that many people successfully program in it casually, without first learning the language in any depth. Caja<sup>2</sup> is a subset of JavaScript we designed to make as little impact as possible on regular JavaScript programming, while still providing object-capability security. The subset is enforced by a static verifier and the insertion of runtime checks into the code. In this section, we provide a brief inaccurate overview of the differences

<sup>1</sup> Thanks to Mark Lillibridge for this formulation.

<sup>2</sup> *Caja*, pronounced “KA-hah”, is Spanish for “box”. With Caja, capabilities attenuate JavaScript authority.

```

function F(x) { this.x_ = x; }
F.prototype.getX = function() {
  return this.x_;
};
F.make = function(x) {
  return new F(x);
};
function test() {
  return new F(3).getX() === 3;
}

```

Figure 1: Caja Functions. `F` is a *constructor*. It can only be initialized and used with `new` and `instanceof`. The function `F.prototype.getX` is a *method*. It can only be called as a method. `F.make` and `test` are *simple functions*. They are not restricted.

between Caja and JavaScript suitable for the casual JavaScript programmer. The rest of this document then accurately goes into more depth.

**Forbidden names.** In Firefox, access to the “`__proto__`” property of an object would grant the authority to create more objects like it, which violates the principle of least authority. Therefore, Caja rejects all names ending with “`__`” (double underscore). This also gives the Caja implementation a place to store its book-keeping information where it is invisible to the Caja programmer.

**Frozen objects.** In JavaScript, all objects are mutable, so passing the same reference to two objects automatically grants them the authority to communicate, which is undesirable. Therefore, Caja adds the ability to *freeze* an object. If an object is frozen, an attempt to set, add, or delete its properties will throw an exception instead. Functions and prototypes are implicitly frozen. In addition, the Caja programmer can explicitly freeze objects to prevent their direct modification. All objects in the default global environment are *immutable*, or transitively frozen.

**No shared global environment.** Caja code is compiled into units of isolation called *modules*; in practice, these are JavaScript functions. A *container* loads the modules and grants them authority by means of references passed as arguments to the module functions. These arguments are called *imports*.

A module that displays the local weather on a webpage should not be able *a priori* to communicate with a module that has access to your bank account. Therefore, each module has its own global environment which inherits from the default global environment, isolating them from each other. On the other hand, a container can allow two chosen modules to communicate by passing a reference to a common mutable object to each module.

**Protected names.** The state of an object that is not part of its public interface should not be read or changed by the outside world. Javascript supports private variables via closures, but this pattern incurs a large memory overhead. Also, using this as the sole encapsulation mechanism for object patterns conflicts with existing JavaScript programming practice. Therefore, Caja enforces the convention that property names ending in “`_`” (single underscore) are protected instance variables. Such names can only appear as property names of “`this`”. As with Smalltalk instance variables or **protected** instance variables in C++, these protected instance variables are visible up and down the inheritance chain within an object, but are not visible outside an object.

**No “this” stealing.** The single-underscore rule above only protects an object’s state from its clients if its clients cannot add methods to it which alias its “`this`”. For example, consider the following constructor:

```

function Cell(value) {
  this.x_ = "secret";
  this.value = value;
}

```

At first glance, there seems to be no way for “x\_” to leak. However, the expression

```
(new Cell(
  function (){
    return this.x_;
  })).value()
```

evaluates to the secret value. Therefore, Caja divides functions into three categories: *simple functions* are those which do not mention “this”. They are first-class and can be used without further restriction. *Constructors* are named functions which mention “this”. *Methods* are anonymous function which mention “this”.

Caja supports the conventional class-like usage of constructors and methods (Figure 1), but prohibits certain other dangerous usage patterns. A constructor can only be called *as a constructor* using `new`, or by a directly derived constructor to initialize a derived instance. An object’s methods can only be called as methods of *that* object, even when calling the method reflectively using `call`, `apply`, or `bind`.

**Sharp knives removed.** The semantics of “with” are even stranger than those of “this”. For example,

```
var o = { x: 4, f: 2 };
with(o) {
  function f() { }
  alert(f); // This displays 2 !
  var x = 3;
}
// Now o.x === 3 !
```

Caja contains no “with” or “eval”. Caja includes a safe JSON library to support the most common use of `eval`—deserializing object literals—and a safe `caja.cajitaEval` for evaluating code in the *Cajita* subset of Caja.

Cajita, which means “little box” in Spanish, is essentially the subset of Caja without “this”. It is far easier to analyze and rewrite than Caja—so

much so that the client-side rewriter `cajitaEval` is feasible—but requires a much different programming style than most JavaScript programmers are accustomed to.

Just as Caja modules receive their authority from the container, `cajitaEval` takes as a parameter an object `imports`. Any free variable appearing in the code passed to `cajitaEval` is considered to be the name of a property of `imports`.

Hopefully, this is all the casual Caja programmer needs to know to get started. Section 2 is a partisan history of access control on the web, in order to motivate the problems Caja addresses. It may safely be skipped. Section 3 explains the problems faced when securing JavaScript, many of which involve the use of “this”.

We then present Caja in two stages. Section 4 presents *Cajita*, the subset of Caja without “this”. For new code, *Cajita* is a reasonably expressive language resembling an object-oriented Scheme. Section 5 then presents the remainder of the Caja language beyond *Cajita*. Caja adds back enough of JavaScript for most old habits and old code to port pleasantly and painlessly. Caja and *Cajita* interoperate without problems. Section 6 briefly surveys related work.

## 2 Identity-centric Epicycles

**erights:** Consider Bruno’s suggestion on simplifying this.

When a document contains live interactive programs, we say it contains *active content*. The computer industry has spent over a billion dollars in failed attempts to support active content. But the success of web apps—themselves a form of active content—demonstrates that this dream was worth pursuing. Unfortunately, web developers today face a maze of complex security mechanisms that have, so far, prevented web apps themselves from supporting active content. To navigate our way out of this maze, we must first see how we got here.

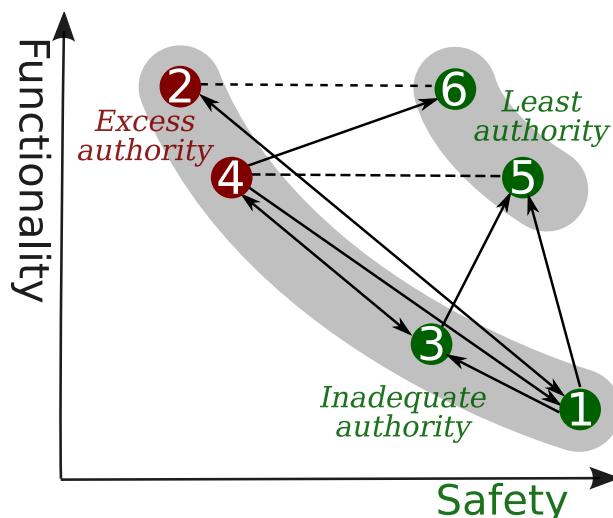


Figure 2: The Evolving Authority of Active Content. Identity-centric access controls have led to thrashing between lost functionality and lost safety. To have both, we need to provide *least authority*: adequate authority for desired functionality without excess authority which invites abuse.

Today’s desktop operating systems all use some form of identity-centric access control [4], in which an installed application runs *as* its user, and so is entrusted with all its user’s authority. Such an application can provide its user all the functionality modern operating systems support, but at the price of being able to do anything its user may do. We depict this situation at 2 on Figure 2. When you run Solitaire, it can delete all your files while playing within the rules of your system, without exploiting any bugs. (For the remainder of this document, we will ignore hazards due to implementation bugs, and explain only hazards due to architectural choices.)

At first, the documents handled by applications were safe passive data 1. Applications first supported active content by running scripts in documents with all of their user’s authority 1→2. Excess authority invites abuse. Simply “reading” a malicious document would allow it to delete all your files. In reaction, installed office applications now




Figure 3: Only Bad Choices. When documents contain scripts, users can disable themselves from getting any work done 1 or enable scripts to destroy all their other work 2.

encourage users to disable scripts (Figure 3) reducing content back to passive data 2→1. The failures of excess authority shown on the upper left thus led to the failures of inadequate authority shown on the lower right.

The web browser is itself an installed application that runs scripts in two contexts. Browser extensions run with all the user’s authority 2. Scripts in web pages run sandboxed, with no authority to the user’s local files. The browser’s *same origin policy*, another layer of identity-centric control [14], provides scripts with the authority to communicate with their site of origin 1→3. Regarding both decisions, the user is helpless. The user has no practical way to grant a script the authority to edit one of the user’s local files, nor can the user deny a script the ability to call home. So long as the user’s valuable assets were local, this model successfully protected the user.

Web apps leverage this success. To the browser, the page on which a web app resides is a document, and the web app itself is simply active content within that document. But to the user, a web app is an application managing yet other documents on the user’s behalf. For example, when the user interacts with webmail, the documents of interest are email messages. Likewise for groups, blogs, chat, docs and spreadsheets, wikis, and more. Let us refer to the documents managed by web apps as *passages*, to distinguish them from the web pages on which they appear.

Since the user can neither grant a web app access to local files nor deny it the ability to call home, the only place a web app could store these passages is on its site of origin. The browser security model protected

the user’s local files from being harmed *or used*. As users shift to using web apps, the assets they value come to be the passages stored at these various origin sites.

