statacons: An SCons-based build tool for Stata

Raymond P. Guiteras
North Carolina State University
Raleigh, N.C., USA
rpguiter@ncsu.edu

Ahnjeong Kim
North Carolina State University
Raleigh, N.C., USA
akim7@ncsu.edu

Brian Quistorff
Bureau of Economic Analysis
Washington, D.C., USA
Brian.Quistorff@bea.gov

Clayson Shumway
North Carolina State University
Raleigh, N.C., USA
cshumwa@ncsu.edu

Abstract. This paper presents `statacons`, an SCons-based build tool for Stata. Because of the integration of Stata and Python in recent versions of Stata, we are able to adapt SCons for Stata workflows without the use of an external shell or extensive configuration. We discuss the usefulness of build tools generally, provide examples of the use of `statacons` in Stata workflows, present key elements of the syntax of `statacons`, and discuss extensions, alternatives, and limitations. Appendices provide installation instructions and recommendations for collaborative workflows.

Keywords: statacons, reproducible research, build tools, scons, Python

1 Build Tools

Transparency and reproducibility are increasingly important norms in empirical science (Wilson et al. 2014; Stodden et al. 2018; Christensen et al. 2019; Orozco et al. 2020). Public posting of data and code, use of version control software, and literate programming are all becoming more common, both in general and with Stata users specifically.

One set of tools that have not received as much attention in Stata are *build tools* (Mokhov et al. 2018). Build tools track *dependencies* throughout a project: which inputs and programs are used to create which outputs, which outputs are subsequently used or combined to create final outputs, etc. A project with a workflow encoded in a build script can be replicated from beginning to end with one command. Just as importantly, when inputs or code change, a build tool can decide exactly which outputs need to be rebuilt, avoiding both errors of omission – failing to update outputs when inputs change – and errors of inclusion – unnecessarily re-running potentially time-consuming code when no relevant inputs have changed.

Stata users already often instinctively create a rudimentary build tool: a "master" do-file that lists all of a project's do-files in order. This is adequate for simple projects, but has limitations in complex workflows. First, a list of do-files provides no information on what produces what, and what the relationships among different inputs and outputs are.[1] This makes it difficult to trace outputs back to their source or debug problems.

---

1. Comments in a master do-file can provide some documentation of inputs and outputs, but drift be-

Second, a master do-file does not handle the tradeoff between errors of omission and errors of inclusion well. Running a master do-file from beginning to end should avoid errors of omission, but is inefficient if some sections of the project take a long time to run. It would not make sense to re-run a computationally intensive estimation or bootstrapping procedure just to adjust the formatting of some graphs. Again, users often intuitively handle these problems by commenting out lines of the master do-file or setting up switches to skip sections of code, and this is usually fine for simple projects.

However, as projects become more complex, an approach like a master do-file can be inadequate. For example, many projects have some or all of the following characteristics: merging multiple input datasets and producing several inter-related datasets for analysis; results of some analyses being used as inputs for others (e.g., estimation results used in subsequent simulations); multiple collaborators working on different aspects of the project simultaneously, using version control software and file-sharing platforms to share code and files; combining Stata tasks with non-Stata tasks, e.g., compiling text, tables and figures into a PDF. In such cases, it is easy to lose track of which parts of a project need to be rebuilt. Similarly, in projects with multiple collaborators (or a single researcher collaborating with her past self), one collaborator may not know about the dependence of one part of the project on another.

Software developers face similar problems, and have developed sophisticated build tools to deal with them. Two ideas are central: first, the relationships between inputs, code and outputs are encoded in a *build script*; second, the build software uses this script and information it collects on the state of the project to manage the build, deciding what is up to date and what needs to be built or rebuilt. The build script itself provides living documentation of the full structure of the project, and the build tool can provide the user with the status of of any particular component, or of all components.

The use of build tools does not appear to be common among Stata users. We suspect the main barriers to adoption are: first, lack of awareness of the existence and value of build tools; second, the startup cost of configuring the tool, system and Stata so that they can work together; third, the startup cost of learning how to use these tools.

In this paper and with `statacons`, we attempt to address all three of these barriers. First, we will show in our examples that `statacons` can be used to manage both simple and complex workflows efficiently. Second, because of the integration of Python and Stata in recent versions (16+), as well as our own adaptations, we are able to make SCons, a popular Python-based build tool, accessible in Stata through our command `statacons`. With a minimal amount of configuration, Stata users can use `statacons` directly in Stata and write the build scripts, `SConstruct`s, directly in Stata's do-file editor, taking advantage of the editor's Python syntax highlighting.[2] Third, our examples in this paper, as well as additional examples in our companion web tutorial and project

---

tween comments and the true content of code is a notorious source of confusion and error (Gentzkow and Shapiro 2014). Furthermore, after an initial setup, virtually all the effort in using a build tool consists in encoding inputs and outputs, so on the margin there is little additional effort involved in using a build tool relative to providing comments in a master do-file.

2. `SConstruct` files written for `statacons` are standard `SConstruct`s and will work without modification in standard SCons, so users who prefer the terminal can continue to work there.

Wiki, can easily be adapted to users' own work.

More specifically, `statacons` offers the following benefits: (a) not requiring changes to existing Stata code or the structure of existing workflows; (b) rebuilding exactly those parts of a project that need to be rebuilt at any given moment, with neither errors of inclusion nor omission; (c) complementing existing workflow practices, including literate programming and use of version control software and file-sharing platforms; (d) allowing easy extensiblity; and (e) operating fully within Stata.

Our paper proceeds as follows: in Section 2, we use a simple example of managing a Stata workflow with `statacons` and an `SConstruct` to introduce basic concepts, vocabulary and mechanics; in Section 3, we provide the syntax of the `statacons` command and `SConstructs`; in Section 4, we present an applied example of using `statacons` to manage the complex workflow of an empirical project (Shumway and Wilson 2021); in Section 5, we describe some of the technical details of our adaptation of SCons to Stata; in Section 6, we discuss extensions, alternative approaches, and limitations; Section 7 concludes. We include a detailed installation guide at the end of the paper. Online Appendices discuss parallel builds, illustrate some additional advanced features, and provide examples of the use of `statacons` in collaborative projects using version control software and shared files. Our project web page[3] provides installation instructions and hosts a companion web tutorial based on Software Carpentry's "Automation and Make" tutorial (Jackson 2016) with several additional worked examples.[4] The `statacons` package is available at our project repository[5], where we have posted complete replication code and data for the Introductory Example and key code from the Applied Example, as well as a project Wiki with examples of additional tools.

2   Introduction by Example

In this section, we use a simple example to introduce the basic concepts, terminology and mechanics of managing a workflow with `statacons`. Replication code and data for this example are provided in `stataconsIntro.zip`, posted to our repository.[6]

Our project has two steps, using the following two two do-files:

```
// dataprep.do
version 16.1

use "inputs/auto-original.dta", clear

generate mpg_sqd = mpg^2
label variable mpg_sqd "Mileage (mpg) squared"

save "outputs/auto-modified.dta", replace


// analysis.do
version 16.1
```

---

3. https://bquistorff.github.io/statacons
4. https://bquistorff.github.io/statacons/swc
5. https://github.com/bquistorff/statacons
6. https://github.com/bquistorff/statacons/raw/main/examples/stataconsIntro.zip

```
use "outputs/auto-modified.dta", clear
twoway scatter price mpg
graph export "outputs/scatterplot.pdf", replace
regress price mpg
eststo linear
regress price mpg mpg_sqd
eststo quadratic
esttab linear quadratic using "outputs/regressionTable.tex", ///
  se r2 replace
```

In the first step, the do-file `dataprep.do` takes an input file, `auto-original.dta`, creates a new variable `mpg_sqd`, which is the square of `mpg`, and saves the resulting dataset as `auto-modified.dta`. In the second step, the do-file `analysis.do` uses `auto-modified.dta` to create a scatterplot of `price` against `mpg`, saved as `scatterplot.pdf`, and regresses `price` on `mpg` and `mpg_sqd`, exporting the regression results to `regressionTable.tex`.

This workflow is represented visually in Figure 1 and can easily be managed by a simple master do-file:

```
// master.do
version 16.1
do "code/dataprep.do"
do "code/analysis.do"
exit
```

When using a build tool, it is useful to think about the workflow differently. Rather than think about a sequence of do-files in a particular order, we think about the *targets* that we want to build and the *dependencies* that we use to create these targets. In this example, the ultimate targets are `scatterplot.pdf` and `regressionTable.tex`. The dependencies for these targets are the *source* code `analysis.do` and the dataset `auto-modified.dta`. The dataset `auto-modified.dta` is itself a target, with dependencies `dataprep.do` and `auto-original.dta`.

This workflow is depicted visually in Figure 2.

To use `statacons`, we encode these relationships in an `SConstruct` file:

```
# SConstruct for Introductory Example
# **** Setup from pystatacons package *****
import pystatacons
env = pystatacons.init_env()
# **** Substance begins        *****
# analysis
cmd_analysis = env.StataBuild(
    target = ['outputs/scatterplot.pdf',
              'outputs/regressionTable.tex'],
    source = 'code/analysis.do'
)
Depends(cmd_analysis, ['outputs/auto-modified.dta'])
```
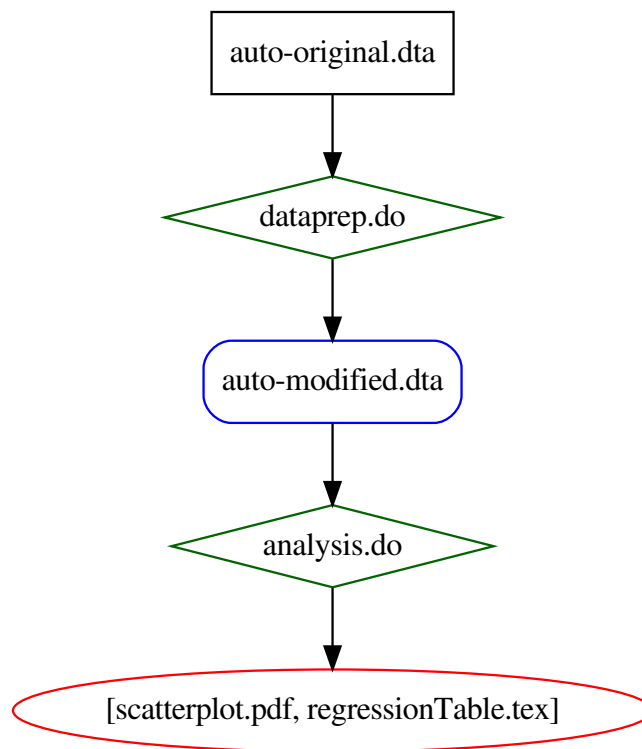
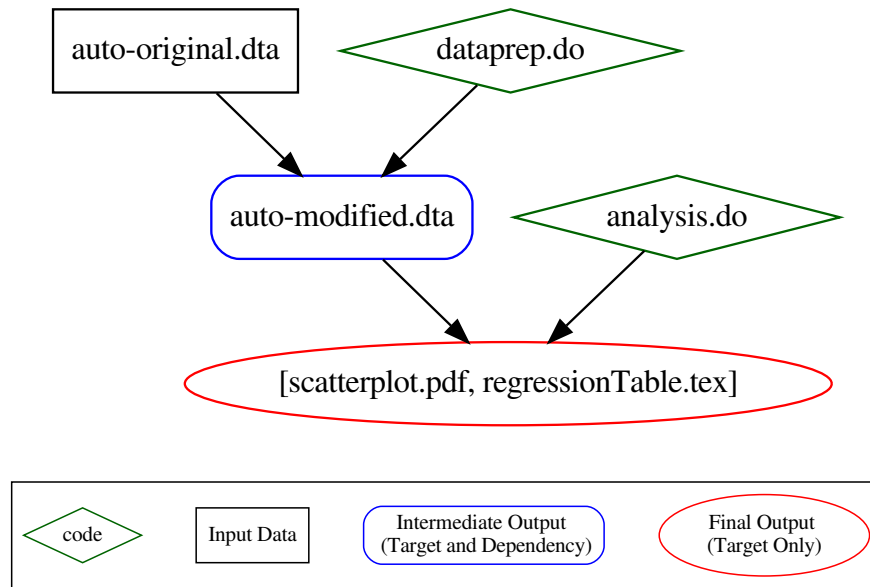Figure 1: Linear Workflow for Introductory Example

Figure 2: Target-Build Workflow for Introductory Example

```
# dataprep
cmd_dataprep = env.StataBuild(
    target = [´outputs/auto-modified.dta´],
    source = ´code/dataprep.do´
)
Depends(cmd_dataprep, [´inputs/auto-original.dta´])
```

Let's consider first the analysis task. We have given this task the name `cmd_analysis`.[7] The *targets* are `scatterplot.pdf` and `regressionTable.tex`. The *source* file for this task is `analysis.do`. The source file is the code we will run to build these targets. The `Depends()` line lists any *dependencies* other than the source file.[8] In this case, this is `auto-modified.dta`. Finally, `env.StataBuild` tells `scons` that this is a task for Stata, to be handled in the way defined by the *Builder method* `StataBuild`. The essence of the `StataBuild` Builder method is that `scons` will build the targets (here, `scatterplot.pdf` and `regressionTable.tex`) by running the source file (here, `analysis.do`) in Stata's batch mode.[9]

---

7. More precisely, `cmd_analysis` is a *NodeList*, where the nodes are the targets. Conveniently, you can use the NodeList elsewhere in the SConstruct to refer to the targets. That is, if we use `cmd_analysis` elsewhere in this SConstruct, scons will understand this to mean the targets `scatterplot.pdf` and `regressionTable.tex`.

8. While not strictly necessary, it is useful to split the source from the other dependencies so there is no ambiguity about what is the code that will be used in the StataBuild environment.

9. Batch mode does not require user input and is often used for automating Stata tasks. See the "Stata batch mode" section of the relevant Getting Started With Stata manual ([GSW] **B.5**, [GSU] **B.3**, [GSM] **B.3**).

The dataprep task, `cmd_dataprep`, is similar and we leave parsing it as an exercise for the reader.

We will now use `statacons` to build our outputs. We will run `statacons` with the option `debug(explain)`, which adds some helpful screen output that tells us what is going on and why.

```
. statacons,  debug(explain)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
scons: building `outputs\auto-modified.dta´ because it doesn´t exist
stata_run(["outputs\auto-modified.dta"], ["code\dataprep.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\dataprep.do".
  Starting in hidden desktop (pid=22460).
scons: building `outputs\scatterplot.pdf´ because it doesn´t exist
stata_run(["outputs\scatterplot.pdf", "outputs\regressionTable.tex"], ["code\ana
> lysis.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\analysis.do".
  Starting in hidden desktop (pid=34912).
scons: done building targets.
```

There are a few key lines from the screen output that we will highlight:

```
scons: building `outputs\auto-modified.dta´ because it doesn´t exist
```

This tells us that `scons` has noticed that `auto-modified.dta` does not exist, and must be created. The way it is created is given in the next two lines:

```
stata_run(["outputs\auto-modified.dta"], ["code\dataprep.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\dataprep.do"
```

The first of these is what is sent to the Stata builder in `scons`: in words, "build the target `auto-modified.dta` using the source `dataprep.do` in the way defined in the Stata builder." The second is what the builder tells the computer to do: run the do-file `dataprep.do` in batch mode.[10]

Having created `auto-modified.dta`, SCons now notices that its next target, `scatterplot.pdf`, does not exist, and needs to be built, which it does by running `analysis.do` in batch mode. SCons then concludes that all targets are up to date, and exits.