To protect their user’s remote passages, web apps employed yet another layer of identity-centric controls, relying on cookies or other forms of authentication to identify their user. But when scripts within these passages ran, they would run within the web page containing the web app serving them, and were thereby authorized to do anything their web app could do on behalf of its user ④. For example, if a webmail application allowed HTML email messages to carry scripts, simply “reading” an incoming email message would allow it to delete your inbox. The ③→④ transition is not a technical change, but a change in where the user’s value resides, and thus a change in the user’s risks. By this dynamic, failures of inadequate authority led to failures of excess authority.

To protect against malicious passages, some web apps do safely provide active content using *iframes*—effectively nested web pages—at the cost of isolating themselves from this content ④→③ [14]. Most web apps *sanitize* HTML content by removing all scripts, reducing content again to passive data ④→①. Existing HTML sanitizers disinfect the patient but leave a corpse. This recapitulates the loss of active content in installed office applications. Some proposals would address these next incremental problems by adding yet another identity-centric epicycle. Can we do better?

If we could start over again, we could use an authorization-centric model such as object-capabilities [1]. The object-capability alternative naturally supports POLA, the principle of least authority, shown in the upper right of Figure 2. An object in an object-capability language can only cause effects by invoking the public interfaces of objects it can reach. An invocation provides references to other objects as arguments, providing the invoked object the least authority needed to carry out these requests [8]. Within these rules, active content would run with exactly the authority explicitly provided by its containing document. Surprisingly, we can gain these benefits simply by applying a milder, non-lethal




Figure 4: Ptolemy’s epicycles. Ptolemy attempted to model the motion of the heavenly bodies using only circles. With each discovery that the model didn’t fit, yet another layer of circle was added to adjust. By contrast, Kepler’s ellipses fit the problem directly, with no need for endless additional layers.

---

sanitizer.

Experience with Java, Scheme, OCaml, Pict, Perl and others demonstrates that existing memory safe languages often already contain an expressive object-capability subset [7, 9, 11, 5, 6, respectively]. We refer to the object-capability subset of JavaScript as *Caja*. This subset is still a general purpose object programming language which JavaScript programmers should find familiar, pleasant, expressive, and easy to learn and use.

```

function Counter() {
  var count = 0;
  return caja.freeze({
    toString: function() {
      return "<counter: " + count + ">";
    },
    incr: function() {
      return count += 1;
    },
    decr: function() {
      return count -= 1;
    }
  });
}

```

Figure 5: A Cajita Counter. Each call to `Counter()` produces a new counter object. Access to a counter provides the authority to read, invoke, or enumerate its properties, all of which are simple functions serving the role of methods. Caja functions are implicitly frozen; the returned object is explicitly frozen; and the instance-state of the object—the count variable—is accessible only as encapsulated state captured by these pseudo-methods. A counter object as a whole, as well as each of its pseudo-methods, are thus proper protected capabilities. Someone with access only to a counter’s `incr` function can increment *that* counter and observe the result, but not do anything else.

---

Some web apps could use the Caja sanitizer to allow active content in their passages ①→⑤. Other web apps could use Caja to overcome the limits of iframes ③→⑤. Browser extensions, which run with their user’s full authority, could make a *powerbox* available to scripts in pages [13, 12, 10, 3]. A web app, on detecting the presence of a powerbox, could offer to edit a local file chosen by the user ④→⑥.

### 3 Subsetting JavaScript

Our starting point is JavaScript as documented in the third edition of the EcmaScript 262 standard [2];

hereafter *ES3*<sup>3</sup>. ES3 code is passed to a Java program known as the the Caja sanitizer, or “cajoler<sup>4</sup>”. The first set of restrictions is enforced by a static verifier. These restrictions mostly involve the use of trailing underscores, where the keyword “`this`” may appear, and the class definition pattern. The second set of restrictions is imposed at runtime. After statically verifying the code, the cajoler rewrites the code, inserting dynamic checks throughout. These involve restricting access to private members, forbidding modification of frozen objects, and so forth. The actual logic of the runtime checks is contained in a runtime library, `caja.js`, that must be loaded by the JavaScript interpreter before loading a Caja module.

The remainder of this document explains the differences between Caja—the JavaScript subset accepted by the Caja sanitizer—and ES3. Other documents will explain the interface between cajoled and uncajoled JavaScript, and Caja’s sanitization of the remaining elements of active web content: HTML, CSS, and the DOM and other APIs provided by browsers to JavaScript. We refer collectively to the subset of these accepted by the Caja sanitizer as *Caja web content*, and to the sanitizer’s corresponding output as *Cajoled web content*.

#### 3.1 The OS analogy

A web app (or any other JavaScript-based embedding application framework) can be written partially in JavaScript and partially in Caja. The web app must load the Caja runtime library, which is written in JavaScript. All untrusted scripts must be provided as Caja source code, to be statically verified and cajoled by the Caja sanitizer. The sanitizer’s output is either included directly in the containing web page or loaded by the Caja runtime.

A loose analogy with machine and operating system architecture may help explain the relationships. In the analogy, the full JavaScript language serves the role of the machine’s full instruction set. JavaScript’s global environment serves the role of physical memory addresses. The I/O-capable objects provided to

---

<sup>3</sup> ES3 is approximately a bit more than JavaScript 1.4 and a bit less than JavaScript 1.5.

<sup>4</sup> We thank Pat Patterson for this term.

JavaScript by a hosting environment, such as the DOM objects provided by the browser, serve the role of devices.

**User-mode.** By a combination of static and dynamic checks, the Caja sanitizer allows only a safe “user-mode” subset of JavaScript. As with user-mode instructions, this subset can compute any computable function, but cannot cause external effects nor sense the outside world.

**Address mapping.** A package of Caja source code to be cajoled together defines a Caja module. All code within the same module shares a global environment, but distinct modules see disjoint global environments. The Caja sanitizer implements this by rewriting free variables as properties of a container-provided “imports” object.

**Context switching.** When Caja object A has a reference to Caja object B, this should enable A to invoke B’s public interface but not access B’s internal state. A and B should both be able to defend their integrity from the other’s possible misbehavior.

**System calls, device drivers.** When a Caja object A invokes an object B written directly in JavaScript, the operations provided by B serve the role of system calls. Caja protects B from A, but A is fully vulnerable to B. When B is a safe wrapper around one of the host’s device-like objects, such as a DOM node, B also serves as a device driver.

A “system call” corresponds to a Caja object invoking a JavaScript object. A web app that is written entirely in JavaScript and provides many services to its Caja objects directly would be like a monolithic kernel. For compatibility with existing JavaScript apps, we support this usage pattern but we don’t recommend it. By analogy with kernel code at the boundary with untrusted code, such JavaScript code needs to maintain delicate invariants that it is easy to get wrong.

The other extreme is analogous to a micro-kernel. The minimal necessary JavaScript code would be the

```
function Point(x, y) {
  return caja.freeze({
    toString: function() {
      return "<" + x + "," + y + ">";
    },
    getX: function() { return x; },
    getY: function() { return y; },
  });
}

var ptA = Point(3, 5);
var ptB = Point(4, 7);
```

Figure 6: A Cajita Point. As a baseline, we first express this simple example in Cajita with no support for inheritance. Other elaborations will show how to support inheritance and various styles of definition in both Cajita and full Caja.

---

app-neutral Caja runtime itself, and a small app-dependent powerbox providing device drivers and initialization. All other services should be Caja objects to be invoked by other Caja objects. Most of the logic of a web app should be structured as such Caja-based services.

### 3.2 JavaScript specific problems

Most of the above remarks would apply equally well were we starting from various other base languages. There are additional issues peculiar to JavaScript that we must deal with. Many of these issues are also software engineering hazards for which JavaScript programmers have developed defensive programming conventions. Where possible, Caja copes with these issues by adapting and enforcing these existing conventions.

**Unconstrained properties.** JavaScript objects contain *properties*, i.e., named fields holding references to other objects. JavaScript specifies that some properties are constrained to be *Protected*, *ReadOnly*, *DontEnum*, or *DontDelete*. Such constraints would help an object protect



```

function PointMixin(self, x, y) {
  self.toString = function() {
    return "<" + self.getX() + "," +
      self.getY() + ">";
  };
  self.getX = function() { return x; };
  self.getY = function() { return y; };
  return self;
}
function Point(x, y) {
  return caja.freeze(PointMixin({}, x, y));
}

```

Figure 7: Cajita Inheritance. In the Cajita inheritance pattern, the equivalent of a non-final class is a function ending with “Mixin” with `self` as its first parameter. The method-like functions can use `self` analogously to the use of `this` in full Caja, in order to refer to the overall object being defined. If the class is non-abstract, it should also have a pseudo-constructor function such as `Point` for making direct instances. This “\*Mixin” function should only be called by these pseudo-constructor functions, such as `WobblyPoint` in Figure 8.

---

itself from its clients, but JavaScript provides no way to express these constraints in the language. Instead, any user-defined object in JavaScript is freely mutated by any other object with access to it.

**Global environment.** All JavaScript code executing within the same JavaScript engine (such as a web page or `iframe`) implicitly share access to the same global environment. Therefore, in JavaScript, objects cannot be isolated from each other.

**Implicit mutable state.** Some base JavaScript objects, such as `Array.prototype`, are implicitly reachable even without naming any global variable names. Even after global environment problems are fixed, the mutability of these objects would prevent isolation.

```

function WobblyPointMixin(self) {
  var super = caja.snapshot(self);
  self.getX = function() {
    return Math.random() + super.getX();
  };
  return self;
}
function WobblyPoint(x, y) {
  var self = PointMixin({}, x, y);
  self = WobblyPointMixin(self);
  return caja.freeze(self);
}

```

Figure 8: Cajita `WobblyPointMixin`. The equivalent of a non-final subclass is a “\*Mixin” function with `self` as its first parameter, where the body calls `caja.snapshot` to make a frozen copy of the partially initialized `self` at that moment, to serve as the conventional `super` for the other functions defined within this scope.

---

**Lack of encapsulation.** To support the “context switching” criterion explained in section 3.1, objects need to be able to encapsulate their private state. JavaScript does provide one such mechanism: lexical variables captured by nested functions. For example, in the following code, the variable `secret` cannot leak or be changed:

```

function makePointFunction(secret) {
  return function(value) {
    return value === secret;
  }
}

```

However, using this as the sole encapsulation mechanism for object patterns conflicts with existing JavaScript programming practice.

**“this” what?** JavaScript’s rules for binding “this” depend on whether a function is invoked by *construction*, by *method call*, by *function call*, or by *reflection*. If a function written to be called in one way is instead called in another way, its

“`this`” might be rebound to a different object or even to the global environment.

**Foreign for/in loops.** Caja has a stated goal of supporting as much legacy code as possible, where it is safe to do so. Nearly all JavaScript libraries use JavaScript’s `for/in` loop to enumerate the names of all an object’s properties, whether inherited or not<sup>5</sup>. As a result, the properties used by the internals of the Caja runtime library, which are hidden to Caja code, need to be skipped by the loop body. Every JavaScript coding style invents its own defensive pattern of additional tests to skip unwanted property names.

Though not part of Caja itself, the Caja distribution includes an “innocent code” transformer that parses JavaScript and surrounds the body of all `for/in` loops with a check that skips properties internal to the Caja library—*i.e.* properties ending in “`__`”, triple underscore.

**Weak static analysis.** Although Caja is less dynamic than JavaScript, we still assume that it is impractical to perform any interesting analysis, such as type inference, both statically and safely. As a result, Caja’s static restrictions can only enforce simple syntactic rules. Remaining restrictions must be enforced by runtime checks.

**Fast path.** For the micro-kernel approach to be attractive, Caja’s extra runtime checks must not cost too much. Frequent operations, such as property access using “`.`” must run close to full speed.

**Uncontrolled language growth.** The ES3 spec allows one to add new dangerous properties to core objects while claiming ES3 compatibility. JavaScript language implementors, platform providers, and standards committees make use of this freedom with unpredictable results. For example, some JavaScript implementations have added dangerous properties, like `eval`, to core objects, like `Object.prototype`. A safe subset

<sup>5</sup> ... unless the property is `DontEnum`, but the JavaScript programmer has no way to express that in his own code.

must deny access to these additional unknown properties. But since these new properties are often `DontEnum`, there isn’t even a reliable way to detect them.

**Browser compatibility.** Web content must work on widely deployed browsers whether or not these browsers strictly conform to the relevant standards. At the time of this writing, the plausible baseline platform is Yahoo!’s list of A-Grade browsers `/citeYahoo:AGrade`. Fortunately, these browsers do conform closely to ES3. Later versions of Caja may specify larger subsets of ES3.

**Multiple worlds.** As with many languages, each instantiation of a JavaScript language world creates a set of primordial objects (like `Object.prototype`) that are global to that world. Unlike other languages, JavaScript is built to support multiple interacting worlds. For example, in the browser environment, a new JavaScript world is created for each `iframe`. An object from one `iframe` can hold a direct reference to an object from another `iframe` of the same origin. This leads to some surprises. Even if `x` holds an array, `x instanceof Array` may evaluate to false because `x` is an instance of the `Array` from a different JavaScript world.

**Silent errors.** In JavaScript, various operations, such as setting a `ReadOnly` property, fail silently rather than throwing an error. Program logic then proceeds along normal control flow paths premised on the assumption that these operations succeeded, leading to inconsistency. To program defensively in the face of this hazard, every assignment would be followed by a “*did it really happen?*” test. This would render programs unreadable and unmaintainable. Where practical, Caja deviates from standard JavaScript by throwing an exception rather than failing silently.

**Object detection.** In JavaScript, reading a non-existent property returns `undefined` rather than throwing an exception. The JavaScript *object*

*detection*<sup>6</sup> pattern relies on this behavior. Since, in this case, the program naturally notices the problem anyway, Caja does not turn this case into a thrown exception.

The above point about “Silent errors” is another reason to avoid the monolithic kernel approach. Web apps in uncajoled JavaScript are vulnerable to any malicious active content that finds a way to provoke a silent error and exploit the resulting inconsistency.

### 3.3 A fail-stop subset

There are four ways that the semantics of cajoled code may differ from those of uncajoled code. First, it may fail to pass the static verifier; second, it may throw an exception at runtime; third, it may return `undefined` when trying to read a protected variable from outside the encapsulating object; and fourth, by hitting one of a few rare corner cases where the semantics have to differ in order to preserve the security properties. We call the last set of deviations “gotchas”, and detail them in sections 4.6 and 5.4. Since (in all cases but the gotchas) the semantics differ only when there is an error, or “failure”, Caja is a *fail-stop subset*<sup>7</sup> of ES3.

A *Caja-compliant* JavaScript program is one which

1. is statically accepted by the Caja sanitizer,
2. does not provoke Caja-induced failures when run cajoled, and
3. avoids these gotchas.

Such a program should have the same semantics whether run cajoled or not.

## 4 Cajita Specification

Before describing Caja and all the Rube-Goldbergian complexity of the semantics of “`this`”, we’ll describe the subset of Caja without “`this`”—a perfectly reasonable and expressive programming language. Caja supports “`this`” in order to ease the porting of old

code. For new code, we recommend sticking to Cajita<sup>8</sup>.

The Caja runtime will provide a safe *eval* operation, `caja.cajitaEval`. For this operation to accept Caja code, the Caja sanitizer would need to be written in JavaScript and included in the Caja download. To minimize download size, `caja.cajitaEval` will instead accept only Cajita code.

To explain the restrictions Cajita imposes, we need some definitions.

**Record.** An object whose prototype’s “`constructor`” property is `Object`, i.e., under normal conditions, an object inheriting directly from `Object.prototype`. Records are normally created using the `{...}` syntax.

**Array.** An object whose prototype’s “`constructor`” property is `Array`, i.e., under normal conditions, an object inheriting directly from `Array.prototype`. Arrays are normally created using the `[...]` syntax.

**JSON Container.** A record or array. These are the non-primitive objects that can be directly expressed in JSON syntax. Note: whenever the word “container” appears unqualified, we are referring to the *module container*, not a JSON container.

**Function.prototype.bind.** Cajita and Caja add the `bind` method to all functions, defined in equation 5 of figure 9. The popular Prototype library, ES3.1, and ES4 all define `bind` in this way.

**Invocation.** A function can be invoked

- as a function (`foo(a...)`),
- as a method (`foo.m(a...)`),
- as a constructor (`new Foo(a...)`), or
- reflectively (by calling its `call`, `apply`, or `bind` methods).

<sup>6</sup> See <http://www.quirksmode.org/js/support.html>.

<sup>7</sup> We thank Dan Rabin for this formulation.

<sup>8</sup> The design of Cajita was inspired by Doug Crockford’s *ADsafe*.

**Simple functions.** A function whose body does not mention “`this`” is a *simple function*. A simple function can be either named or anonymous. Simple functions are first-class—they can be stored in variables and passed around freely, just like any other value.

**Frozen.** If an object is *frozen*, any attempt to directly assign to its properties, add new properties to it, or delete its properties causes an exception to be thrown. Frozen is a shallow restriction: Frozen objects can retain and provide non-frozen objects. (Imagine a frozen surface covering a liquid lake.) In Cajita and Caja, functions are implicitly frozen once they’ve been initialized. The Caja runtime library additionally provides an explicit operation for freezing JSON containers: “`caja.freeze(obj)`”.

**Immutable.** If an object is *immutable*, then it is frozen, and all objects it has access to are themselves immutable. Shared access to an immutable object does not provide a communication channel, and so does not endanger isolation. With the exception of `Math.random` and `Date`, all objects that are globally or implicitly accessible to all Caja programs are immutable. We discuss these exceptions in section 4.5.

## 4.1 Cajita regularities

The regularities in Figure 9 apply when calling simple functions, whether the calling code is in Cajita or Caja. When calling other functions, only the weaker *Caja regularities* shown in Figure 12 apply. The regularities in both sections are often stronger than ES3, but are all within a fail-stop subset of ES3.

- Equation (1) of Figure 9 states that the `new` keyword does not change the meaning of calling a simple function. This holds only for simple functions that explicitly `return` a value. As in uncajoled JavaScript, if a simple function instead implicitly returns, it will return `undefined` when called without `new`, but will return a useless object when called with `new`.

$$F(\dots) \equiv \text{new } F(\dots) \quad (1)$$

$$\equiv F.\text{call}(v, \dots) \quad (2)$$

$$\equiv F.\text{apply}(v, [\dots]) \quad (3)$$

$$\equiv F.\text{bind}(v)(\dots) \quad (4)$$

$$F(\dots_1, \dots_2) \equiv F.\text{bind}(v, \dots_1)(\dots_2) \quad (5)$$

$$x.m \equiv \text{true} \ \&\& \ x.m \quad (6)$$

$$(x.m)(\dots) \equiv (\text{true} \ \&\& \ x.m)(\dots) \quad (7)$$

$$\{\dots\} \equiv (\text{function}())\{\dots\}() \quad (8)$$

$$(\dots) \equiv (\text{function}()\{\text{return } \dots\})() \quad (9)$$

Figure 9: Cajita Regularities. Given that `F` is a simple function, `x.m` holds a simple function, and `v` is an expression with no effects and stable value (such as a variable reference), then most of these equivalences hold in Caja as well as Cajita. Equation (8) holds in general only in Cajita. See section 4.1 for further qualifying conditions.