Incidentally, note that the order in which the tasks appear in the SConstruct is not important. Even though the SConstruct lists the analysis task first, the logic of the build requires that SCons perform the dataprep task first.[11]

---

10. Exactly where the Stata executable is on your machine will differ from user to user. We have written `statacons` to find it automatically for most standard setups, and allow users to adjust the default in configuration files. We discuss configuration files, by default named `config_project.ini` and `config_local.ini`, at greater length in Section 3.4.

11. Of course, the way the SConstruct is written affects how easy it is for a *human* reader to understand it. In some cases, it may be easier to follow if the main outputs are listed first, as we have done here. In other cases, a "chronological" ordering may be preferable.

Now if we ask `statacons` to rebuild, it will check whether the dependencies of any target have changed since the previous build.

```
. statacons, debug(explain)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
scons: `.´ is up to date.
scons: done building targets.
```

`statacons` does this by calculating file *signatures* – a compact string that will not change if the file contents do not change, and almost certainly will change if the file contents change – for each target and dependency, and comparing these signatures to those from the previous build, which it has saved in a database called `.sconsign.dblite`.[12] The `stataconsign` command (a Stata call to the SCons command `sconsign`) will print the content of the.`sconsign.dblite`.

As an example, here is the the current state of the `.sconsign.dblite` database for this Introductory Example:

```
. stataconsign
=== .:
SConstruct: None 1651866262 581
=== code:
analysis.do: af3fe510ee8586a10401fe8e9b4fef49 1652884882 317
dataprep.do: 57d511b57ba3f9d41a1c488dab10b3ff 1652884882 192
=== inputs:
auto-original.dta: 74:12(71728):3831085005:1395876116:17340132 1640899314 12765
=== outputs:
auto-modified.dta: 74:13(15616):2430311699:721316426:3189221131 1652884885 13701
        code/dataprep.do: 57d511b57ba3f9d41a1c488dab10b3ff 1652884882 192
        inputs/auto-original.dta: 74:12(71728):3831085005:1395876116:17340132 16
> 40899314 12765
        947c32d25892d9cb39480bf147f78ed3 [stata_run(target, source, env)]
regressionTable.tex: bb4d5242836537f8cb562435b9e17cb8 1652884889 931
        code/analysis.do: af3fe510ee8586a10401fe8e9b4fef49 1652884882 317
        outputs/auto-modified.dta: 74:13(15616):2430311699:721316426:3189221131
> 1652884885 13701
        947c32d25892d9cb39480bf147f78ed3 [stata_run(target, source, env)]
scatterplot.pdf: 5e43c38cd77e29c7235c2866a448742f 1652884889 65452
        code/analysis.do: af3fe510ee8586a10401fe8e9b4fef49 1652884882 317
        outputs/auto-modified.dta: 74:13(15616):2430311699:721316426:3189221131
> 1652884885 13701
        947c32d25892d9cb39480bf147f78ed3 [stata_run(target, source, env)]
```

Examining the database provides insight into how SCons and `statacons` work. For

---

12. Technically there are different types of functions: "hashes", which provide some cryptographic guarantees, and "checksums", which do not. In the text, we use the more generic term "signature" for all these functions.

the code (`analysis.do` and `dataprep.do`), the database is straightforward. Each is listed with three elements: a signature, a timestamp, and the file length. Similarly, for the input dataset `auto-original.dta`, the dataset entry consists of the same three items, although the signature is calculated differently, as we discuss in Section 5.

For *generated* files, i.e., files that are built by `statacons`, the database entries are more complex. As an example, consider the line for `auto-modified.dta`. Recall that this is the output of the dataprep task. The first line in the database entry for `auto-modified.dta` has the same structure as those of the code and input data: a signature, a timestamp and the file length. What is different is that the entry for `auto-modified.dta` has three additional lines: one for each of the two dependencies (`auto-original.dta` and `dataprep.do`) and one for the build action `stata_run`. The lines for the dependencies just replicate their own database entries. The final line is a signature for the build action. These are the three items that SCons needs to consider when deciding whether to rebuild `auto-modified.dta`. If any of the dependencies or if the build action have changed, SCons will decide to update `auto-modified.dta`. If none have changed, SCons will know that it does not need to rebuild. The remaining entries are all for derived files and all have the same structure: one line for the file itself, one line for each dependency and one for the build action.

In this case, `statacons` determines that no dependencies have changed since the last build, so nothing needs to be done.

Now, let's see what `statacons` does when we make a small modification to `analysis.do` only. We add a title to the scatterplot, changing

```
twoway scatter price mpg
```

to

```
twoway scatter price mpg, title("Price vs. MPG")
```

What happens when we run `statacons`?

```
. statacons, debug(explain)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
scons: rebuilding `outputs\scatterplot.pdf´ because `code\analysis.do´ changed
stata_run(["outputs\scatterplot.pdf", "outputs\regressionTable.tex"], ["code\ana
> lysis.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\analysis.do".
  Starting in hidden desktop (pid=17496).
scons: done building targets.
```

`statacons` sees that we need to rebuild our target `scatterplot.pdf`, because its source file `analysis.do` has changed. `statacons` also sees that there have been no changes to the dependencies for `auto-modified.dta`, so it does not need to rebuild that.

Next, let's explore what happens when we make an edit to `dataprep.do`. We add a variable `mpg_cub`, the cube of `mpg`:

```
generate mpg_cub = mpg^3
label variable mpg_cub "Mileage (mpg) cubed"
```

And we rebuild using `statacons`:

```
. statacons, debug(explain)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
scons: rebuilding `outputs\auto-modified.dta´ because `code\dataprep.do´ changed
stata_run(["outputs\auto-modified.dta"], ["code\dataprep.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\dataprep.do".
  Starting in hidden desktop (pid=22936).
scons: rebuilding `outputs\scatterplot.pdf´ because `outputs\auto-modified.dta´
> changed
stata_run(["outputs\scatterplot.pdf", "outputs\regressionTable.tex"], ["code\ana
> lysis.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\analysis.do".
  Starting in hidden desktop (pid=21064).
scons: done building targets.
```

First, `statacons` sees that `dataprep.do`, the source for `auto_modified.dta`, has changed, so `auto_modified.dta` must be rebuilt. Then, after rebuilding `auto_modified.dta`, `statacons` sees that it has changed, and since `auto_modified.dta` is a dependency for our ultimate targets `scatterplot.pdf` and `regressionTable.tex`, these targets must be re-built.[13]

Finally, let's make a small edit to `dataprep.do` that does not change the content of `auto-modified.dta`:

```
// this is an edit to dataprep.do that does not change auto-modified.dta
```

The result of rebuilding with `statacons` may be surprising:

```
. statacons , debug(explain)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
```

---

13. Careful readers of the underlying do-files might note that `analysis.do` does not actually use the variable `mpg_cub`. However, `statacons` does not know that – it only sees that a dependency has changed and, since the targets potentially could change, they must be rebuilt. See Section 4 for a strategy to avoid this sort of quasi-false positive.

```
scons: rebuilding `outputs\auto-modified.dta´ because `code\dataprep.do´ changed
stata_run(["outputs\auto-modified.dta"], ["code\dataprep.do"])
Running: "C:\Program Files\Stata16\StataMP-64.exe" /e do "code\dataprep.do".
  Starting in hidden desktop (pid=31036).
scons: done building targets.
```

This highlights an important way in which `scons`, and therefore `statacons`, differs from the classic and widely used build tool `make` and many of its derivatives.[14] As we would expect, `statacons` sees that `dataprep.do` has changed, so `auto-modified.dta` must be rebuilt. However, notice that `statacons` does *not* rebuild `scatterplot.pdf` and `regressionTable.tex`. This is because, while `auto-modified.dta` has been re-built, nothing about it has *changed* other than its timestamp. `statacons` will re-build a target only if a dependency has changed. When deciding whether to re-build `scatterplot.pdf` and `regressionTable.tex`, `statacons` computes a signature of `auto-modified.dta` that depends only on the content of the dataset, not on its timestamp. Then `statacons` will check whether this signature has changed since the previous time the target was built. Since none of the content of `auto-modified.dta` has changed, its signature is unchanged, and `statacons` sees that there is no need to rebuild any more targets.

3   Syntax

The syntax of `statacons` follows that of `scons`. For the most commonly used options, we have adapted SCons syntax into a style more familiar to Stata users. These options are listed below. We also allow users to use the standard SCons syntax, although not a mix of the two in the same command.[15]

The full syntax of `scons` is too long to replicate here, but can be found in the SCons manual (SCons Development Team 2021a). Here, we present the basic syntax and several of the most useful options, which we have made available in Stata syntax. For options not listed here, the user will need to use standard SCons syntax. More in-depth and comprehensive discussion can be found in the SCons User Guide (SCons Development Team 2021b).

The syntax of `statacons` is:

<u>statacons</u> $\big[$targets$\big]\big[$, options(values) $\big]$

The SCons equivalent is:

<u>statacons</u> $\big[$option$\big[$=value$\big]\big]\big[$targets$\big]$

By default, `statacons` will look in the current directory for an `SConstruct` file, consider all targets described by that file, and then build or re-build targets when needed.

---

14. For example, GNU Make: https://www.gnu.org/software/make/.
15. Invoking SCons from a terminal prompt requires SCons syntax.

3.1 Selecting targets

The user can select a single target or a subset of targets by specifying them on the command line. In the example from the previous section, to rebuild only `auto-modified.dta` but not the other targets, we would use:

. `statacons outputs/auto-modified.dta`

The SCons equivalent is identical.

Alternatively, the user can edit the `SConstruct` file to specify the default targets using the `Default()` function. See the SConstruct Syntax subsection below.

If no targets are specified on the command line and no `Default()` function given in the `SConstruct`, then `statacons` will consider all targets in the `SConstruct` in the current directory or sub-directories.

3.2 Options

Standard `scons` options

We discuss several of the most commonly used options. For a complete list of options, see the SCons manual (SCons Development Team 2021a).

`clean` will *clean*, i.e., delete, specified targets. Use this option with care, since by default it will remove all targets defined in the `SConstruct`.

   SCons equivalent: `-c`, `--clean`, `--remove`

`debug(type[,type...])` will print additional information to the screen to help *debug* the build process. What information is printed depends on the `type` or `type`s specified. The most useful `type` is `explain`, which will lead `statacons` to print an explanation for each target it selects to build or re-build – typically, which dependency has changed. Note that explanations are not given for targets that are not built.

   SCons equivalent: `--debug=type[,type...]`

`directory(directory)` change directory to `directory` before starting build

   SCons equivalent: `-C directory`, `--directory=directory`

`dry_run` will put SCons into *no execute* mode. SCons will attempt to determine what targets it would build and print the commands it would execute, but will not actually execute these commands.

   SCons equivalent: `-n`, `--no-exec`, `--dry-run`, `--just-print`

   It is important to note that because SCons by default rebuilds targets only when dependencies have changed, SCons may not be able to determine whether it will need to rebuild "downstream" targets. For example, suppose `A.dta` is a dependency of `B.dta` and `B.dta` is a dependency of `C.dta`. When `A.dta` changes, SCons knows

it must rebuild `B.dta`. However, since in *no execute* mode SCons does not actually rebuild `B.dta`, it does not know whether `B.dta` will *change*. As a result, it does not know whether it will have to rebuild `C.dta` and will not print any commands for rebuilding `C.dta`. In general, we can think of the default set of commands printed in *no execute* mode as the smallest set that will be executed.

`file(file)` or `sconstruct(file)` allows the user to specify *file* as the initial SConstruct file for `statacons` to read. By default, `statacons` will look for a file named `SConstruct` in the current directory.

SCons equivalent: `-f file`, `--file=file`, `--sconstruct=file`

`help` prints the full SCons help message to the screen. The user can define an additional help message in the `SConstruct` using the `Help` function, see the SCons manual for details. (`statacons, help` is distinct from `help statacons`, as the latter gives the help file for the Stata command, not the SCons program.)

SCons equivalent: `-h, --help`

`q` do not print SCons status messages ("quiet")

SCons equivalent: `-Q`

`silent` do not print messages about SCons actions nor SCons status messages ("silent"). Implies `q`.

SCons equivalent: `-s, --silent, --quiet`

`tree(type[,type...])` prints a *dependency tree*, which is useful for debugging or checking on the status of a build, or just for visualizing and understanding the workflow. The part of the tree printed and the format depend on the `type` or `type`s specified:

`all`: print the entire tree

`derived`: print only "derived" (i.e., target) files. Omit files that are not targets, e.g., code or input data.

`linedraw`: uses unicode characters to draw the tree rather than the default ascii. Some users may find this easier to read on the screen.

`status`: prints the current status of each item on the tree. This is useful for understanding the status of a build, since it will tell you which files are current ("C"), which files exist ("E"), and several other status categories.

`prune`: makes the tree easier to read by not repeating dependencies for files that have already been described by the tree. Without this option, `tree` prints a lot of redundant information. With the `prune` option invoked, `tree` will indicate that a file's dependencies have already been listed by enclosing the file name in [square brackets].

SCons equivalent: `--tree=type[,type...]`

We most commonly invoke `tree` as `tree(status,prune)`. For our introductory

```
. statacons, tree(status,prune) dry_run silent
 E         = exists
  R        = exists in repository only
   b       = implicit builder
   B       = explicit builder
    S      = side effect
     P     = precious
      A    = always build
       C   = current
        N  = no clean
         H = no cache
[E b   C  ]+-.
[E b   C  ]  +-code
[E     C  ]  | +-code\analysis.do
[E     C  ]  | +-code\dataprep.do
[E b   C  ]  +-inputs
[E     C  ]  | +-inputs\auto-original.dta
[E b   C  ]  +-outputs
[E B P C  ]  | +-outputs\auto-modified.dta
[E     C  ]  | | +-code\dataprep.do
[E     C  ]  | | +-inputs\auto-original.dta
[E B P C  ]  | +-outputs\regressionTable.tex
[E     C  ]  | | +-code\analysis.do
[E B P C  ]  | | +-[outputs\auto-modified.dta]
[E B P C  ]  | +-outputs\scatterplot.pdf
[E     C  ]  |    +-code\analysis.do
[E B P C  ]  |    +-[outputs\auto-modified.dta]
[E     C  ]  +-SConstruct
```

Figure 3: SCons tree from our Introductory Example

example, this produces the output shown in Figure 3.