- 
- Equation (9) holds in Cajita and Caja when the left-hand side does not mention `arguments` freely.
  - Equation (8) holds only in Cajita, and only when the left hand side does not contain a free break, a free continue, or a return statement.
  - When calling the `call`, `apply`, or `bind` method of a simple function, the first argument is ignored: a simple function cannot contain the keyword `this` and thus has no way to refer to that argument.
  - The `apply` method differs from `call` only in packaging all arguments together into a list.
  - A single-argument `bind` of a simple function returns a function with equivalent invocation behavior—a function that behaves the same, whether called as a function, as a constructor, as a method, or reflectively.
  - When `bind` has additional arguments, it returns

a new function representing  $F$  carried over these additional arguments.

- In JavaScript, when the left operand of an `&&` expression evaluates to a “truthy” value—that is, any value  $x$  such that `Boolean(x) === true`—the `&&` expression as a whole evaluates to the value of its right operand. Therefore, you might expect Equation (6) to hold in general. The next item sheds light on why it does not hold in JavaScript when the value of the right operand is a non-simple function.
- In JavaScript, when the value of a property `x.m` is a function, the expression `(x.m)(...)` binds `this` to `x`, whereas `(true && x.m)(...)` binds `this` to the global scope. In a web browser, the global scope is reified as the object `window`, so the browser calls `m` as a method on `window` instead of on `x`. Fortunately, when `m` is a simple function, these two forms of invocation have the same meaning.

## 4.2 Common static restrictions

Any source code statically accepted by the Caja sanitizer is a *legal Caja program*. A legal Caja program satisfying additional static restrictions is also a *legal Cajita program* and will be accepted by the Cajita sanitizer. A Caja-compliant JavaScript program that is also a legal Cajita program is a *Cajita-compliant JavaScript program*—it will have the same semantics whether uncajoled, cajoled by the Caja sanitizer, or cajoled by the Cajita sanitizer.

The static restrictions immediately below apply to both Caja and Cajita. This is followed by the additional static restrictions specific to Cajita.

**Stable language.** Virtually any input which should be statically rejected by ES3 is forbidden, even if it would be allowed by a target browser or later JavaScript specifications. This includes any use of keywords reserved in ES3. But we reserve the right to include *de-facto extensions* to ES3 as explained below.

```
function Brand() {
  var flag = false, payload = null;

  return caja.freeze({
    seal: function(payloadToSeal) {
      function box() {
        flag = true;
        payload = payloadToSeal;
      }
      box.toString = function() {
        return "(box)";
      };
      return box;
    },
    unseal: function(box) {
      flag = false;
      payload = null;
      try {
        box();
        if (!flag) { throw ...; }
        return payload;
      } finally {
        flag = false;
        payload = null;
      }
    }
  });
}
```

Figure 10: Rights Amplification. Each brand has a `seal` and `unseal` function, acting like a matched encryption and decryption key. Sealing an object returns a sealed `box` that can only be unsealed by the corresponding `unseal` function. The implementation technique shown here is due to Marc Stiegler.

**De-facto extensions.** As we identify widely supported extensions of ES3 that we can accept as input, but still cajole to conforming ES3 on output, we may add these to Caja. For example, we are currently considering allowing backslash as a line continuation character, since this is allowed by virtually all JavaScript implementations and can be trivially cajoled to correct ES3.

**Without “with”.** The “with” keyword is forbidden. Because of the scope confusion it causes, “with” is a widely hated and avoided feature that would be a lot of trouble to support safely.

**Beware unicode.** Cajita and Caja accept unicode characters only in string literals. Some of these create parsing problems on some widely deployed JavaScript platforms. Prohibiting these protects against some character-encoding attacks. We expect to relax this restriction once we know how to do so safely.

**Forbidden names.** An identifier ending with a double underscore is forbidden, either as a variable name or a property name. We reserve the triple underscore for use by the sanitizer’s cajoled output and by the Caja runtime. Firefox reserves the double underscore for itself.

**“new” is ok.** Since Cajita does not have this, constructors, nor prototypes, `new` isn’t needed purely within Cajita. But since Cajita code must interoperate smoothly with Caja and uncajoled JavaScript code, `new` is considered a valid part of Cajita.

**No assignment to imports or declared functions.**

Variables used freely in Caja code refer to properties of the `IMPORTS_` object, and assignment to these properties is statically rejected. Declared function names may not be reassigned. For example, the following code is illegal in Cajita and Caja:

```
function foo()
foo = 3;
```

**No deleting variables.** Allowing variables to be deleted prevents static scope analysis, so Cajita and Caja both prohibit it. Properties of objects, however, may still be deleted.

### 4.3 Cajita-only static restrictions

The following features are present in Caja in order to accommodate old code, rather than to enhance expressiveness. Since Cajita is for new code, in order

```
function Mint() {
  var brand = Brand();
  return function Purse(balance) {
    caja.enforceNat(balance);
    function decr(amount) {
      caja.enforceNat(amount);
      balance =
        caja.enforceNat(balance-amount);
    }
    return caja.freeze({
      getBalance: function() {
        return balance; },
      makePurse: function() {
        return Purse(0); },
      getDecr: function() {
        return brand.seal(decr); },
      deposit: function(amount,src) {
        var box = src.getDecr();
        var decr = brand.unseal(box);
        var newBal =
          caja.enforceNat(balance+amount);
        decr(amount);
        balance = newBal;
      }
    });
  };
}
```

Figure 11: The MintMaker Example. Calling `Mint()` creates a `Purse` function for making purses holding new transferable units of a distinct “currency”. Given two purses of the same currency, one can transfer money between them, but one can’t violate conservation of currency.

---

to minimize the download size of the Cajita sanitizer, as well as to simplify the semantics of Cajita considered on its own, these features are absent from Cajita. Code containing these features is not legal Cajita.

**“this”.** The central difference between Caja and Cajita is that only Caja includes “this”.

**Protected names.** In Caja, an *protected name* is a property name ending in “\_” (a single under-

bar). Such names are used for encapsulation in Caja but are prohibited in Cajita. Cajita’s only encapsulation mechanism is lexical scoping.

**Prototypes.** In Caja and Cajita, if “Foo” is a function name, then static properties of the function can be initialized until the first time the function is used. Cajita prohibits access to Foo’s “prototype” property, and so prevents use of JavaScript’s prototype inheritance within Cajita.

**“instanceof”.** Without “this” and prototypes, Cajita has no need for `instanceof`. Rather, JavaScript’s `typeof` is almost an adequate type discriminator for Cajita. But Cajita still needs a way to distinguish records from arrays. We could allow the conventional `x instanceof Array` expression, but it does not work correctly when `x` is an array from another cajoled JavaScript world, such as another iframe with the same origin. Instead we provide `caja.isArray(x)` as a correct alternative.

**Literal RegExp syntax.** In JavaScript implementations, the literal pattern syntax is often optimized into a static object with mutable state, violating isolation. The Caja sanitizer cajoles the `/pattern/` syntax to `new RegExp("pattern")`. In Cajita, the second form must be written explicitly.

**for/in loops.** Because of the confusing semantics of JavaScript’s `for/in` loops, these are absent from Cajita. Instead, Cajita code should enumerate the properties of `obj` by doing

```
caja.forEach(obj,function(v,k){...});
```

This code will reliably give the same results whether run cajoled or not. It will enumerate only the non-inherited publicly Caja-visible property value / property name associations of `obj`. If `caja.isArray(obj)`, then `k` will enumerate successive indexes into the array.

**Semicolon insertion.** JavaScript will insert semicolons automatically in certain situations involving newlines. Code which parses correctly but

differently if semicolons are automatically inserted is not legal Cajita. For example, due to the newline at the end of the first line, this code

```
x = a
  + b;
```

has two different correct parses: `x = a + b;` and `x = a; +b;`.

Therefore, code which has a newline between `x = a` and `+ b` is not legal Cajita. It should be replaced by `x = a + b;` which is unambiguous.

**Block-breaking scopes.** Cajita variable names are visible only according to the intersection of ES3’s scoping rules and conventional Java-like block-level lexical scoping. This is essentially lexical scoping from the point of introduction with the restriction that a function cannot contain two definitions of the same variable name, even in two separate blocks. If JavaScript scope analysis and conventional block-level lexical scope analysis would disagree on the variable bindings of a given piece of code, then that code is not legal Cajita.

**Coercing equality.** JavaScript’s coercing rules for the “==” and “!=” operators are complex, accident prone, and not even transitive. Cajita only includes the equality operators “===” and “!==”.

## 4.4 Cajita dynamic restrictions

The following restrictions apply to both Caja and Cajita.

**Frozen Functions.** An anonymous simple function is implicitly frozen. A named simple function may be initialized, but is implicitly frozen immediately before its first non-initializing use or escaping occurrence. For example, the assignment to `box.toString` in Figure 10 will succeed, because it occurs before `box` is implicitly frozen by the following `return` statement. Initializing assignments can thus be considered declarative initializations rather than mutations.

*Claim: No Caja program can cause a Caja-observable mutation of a function or of any object Caja considers frozen.*

## 4.5 Modules

The output of the Caja sanitizer consists almost entirely of a JavaScript function called a *module*. The *bound* variables of a module are those that appear as the names of functions declared in the Caja code or in a `var` declaration. The *free* variables are the variables that are not bound.

One of the arguments of every module function is named `IMPORTS___`, and the Caja sanitizer rewrites all free variables to be properties of `IMPORTS___`. For example, the cajoling process rewrites

```
var list = new Foo(6);
```

to (approximately)

```
var Foo = ___readPub(IMPORTS___, 'Foo');  
var list = new (___asCtor(Foo))(6);
```

From the container’s perspective, this effectively reifies the module’s global scope. Note that the module itself does not necessarily have a means of obtaining a reference to `IMPORTS___`.

If a container wants to allow communication between two modules, it can provide a mutable object as a property of `IMPORTS___`, say,

```
IMPORTS___channel = {};
```

Then if one module sets a property of `channel` in its code, the other can read the property and vice-versa:

```
// In module 1:  
channel.message = "Hi there!";
```

```
// In module 2:  
alert(channel.message);  
// Displays "Hi there"
```

Similarly, if the container wishes to grant a module reified access to its `IMPORTS___`, it can set, for example, `IMPORTS___global = IMPORTS___`.

*Claim: Two separate module instances, even if they instantiate the same module, are isolated from each*

*other unless they have been granted references to the same mutable object.*

Here are some restrictions that the `IMPORTS___` object must have in order to preserve Caja’s security properties.

**eval** The whole point of the cajoler and runtime library is to enforce the Caja restrictions on JavaScript. The `eval` method would allow arbitrary JavaScript to be executed, so it’s important that a Caja module never gets a reference to the `eval` method.

Instead, “`caja.cajitaEval`” will evaluate Cajita source code (text or AST). The `cajitaEval` function will take an `imports` object, and free variables in the source will be bound to properties of `imports`.

**Function** The JavaScript `Function` constructor is absent for the same reason as `eval`.

**Restricted reflection** Allowing access to the `constructor` property of prototypical objects and functions would grant the authority to create more objects like them, which violates the principle of least authority. Therefore, the built-in `constructor` property is absent from Caja.

The `prototype` property of functions can only be used in the limited<sup>9</sup> ways shown in Figure 19.

The `call`, `apply`, and `bind` methods of functions cannot be replaced or overridden.

*Claim: The restrictions stated in this document together make the `Function` object unreachable from Caja programs.*

**new Date()** Nearly all of the members of the global environment are immutable. However, in JavaScript, “`new Date()`” gives ambient access to the current date and time, in violation of object-capability rules as well as dependency injection discipline. `Date` is therefore a member of the global environment which is not actually

<sup>9</sup> See <http://code.google.com/p/google-caja/issues/detail?id=346> for details of the attack enabled by unrestricted access.



immutable. Further, this ambient access to the current time provides a timing channel, further impeding any attempts to stem the leakage of bits over covert channels. Nevertheless, despite these concerns, because it provides only a read-only channel for sensing the world, Caja provides the JavaScript `Date` constructor to Caja programs.

**Math.random()** The JavaScript `Math.random` method is not even read-only. The ES3 standard places no obligations regarding quality of the randomness produced. In particular, an implementation could conform to ES3 and still leak to a given caller of `Math.random()` the ability to infer how many previous times it had been called. Nevertheless, Caja provides the JavaScript `Math.random` method to Caja programs. We recommend that JavaScript platform providers provide good enough randomness that this method doesn't serve as an information channel between otherwise-isolated modules.

## 4.6 Cajita gotchas

Caja seeks to define a fail-stop subset of ES3, as explained in section 3.3. However, it falls short of this goal in several minor ways. To write a correct program that executes correctly whether run cajoled or uncajoled, it should avoid these gotchas. In this section, we enumerate those gotchas relevant to the Cajita subset of Caja.

**Snapshot “arguments”.** In ES3, if `x` is the *i*'th parameter of a function, assignments to `x` are visible as changes to `arguments[i]` and vice versa. In Caja, if “arguments” is mentioned, it is bound to a proper array snapshot of the arguments list when the function was entered, *not* an array-like object. In order for Caja to be a fail-stop subset of ES3, a future version of the Caja sanitizer will statically disallow assignments to any parameter variable within a function that mentions “arguments”. But in the initial Caja implementation, this minor gotcha remains.

**Absent ReferenceError.** In ES3, when a reference

to an undefined variable is evaluated as an expression, a `ReferenceError` is thrown. The Caja sanitizer cajoles a reference to an undefined variable into a reference to a property of `IMPORTS_`. Given the current cajoling rules, a reference to an undefined outer variable will evaluate to `undefined` rather than throwing a `ReferenceError`.

**Top-level variable declaration is not an import.**

In JavaScript, free variables and declared variables at the top-level are both properties of the global object. In Caja, even though free variables are rewritten to properties of imports, declaring a variable with `var` does not add a property to the `IMPORTS_` object.

## 5 Caja Specification

Whereas Cajita is a small subset of JavaScript meant to support new code, Caja is a large subset of JavaScript meant to ease the porting of old JavaScript code and practices. Cajita is small enough that its security properties can be understood. Caja seeks to accept as large a subset of JavaScript as is practical without losing the security properties provided by Cajita. In this section, we explain only the remaining elements of Caja beyond the elements of Cajita already explained.

To explain the remaining elements of Caja, we need some additional definitions.

**Constructed object.** An object defined by Caja code that's not a JSON container and not a function must have been constructed by calling “new” on a function other than `Array` or `Object`.

**Prototypical objects.** As in ES3, a constructed object's implicit *prototype*—the object it directly inherits from—is the value of the `prototype` property of the function which constructed it (which must have been called with “new”). In Caja, these *prototypical objects* are not first-class. When a function is implicitly frozen, so is its `prototype`. Until then, both it and its `prototype` may be initialized.

**Constructors.** A named function whose body mentions “`this`” is a *constructor*.

**Methods.** An anonymous function whose body mentions “`this`” is a *method*. A method definition may appear in one of two constructions. The first is as a parameter in the member map in a call to `caja.def`. For example, `getX` and `setX` are methods in the following code:

```
function Foo(x) { this.x_ = x; }
caja.def(Foo, Object, {
  getX:function () { return this.x_; },
  setX:function () { this.x_ = x; }
});
```

The second is as the right-hand side of an assignment to a property of a constructor’s prototype, before the constructor’s first use. For example:

```
function Foo(x) { this.x_ = x; }

// These are allowed
Foo.prototype.getX = function () {
  return this.x_;
};
Foo.prototype.setX = function (x) {
  this.x_ = x;
};

// First reference to Foo
var f = new Foo(0);

// This is no longer allowed
Foo.prototype.add3 = function (x) {
  return x + 2;
};
```

In both these cases, the function is marked as a method and is added as a property of object bound to the constructor’s `prototype` property. Such a function is called an *unattached method*, in contrast to an *attached method*, explained below. Direct references to unattached methods should never be accessible to Caja code.

$$\begin{aligned}
x.m(\dots) &\equiv (\text{true}\&\&x.m).call(x,\dots) &(10) \\
&\equiv x.m.call(x,\dots) &(11) \\
&\equiv x.m.apply(x,[\dots]) &(12) \\
&\equiv x.m.bind(x)(\dots) &(13) \\
x.m.bind(x) &\equiv (\text{true} \&\& x.m).bind(x) &(14) \\
x.m(\dots_1,\dots_2) &\equiv x.m.bind(x,\dots_1)(\dots_2) &(15) \\
(\dots) &\equiv (\text{function}()\{ &(16) \\
&\quad \text{return} \dots\}).call(\text{this}) \\
\{\dots\} &\equiv (\text{function}()\{ &(17) \\
&\quad \dots\}).call(\text{this})
\end{aligned}$$

Figure 12: Caja Regularities. In Caja, given that `x.m` is associated with either a simple function or a method, then these equivalences hold. See section 5.1 for qualifying conditions.

When a method like `getX` is read as a property of an object `o`, it returns instead an *attached method*, a wrapper around the unattached method that stores `o`; when it is invoked as a method on some object `o2`, the wrapper first verifies that `o === o2`. If the two are not equal, then the wrapper throws an exception.

An *inline method* is an anonymous function mentioning “`this`” that is immediately invoked using `call`, `apply`, or `bind` with “`this`” as the first parameter. Inline methods are a means of achieving true block scoping in Caja. See section 5.1 for more details.

## 5.1 Caja regularities

The regularities shown in Figure 12 apply when Caja code calls any Caja function *other than* a constructor. These regularities are often stronger than ES3, but are all within a fail-stop subset of ES3.

- The code on the left of Equation (10) of Figure 12 calls `x.m` as a method on `x`. The code on the right first extracts the value of `x.m`. When

`x.m` is a method, the extracted value is an *attached method* whose attachment is `x`. When `x` is an expression with no effects and a stable value (such as a variable reference), the code on the right then calls the attached method’s `call` method with its attachment and the original arguments. These two calls are equivalent.