Exporting the output of `tree` to a text file[16] will produce output similar to that of the BITSS / Social Science Reproduction Platform Diagram Builder (BITSS 2020).

Custom `statacons` **options**

All of the options above are standard to SCons. We have written several options specifically for `statacons`.[17]

`assume_built("target")` will instruct the Stata builder to skip a task if all of its targets are listed and then to mark those targets as up-to-date. This can be a colon-separated list of file patterns.

  SCons syntax: `--assume-built="target"`

`assume_done("filename.do")` will instruct the Stata builder to skip the given do-file

---

16. For example, in a Unix-like terminal, enter `scons -ns --tree=prune > outputFile.txt`

17. Standard SCons (i.e., invoked as `scons` from the command line) will process these commands as long as you have installed the Python package `pystatacons` and your `SConstruct` imports `pystatacons`. SCons-style syntax will be required.

in the current build, but mark the associated target(s) as up-to-date. This is useful if we know that a target is up to date but SCons does not recognize it as such, for example: (1) if we have just built `target.tex` by running `filename.do` directly in Stata rather than through `scons` or `statacons`; (2) if we have only made edits that we know will not result in any changes to `target.tex`, such as adding comments in `filename.do`. Works even if the task for this do-file includes parameters or a non-`do` file command. See discussion in Section 6.3, "Limitations."

SCons syntax: `--assume-done="`*filename.do*`"`

`assume_done("`*file name1.do*`:`*file name2.do*$\big[$`:` ... $\big]$`")` will split by the colon and skip each file. Each component can be a file pattern, as seen in the next example.

SCons syntax: `--assume-done="`*file name1.do*`:`*file name2.do*$\big[$`:` ... $\big]$"

`assume_done(*)` will skip all do-files in the current build and mark all their targets as up-to-date.

SCons syntax: `--assume-done=*`

- `config_file(CONFIG_FILE)`. By default, `statacons` will look for `config_project.ini` and `config_local.ini` in the same directory as the `SConstruct`. If these files do not exist or they do not specify several core parameters, such as the location of the Stata executable, defaults based on typical configurations for the user's operating system, version and edition of Stata, will be used. Setting `config_file(CONFIG_FILE)` will instruct the Stata builder to use CONFIG_FILE as the configuration file. See Section 3.4 for more on default settings and configuration files. The user can specify multiple files as FILE1:FILE2:...:FILEN. In case of conflicts between files, the file that appears later in FILE1:FILE2:...:FILEN will take precedence.

    SCons syntax: `--config-file=CONFIG_FILE`

- `show_config` will print the current values of all configuration options, whether chosen by default or specified in a configuration file. `statacons, show_config` activates the `dry_run` and `silent` options, so `statacons` will not build any files or report anything about the state of the build.

    SCons syntax: `--show-config`

3.3 SConstruct Syntax

Basic SConstruct Recipe

The basic recipe for encoding a target built by Stata into an `SConstruct` file is:

```
task_name = env.StataBuild(
      target = ['path/to/target1.ext', 'path/to/target2.ext'],
      source = 'path/to/dofile.do'
)
```

```
Depends(task_name, ['path/to/dependency1.ext',
                    'path/to/dependency2.ext']
)
```

In Section 2, we described each of the elements: *builder* (here, `StataBuild`), *target*, *source*, and *dependencies*. `target` and `source` are both required.[18]

Additional Options for SConstruct Recipe

We have added some options to the basic recipe:

```
task_name = env.StataBuild(
     target = ['path/to/target1.ext', 'path/to/target2.ext'],
     source = 'path/to/source.ext',
     file_cmd = "command",
     params = 'arguments or options',
     depends = ['path/to/dependency1.ext', 'path/to/dependency2.ext']
)
```

The additional options are as follows:

- `file_cmd`: the command that SCons should pass to Stata's batch mode. The default is `do`, but the user can specify anything that Stata can accept as a command, e.g., `dyndoc`, `markdown`, etc.

- `params`: arguments or options that should follow the `source` in the call to Stata batch mode. For example, a call to `markdown` might specify

  ```
  params = ', saving(myfile.html) replace'
  ```

- `depends`: an alternative to using the `Depends()` function

We include several additional examples of these options in our web tutorial.[19]

Defining Lists of Files in SConstructs

As projects grow in size, it may become unwieldy to write out long lists of dependencies. Furthermore, if different targets have overlapping sets of dependencies, writing the same dependencies in different places is inefficient and can lead to errors. We can do better by defining variables in our SConscript files using standard Python syntax and some special SCons functions. As a preview of the Applied Example in Section 4 below, we illustrate two of these methods.

---

18. For managing tasks that do not build specific targets, see "Alias and Phony Targets" below.
19. https://bquistorff.github.io/statacons/swc.html, especially the "Paramters" lesson.

One of the targets, `ctyQtrWages.dta`, uses as inputs five Excel files named `wagedata1.xlsx`, ..., `wagedata5.xlsx`. We could write these out explicitly in the `Depends()` line, but instead we will demonstrate two alternatives. The first is the SCons `Glob()` function:

```
wageDataFiles = Glob("inputs/wagedata[1-5].xlsx")
Depends(wage, wageDataFiles)
```

where `[1-5]` means "any element of the sequence 1, 2, ..., 5."[20] SCons will search the directory `inputs` for files matching the pattern and will define `wageDataFiles` to be the set of files that match.

A more flexible alternative is to use Python's "list comprehension" method as follows:[21]

```
wageDataFiles = ["inputs/wagedata"+e+".xlsx"
                 for e in ['1','2','3','4','5']]
Depends(wage, wageDataFiles)
```

While this method is not as concise as `Glob()`, it has a few advantages. First, it allows for more complex structures and patterns. For example, a sequence of ISO country codes would begin `'AFG'`,`'ALA'`,`'ALB'`. Second, it can define lists of targets that may not yet exist, while `Glob()` is limited to matching existing files. Third, it is more precise. As a concrete example, suppose the file `wagedata1.xlsx` is missing from the `dataprep/inputs` directory. If we write out all our dependencies explicitly or use the list comprehension method above, SCons will return an error, which is useful. If we use `Glob("inputs/wagedata[1-5].xlsx")`, SCons will think that the dependencies are `wagedata2.xlsx`, ..., `wagedata5.xlsx`, i.e., the files that are present, and will proceed with the build.[22]

See the companion web tutorial[23] for more examples of using Python coding to write more concise `SConstruct`s.

SConstruct Functions

We discuss only a few of the more useful functions available in SConstruct files. See the SCons manual or the SCons User Guide for a comprehensive list and more details on implementation. We also note a few additional options that the `pystatacons` package provides.

---

20. `Glob()` allows unix shell pattern matching, i.e., `*` (match everything), `?` (match single character), `[seq]` (match any singe character in the sequence `seq`), and `!seq` (match any single character not in the sequence `seq`). See the SCons User Guide for more details.
21. See https://docs.python.org/3/tutorial/datastructures.html for more information on list comprehension and other Python data structures.
22. Despite our concerns about `Glob()`, we have used it in our Applied Example for the sake of providing an example. See the definition of the `industryCompFiles` variable in `dataprep/SConstruct`.
23. https://bquistorff.github.io/statacons/swc.html, especially the "Variables" lesson.

Note that some of these functions (e.g., `SetOption()`), will require the use of *nix-like standard SCons syntax rather than our adaptation to Stata syntax.

`AlwaysBuild(target [,target...])`: instructs SCons that the given targets should always be rebuilt when they are specified, even if they are up to date. Note that a target must be specified either in the command line or as part of the default set – that is, `AlwaysBuild()` neither implies nor is implied by `Default()`.

`Default()`: allows the user to specify which targets SCons will examine by default. In our introductory example, we would set a default of rebuilding only `auto-modified.dta` but not the other targets by adding

    Default('outputs/auto-modified.dta')

to the `SConstruct`. If `Default()` is not specified, SCons will examine all targets in current directory or subdirectories unless otherwise specified on the command line.

`Decider(function)`: determine how SCons will decide whether a target is up to date.

`"content"`: the default. For each dependency of a target, SCons will calculate a signature and compare the result against the last time the target was built.[24] This means that a target will be rebuilt only if one if its dependencies have *changed*. So, suppose `A.dta` is a dependency of `B.dta` and `B.dta` is a dependency of `C.dta`. Now, suppose that `A.dta` changes, but when `B.dta` is rebuilt, it does not change. Even though `B.dta` is *newer* than `C.dta`, SCons will not rebuild `C.dta`.

`"timestamp-newer"`: SCons will rebuild a target if any of its dependencies are newer than the target itself (similar to the classic build tool `make`). Suppose `A.dta` is a dependency of `B.dta` and `B.dta` is a dependency of `C.dta`. Now, suppose that `A.dta` changes, but when `B.dta` is rebuilt, it does not change. Even though `B.dta` has not changed, it is newer than `C.dta`, so SCons will rebuild `C.dta`. Because it does take time for SCons to calculate signatures, `timestamp-newer` may be faster despite some unnecessary rebuilds if there are relatively many input files (and therefore many signatures to calculate) and relatively low computation time.

`"timestamp-match"`: SCons will rebuild a target if any of its dependencies have a timestamp different (newer or older) than the last time the target was built. If all dependencies' timestamps are the same as the last time the target was built, the target will be considered up to date and will not be rebuilt.

`"content-timestamp"`: rebuild a target only if a dependency's timestamp *and* signature have changed since the last time the target was rebuilt, that dependency is considered up to date. This is similar to `content`, except it skips checking signatures for dependencies whose timestamps are unchanged and therefore may run faster.

`"content-timestamp-newer"`: rebuild a target only if a dependency's signature

---

24. Rather than recalculating the signature, SCons will save time by using a cached value of the signature if (1) a cached value is available from previous builds, (2) the modification timestamp has not changed, and (3) the timestamp is at least as old in seconds as the `max-drift` parameter (default is two days, but is configurable).

has changed since the last time the target was rebuilt *and* any of its dependencies are newer than the target itself. Similar to the `assume_done` and `assume_built` command-line options defined above, this means that SCons will not re-run Stata code that has already been run outside of SCons. Unlike `assume_done` and `assume_built`, it only needs to be specified once in the SConstruct, not specified at the command line each time. `"content-timestamp-newer"` is defined by `pystatacons`. To use it, include `Decider(pystatacons.decider_str_lookup["content-timestamp-newer"])` in your SConstruct after `pystatacons` has been imported.

`Ignore(`*target, dependency*`)`: tell SCons to ignore the file *dependency* when deciding whether to rebuild *target*.

`Ignore(`*directory, target*`)`: remove *target* from the set of default targets. The first argument of `Ignore()` must be the *directory* where *target* would be rebuilt. SCons will still build *target* if (1) *target* is specified on the command line or (2) *target* is a dependency of another target and needs to be rebuilt.

`Requires(`*target, prerequisite*`)`: specify that *target* should be built after *prerequisite* even if *prerequisite* is not a dependency for *target*. This is useful if the user wishes to impose some order on the build even if that order is not required by the build logic. For example, the user may want to move longer-running tasks later in the build so she can inspect the output of shorter-running tasks while the others run.

`SetOption(`*name, value*`)`: allows some command-line options to be set in the SConstruct. For example, `SetOption('silent',1)` is equivalent to `silent` on the command line. See the User Manual for the options that can be set by `SetOption()`. The command line over-rides values set through `SetOption()`.

Alias and Phony Targets

The `Alias()` function allows users to give names to targets or groups of targets. This can make it easier to build only a subset of targets, or to refer to groups of targets in an SConstruct.

In our Introductory Example, suppose we wanted a shorthand for the dataprep process so that, rather than call `outputs/auto-modified.dta` from the command line, we could use just `dataprep` which is shorter and easier to remember. We would code:

```
dataprep_Targets = ['outputs/auto-modified.dta']
Alias('dataprep',dataprep_Targets)
```

into our SConstruct and then, if we wanted to build only the dataprep targets, call

```
statacons dataprep
```

instead of

```
statacons outputs/auto-modified.dta
```

For additional uses of `Alias()` to create a convenient way to reference groups of files,

see `SConstructWithAlias` and other SConstructs from the Introductory Example.[25]

A second use of `Alias()` is to incorporate commands associated with a project that do not build files: configuring the local environment, communicating with a server (e.g., publishing a package), or displaying information in the terminal. For convenience, these can be stored in your SConstruct by assigning them to "Phony" targets (targets that are not real files) using Alias. For example, if we put the following line in our SConstruct, we could run the shell command `dir` whenever we run SCons with the target show_dir.

```
env.AlwaysBuild(env.Alias("show_dir", action="dir"))
```

As show_dir is not a real target with an associated build recipe, we encase `Alias()` in `AlwaysBuild()`. The command `dir` could be replaced with a Python function (potentially user-defined) to allow for more options, as we will see in Section 6.1.

3.4   Default Settings and Configuration Files

We have incorporated what we think are sensible defaults into `statacons`. These include:

- finding the Stata executable automatically by searching the user's `PATH` variable and other default locations;

- using our custom `complete_datasignature.ado` to compute signatures for `.dta` that depend only on their content, not their embedded timestamp. (See discussion in Section 5);

- deleting the batch-mode generated log-files after a do-file is successfully executed in batch mode, but retaining it in the default directory if the do-file fails.

By setting these defaults, we reduce the need for users to spend time figuring out how to configure the tool, which we suspect has been a barrier to adoption of build tools previously. However, some users may wish to choose other options or have unusual setups. Therefore, we allow users to change their configurations in two files, `config_project.ini` (for settings common to all users in a given project) and `config_local.ini` (for user-specific settings). We provide templates for both with suggested values in the project files installed with `statacons`. In case of conflicts between `config_project.ini` and `config_local.ini`, `config_local.ini` will take precedence.

These configuration files follow a simplified INI format readable by Python's standard library[26]. They are simple text files that are easily read and edited. Each file stores key-value pairs organized under one or more headers.

```
[header1]
```

25. https://github.com/bquistorff/statacons/raw/main/examples/stataconsIntro.zip
26. See https://docs.python.org/3/library/configparser.html for full details

```
key1: value 1
key2: value 2

[header2]
key3: value 3
```

As an example, the user may wish to retain the log-files created by batch-mode Stata in a folder called `logs`, which can be done in `config_local.ini` with

```
[SCons]
success_batch_log_dir: ./logs/
```

As a second example, by default SCons operates relative to the directory where the `SConstruct` is found, so this is the directory from which Stata will run in batch mode. This is a reasonable default, but some users may wish for Stata to start in a different directory. We allow this through a `stata_chdir` option in `config_project.ini`. See the `config_project.ini` file provided with this package for the available options.

We provide some additional examples of the use of configuration files elsewhere in this paper, such as telling SCons where to find shared folders in Dropbox. Users can specify other configuration files using the `config_file()` option listed in the Custom `statacons` options section above. As mentioned there, in the case of conflicts between config files, the file listed later in `config_file()` will take precedence.

3.5   Other Advanced Features

We conclude this section by briefly mentioning a few advanced features of SCons. We have additional discussion and examples in the Appendices. This is not a complete list nor a comprehensive guide; rather, we aim to familiarize readers with some concepts and vocabulary they may encounter while using the tool.