- The `apply` method differs from `call` only in packaging all arguments together into a list.
- Binding an attached method to its attachment yields a conventional bound method—a simple function of the remaining arguments which calls the original method as a method on its attachment.
- When `bind` has additional arguments, it returns a new function representing `F` carried over these additional arguments.
- Equation (16) holds when the expression does not mention `arguments`.
- Equation (17) holds when no `break`, `continue`, or `return` appears freely in the body and no variable defined in the body has the same name as a variable in scope outside the body.

## 5.2 Caja static restrictions

Any source code statically accepted by the Caja sanitizer is a *legal Caja program*. The following syntactic explains why a program may instead be statically rejected.

**Protected properties.** A property name ending in a single underscore may be used only to name protected properties. It may appear as a property name of `“this”`.

**Constructor names.** A Caja constructor can only be called as a constructor using `new`, in order to instantiate a direct instance, or reflectively using `super`. Caja adds a `super` property to constructors to refer to their superclass. For an example, see the definition of the `WobblyPoint` function in figure 15.

```
function Point(x, y) {
  this.x_ = x;
  this.y_ = y;
}
Point.prototype.toString = function() {
  return "<" + this.getX() + "," +
    this.getY() + ">";
};
Point.prototype.getX = function() {
  return this.x_;
};
Point.prototype.getY = function() {
  return this.y_;
};

var ptC = new Point(3, 5);
var ptD = new Point(4, 7);
```

Figure 13: A Caja Point. The point example, written in this common class-like pattern of Javascript programming, is valid Caja. `Point` is frozen by its first use, after which neither `Point` nor `Point.prototype` can be further initialized.

---

Like a named simple function, a constructor and its `prototype` property may be initialized—that is, Caja code may add properties to a constructor and its `prototype`—but are implicitly frozen on first use.

**Methods.** To avoid the confusions regarding `“this”`, Caja methods may only appear in the positions marked `“member”` in the online documentation. Methods may thus be used to initialize properties of prototypes.

Although constructors are normally frozen and the `“prototype”` property of functions is generally not accessible, we allow the patterns shown in Figures 13 and 14 for declaring a constructor, initializing it, and initializing its prototype.

If the first argument to `“caja.def”` is a function name, this is considered an initializing use, and so does not implicitly freeze that function.

```

function Point(x, y) {
  this.x_ = x;
  this.y_ = y;
}
caja.def(Point, Object, {
  toString: function() {
    return "<" + this.getX() + "," +
           this.getY() + ">";
  },
  getX: function() { return this.x_; },
  getY: function() { return this.y_; }
});

```

Figure 14: A Brief Caja Point. Caja also accepts this more compact pattern for initializing a top-level prototype all at once.

### 5.3 Caja dynamic restrictions

The following additional dynamic restrictions are relevant to Caja code.

**Frozen prototypes.** In Caja, until a function is frozen, both it and its **prototype** property may be initialized. When a function is frozen, so is the value of its **prototype** property. Therefore, only the instances at the leaves of the JavaScript inheritance tree may remain unfrozen. Initializing assignments to a function's **prototype** can thus be considered declarative initializations rather than mutations.

*Claim: No Caja program can cause a Caja-observable mutation of a prototypical object.*

**Well formed inheritance.** JavaScript provides an interesting set of primitives for building non-standard inheritance arrangements. Many of these arrangements will break assumptions in other code. In practice, these primitives are used in a particular arrangement in which, for example, for all functions  $F$ ,  $F.prototype.constructor === F$ . Caja allows only this classical inheritance pattern, so that Caja code and the Caja implementation can rely on it.

**Shape change.** When one adds or deletes properties of an object, we can describe this as changing the shape of the object. Of course, no one can change the shape of a frozen object. Anyone with access to a non-frozen JSON container may freely change its shape. A constructed object can directly change its own shape, by assignment or **delete** using **this**. Clients of a constructed object cannot directly change its shape. But since a constructed object can directly change its own shape, it can provide methods enabling its clients to ask it to change its shape. In other words, a constructed object has control of its own shape.

Adding a property that overrides an inherited property is considered a shape change, so only a constructed object may do this directly for itself. If a constructed object does create a public, non-inherited property, its clients can directly assign to it.

**Non-reflective constructors** Any attempt to call a constructor's **call**, **apply**, or **bind** methods must fail, except for the statically exempted use of **call** mandated in section 5.2 for derived constructors.

**Attached methods** Any attempt to obtain a method as a value will instead yield an *attached method*. If  $x.foo(...)$  would directly call a method, then  $x.foo$  will return that method as attached to  $x$ . An attached method can only be invoked by calls that bind its **this** to its attachment, whether called as a method on its attachment, or called reflectively by providing its attachment as the first argument of **call**, **apply**, or **bind**. Since calling an attached method either fails or acts like calling the original method, an attached method behaves within a fail-stop subset of the behavior associated with the original method.

### 5.4 Caja gotchas

Caja seeks to define a fail-stop subset of ES3, as explained in section 3.3. However, it falls short of this

```

function WobblyPoint(x, y) {
  WobblyPoint.super(this, x, y);
}
caja.def(WobblyPoint, Point, {
  getX: function() {
    return Math.random() +
      WobblyPoint.super.prototype.
        getX.call(this);
  }
});

```

Figure 15: A Caja Subclass. `caja.def` supports classical inheritance. The second argument serves as a “superclass”. The third argument provides instance members including methods. A fourth optional argument (unshown) provides static members. The `super` call within the `WobblyPoint` constructor asks the “superclass” constructor to do its part in initializing the new instance. We have not yet decided on the form of “`super`” method calls; the `...super.prototype...` syntax above is one proposal we are considering.

---

goal in several minor ways. To write a correct program that executes correctly whether run cajoled or uncajoled, it should avoid these gotchas. In this section, we enumerate those remaining gotchas relevant specifically to Caja.

**Bare for/in loops.** More properties are visible and enumerable to uncajoled programs than cajoled programs. To write a program which will see the same properties whether run cajoled or not, write the following instead:

```

for (var k in obj) {
  if (caja.canEnumPub(obj,k)) {
    ...k...obj[k]...
  }
}

```

This conditional does not affect the behavior of cajoled programs, so programs that only need to run cajoled can safely leave it out.

Using `canEnumOwn` instead will further restrict the enumeration to non-inherited properties, as is typically desired. The same effect can still be obtained more compactly using Caja’s `caja.forEach` construct as explained in section 4.3.

**Isolated RegExps.** ES3 specifies that a literal regular expression pattern corresponds directly to a single mutable `RegExp` object. Caja, as well as the Internet Explorer version of JavaScript (JScript), instead create a new `RegExp` on each evaluation of a literal pattern, avoiding the implicit sharing of mutable state. For any program already compatible with JScript, this is not an issue.

**Permissive constructors.** In JavaScript, if a constructor is stored in an object’s property, and that property is then invoked as a method of the object (without using `new`), the constructor would run with its `this` bound to that object, which in Caja would violate that object’s encapsulation. Even worse, in JavaScript, if a constructor is called as a function, its `this` would be bound to the global object—which would be a fatal escalation of privilege.

In order for Caja to be both safe and a fail-stop subset of JavaScript, these cases should fail. Instead, in the initial Caja implementation, in these cases the constructor may instead act as if called with `new`. This is safe, but it silently diverges from JavaScript behavior.

**Attachment breaks identity.** Figure 6 and Figure 13 each instantiate two points. Both are in Caja-compliant JavaScript—they work correctly whether cajoled or not. After Figure 6, which is also Cajita-compliant, `ptA.getX===ptB.getX` will always be `false`. Whether cajoled or not, each point instance returns its own unique `getX` function.

By contrast, `ptC.getX===ptD.getX` will be `true` if Figure 13 is run uncajoled, but `false` if cajoled. In uncajoled JavaScript, both operands return the `Point.prototype.getX` method itself. When cajoled, the left operand returns the

```

function Shadow(model) {
  this.state_ = model.getState();
  var listener = (function(newState) {
    this.state_ = newState;
  }).bind(this);
  model.addStateListener(listener);
}
Shadow.prototype.getState = function() {
  return this.state_;
};

```

Figure 16: A Caja Inline Method. The anonymous function above mentions “`this`”, and so is a form of method, but it is not used to initialize a property of a shared prototype. Such an *inline method* may appear only as the receiver of a `call`, `apply`, or `bind` call whose first argument is “`this`”. The above listener, when invoked, runs in the lexical scope in which it was created, including the binding of “`this`”.

---

method as attached to `ptC` whereas the right operand returns the method as attached to `ptD`. This difference in object identity is a genuine Caja gotcha. Caja-compliant programs should avoid testing the object identity of methods. Cajita-compliant programs need not worry.

**Exceptions break identity.** The `Error` class exposes too much authority, so instances of `Error` are caught and replaced with frozen records containing the relevant information.

## 6 Related Work

### 6.1 Browser Shield

**all:** [Write this section](#)

### 6.2 ADsafe

**all:** [Write this section](#)

### 6.3 Jacaranda

**all:** [Write this section](#)

## 7 Conclusions

**all:** [Write this section](#)

## 8 Acknowledgements

We thank Dirk Balfanz, Bruno Bowden, Jon Bright, Andrea Campi, Doug Crockford, Jed Donnelley, Brendan Eich, David-Sarah Hopwood, Ken Kahn, Adam Langley, Marcel Laverdet, Kevin Reid, Graham Spencer, Marc Stiegler, and David Wagner for various comments and suggestions.