The first is the `SConscript`, which allows users to split a long `SConstruct` into more manageable parts, i.e., a *hierarchical build*. See Appendix A.1.

This second is SCons's ability to manage *parallel builds*, conducting tasks simultaneously rather than sequentially when the logic of the workflow permits. SCons parallel builds are well-suited for a relatively small number of relatively slow tasks. For tasks with a relatively high number of relatively quick tasks, like bootstrapping or simulation, the approach of `parallel` (Vega Yon and Quistorff 2019) is preferable. We provide an application of parallel builds to our Applied Example in Appendix A.2. Parallel builds are currently only available in `scons` (i.e., from a terminal), not in `statacons`.

The third is the SCons *cache*, which allows storing and sharing of derived files. We describe the use of the cache in Appendix B in the context of collaborative workflows.

4   Applied Example

This section describes the use of `statacons` in an empirical project (Shumway and Wilson 2021). Key build scripts and do-files are provided in `appliedExample.zip`, posted to our repository.[27]

For our purposes, this project has two main parts: dataprep, in which raw data are cleaned and merged, and new variables are created for analysis; and analysis, in which the analysis of these data is carried out. The value of a build tool is apparent from Figures 4, 5, 6, and 7: rather than trying to remember all these interdependencies, we encode them in SConstructs. Consulting an SConstruct and the output of `tree()` is a more reliable way to document interdependencies, to orient a new collaborator, remind the user's future self of how a project works after some time away, trace errors, etc., than depending on memory. The SConstructs and output of `tree` for both the dataprep and analysis tasks are provided in `appliedExample.zip` on the project repository.
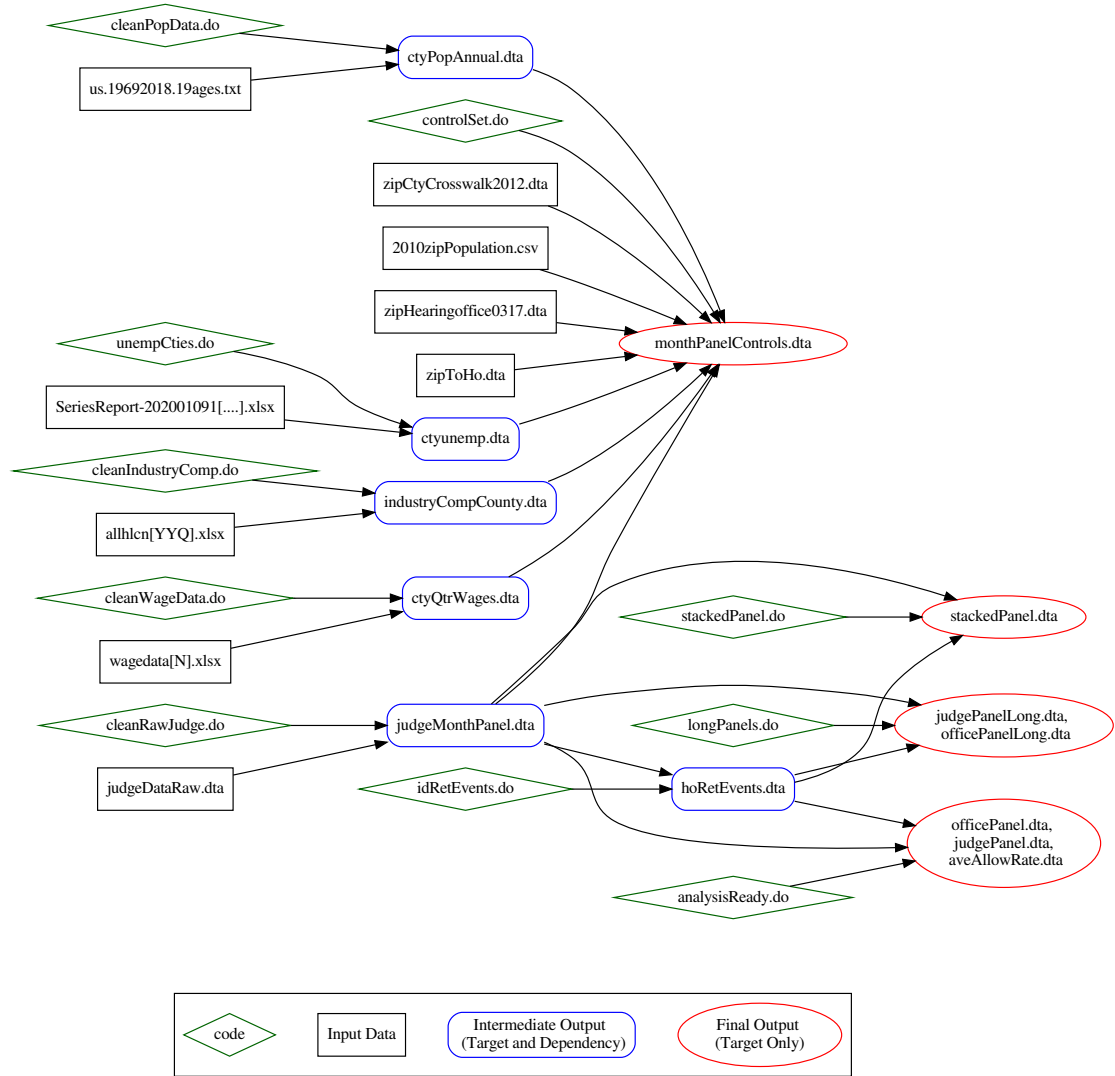
We have separated the two main parts of the project, data cleaning and data analysis, into two directories, each with its own `SConstruct`. These two directories are linked by a third, which we have called `transferDataprep`. The phony target `post_outputs` in the dataprep SConstruct copies the main outputs from `dataprep/outputs` to `transferDataprep`, while the phony target `pull_inputs` in the analysis SConstruct copies these from `transferDataprep` to `analysis/inputs`. See Section 6.1 for more details and an extension to shared files.

One of the estimation procedures, the dynamic treatment effects estimator of Sun and Abraham (2021), takes over two hours to run even on a fast desktop. We have written our code and SConstruct to avoid repeating this estimation when it is not necessary. The relevant section of the SConstruct is shown in Figure 8 along with a diagram.

First, this task uses only a subset of the variables and observations in `judgePanel.dta`. While we do need to re-run the estimation if the data used have changed, we would not want to re-run the estimation if the only variables or observations to change were not used in this task. The first task, `sunAbrahamSelect`, creates `sunAbrahamEstimationSample.dta`, which keeps only the variables and observations of `judgePanel.dta` needed for this task. So, suppose `judgePanel.dta` has changed. SCons will rebuild `sunAbrahamEstimationSample.dta`. If `sunAbrahamEstimationSample.dta` changes, then SCons will go on to the estimation step in `sunAbrahamEst`, since `sunAbrahamEstimationSample.dta` is the dependency for that task. However, if `sunAbrahamEstimationSample.dta` has not changed, then SCons will not repeat `sunAbrahamEst`.
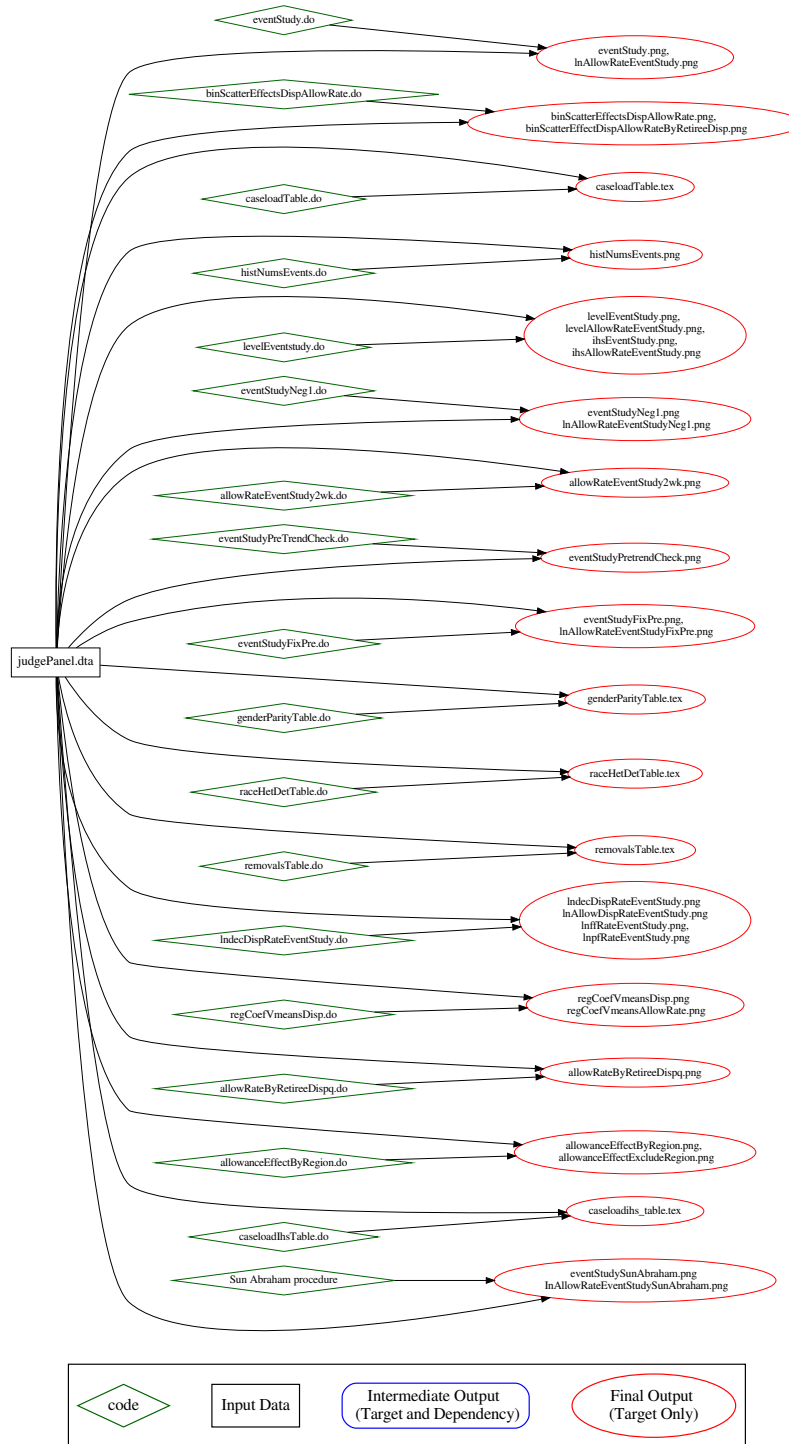
Similarly, the ultimate outputs are two graphs of estimated treatment effects. Suppose we wanted to change some aspect of the graphs, e.g., the axis titles. We would not want to have to repeat the full estimation. On the other hand, we do want to re-create these graphs whenever the underlying estimates change. We handle this by separating the estimation and the creation of graphs into distinct steps. The task `sunAbrahamEst`

---

27. https://github.com/bquistorff/statacons/raw/main/examples/appliedExample.zip

Figure 4: Workflow for Applied Example: Dataprep

Notes: brackets represent a series of similarly-named files, e.g., `wagedata[N].xlsx` represents `wagedata1.xlsx`, ..., `wagedata5.xlsx`. We have omitted file paths for brevity.

Figure 5: Workflow for Applied Example: Analysis (Part 1)
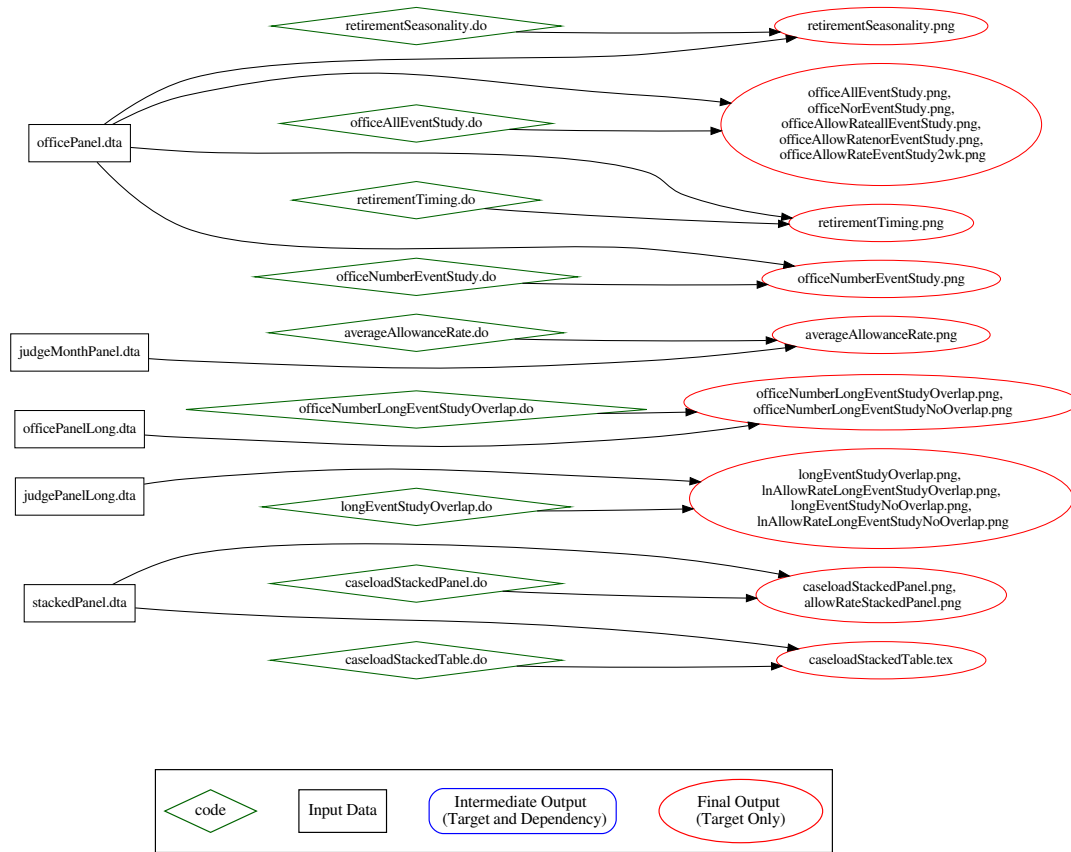
Notes: We have omitted file paths for brevity. The workflow for the Sun-Abraham event study procedure is simplified, see Figure 8 for the complete workflow.

Notes: We have omitted file paths for brevity.

Figure 6: Workflow for Applied Example: Analysis (Part 2)

Notes: We have omitted file paths for brevity.