## A Tables

## References

- [1] J. B. Dennis and E. C. V. Horn. Programming Semantics for Multiprogrammed Computations. Technical Report MIT/LCS/TR-23, M.I.T. Laboratory for Computer Science, 1965.
- [2] ECMA. *ECMA-262: ECMAScript Language Specification*. ECMA (European Association for Standardizing Information and Communication Systems), Geneva, Switzerland, third edition, Dec. 1999.
- [3] I. K. S. L. Garfinkel. Bitfrost: the One Laptop per Child Security Model. *Symposium On Usable Privacy and Security*, 2007.
- [4] A. H. Karp. Authorization-based access control for the services oriented architecture. *c5*, 0:160–167, 2006.
- [5] M. Košík. Backwater Operating System, 2007. `altair.dcs.elf.stuba.sk:60001/mediawiki/upload/2/2b/Backwater.pdf`.

| Caja expression                      | cajoles to ES3 code equivalent to                                       |      |
|--------------------------------------|---|------|
| <code>with</code>                    | <code>/* rejected in all positions */</code>                            | (18) |
| <code>local_</code>                  | <code>/* rejected in all positions */</code>                            | (19) |
| <code>glob_</code>                   | <code>/* rejected in all positions */</code>                            | (20) |
| <code>glob</code>                    | <code>____.readPub(IMPORTS____, "glob")</code>                          | (21) |
| <code>this.p_</code>                 | <code>/* rejected in all positions */</code>                            | (22) |
| <code>foo.p_</code>                  | <code>/* rejected in all positions */</code>                            | (23) |
| <code>this.p</code>                  | <code>____.readProp(this, "p")</code>                                   | (24) |
| <code>foo.p</code>                   | <code>____.readPub(foo, "p")</code>                                     | (25) |
| <code>this[bar]</code>               | <code>____.readProp(this, bar)</code>                                   | (26) |
| <code>foo[bar]</code>                | <code>____.readPub(foo, bar)</code>                                     | (27) |
| <code>bar in this</code>             | <code>____.canReadProp(this, bar)</code>                                | (28) |
| <code>bar in foo</code>              | <code>____.canReadPub(foo, bar)</code>                                  | (29) |
| <code>for (key in this) {...}</code> | <code>for (key in this) {if (____.canEnumProp(this, key)) {...}}</code> | (30) |
| <code>for (key in foo) {...}</code>  | <code>for (key in foo) {if (____.canEnumPub(foo, key)) {...}}</code>    | (31) |
| <code>this.p = baz</code>            | <code>____.setProp(this, "p", baz)</code>                               | (32) |
| <code>foo.p = baz</code>             | <code>____.setPub(foo, "p", baz)</code>                                 | (33) |
| <code>this[bar] = baz</code>         | <code>____.setProp(this, bar, baz)</code>                               | (34) |
| <code>foo[bar] = baz</code>          | <code>____.setPub(foo, bar, baz)</code>                                 | (35) |
| <code>delete this.p</code>           | <code>____.deleteProp(this, "p")</code>                                 | (36) |
| <code>delete foo.p</code>            | <code>____.deletePub(foo, "p")</code>                                   | (37) |
| <code>delete this[bar]</code>        | <code>____.deleteProp(this, bar)</code>                                 | (38) |
| <code>delete foo[bar]</code>         | <code>____.deletePub(foo, bar)</code>                                   | (39) |
|                                      |   | (40) |

Figure 17: Cajoling Property Access. Under the assumption that the Caja runtime environment is as specified, the Caja sanitizer generates cajoled Javascript equivalent to that specified above, but inlined and optimized where possible. The meaning of sanitizing is thereby determined by the specification of these entry points into the Caja runtime library. Where we show cajoled code apparently duplicating an expression, the Caja sanitizer instead introduces temporary variables as needed so that each expression evaluates exactly as many times and in the same order as in the original.

| Caja expression                          | cajoles to ES3 code equivalent to  |      |
|--|--|------|
| <i>/* caja module body */</i>            | <code>___loadModule(function(___ , IMPORTS___) {</code>                            | (41) |
|  | <code>  <i>/* cajoled module body */</i></code>                                    |      |
|  | <code>  });</code>   |      |
| <code>this.m(a...)</code>                | <code>___callProp(this, "m", [a...])</code>  | (42) |
| <code>foo.m(a...)</code>                 | <code>___callPub(foo, "m", [a...])</code>  | (43) |
| <code>this[bar](a...)</code>             | <code>___callProp(this, bar, [a...])</code>  | (44) |
| <code>foo[bar](a...)</code>              | <code>___callPub(foo, bar, [a...])</code>  | (45) |
| <code>new foo(a...)</code>               | <code>new (___asCtor(foo))(a...)</code>  | (46) |
| <code>foo(a...)</code>                   | <code>___asSimpleFunc(foo)(a...)</code>  | (47) |
| Methods                                  |  |      |
| <code>function(a...) {...this...}</code> | <code>___method(function(a...) {...this...})</code>                                | (48) |
| Simplefunctions                          |  |      |
| <code>function F(a...) {...}</code>      | <code>___primFreeze(___simpleFunc(function F(a...) {...}))</code>                  | (49) |
| <code>function(a...) {...}</code>        | <code>___primFreeze(___simpleFunc(function(a...) {...}))</code>                    | (50) |
| <code>arguments.callee</code>            | <i>/* rejected */</i>  | (51) |
| <code>...arguments...</code>             | <code>...args_____</code>  | (52) |
|  | <code>var args___ = ___args(arguments);</code> <i>/* move to function start */</i> |      |
| <code>/pattern/</code>                   | <code>new RegExp("pattern")</code>   | (53) |
| <code>/pattern/flags</code>              | <code>new RegExp("pattern", "flags")</code> <i>/* where flags is [igm] * */</i>    | (54) |

Figure 18: Cajoling Callers and Callees. A cajoled Caja module can be loaded/evald once, creating an anonymous plugin-maker function. Each time a plugin-maker is called, it makes a new confined plugin. The use of a terminal “;” is shorthand for testing whether the matching expression is evaluated for effects only, not for its value.



| Caja expression                         | Special cases for function names and methods                                      |      |
|---|---|------|
|   | Initializes, doesn't freeze Foo   |      |
| Foo.prototype.m = member;               | ----.setMember(Foo,"m",member);   | (55) |
| Foo.prototype = {...:member,...};       | ----.setMemberMap(Foo,{...:member,...});  | (56) |
| Foo.m = ...                             | ----.setPub(Foo,"m",...)  | (57) |
| caja.def(Foo,Base)                      |   | (58) |
| caja.def(Foo,Base,{...:member,...},...) |   |      |
|   | An inner method within a method or constructor                                    |      |
| member                                  | ----.attach(this,member)  | (59) |
|   | Freezes Foo to prevent further initialization                                     |      |
| new Foo(...)                            |   | (60) |
| caja.def(Derived, Foo, ...)             |   |      |
| ...Foo...                               | .....primFreeze(Foo)...   | (61) |
| ... instanceof Foo                      | <i>allow, whether Foo is frozen or not</i>  | (62) |
| Foo = ...                               | <i>reject assignment to a function name</i>                                       | (63) |
| var Foo = ...                           | <i>reject conflicting initialization as well</i>                                  | (64) |
|   | Can only happen if Foo is already frozen  |      |
| Foo.call(this,...);                     | <i>Only at start of Derived,<br/>and only if the remaining args have no this.</i> | (65) |
| Foo.prototype.m                         | <i>Only within methods of Derived</i>   | (66) |
|   | ----.attach(this, Foo.prototype.m)  |      |

Figure 19: Cajoling Special Cases. When Foo is the name of a named function or a constructor, then these special cases are checked before the general cajoling rules. At the member positions above, either normal expressions or methods may appear.

| Methods of ___            | method body   |
|---------------------------|---|
| enforce(test,complaint)   | if (test) { return true; }<br>throw new CajaRuntimeError(complaint);  |
| canRead(obj,name)         | return !!obj[name+"_canRead___"];   |
| canEnum(obj,name)         | return !!obj[name+"_canEnum___"];   |
| canCall(obj,name)         | return !!obj[name+"_canCall___"];   |
| canSet(obj,name)          | return !!obj[name+"_canSet___"];  |
| canDelete(obj,name)       | return !!obj[name+"_canDelete___"];   |
| allowRead(obj,name)*      | obj[name+"_canRead___"] = true;   |
| allowEnum(obj,name)*      | allowRead(obj,name);<br>obj[name+"_canEnum___"] = true;   |
| allowCall(obj,name)*      | obj[name+"_canCall___"] = true;   |
| allowSet(obj,name)*       | enforce(!isFrozen(obj),...);<br>allowEnum(obj,name);<br>obj[name+"_canSet___"] = true;  |
| allowDelete(obj,name)*    | enforce(!isFrozen(obj),...);<br>obj[name+"_canDelete___"] = true;<br>/*other bookkeeping yet to be determined*/   |
| hasOwnProp(obj,name)      | /*like the original: obj.hasOwnProperty(name)*/   |
| isJSONContainer(obj)      | var constr = directConstructor(obj);<br>return constr === Object    constr === Array;   |
| isFrozen(obj)             | return hasOwnProp(obj,"___FROZEN___");  |
| primFreeze(obj)*          | for (k in obj) {<br>if (endsWith(k,"_canSet___")  endsWith(k,"_canDelete___")) {<br>obj[k] = false; }}<br>obj.___FROZEN___ = true;<br>if (typeof obj === "function") {<br>primFreeze(obj.prototype); }<br>return obj; |
| method(constr, meth)      | enforce(typeof constr === "function",...);<br>enforce(typeof meth === "function",...);<br>meth.___METHOD_OF___ = constr;<br>return primFreeze(meth);  |
| allowMethod(constr,name)* | method(constr,constr.prototype[meth]);<br>allowCall(constr,name);   |