Figure 7: Workflow for Applied Example: Analysis (Part 3)

```
# Sun-Abraham event study estimation methodology
sunAbrahamSelect = env.StataBuild(
    source = "code/sunAbrahamSelect.do",
     target = [´outputs/dta/sunAbrahamEstimationSample.dta´]
)
Depends(sunAbrahamSelect, [´inputs/judgePanel.dta´])

# do estimation
sunAbrahamEst = env.StataBuild(
    source = "code/sunAbrahamEstimation.do",
    target = [´outputs/dta/sunAbrahamResults.dta´]
)
Depends(sunAbrahamEst, [´outputs/dta/sunAbrahamEstimationSample.dta´])

# plot results
sunAbrahamFigs = env.StataBuild(
    source = "code/sunAbrahamGraph.do",
    target = [´outputs/figures/eventStudySunAbraham.png´,
        ´outputs/figures/lnAllowRateEventStudySunAbraham.png´]
)
Depends(sunAbrahamFigs, [´outputs/dta/sunAbrahamResults.dta´])
```
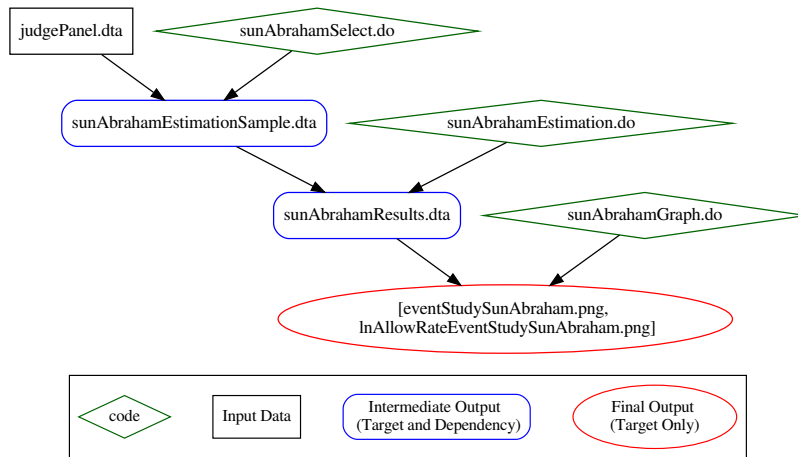
Figure 8: Separation of Concerns in Applied Example

handles the estimation, saving the estimation results in `sunAbrahamResults.dta`.[28]
Then the task `sunAbrahamFigs` handles producing the graphs. If we edit `sunAbrahamGraphs.do`
to change some aspect of the graphs, SCons will not see any need to rebuild `sunAbrahamEst`.
On the other hand, if something about the estimation changes such that `sunAbrahamResults.dta`
changes, SCons will automatically rebuild the graphs.

This division of tasks is known as *separation of concerns* in computer science and is in
our view a key distinction between working with a build tool and working in the literate
programming paradigm described in Section 6.2. The virtue of literate programming
is that everything is in one piece of code, which is certainly convenient and simple. A
build tool workflow in which discrete tasks are separated will save computation time.
Of course, there is a cost to creating separate do-files, so this tradeoff may not be
worthwhile for simple, quick tasks.

Finally, we can further reduce computation time by using SCons's ability to manage
parallel builds. We discuss this in Section A.2 of the Appendix.

5   Technical Details

We note here some technical details of how our package works.

First, as we have discussed above, SCons's default behavior is to look for changes
in the content of dependencies when deciding whether to rebuild a target. SCons does
this by computing a *signature* of each dependency. However, because Stata embeds
a timestamp in `.dta` files, `.dta` files created at two different times will have different
signatures even if the contents are identical. This will lead to unnecessary rebuilds of
target files. Stata's `datasignature` command can create signatures that depend only
on the values of the data in the `.dta` file, not the timestamp, but `datasignature`
signatures do not depend on variable labels or value labels.

We create an ado program `complete_datasignature.ado` that augments Stata's
datasignature with a signature of the dataset's metadata (variable formats, variable
and value labels, notes and characteristics) but not the timestamp. We then patch
SCons so that it uses our `complete_datasignature.ado` instead of its default signature
algorithm. In the `config_project.ini` file, the user has several options:

- `Strict`: include data and all metadata in the signature (variable formats, variable
  and value labels, notes and characteristics)

- `LabelsFormatsOnly`: include data, variable formats, and variable and value labels
  in the signature (do not include dataset labels, notes, characteristics)

---

28. Here we have saved estimation results into a `.dta` dataset. An alternative is to use Stata's `estimates`
    `save` command to save results to a `.ster` file. The user-written `estwrite` and `estread` (Jann 2005)
    improve on `estimates save` by allowing users to store multiple estimation results in a single `.sters`
    file and then refer to individual estimates by model name. Note that as of version `1.2.5`, `estwrite`
    has an option `reproducible` that will eliminate the timestamps previously embedded in `.sters`
    files. As of this writing, `.ster` files produced by `estimates save` contain an embedded timestamp.

- `DataOnly`: use Stata's standard `datasignature`, which depends on data only (does not depend on any metadata or the embedded timestamp)

- `False`: use SCons's default signature, which will depend on the embedded timestamp

In addition, we provide some control over the use of `datasignature`'s `fast` mode. The fast mode speeds up the calculation of the file signature, but is not *machine-independent*: users on different machines may produce different signatures. This is not so important for solo users working on a single machine, but for collaborative workflows (or a single user working on different machines) that save time by sharing built files in a common directory, using the `fast` option may lead `statacons` to incorrectly believe that a build is out of date. We make `fast` the default unless the `CacheDir()` option is specified (indicating a collaborative workflow, see Appendix B), in which case the machine-independent but slower method is the default. Users can override these defaults, for example to use the slow option for a solo project, through the `dta_sig_mode` option in the `config_project.ini` configuration file.

Our next set of issues deals with SCons directing Stata to run in batch mode. First, SCons needs to know if the do file finished successfully to know if targets should be considered built. This would typically be communicated via the program's exit status, but Stata always returns a normal exit even if the script it runs encountered an error. Our Stata builder therefore parses the end of the Stata batch-mode log file to see if there was error and communicate this back to SCons. This requires that batch-mode Stata produce plain-text log-files, so we do not allow the user to change the options with which batch-mode Stata is called.

Second, Stata typically uses the do-file name as the name of the batch-mode log-file, but we may need to change this for a variety of reasons. If two commands would result in the same log file name (e.g., we call the same do-file with different parameters or we have multiple non-`do` commands), then they can't be run in parallel because Stata will fail to start when the second instance can't open the log file. Therefore, for all commands aside from plain do commands (with no parameters), we have Stata run a generated do-file that then runs the user's command. We name this generated file so that its resulting log file will be unique (we take the original command's typical log-file name and append a short hash of the text of the command).

Finally, although SCons is a Python tool, and Stata 16+ has a Python environment, several hurdles need to be overcome to use SCons from inside Stata. First, Stata sets a handler for several program "signals" for Python at startup. Later, SCons temporarily sets new signal handlers. When SCons is done, though, it tries to reinstate the original handlers in a way that does not deal correctly with the handlers Stata set from outside Python. We therefore patch the SCons function that reinstates the signal handlers to deal with this case and avoid a crash. Second, Stata changes the `stdout` and `stderr` communication channels for the Python process at startup. Later, Scons temporarily changes these, but does not restore them correctly when finished. We patch this restore function so that output can continue to be displayed in the Stata window. Third,

Stata's Python environment is kept intact between Python script calls. SCons, however, is typically only able to be called from a clean environment. We therefore remove references to SCons modules before starting another run.

6 Extensions, Alternative Approaches and Limitations

6.1 Extensions

There are several ways to add to the basic functions we have described to this point. In increasing order of difficulty, these are: use additional build tools already supplied with SCons; use built-in SCons commands; execute Python commands and utilities; add user-written SCons build tools; write a new builder. We provide a few examples of the first three, and provide references for those interested in the last two.

Additional Built-in Build Tools: Compiling LaTeX with SCons

SCons comes with a builder for LaTeX documents. As long as the user has a TeX distribution (e.g., TexLive, MiKTeX, etc.) installed and known to the system, this will work more or less out of the box. Conveniently, SCons will automatically scan the `.tex` file for dependencies, so it is not necessary to code dependencies explicitly into the SConstruct. By default, these implicit dependencies include `.tex` files that have been `\include{}`ed or `\imported{}`ed, graphics files (`.pdf`, `.eps`, etc.) included through `\includegraphics{}`, `.bib` files, among others. The user may need to adjust some construction variables to get the document to compile in the desired way. In Appendix C, we add compiling a `.tex` file into a `pdf` to our Introductory Example.

All tools built into SCons are listed in the SCons User Guide.[29]

Built-in SCons Commands: "Install" (Copy) Built Files Across Directories

SCons has several built-in utilities that can be convenient tools. See Sections 11 and 12 of the SCons User Guide for a full list. One simple and especially useful tool, `Install()`, copies files across directories. We use `Install()` extensively in the code for this paper.[30]

First, in our Introductory Example, we use `Install()` to post our table, graph and compiled PDF to Overleaf for use in this paper. We need to tell statacons where it should copy these files, i.e., the Overleaf directory. In general, Overleaf files will reside in `Dropbox/apps/Overleaf/ProjectName`, where in this case our ProjectName is `SJ-BuildTool`. If there is only one user, or if all users have the same path to Dropbox,[31] we can define this directly in the SConstruct. For a Windows user, this would be:

---

29. As of Version 4.3.0 (December 2021), this is "Appendix C. Tools,", https://scons.org/doc/production/HTML/scons-user/apc.html.

30. All code is available on our project website, and the project files included with statacons contain templates for the configuration files discussed below, with extensive examples.

31. For example through a symbolic link.

```
# define Overleaf directory
overleaf_dir = 'C:/Users/UserName/Dropbox/apps/Overleaf/SJ-BuildTool'
```

If there are multiple users with different paths to Dropbox, each path must[32] be defined in the user's own `config_local.ini` configuration file:

```
# define Overleaf directory
overleaf_dir: C:/Users/UserName/Dropbox/apps/Overleaf/SJ-BuildTool
```

and then, in the SConstruct, we read in this variable:

```
# read in overleaf directory from config_local.ini
overleaf_dir = env['CONFIG']['Project']['overleaf_dir']
```

Having obtained the target location, we can add our `Install()` command to the SConstruct:

```
# transfer files to overleaf
env.Install(overleaf_dir, [cmd_analysis, pdf_output])
env.Alias('post_to_overleaf', overleaf_dir)
```

`Install()` takes two arguments: the first is the directory to which the files should be copied; the second is the list of files to be copied. Here, the first argument is the path to the Overleaf folder, as defined in the `overleaf_dir` variable. The second argument, `[cmd_analysis, pdf_output]`, tells SCons to copy the targets of the tasks `cmd_analysis` and `pdf_output`. Finally, in the last line, we use `Alias()` to give this task a convenient name. Now we can post our output files to overleaf with

```
statacons post_to_overleaf
```

Our second application of `Install()` is in our Applied Example.[33] There are two components to this project, dataprep and analysis. We want the outputs of dataprep to become inputs for analysis. A single user working alone could simply write code in analysis to use output files from dataprep directly, or could write a single `Install()` task in the dataprep SConstruct to push from `dataprep/outputs` to `analysis/inputs`, or in the analysis SConstruct to pull from `dataprep/outputs` to `analysis/inputs`. However, these strategies may not work well in a collaborative project, since a user in analysis would not want datasets to be changed without warning, and a user in dataprep would not want work in progress to be pulled into analysis. At the cost of slightly increasing complication, we set up a more deliberate workflow.

---

32. Adept Python coders may be able to program their way around this in the SConstruct by accessing the operating system environment. As we note in the Conclusion, we encourage users who develop enhancements to `statacons` to make them available publicly, and we provide a project Wiki to host.

33. Code posted at https://github.com/bquistorff/statacons/raw/main/examples/appliedExample.zip

First, we add a `transferDataprep` directory on the same level as the `dataprep` and `analysis` directories.

Second, we add an `Install()` task to the dataprep SConstruct that the user can select to push finished outputs from `dataprep/outputs` to `transferDataprep`.

```
# post outputs to ../transferDataPrep folder
env.Install('../transferDataPrep', outputs)
env.Alias('postOutputs', '../transferDataPrep')
```

To avoid transferring incomplete work, we want the `postOutputs` task to run only when it is explicitly called (`statacons postOutputs` at the Stata prompt or `scons postOutputs` in a terminal). In this case, `postOutputs` does not run automatically because its "target", the `transferDataPrep` directory, is not in the `dataprep` directory – by default, SCons will look for targets only inside the SConstruct's directory and subdirectories.

However, if the target directory is in `dataprep` or its subdirectories, we will need to remove `postOutputs` from the default targets. There are two ways to do this. The first is to set the `Default()` function explicitly and not include the transfer task among the defaults. This may be cumbersome if there are many targets. The second way is to use the `Ignore()` function in the SConstruct to tell SCons not to build `postOutputs` unless explicitly called from the command line. The code in the SConstruct would be

```
# post outputs to ./transferDataPrep folder
env.Install('./transferDataPrep', outputs)
env.Alias('postOutputs', './transferDataPrep')
# since ./transferDataPrep is a subdirectory, need
#  to Ignore postoutputs so it is not built by default
Ignore(postOutputs,'./transferDataPrep')
```

Third and finally, we set up a similar task in the analysis SConstruct to pull data from `transferDataprep` into `analysis/inputs`. In this example, we demonstrate the use of Python utilities, here `shutil.copytree`:[34]

```
# get inputs from folder "judgesExample/transferDataPrep"
# relative path from judgesExample/analysis is ../transferDataPrep
# need to specify alias pullInputs at command line
def copy_inputs(src = "../transferDataPrep", dst="inputs", **kwargs):
    import shutil
    shutil.copytree(src, dst, dirs_exist_ok=1)
env.AlwaysBuild(env.Alias("pullInputs", action=copy_inputs))
```

While we have made the `transferDataprep` directory local in this example, it could be a shared folder such as on Dropbox or Google Drive. Each user would just have to

---

34. The `dirs_exist_ok` option of `shutil.copytree` requires Python 3.8 or higher.

specify the right path in their `config_local.ini` configuration file. See Appendix B on collaborative workflows for additional comments.

Python Commands – Download and Unzip Archives

It is straightforward to use Python utilities in SCons. As an example, the data for our Applied Example are stored on Dropbox in a zip file. In `dataprep/SConstruct`, we include instructions to download and unzip the `.zip` file using Python utilities.

```
def url_download(fname = "judgesExampleInputs.zip",
     url="https://dl.dropboxusercontent.com/s/abc123/judgesExampleInputs.zip",
     **kwargs):
    import urllib.request
    urllib.request.urlretrieve(url, fname)
env.AlwaysBuild(env.Alias("dl_orig_data", action=url_download))

env.AlwaysBuild(env.Alias("unzip_inputs",
    action="python -m zipfile -e judgeExampleInputs.zip .")
)
```

Since these actions do not have targets, their aliases must be specified in the call to `statacons`, e.g.,

```
statacons dl_original_data
statacons unzip_inputs
```

Other Types of Extensions

There are many other ways to add to SCons. We list a few resources here. First, there are several user-written Tools implemented as Python packages.[35] Second, the user can write Commands, Builders or Tools. See the SCons User Guide for documentation.[36] A nice step-by-step example is provided by Wichman (2021).