Figure 20: Hidden Attributes. These methods handle the concrete representations of object and property attributes. Only methods marked with a \* should be called by JavaScript code during initialization of the embedding app to express taming decisions. All objects that are reachable from the ES3 shared environment should be frozen, so that the shared environment is transitively read-only to all Caja code.

| Methods of ___                      | method body  |
|-------------------------------------|--|
| <code>canReadProp(self,name)</code> | <code>if (endsWith(name,"_")) { return false; }<br/>return canRead(self,name);</code>  |
| <code>readProp(self,name)</code>    | <code>return canReadProp(self,name) ? self[name] : undefined;</code>   |
| <code>canReadPub(obj,name)</code>   | <code>if (endsWith(name,"_")) { return false; }<br/>if (canRead(obj,name)) { return true; }<br/>if (!isJSONContainer(obj)) { return false; }<br/>if (!hasOwnProp(obj,name)) { return false; }<br/>allowRead(obj,name); /*memoize*/<br/>return true;</code> |
| <code>readPub(obj,name)</code>      | <code>return canReadPub(obj,name) ? obj[name] : undefined;</code>  |
| <code>canEnumProp(self,name)</code> | <code>if (endsWith(name,"_")) { return false; }<br/>return canEnum(self,name);</code>  |
| <code>canEnumPub(obj,name)</code>   | <code>if (endsWith(name,"_")) { return false; }<br/>if (canEnum(obj,name)) { return true; }<br/>if (!isJSONContainer(obj)) { return false; }<br/>if (!hasOwnProp(obj,name)) { return false; }<br/>allowEnum(obj,name); /*memoize*/<br/>return true;</code> |
| <code>canSetProp(self,name)</code>  | <code>if (endsWith(name,"_")) { return false; }<br/>if (canSet(self,name)) { return true; }<br/>return !isFrozen(self);</code>   |
| <code>setProp(self,name,val)</code> | <code>enforce(canSetProp(self,name),...);<br/>allowSet(self,name); /*grant*/<br/>return self[name] = val;</code>   |
| <code>canSetPub(obj,name)</code>    | <code>if (endsWith(name,"_")) { return false; }<br/>if (canSet(obj,name)) { return true; }<br/>return !isFrozen(obj) &amp;&amp; isJSONContainer(obj);</code>   |
| <code>setPub(obj,name,val)</code>   | <code>enforce(canSetPub(obj,name),...);<br/>allowSet(obj,name); /*grant*/<br/>return obj[name] = val;</code>   |
| <code>deleteProp(self,name)</code>  | <code>enforce(canDeleteProp(self,name),...);<br/>/*XXX Bookkeeping yet to be determined*/<br/>return enforce(delete self[name],...);</code>  |
| <code>deletePub(obj,name)</code>    | <code>enforce(canDeletePub(obj,name),...);<br/>enforce(isJSONContainer(obj),...);<br/>/*XXX Bookkeeping yet to be determined*/<br/>return enforce(delete obj[name],...);</code>  |
| <code>args(original)</code>         | <code>return primFreeze(Array.prototype.slice.call...(original,0));</code>   |

Figure 21: Property Access. The calls to `allowRead` and `allowEnum` merely memoize a query result. The calls to `allowSet` track the implications of side effects.

| Global ES3 non-constructor      | Property            | Taming       |
|---------------------------------|---------------------|--------------|
| <code>NaN</code>                |                     | ok           |
| <code>Infinity</code>           |                     | ok           |
| <code>undefined</code>          |                     | ok           |
| <code>eval</code>               |                     | hidden       |
| <code>parseInt</code>           |                     | ok           |
| <code>parseFloat</code>         |                     | ok           |
| <code>isNaN</code>              |                     | ok           |
| <code>isFinite</code>           |                     | ok           |
| <code>decodeURI</code>          |                     | ok           |
| <code>decodeURIComponent</code> |                     | ok           |
| <code>encodeURI</code>          |                     | ok           |
| <code>encodeURIComponent</code> |                     | ok           |
| <code>Math</code>               |                     | ok           |
|                                 | <code>random</code> | callable*    |
|                                 | all others in ES3   | ok, callable |

Figure 22: Taming ES3 Global Non-Constructors. Except for `eval`, all non-constructors specified by ES3 are visible in Caja’s outer environment as immutable objects. Note that `Math.random` is not actually immutable, and therefore neither is `Math` nor Caja’s outer environment itself. We allow it anyway for reasons explained in the text.

- [6] B. Laurie. Safer Scripting Through Pre-compilation. *Security Protocols 13*, LNCS 4631, 2004.
- [7] A. M. Mettler and D. Wagner. The Joe-E Language Specification (draft). Technical Report UCB/EECS-2006-26, EECS Department, University of California, Berkeley, March 17 2006.
- [8] M. S. Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, Baltimore, Maryland, USA, May 2006.
- [9] J. A. Rees. A Security Kernel Based on the Lambda-Calculus. Technical report, Massachusetts Institute of Technology, 1996.
- [10] M. Seaborn. Plash: The Principle of Least Authority Shell, 2005. [plash.beasts.org/](http://plash.beasts.org/).
- [11] M. Stiegler. Emily, a High Performance Language for Secure Cooperation, 2006. [skyhunter.com/marcs/emily.pdf](http://skyhunter.com/marcs/emily.pdf).
- [12] M. Stiegler, A. H. Karp, K.-P. Yee, and M. S. Miller. Polaris: Virus Safe Computing for Windows XP. Technical Report HPL-2004-221, Hewlett Packard Laboratories, 2004.
- [13] D. Wagner and E. D. Tribble. A Security Analysis of the Combex DarpaBrowser Architecture, Mar. 2002. [combex.com/papers/darpa-review/](http://combex.com/papers/darpa-review/).
- [14] H. J. Wang, X. Fan, C. Jackson, and J. Howell. Protection and communication abstractions for web browsers in MashupOS. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles (SOSP’07)*. ACM, Oct. 2007.

| Global ES3 constructor | Property                | Taming         |
|------------------------|-------------------------|----------------|
| constructor            |                         | default ctor   |
| constructor.prototype  |                         | hidden         |
|                        | constructor             | hidden         |
|                        | toString                | default method |
|                        | toLocaleString          | default method |
|                        | valueOf                 | default method |
| instances              | length                  | default ok     |
|                        | /*stringified numbers*/ | default ok     |
| Object.prototype       | hasOwnProperty          | handled        |
|                        | isPrototypeOf           | method         |
|                        | propertyIsEnumerable    | handled        |
|                        | freeze_                 | added method   |
| Function               |                         | hidden         |
| Function.prototype     |                         | hidden         |
|                        | apply                   | handled        |
|                        | call                    | handled        |
|                        | bind                    | added method   |
| instances              | prototype               | hidden         |
|                        | length                  | ok             |
| Array.prototype        | concat                  | method         |
|                        | join                    | method         |
|                        | pop                     | handled        |
|                        | push                    | handled        |
|                        | reverse                 | handled        |
|                        | shift                   | handled        |
|                        | slice                   | method         |
|                        | sort                    | handled        |
|                        | splice                  | handled        |
|                        | unshift                 | handled        |
| String                 | fromCharCode            | callable       |
| String.prototype       | match                   | handled        |
|                        | replace                 | handled        |
|                        | search                  | handled        |
|                        | split                   | handled        |
|                        | all others in ES3       | ok, method     |

Figure 23: Taming ES3 Global Constructors, Part 1. The first section above shows the taming decisions that apply by default to global ES3 constructors, their prototypes, and their instances, unless stated otherwise in a specific table entry. A *stringified number* is any `x` for which `x === String(Number(x))`. A *handled* method acts differently when called by cajoled *vs.* uncajoled code. Handled mutating methods like `Array.pop` obey Caja’s mutability constraints.

| Global ES3 constructor  | Property             | Taming   |
|-------------------------|----------------------|----------|
| <b>Boolean</b>          |                      | ctor     |
| <b>Number</b>           | MAX_VALUE            | ok       |
|                         | MIN_VALUE            | ok       |
|                         | NaN                  | ok       |
|                         | NEGATIVE_INFINITY    | ok       |
|                         | POSITIVE_INFINITY    | ok       |
| <b>Number.prototype</b> | toFixed              | method   |
|                         | toExponential        | method   |
|                         | toPrecision          | method   |
| <b>Date</b>             |                      | ctor*    |
|                         | parse                | callable |
| <b>Date.prototype</b>   | UTC                  | callable |
|                         | to*String all in ES3 | method   |
|                         | get* all in ES3      | method   |
|                         | set* all in ES3      | handled  |
| <b>RegExp.prototype</b> | exec                 | handled  |
|                         | test                 | handled  |
| instances               | source               | ok       |
|                         | global               | ok       |
|                         | ignoreCase           | ok       |
|                         | multiline            | ok       |
|                         | lastIndex            | ok       |
| <b>Error.prototype</b>  | name                 | ok       |
|                         | message              | ok       |
| <b>*Error</b>           | all in ES3           | ok       |
| <b>*Error.prototype</b> | all in ES3           | ok       |

Figure 24: Taming ES3 Global Constructors, Part 2. The **Date** constructor itself gives ambient read-only access to the current time, and is therefore not immutable. We allow it anyway for reasons explained in the text.