6.2   Alternative Approaches

Alternative Build Tools

We chose to work in SCons for two main reasons. First, because it is a Python tool we could integrate it into Stata without requiring the use of a shell. Second, it provides the fundamental desirable features of build systems, as articulated in Mokhov et al. (2018): it is "minimal", in that it only rebuilds when dependencies have changed, and

---

35. See https://github.com/SCons/scons/wiki/ToolsIndex for a list.
36. As of Version 4.3.0 (December 2021), the relevant sections are 17, "Extending SCons: Writing Your Own Builder", and 18, "Not Writing a Builder: the `Command` Builder."

only runs each task once, and it provides "early-cutoff optimization," knowing when a newly built object has not changed and then ceasing to rebuild downstream targets. There are many alternatives to SCons; in this section we list several and discuss what we see as the advantages and disadvantages relative to SCons.

The best-known build tool is `make`. SCons has a few advantages relative to `make`. First, `make` rebuilds targets when any dependency's timestamp is newer than the target, which leads to unnecessary rebuilds if a dependency's content has not changed.[37] Second, `make` uses a custom syntax that is typically difficult for beginners (it uses many special characters and is unforgiving of rules for which whitespace characters to use). Third, `make` requires the use of a Unix shell (or emulator like Terminal on Macintosh or Windows Subsystem for Linux) to be used interactively (using `shell make` in Stata will not show you the output). One advantage of `make` over Scons is that it is somewhat easier to code straightforward commands into a script without having to create a build environment. A second advantage is the `makefile2graph` utility, which produces graphs of dependency trees. We are not aware of a similar utility for SCons.

There are many other build tools, such as `cmake`, some of which have support for using content signatures instead of file timestamps, but most are not in Python and so can not be used interactively inside of Stata. Among Python-based build tools, we chose to work in `scons` because its widespread use means there are active discussion and support communities.

The user-written `project` command in Stata (Picard 2013) is an ingenious pure Stata approach: dependencies are encoded into a project's do-files and the `project` program maintains a database of signatures to decide which outputs to rebuild. We prefer working in SCons for a few reasons, including integration with other tools, configurability, and access to advanced option such as the cache and parallel jobs, and that it does not require making changes to existing do-files.

Literate Programming

A more popular approach to reproducible research in Stata has followed the *literate programming* paradigm, in which text and code are combined in a single document. There are several excellent options in this paradigm, including built-in Stata commands such as `dyndoc` and `putdocx`, "notebooks" like the Jupyter notebooks popular in Python and now accessible in Stata 17, and user-written commands like `texdoc` (Jann 2016), `webdoc` (Jann 2017), `markdoc` (Haghish 2016), and `markstat` (Rodríguez 2017).

The literate programming approach is appealing and is certainly useful in many cases.[38] A literate programming script does provide a degree of self-documentation since input and output file names are present in the same file as the text. However, literate programming has limitations similar to those of a a master do-file: as projects grow more complex and computationally intensive, either all analyses must be repeated

---

37. SCons can replicate this behavior if desired, see the `Decider()` function mentioned in Section 3.3.
38. For example, we used `texdoc` with `statacons` to produce Section 2 of this paper, `dyndoc` with `statacons` to produce the companion tutorial website, and `markdoc` to produce our help files.

every time the document is produced, or the user must make some manual decisions about which parts to refresh and which to skip. See our discussion of *separation of concerns* in Section 4. Fortunately, a build tool such as `statacons` can complement a literate programming tool by automating some of these decisions. For example, our web tutorial consists of several web pages, each produced by a corresponding `dyndoc` file. We have written an SConstruct to manage this process, so that we rebuild a web page if and only if the inputs have changed. This code is available on our Wiki page.[39]

6.3   Limitations

The main limitation we have found in working with SCons is that SCons only recognizes builds it has performed itself. That is, SCons determines which targets need to be rebuilt by comparing the current state of the project to its state after the last time a project was built *by SCons*, as recorded in the `sconsign.dblite` database.[40]

As an example, suppose that in our Introductory Example, the user has edited `analysis.do` to change an axis title in the target `scatterplot.pdf`. The user runs her code from the Stata command window in the usual way (`do code/analysis.do`) and is satisfied with the result. In reality, the targets `scatterplot.pdf` and `regressionTable.tex` are now up to date. However, SCons will not realize this – SCons will compute a signature of `analysis.do`, see that it has changed since its signature was last recorded in `sconsign.dblite`, and believe that the targets need to be rebuilt.

In `make`, if the user knows that a build is up-to-date, the `--touch` option will change targets' timestamps to the present. Since `make` rebuilds only when dependencies are *newer* than the target, `make` will no longer rebuild the target.

SCons does not have a similar option, so we have added the `assume_built` and `assume_done` options to the `statacons` command, in which the user can designate certain targets that do not need to be rebuilt or do-files that do not need to be re-run. If the already built targets all have newer timestamps, then our provided "content-timestamp-newer" `Decider()` can also be used. See the respective entries in Section 3.2.

A second, related limitation is that setting up a collaborative project, for example keeping code under version control (e.g., Git) and sharing files through a service like Dropbox, requires a bit of care. Our preferred approach is a shared SCons *cache*. This is a powerful tool provided by SCons but there is a learning curve. We describe this method and some alternatives in Appendix B.

Finally, as mentioned in Section 6.2, it is easier to implement one-off tasks with `make` than with SCons.[41] We try to reduce this limitation by providing a few examples that users can adapt, but some familiarity with Python coding and some trial-and-error will likely be required to extend these examples to other tasks.

---

39. https://github.com/bquistorff/statacons/wiki/Statacons-and-literate-programming.
40. This applies for all settings of the `Decider()` function, including `timestamp-newer`.
41. At least, it is easier if the user is already comfortable using a shell and shell commands.

7   Conclusion

In this paper, we have provided introductions to build tools and `statacons`, a build tool for Stata we have adapted from SCons. `statacons` runs directly in Stata without an external shell, can be used with a minimum of initial configuration, and does not require modification of existing code. While there is a startup cost to integrating `statacons` into a project, and it requires some discipline to keep SConstructs up to date, the benefits – reliable reproducibility and saved computation time – are considerable.

We recommend that users start small: while retrofitting a large, existing project may prove frustrating for new users, using `statacons` and a build-tool mindset from the beginning of a new project can promote better code by encouraging breaking down tasks into parts, simplifying the scope of individual pieces of code, and separation of concerns. In debugging problems, starting small is also advised – first, run code in Stata to see if the problem is with the code rather than the build tool, then build one or a few targets at a time until the problem is isolated, examine the information provided by the `debug`, `tree` and `show_config` options and the `stataconsign` function.[42]

It is our hope that this paper will encourage the use of build tools in Stata. One barrier is that there are few worked examples that are simple enough for the novice to understand easily but also show some of the power of the tool. We provide a few such examples in this paper and the accompanying material. We hope that other users will build on this work by providing additional annotated examples and by extending the scope of tasks that can be performed. We encourage users to share these advances publicly, and we provide a project Wiki to host such contributions.

Acknowledgements

8   References

BITSS. 2020.   Guide for Advancing Computation Reproducibility in the Social Sciences.   Berkeley Initiative for Transparency in the Social Sciences. https://bitss.github.io/ACRE.

Christensen, G., Z. Wang, E. L. Paluck, N. Swanson, D. J. Birke, E. Miguel, and R. Littman. 2019. Open Science Practices are on the Rise: The State of Social Science (3S) Survey. preprint, MetaArXiv. https://osf.io/5rksu.

Gentzkow, M., and J. M. Shapiro. 2014.   Code and Data for the Social Sciences: A Practitioner's Guide.   University of Chicago mimeo. http://web.stanford.edu/ gentzkow/research/CodeAndData.pdf.

---

42. The do-file `debugging-checklist.do` included with the `statacons` package may be useful.

Haghish, E. F. 2016. Markdoc: Literate Programming in Stata. *The Stata Journal* 16(4): 964–988. https://doi.org/10.1177/1536867X1601600409.

Jackson, M. 2016. Software Carpentry: Automation and Make. http://swcarpentry.github.io/make-novice/.

Jann, B. 2005. ESTWRITE: Stata module to store estimation results on disk. Statistical Software Components, Boston College Department of Economics. https://github.com/benjann/estwrite.

———. 2007. Making Regression Tables Simplified. *The Stata Journal* 7(2): 227–244. https://doi.org/10.1177/1536867X0700700207.

———. 2016. Creating LaTeX Documents from within Stata using Texdoc. *The Stata Journal* 16(2): 245–263. http://journals.sagepub.com/doi/10.1177/1536867X1601600201.

———. 2017. Creating HTML or Markdown Documents from within Stata using Webdoc. *The Stata Journal* 17(1): 3–38. http://journals.sagepub.com/doi/10.1177/1536867X1701700102.

Mokhov, A., N. Mitchell, and S. Peyton Jones. 2018. Build systems à la carte. *Proceedings of the ACM on Programming Languages* 2(ICFP): 1–29. https://dl.acm.org/doi/10.1145/3236774.

Orozco, V., C. Bontemps, E. Maigné, V. Piguet, A. Hofstetter, A. Lacroix, F. Levert, and J.-M. Rousselle. 2020. How to Make a Pie: Reproducible Research for Empirical Economics and Econometrics. *Journal of Economic Surveys* 34(5): 1134–1169. http://onlinelibrary.wiley.com/doi/abs/10.1111/joes.12389.

Picard, R. 2013. PROJECT: Stata module providing a set of tools to build and manage a Stata project. Statistical Software Components, Boston College Department of Economics. https://ideas.repec.org/c/boc/bocode/s457685.html.

Rodríguez, G. 2017. Literate Data Analysis with Stata and Markdown. *The Stata Journal* 17(3): 600–618. http://journals.sagepub.com/doi/10.1177/1536867X1701700304.

SCons Development Team. 2021a. SCons 4.3.0 Man page. https://scons.org/doc/4.3.0/PDF/scons-man.pdf.

———. 2021b. SCons 4.3.0 User Guide. https://scons.org/doc/4.3.0/PDF/scons-user.pdf.

Shumway, C., and R. Wilson. 2021. Workplace Disruptions, Judge Caseloads, and Judge Decisions: Evidence from SSA Judicial Corps Retirements. Working Paper, forthcoming, Journal of Public Economics. https://go.ncsu.edu/shumway.wilson.judges.

Stodden, V., J. Seiler, and Z. Ma. 2018. An empirical analysis of journal policy effectiveness for computational reproducibility. *Proceedings of the National Academy of Sciences* 115(11): 2584–2589. http://www.pnas.org/lookup/doi/10.1073/pnas.1708290115.

Sun, L., and S. Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2): 175–199. https://www.sciencedirect.com/science/article/pii/S030440762030378X.

Vega Yon, G. G., and B. Quistorff. 2019. parallel: A command for parallel computing. *The Stata Journal* 19(3): 667–684. https://doi.org/10.1177/1536867X19874242.

Wichman, M. 2021. ToolsForFools. https://github.com/SCons/scons/wiki/ToolsForFools.

Wilson, G., D. A. Aruliah, C. T. Brown, N. P. Chue Hong, M. Davis, R. T. Guy, S. H. D. Haddock, K. D. Huff, I. M. Mitchell, M. D. Plumbley, B. Waugh, E. P. White, and P. Wilson. 2014. Best Practices for Scientific Computing. *PLoS Biol* 12(1): e1001745. http://dx.doi.org/10.1371/journal.pbio.1001745.

About the authors

Guiteras is Assistant Professor in the Department of Agricultural and Resource Economics and Core Faculty in the Global WaSH Research Cluster at North Carolina State University.

Kim is a Ph.D. student in Economics at North Carolina State University.

Quistorff is a Research Economist in the Office of the Chief Economist at the Bureau of Economic Analysis.

Shumway is a Ph.D. student in Economics at North Carolina State University.

The authors are listed alphabetically. All views expressed in this paper are our own and do not necessarily reflect those of our employers or any other institutions with which we are affiliated.

Quistorff is the corresponding author, though for concerns regarding the package (additional information, questions, suggestions, and bug reporting) please use our project repository: https://github.com/bquistorff/statacons/.

For the latest version of the software and installation instructions, go to the project web page: https://bquistorff.github.io/statacons.

For previous versions of the software and installation instructions, go to our project archive https://osf.io/gbh4m/.

Installation Guide

Requirements

These installation instructions correspond to `statacons` v3.0.0, released with this version of the paper. For up-to-date installation instructions, corresponding to the latest version of `statacons`, see our project web page: https://bquistorff.github.io/statacons.

The following are required for `statacons`:

- Stata 16.0 or later
- Python 3.5 or later[43]
- SCons 4.2 or later

Step 1: Python Setup

We provide the basic steps for installing Python and configuring Stata for use with Python. For more details, we recommend a series of posts by Chuck Huber on the Stata Blog.[44]

Python commands are issued at a command prompt in a terminal window. Exactly which terminal will depend on your operating system and implementation of Python. Some popular examples include

- Windows
  - Python: Windows Command Prompt ("`cmd.exe`") or Powershell
  - Anaconda: Anaconda Powershell Prompt or Anaconda Prompt
- macOS: Terminal
- Unix-like (e.g., Linux): terminal

Step 1.1: Install Python

Install Python 3.5 or later (3.8 or later recommended).

Step 1.2: Configure Stata to work with Python

Step 1.2.1: Locate Python executable

In Stata, enter

---

43. Some of the advanced options in our Applied Example require Python 3.8 or later.
44. https://blog.stata.com/tag/python/, see especially posts 1 (Huber, C., "Stata/Python integration part 1: Setting up Stata to use Python") and 3 (Huber, C., "Stata/Python integration part 3: How to install Python packages").

```
python search
```

to locate the Python executable. If your system has both Python 2 and Python 3 installed, be sure to use the Python 3 executable.

Step 1.2.2: Set path to Python executable in Stata

Stata will use the default Python, unless configured before first use. To do so, in Stata, enter

```
python set exec pyexecutable , permanently
```

where *pyexecutable* is a path found in Step 1.2.1.

If Stata is started in a specific Anaconda environment, we provide a simple command in Stata to query the environment and set the Python path accordingly.

```
set_python_exec_env
```

Step 1.2.3: check setup

In Stata, check your setup by entering

```
python query
```

Step 1.3: Install Python packages scons and pystatacons

`pip` is the standard way to install Python packages. In the terminal, enter

```
  pip install scons
  pip install pystatacons
```

To check that these are installed and recognized by Stata, open Stata and enter

```
  python which SCons // case-sensitive
  python which pystatacons
```

Step 2: statacons installation and setup

Step 2.1: Install statacons Stata package

First, check your `adopath`. By default, `statacons` will be installed to your `PLUS` directory. Users working in collaboration with others may wish to change the `PLUS` and `PERSONAL` directories to be local to the project so that all users will be using the same version of user-written do-files. See the `profile_template.do` we have included in our project files for what we recommend.

Second, `statacons` requires both the core program files and a key set of project files. You may wish to install these separately. To install the main program use:

```
net install statacons, from(https://raw.github.com/bquistorff/statacons/main/).
```

The core program files are:

```
complete_datasignature.ado
statacons.ado
staconsign.ado
runscons.py
sconsign-script.py
sconstruct_fns.py
```

To check that `statacons` is installed, enter `which statacons` in Stata.

Step 2.2: Install `statacons` project files

To install the project files, navigate to the root of your project directory and then run

```
// Change to project's root directory
cd path/to/your/project
// set other to the current directory, i.e. project's root directory
net set other .
// get project files
net get statacons, from(https://raw.github.com/bquistorff/statacons/main/)
unzipfile project_files
rm project_files.zip
```

The project files packages in `project_files.zip` are

```
SConstruct
config_project.ini
utils/config_local_template.ini
utils/profile_template.do
utils/debugging-checklist_template.do
utils/.gitignore_template
```

The only project file required is `SConstruct`. You may also wish to use the template files; if so, copy them into the main project folder (removing the `_template` suffix).

Step 2.3: Configuration

`statacons` is intended to work without additional user configuration, for example by automatically detecting the location of the Stata executable. However, some configura-

tion may be required. For example, if you have more than one version of Stata installed, `statacons` may not find your preferred version first.

Our default is to use two configuration files, `config_project.ini`, which sets parameters that are common to all users working on a given project, and `config_local.ini`, which sets parameters specific to the local user. The files `config_project.ini` and `config_local_template.ini` provided with the Stata project files give several examples of how to set commonly-used parameters. See Section 3.4 for more information on these configuration files and how to read from them in SConstructs.

The default is that these configuration files, if used, will be stored in the same directory as the SConstruct. The command line option `config_file()` can instruct SCons to use a custom location or filename instead.

Users can specify other locations and names for configuration files using the `config_file()` option, see Sections 3.2 and 3.4.

To check the configuration of `statacons`, enter `statacons, show_config` in Stata. The do-file `debugging-checklist.do` included in the `utils` folder may also be useful in checking that components are in the right places.

Step 2.4 (optional): Setup for Version Control

If you are using version control software such as git, here are a few recommendations:

- set your ado-path so that `statacons` and other packages are installed in the project folder. This will ensure that all users have the same version of all ado-files. See `profile_template.do` in the `utils` folder

- files to keep under version control: `SConstruct`, `config_project.ini`

- files to *exclude* from version control: `config_local.ini`, `debugging-checklist.do`, all generated files, including SConsign files (e.g., `.sconsign.dblite`) and batch-generated `.log`-files. See `.gitignore_template` in the `utils` folder for what we recommend.

See Appendix B for more details on working with version control.

Step 3: Test run with Introductory Example

Replication code and data for our Introductory Example as well as the extensions to that example are provided in `stataconsIntro.zip`, posted to our repository.[45] To run the example, unpack the zip archive and follow the steps in Section 2 of the paper.

---

45. https://github.com/bquistorff/statacons/main/raw/examples/stataconsIntro.zip

A   Other Advanced Features

A.1   SConscripts and Hierarchical Builds

SConscripts allow dividing a project into parts. Figure A1 shows how we can divide our Introductory Example into separate `SConscripts` for dataprep and analysis. Note the `Export()` and `Import()` statements needed to pass environments and variables across the `SConstruct` and the `SConscripts`.

In addition to breaking up large `SConstructs` into more manageable, readable parts, using `SConscripts` allows assigning different options or parameters to different parts of the project.

A.2   Parallel Builds

SCons can manage *parallel builds*, conducting tasks that are not interdependent simultaneously by dividing the computer's processors among the tasks. This can reduce computation time relative to running tasks sequentially.

At present, parallel builds are only supported in SCons itself, i.e., running `scons` from a terminal, not in `statacons`.

Using parallel builds without any modification to code can reduce computation time. Furthermore, writing code with parallel builds in mind can unlock additional gains. In our Applied Example, the event study methodology of Sun and Abraham (2021) is conducted for four different variables. Each estimation is independent of the others, so the four can be conducted in parallel. Figure A2 provides the `SConscript` for this process. We use an SConscript so that we can set the number of separate jobs for these tasks without affecting the rest of the build. Note the `SetOption('num_jobs',2)` line in the SConscript. We use 2 as a minimal working example, users with more processors may wish to allow more jobs.

We have included the do-files for this parallel build in our Supplemental Materials online. The basic idea is that (1) for each outcome variable, we (a) create a dataset with that variable and the relevant controls and identifiers, (b) perform the estimate for that outcome and save the results in a dataset, then (2) once the estimation is complete for all four variables, we combine the estimation results for the four outcome variables and present the results.

The gains from parallel builds will depend on many factors, including the nature of the task, the hardware, version of Stata, file size, and the system I/O throughput. For example, User A with 8 processors but a 4-core version of Stata MP will likely experience greater *relative* gains by adding a second, parallel job than User B with 8 processors and 8-core Stata MP: User A was only using 4 processors previously, and now can use all 8 simultaneously, still applying 4 to each task; User B was previously devoting 8 processors to each task sequentially, and now will devote 4 processors to each

task simultaneously.[A1] Similarly, tasks that use only a small fraction of the available RAM can run in parallel without much cost to the speed of the individual tasks, while tasks that are demanding of available RAM may not be sped up as much when run in parallel. Finally, as noted in the paper, `parallel` (Vega Yon and Quistorff 2019) is a better option for bootstrapping or simulation.

Some care should be taken in writing code in these cases – while there is no issue with multiple Stata processes *accessing* the same dataset at once, it is important that parallel tasks not *write* to the same dataset. This just requires a bit of caution that no output filenames are repeated across jobs.

A1. For some rough guidance on the types of tasks more or less likely to benefit from parallel builds, see the "Stata/MP Performance Report" (https://www.stata.com/statamp/report.pdf) for data on speed gains as the number of cores increases for different commands. Note that the gains as the number of cores increases is *negatively* related to gains from parallel builds if a fixed number of cores are being divided among tasks.

```
#--------- SConstruct file ----------
# **** Setup from pystatacons package *****
import pystatacons
env = pystatacons.init_env()

# **** Substance begins        *****
Export('env') # allow SConscripts to import environment
SConscript('SConscript-dataprep')
SConscript('SConscript-analysis')




# -------- SConscript-dataprep ------------
# get environment from SConstruct
Import('env')

dataprep_Targets = ['outputs/auto-modified.dta']
Alias('dataprep',dataprep_Targets)

cmd_dataprep = env.StataBuild(
    target = dataprep_Targets,
    source = 'code/dataprep.do'
)
Depends(cmd_dataprep, ['inputs/auto-original.dta'])

# allow SConscript-dataprep to access dataprep_Targets variable
Export('dataprep_Targets')




# -------- SConscript-analysis ------------
# get environment from SConstruct
Import('env')
# get dataprep_Targets variable exported by SConscript-dataprep
Import('dataprep_Targets')

analysis_Targets = ['outputs/scatterplot.pdf',
               'outputs/regressionTable.tex']
Alias('analysis',analysis_Targets)

cmd_analysis = env.StataBuild(
    target = analysis_Targets,
    source = 'code/analysis.do'
)
Depends(cmd_analysis, dataprep_Targets)
```

Figure A1: An SConstruct file with two SConscript files

```
# ** Setup
import pystatacons
env = pystatacons.init_env()

# use an sconsign database specific to this SConstruct
SConsignFile("sconsignParallel")

# set up parallel
# set default maximum of 2 jobs
#SetOption(´num_jobs´,2)
# or set from command line, e.g.,
# scons -j 2 --file=SConstructParallel

# selection:
#  pare down dataset to variables and observations used
#  create one dataset per outcome variable
select_parallel = env.StataBuild(
   source = "code/sunabraham_select_parallel.do",
    target = [´outputs/dta/sunAbrahamSample_jallowrate.dta´,
              ´outputs/dta/sunAbrahamSample_awards.dta´,
              ´outputs/dta/sunAbrahamSample_decisions.dta´,
              ´outputs/dta/sunAbrahamSample_dispositions.dta´])
Depends(select_parallel, [´inputs/judgePanel.dta´])

# estimation
#  one task per outcome variable

cmds={}
for outcome in ["awards","decisions","jallowrate","dispositions"]:
    cmds[outcome]=env.StataBuild(
        source = "code/sunabraham_estimation_"+outcome+".do",
        target = [´outputs/dta/sunAbrahamResults_´+outcome+´.dta´],
        depends = [´outputs/dta/sunAbrahamSample_´+outcome+´.dta´]
    )


#  plot results
figs_parallel = env.StataBuild(
    source = "code/sunabraham_graph_parallel.do",
    target = [´outputs/figures/sunabraham_parallel.png´,
              ´outputs/figures/lnallowrate_sunabraham_parallel.png´])
Depends(figs_parallel,
    [´outputs/dta/sunAbrahamResults_awards.dta´,
     ´outputs/dta/sunAbrahamResults_jallowrate.dta´,
     ´outputs/dta/sunAbrahamResults_decisions.dta´,
     ´outputs/dta/sunAbrahamResults_dispositions.dta´])
```

Figure A2: SConstruct: Parallel Jobs in Applied Example

B   Collaborative Workflows

The Problem

In this Appendix, we discuss a few models for collaborative workflows. The main diffi-
culty with collaboration is that, as we mention in the Limitations section (6.3), SCons
only understands builds it has done itself. That is, SCons compares the current state of
targets and dependencies to their state as recorded in the `.sconsign.dblite` database
the last time SCons ran the build.

As an example, suppose that Users A and B are collaborating on the Introductory
Example of Section 2. User A is working primarily on the dataprep task and User B is
working primarily on analysis. Imagine for the sake of this example that the dataprep
task takes a very long time to run. User A has made some edits to `dataprep.do` and
updated `auto-modified.dta`. User A shares these two files with User B.

The issue is that User B's SCons will look at `dataprep.do`, compare its current
status to its status the last time `auto-modified.dta` was built *as recorded in User B's*
`.sconsign.dblite`, see that `dataprep.do` is different, and conclude that `auto-modified.dta`
needs to be rebuilt. User B could use the `assume_built()` or `assume_done()` options we
have added to SCons to tell SCons that in fact `auto-modified.dta` is up to date, but
in more complex projects this sort of mental dependency-tracking is prone to failure.

The underlying problem is that there are three competing values: we want builds
to be complete (all targets up-to-date), we do not want to duplicate work, and we
want to automate the build and avoid manual workarounds, which are error-prone. Our
preferred approach is to use the SCons *cache* to share built files. While this requires a
small amount of effort to set up and some care to maintain, it satisfies all three values
without introducing too much additional complexity. We will describe this approach
below, then mention a few alternatives and their limitations. Our project Wiki contains
a simple worked example of a collaborative project using the SCons cache along with
GitHub.[B1]

SCons Cache

Our preferred approach is to use the SCons *cache* to share derived files. See Figure B1
for an example SConstruct. In the SConstruct, we designate a shared folder to store
the cache:

```
CacheDir('path/to/cache')
```

Rather than hard-code the path into the SConstruct, the path can be coded into
a configuration file (e.g., `config_project.ini`, `config_local.ini`, or other specified
through the `config_file()` option), which the SConstruct can then read. Here, we use a
local `scons_cache` folder so that this can be a self-contained example, but in practice this
may be a shared drive or a folder on a file-sharing service. See `config_local_template.ini`

---

*B*1. https://github.com/bquistorff/statacons/wiki/Collaboration-using-GitHub-and-the-SCons-cache

for several examples.

SCons will then store built files in the cache and, before building targets, check to see whether an up-to-date version is available in the cache. See the SCons User Guide for details on how this works.[B2]

We have found that this approach handles the three competing goals of a build tool (complete builds; reduce unnecessary builds; automaticity instead of manual intervention) well, with a few caveats.

The main caveat is that by default SCons will update the cache whenever a target is newly built, which is not desirable while writing and testing code since it will lead to filling the cache with unnecessary files, and other users will be repeatedly downloading these unnecessary files. It is good practice to update the cache only when you are sure the work is complete. There are two ways to do this.

The first way is to use the command-line option `cache_readonly` to tell SCons to check and read from the cache but not update the cache with any built files, or `cache_disable` to tell SCons to ignore the cache entirely. Once you have successfully completed your build, use the option `cache_force` to force SCons to update the cache. This will also make it less likely that two users will attempt to update the same cached file at the same time.

The second way is to edit the SConstruct so that the default is to have the cache disabled, and a command-line option is required to read to or write to the cache. We provide an example on our project Wiki.[B3]

A few additional caveats. First, the cache can get large and periodically need to be cleaned and rebuilt.[B4] Second, when using the cache, SCons calculates build signatures regardless of the Decider method chosen, so there are no time savings from using, for example, `content-timestamp`. Of course, this does not affect those using the default `content`. Finally, in its own storage, the cache uses file signatures as filenames, instead of the cached file's own name, which is slightly inconvenient if you need to find and examine the cached version of a particular file. The `cache_debug` option will provide information on the files being used, and note that the file will be stored in a subfolder corresponding to the first two characters of the cached filename. For example, a cached file `d7581eba1ac59ad5f92b2fceba04477f` will be in subfolder `D7`.

We list some useful command-line options for the cache below. Using these options together with `dry_run` can be instructive to get a preview of how SCons will interact with the cache.

We list the Stata-style syntax first, then the standard SCons syntax. Recall that either syntax is allowed but the two cannot be mixed in a single command.

B2. As of SCons 4.3.0 (December 2021), this is Section 22, "Caching Built Files." See especially the first footnote in that section, which describes the information SCons stores and uses to determine whether a target is up to date.

B3. https://github.com/bquistorff/statacons/wiki/Cache-with-cache-disabled-as-default.

B4. Rebuilding the cache does not require rebuilding the project, just erasing the cache and then calling SCons with the option `cache_force`.

`cache_debug(-)` prints information on the cache files being used to the screen

    SCons equivalent: `--cache-debug=-`

`cache_debug(file)` saves information on the cache files being used to the file `file`

    SCons equivalent: `--cache-debug=file`

`cache_disable` do not use files from the cache, do not write files to the cache

    SCons equivalent: `--cache-disable`, `--no-cache`

`cache_force` write all derived files to the cache, whether they are built in this call to SCons or have previously been built. This overrides the default, which is to write only files that are built or rebuilt files in the current call to SCons.

    SCons equivalent: `--cache-force`, `--cache-populate`

`cache_readonly` use files from the cache but do not write files to the cache

    SCons equivalent: `--cache-readonly`

`cache_show` by default, SCons will print "`retrieved file from cache`" to the screen when it pulls a file from the cache rather than building it. With the `cache_show` option, SCons will instead print what it *would have done* to build the file if it had not used the cache.

    SCons equivalent: `--cache-show`

Other Options

Here are a few alternatives, along with what we see as their main drawbacks.

Shared `.sconsign.dblite` database

Since SCons uses the `.sconsign.dblite` database to decide which targets need to be rebuilt, it is tempting to use a shared database, whether through a file-synchronization service like Dropbox or version-control system like Git or GitHub. However, this is unlikely to work well. Since the `.sconsign.dblite` is a single, binary file, version conflicts or file corruption can arise even if users are working on different aspects of a project. Furthermore, unless all users have all their code, targets and dependencies synchronized at all times, a shared `.sconsign.dblite` database will cause SCons to get very confused, since it will see discrepancies between files and database entries every time any user updates any file. This might be approximately workable if all users are synced to the same shared folder, e.g., a shared Dropbox folder, but this is not possible if one is using more full-featured source code and collaboration management solutions such as git.

Marking files as built

The `assume_built`/`assume_done` options we have added for `statacons` can help in one-off cases. See 3.2 for their description. In the example above, User B could update to User A's latest `dataprep.do` and `auto-modified.dta`, then run

```
statacons , assume_built('outputs/auto-modified.dta')
```

SCons will then skip rebuilding `auto-modified.dta`, record the current file signature as in `.sconsign.dblite`, and move on to rebuilding the final outputs.

  While this approach can be useful, it relies on users accurately keeping track of what is up to date and what is not, and making manual choices in builds, both of which build tools are intended to make unnecessary. For example, we have found `assume_built(*)` to be useful if a minor change to the builder has caused SCons to think all files need to be re-built when in fact none of them do. Again, though, we have to be sure that the change to the builder will not result in any substantive change.

Additional Comments on Collaborative Workflows

Version Control If you are using a version control system like Git or SVN, we recommend that you do *not* keep `config_local.ini` or `.sconsign.dblite` under version control. In the case of `config_local.ini`, this is simply because the purpose of this file is to account for things that will vary across users. In the case of `.sconsign.dblite`, see our discussion under "Shared `.sconsign.dblite` database" above. There are a few additional Python-generated files (e.g., `__pycache__`) that should not be kept under version control. In Git, this is handled with a `.gitignore` file. In the replication archive we provide for our Introductory Example, we have included an example `.gitignore` that excludes the files above. In addition, we recommend excluding all generated files. That is, as a rule, we keep only code (in the broad sense of things a human has typed, as opposed to things generated by the computer) under version control. This is generally considered good practice but is essential when using the SCons cache, since the cache is already essentially handling version control for generated files, and having both SCons and the version control software attempting the same task will lead to conflicts.

  We do recommend keeping user-written programs under version control to ensure that all users have the same version of these programs. In our Introductory Example, we have kept user-written `estout` (Jann 2007) in the folder `code/ado/plus`. Our `profile.do` sets Stata's `adopath` to instruct Stata to look in this folder (and subfolders). Similarly, we keep `config_project.ini`, the SConstruct, `statacons.ado` and the other associated files under version control, again so that the build will be consistent across users. The downside of this practice is that you will need to maintain separate copies of the same programs for each project.

Shared SCons cache If you are using Dropbox or a similar service to share your SCons cache, make sure that you have offline access to the folder containing the cache.[B5] Using the "Smart Sync" or "Streaming Access" options that store files online and then fetch them when requested is likely to cause several problems. First, Python may not find folders that are only available in the cloud. Second, this will slow down the process of retrieving cached files. Third, you may not be able to access certain recently cached files when working offline.

Finally, as of April 2022 we do not recommend Google Drive for sharing caches. Google Drive will sometimes be able to discern the type of cached files and will add a file extension to the filename in the cache. This happens especially frequently with `.pdf`s. This will confuse SCons and prevent it from successfully retrieving that file from the cache. We have not found this to occur with Dropbox.

---

*B*5. As of April 2022: in Dropbox, either turn off "Smart Sync" to make all of your Dropbox content available offline, or use "Selective Sync" to keep your cache folders available offline. While we currently do not recommend Google Drive for shared caches (see the next paragraph), if you do use Google Drive, choose either "Mirror files" to make *all* of your Google Drive content available offline, or follow the instructions under "Stream files – Choose specific files and folders to make available offline" to keep your cache folders available offline.

```
# **** Setup from pystatacons package *****
import pystatacons
env = pystatacons.init_env()

# use an sconsign database specific to this SConstruct
SConsignFile(".sconsignCache")

# set path to cache directory
# shared cache dir
# hard-coded -- if hard-coding path with spaces, *must* enclose in quotes
# dropbox example
#CacheDir(´C:/Users/Username/Dropbox/SJ-BuildTools/simpleCacheShare´)
# local folder example
#CacheDir(´./scons_cache´)
# from config_local.ini -- in config, do *not* enclose path with spaces in quotes
CacheDir(env[´CONFIG´][´Project´][´cache_dir´])

# **** Substance begins        *****

# analysis
cmd_analysis = env.StataBuild(
    target = [´outputs/scatterplot.pdf´,
              ´outputs/regressionTable.tex´],
    source = ´code/analysis.do´
)
Depends(cmd_analysis, [´outputs/auto-modified.dta´])

# dataprep
cmd_dataprep = env.StataBuild(
    target = [´outputs/auto-modified.dta´],
    source = ´code/dataprep.do´
)
Depends(cmd_dataprep, [´inputs/auto-original.dta´])
```

Figure B1: SConstruct for Introductory Example Using SCons Cache
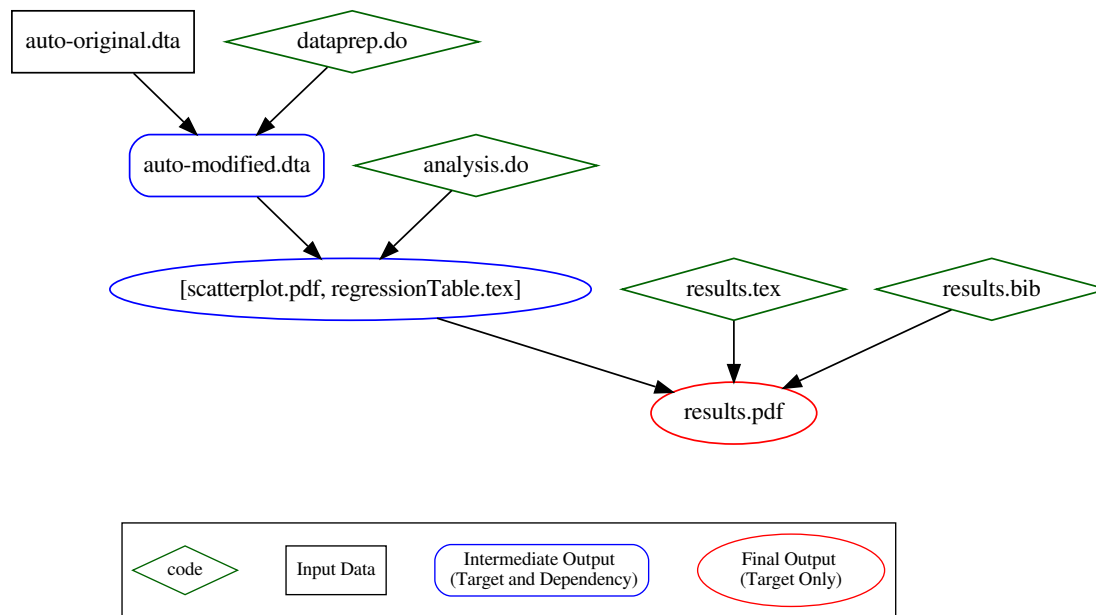
C Compiling a PDF from the Introductory Example

Figure C1: Workflow for Compiling PDF from Introductory Example

```
# SConstructCompilePDF

# **** Setup from pystatacons package *****
import pystatacons
env = pystatacons.init_env()

Help("""
To compile analysis targets only:
    ´statacons, file(SConstructCompilePDF)´
    ´scons --file=SConstructCompilePDF´
To clean only analysis targets (not compiled PDF)
    ´statacons, file(SConstructCompilePDF) clean´
    ´scons --file=SConstructCompilePDF -c´
To compile PDF:
    ´statacons compilePDF, file(SConstructCompilePDF)´
    ´scons --file=SConstructCompilePDF compilePDF´
To clean all targets plus all latex helper files (aux, bbl, etc)
    ´statacons compilePDF, file(SConstructCompilePDF) clean´
    ´scons --file=SConstructCompilePDF compilePDF -c´
To clean only latex helper files (aux, bbl, etc)
    ´statacons cleanPDFhelp, file(SConstructCompilePDF) clean´
    ´scons --file=SConstructCompilePDF cleanPDFhelp -c´
""", append=True)

# **** Substance begins        *****

# dataprep

dataprep_Targets = [´outputs/auto-modified.dta´]
Alias(´dataprep´,dataprep_Targets)

cmd_dataprep = env.StataBuild(
    target = dataprep_Targets,
    source = ´code/dataprep.do´,
    depends = [´inputs/auto-original.dta´])

# analysis

analysis_Targets = [´outputs/scatterplot.pdf´,
                ´outputs/regressionTable.tex´]
Alias(´analysis´,analysis_Targets)

cmd_analysis = env.StataBuild(
    target = analysis_Targets,
    source = ´code/analysis.do´,
    depends = dataprep_Targets)
Default(´analysis´)

# compile PDF - SCons knows how to scan tex files for dependencies
env.Tool("pdftex")
env.AppendUnique(PDFLATEXFLAGS=´-quiet´)
pdf_output = env.PDF(target=´outputs/results.pdf´, source=´code/results.tex´)
Alias(´compilePDF´,´outputs/results.pdf´)

# extra files generated by pdflatex to clean
pdfToClean = ["outputs/results"+e for e in [´.synctex.gz´,´.fls´,´.log´,´.aux´,´
> .bbl´,´.out´,´.dvi´,´.blg´]]
Clean(´cleanPDFhelp´,pdfToClean)
```

Figure C2: SConstruct for Compiling PDF from Introductory Example

```
. statacons compilePDF, file(SConstructCompilePDF)  ///
>     debug(explain) tree(status,prune)
scons: Reading SConscript files ...
Using ´LabelsFormatsOnly´ custom_datasignature.
Calculates timestamp-independent checksum of dataset,
  including variable formats, variable labels and value labels.
Edit use_custom_datasignature in config_project.ini to change.
  (other options are Strict, DataOnly, False)
scons: done reading SConscript files.
scons: Building targets ...
scons: building `outputs\results.pdf´ because it doesn´t exist
cd /D outputs && pdflatex -interaction=nonstopmode -recorder -quiet results.tex
cd /D outputs && bibtex results
cd /D outputs && pdflatex -interaction=nonstopmode -recorder -quiet results.tex
cd /D outputs && pdflatex -interaction=nonstopmode -recorder -quiet results.tex
 E        = exists
  R       = exists in repository only
   b      = implicit builder
   B      = explicit builder
    S     = side effect
     P    = precious
      A   = always build
       C  = current
        N = no clean
         H = no cache
[E b      ]+-compilePDF
[E B   C ]  +-outputs\results.pdf
[E     C ]      +-code\results.tex
[E     C ]      +-code\results.bib
[E B P C ]      +-outputs\regressionTable.tex
[E     C ]      | +-code\analysis.do
[E B P C ]      | +-outputs\auto-modified.dta
[E     C ]      |    +-code\dataprep.do
[E     C ]      |    +-inputs\auto-original.dta
[E B P C ]      +-outputs\scatterplot.pdf
[E     C ]        +-code\analysis.do
[E B P C ]         +-[outputs\auto-modified.dta]
scons: done building targets.
```

Figure C3: SCons output from Compiling PDF from Introductory Example

Results from Introductory Example

Compiled by `statacons`

Figure 1 is a scatterplot.
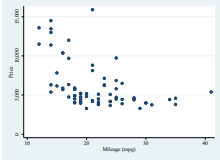Table 1 is a table of regression results created using `estout` [Jann, 2007]

1

Figure 1: A scatterplot.

	(1)	(2)
	price	price
mpg	-238.9***	-1265.2***
	(53.08)	(289.5)
mpg_sqd		21.36***
		(5.939)
_cons	11253.1***	22716.5***
	(1170.8)	(3366.6)
N	74	74
R^2	0.220	0.340

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: A table of regression results

2

References

Ben Jann. Making Regression Tables Simplified. *The Stata Journal*, 7(2):227–244, June 2007. doi: 10.1177/1536867X0700700207.

3

Figure C4: Compiled PDF from Introductory Example

D Version History

Most recent version of `statacons`: project repository https://github.com/bquistorff/statacons.

Most recent version of the paper: https://osf.io/preprints/metaarxiv/qesx6/.

Previous versions of `statacons`: https://osf.io/gbh4m/, see "Components" section .

Previous versions of the paper: https://osf.io/preprints/metaarxiv/qesx6/, click the "Download previous versions" button below the paper image, select desired version of the paper.

v3.0.0 May 2022 (this version, as resubmitted to *The Stata Journal*)

> Citation: Guiteras, Raymond, Ahnjeong Kim, Brian Quistorff and Clayson Shumway, "statacons: An SCons-based build tool for Stata," CEnREP Working Paper 22-001, May 2022, https://go.ncsu.edu/cenrep-wp-22-001. Under review at the Stata Journal.
>
> Paper archive: https://doi.org/10.31222/osf.io/qesx6
>
> Install `statacons`: `net install statacons,`
> `from(https://raw.github.com/bquistorff/statacons/v3.0.0/)`
>
> Installation instructions: https://github.com/bquistorff/statacons/tree/v3.0.0
>
> Code archive: https://osf.io/k4tqc/

v2.0.0 January 19, 2022

> Citation: Guiteras, Raymond, Ahnjeong Kim, Brian Quistorff and Clayson Shumway, "statacons: An SCons-based build tool for Stata," CEnREP Working Paper 22-001, January 2022, https://go.ncsu.edu/cenrep-wp-22-001. Under review at the Stata Journal.
>
> Paper archive (v2): https://doi.org/10.31222/osf.io/qesx6, click the "Download previous versions" button below the paper image, select "Version 2."
>
> Install `statacons`: `net install statacons,`
> `from(https://raw.github.com/bquistorff/statacons/v2.0.0/)`
>
> Installation instructions: https://github.com/bquistorff/statacons/tree/v2.0.0
>
> Code archive: https://osf.io/n42mc/

v1.0.0 January 4, 2022 (as submitted to *The Stata Journal*)

> Citation: Guiteras, Raymond, Ahnjeong Kim, Brian Quistorff and Clayson Shumway, "statacons: An SCons-based build tool for Stata," CEnREP Working Paper 22-001, January 2022, https://go.ncsu.edu/cenrep-wp-22-001. Under review at the Stata Journal.
>
> Paper archive (v1): https://doi.org/10.31222/osf.io/qesx6, click the "Download previous versions" button below the paper image, select "Version 1."
>
> Install `statacons`: `net install statacons,`
> `from(https://raw.github.com/bquistorff/statacons/v1.0.0/)`

Installation instructions: https://github.com/bquistorff/statacons/tree/v1.0.0

Code archive: https://osf.io/cmh56/